Rebuttal to "Comment on 'A theoretical upper limit for offshore wind energy extraction' by Simão Ferreira et al. (2026)": reproducibility analysis, audit of the vdLW implementation, and clarification of finite wind-farm correction methodologies

Jens Nørkær Sørensen[1], Gunner Chr. Larsen[2], and Carlos Simão Ferreira[3]

[1]Section of Aero- and Fluid Dynamics, DTU Wind and Energy Systems, Nils Koppels Allé, bldg. 403, 2800 Lyngby, Denmark
[2]Response, Aero-elasticity, Control and Hydrodynamics, DTU Wind, Frederiksborgvej 399, Roskilde, 4000, Denmark
[3]Faculty Aerospace Engineering, Delft University of Technology, Kluyverweg 1, Delft, 2629 HS, The Netherlands

Correspondence: Jens Nørkær Sørensen (jnso@dtu.dk)

Abstract.

This rebuttal addresses the claims raised in the comment paper by van der Laan and Watson (vdLW) regarding the reproducibility, interpretation, and validity of the finite wind-farm implementation used in our original work on the theoretical limits of offshore wind energy (Simão Ferreira, Larsen, and Sørensen (2026)). The original formulation derives a theoretical upper limit for offshore wind-farm performance from atmospheric boundary-layer momentum-balance principles, where the infinite wind-farm equilibrium solution defines the asymptotic theoretical limit, and finite wind-farm effects are represented through a physically motivated correction accounting for turbines exposed to free inflow conditions near wind-farm edges. We demonstrate that the primary source of disagreement reported by vdLW originates from implementation errors, altered physical inputs, and methodological inconsistencies in the vdLW reproduction of the Simão Ferreira, Larsen, and Sørensen (SLS) results and implementation of their proposed finite wind-farm correction methodology, rather than from deficiencies in the underlying theoretical framework itself. After correcting the identified Type 1 implementation errors and restoring the published SLS inputs and assumptions, the vdLW implementation reproduces the published SLS results, thereby resolving the reproducibility issue itself. A second comparison is then performed using the alternative finite wind-farm correction methodology proposed by vdLW. Even before correcting additional Type 2 preprocessing and geometry-processing inconsistencies associated with this alternative methodology, agreement with the original SLS results remains strong and improves further once selected Type 2 errors are mitigated. This rebuttal further clarifies that the proposed limit is a theoretical limit derived from a physically based atmospheric momentum-equilibrium model, conceptually analogous to the Betz limit in rotor aerodynamics. The finite wind-farm correction does not invalidate the asymptotic limit, but instead provides a physically motivated extension for representing finite-size effects and the transition between free inflow conditions near wind-farm edges and fully developed flow conditions within the wind-farm interior. Importantly, the alternative methodology proposed by vdLW itself ultimately remains consistent with the SLS theoretical framework once implementation issues are corrected. Operational offshore wind-farm data represented using the corrected vdLW methodology continue to cluster below the SLS theoretical limit

curve and close to the 90% theoretical limit curve. The present comparison therefore demonstrates that the discussion raised by vdLW primarily concerns refinement of finite wind-farm correction methodologies rather than invalidation of the underlying asymptotic theoretical framework. Questions regarding national deployment strategies, energy policy scenarios, or broader system-level interpretations raised by vdLW are outside the scope of the present rebuttal, which focuses exclusively on the scientific validity, reproducibility, and physical interpretation of the analytical framework. The present work further highlights that finite wind-farm correction methodologies remain an important and evolving research topic, particularly for highly complex wind-farm clusters and interacting neighboring wind farms. In this regard, the effort by vdLW toward more systematic and automated finite wind-farm correction methodologies is valuable and has provided useful inspiration for future refinement of the framework. Overall, the corrected vdLW results reinforce the robustness, reproducibility, and physical consistency of the SLS theoretical limit framework and remain fully consistent with the Wind Farm Wind Factor formulation and its validation against operational data from 72 offshore wind farms representing more than 420 cumulative operational years.

1 Introduction and structure of the rebuttal

In their comment submitted to Wind Energy Science, van der Laan and Watson (vdLW) present a strong criticism of our article "A theoretical upper limit for offshore wind energy extraction" by Simão Ferreira, Larsen and Sørensen (2026) (hereafter referred to as SLS). In their comment, they argue that the derived theoretical limit is not a theoretical limit, but rather a model limit strongly dependent on the finite wind-farm correction methodology. Furthermore, they claim that the model is validated using an ad hoc and undocumented approach for application to finite wind farms. Unfortunately, several of the results presented in the comment are based on misreadings, conceptual misunderstandings regarding the use of our model, and misinterpretations of the original paper.

We are nevertheless grateful for the effort by vdLW to reproduce, analyze, and challenge the SLS framework. Critical analysis and independent reproduction attempts are essential parts of the scientific process. However, when auditing the vdLW implementation, code, and reproduced results, we identified several implementation errors and undocumented modifications to the original SLS methodology that materially affected the reproduced results.

Importantly, once the most significant implementation errors are corrected, the results of the vdLW methodology are found to be in very good agreement with those of SLS. The comparisons presented in the following sections demonstrate that the vdLW methodology, despite its identified limitations, remains fully consistent with the SLS theoretical framework and therefore reinforces rather than invalidates the central conclusions of SLS.

Most importantly, the present discussion does not challenge the existence or validity of the underlying asymptotic theoretical limit itself. Rather, the discussion concerns how best to represent the deviation of real finite wind farms from the asymptotic infinite-farm equilibrium solution through suitable finite wind-farm correction methodologies. In this sense, the work by vdLW ultimately reinforces the robustness of the theoretical limit framework, as even their alternative methodology converges toward results that remain in full agreement with the SLS limit once implementation issues are corrected.

We therefore view the present discussion not as a challenge to the limit itself, but as part of the natural scientific process of refining and improving finite wind-farm correction methodologies. In this regard, several aspects of the effort by vdLW are valuable and inspiring, particularly their attempt to develop a more systematic and automated finite wind-farm correction procedure. The audit process and comparison of methodologies also helped us reflect upon some of our own assumptions and possible sources of bias, providing useful direction for future refinement of the framework.

The extent and impact of the implementation issues identified during the audit made the analysis extensive. Therefore, in the main body of the present rebuttal we focus on the principal scientific points and key reproducibility results, while the detailed audit of the vdLW implementation and responses to specific comments are presented in the Appendix. This structure is adopted to avoid overwhelming the reader with excessive detail in the main discussion.

In the first core of the text we present:

- A short explanation of the method by SLS
- The reproducibility analysis and key results of the audit of the vdLW implementation
- Conclusions

In the Appendix part we present

- Responses to selected vdLW comments not addressed in the main text
- Audit of the vdLW implementation and methodology: main errors and their origin
- Validation examples of the methodology adopted in the work of SLS

2 A short explanation of the method by SLS

The method is already explained in both the comment and the original paper. Hence, we will keep the description brief. The basis of the method is the derivation of the following equation for the maximal capacity factor of an infinite wind farm,

$$CF_{\max} = \phi_{\mathrm{WF}}^{-3}\Gamma\left(1+\frac{k}{k}\right)^{-3}\Gamma_{\mathrm{ic}}\left(\frac{3+k}{k},\phi_{\mathrm{WF}}^{k}\Gamma\left(1+\frac{k}{k}\right)^{k}\right)+e^{-\phi_{\mathrm{WF}}^{k}\Gamma\left(1+\frac{k}{k}\right)^{k}}, \tag{1}$$

where

$$\phi_{\mathrm{WF}} = \frac{U_r}{U_\infty \epsilon},$$

with U_r being the rated wind speed, λ and k the Weibull parameters, Γ the Gamma function, and ϵ the ratio between the average wake velocity and the free wind speed at hub height.

The expression for ϵ is derived by assuming a momentum balance between the thrust exerted by the wind farm on the atmospheric boundary layer and the change in the velocity profiles in the boundary layer under neutral atmospheric conditions.

The expression contains most parameters characterizing the atmospheric boundary layer and the rotor, such as thrust coefficient, average distance between the turbines, hub height, roughness length, and geostrophic wind speed.

The equation is derived assuming a thrust and a power curve for an optimum operating wind turbine, combined with a Weibull distribution for the site mean wind speed. Hence, within the approximations and assumptions of the model, we expect Eq. (1) to represent the optimum capacity factor for an infinitely large wind farm, i.e., a wind farm that is so large that the difference in power production between turbines operating along the edges of the farm and those operating in the farm interior is negligible compared to the total power production of the farm.

Ignoring the impact of the Weibull shape factor k, which can essentially be taken as constant, the capacity factor only depends on ϕ_{WF}. Hence, the maximum capacity factor of any "infinitely" large wind farm is uniquely determined analytically for any given value of ϕ_{WF}. Interestingly, setting $\epsilon = 1$ gives the maximum capacity factor for a solitary wind turbine, meaning that Eq. (1) expresses two different reference cases: one for a wind turbine located in the interior of a wind farm ($\epsilon < 1$) and one subject to free wind at the edge of the wind farm ($\epsilon = 1$).

2.1 Theoretical limit and finite-size correction

A central criticism raised in the comment paper is the claim that the proposed limit is not a theoretical limit, but merely the limit of an analytical model because application to real wind farms requires a finite wind-farm correction. We believe this interpretation reflects a misunderstanding of the role of the infinite wind-farm equilibrium solution and of how finite-size corrections are commonly treated in aerodynamic theory.

As explained above, the theoretical limit proposed by SLS is rigorously defined for the asymptotic case of an infinitely large wind farm operating under horizontally homogeneous equilibrium conditions within the atmospheric boundary layer, as well as for the limiting case of a solitary wind turbine operating under free-stream conditions at the edge of the farm. The finite wind-farm correction does not redefine this asymptotic limit, but rather provides an engineering extension for estimating the performance of real finite wind farms containing turbines exposed to partially freer inflow conditions. In practice, this corresponds to an equivalent value of ϵ lying between the fully developed infinite-farm case and the free-stream edge-turbine case $\epsilon = 1$.

This distinction is analogous to classical actuator-disk momentum theory and the Betz limit. The Betz limit is rigorously derived for an idealized actuator disk under asymptotic assumptions ignoring the influence of the azimuthal velocity component, while real wind turbines require additional engineering corrections accounting for finite-blade effects, such as the Prandtl tip-loss correction, and finite tip speed ratios. Importantly, the existence of finite-blade corrections does not imply that the Betz limit ceases to be a theoretical limit. Rather, the asymptotic actuator-disk solution remains the underlying theoretical reference, while the engineering corrections provide practical methods to estimate how real finite systems depart from the idealized asymptotic case while still obeying the same underlying momentum-extraction physics.

A key point, and one of the central misunderstandings in the vdLW interpretation, is that not all equivalent values of ϵ are physically admissible. The physically realistic range of finite wind-farm corrections is in fact relatively narrow. Thus, similarly to finite-blade corrections in rotor aerodynamics, finite-farm corrections refine the practical implementation of the model but do

not invalidate or redefine the underlying asymptotic limit itself. Therefore, the SLS limit plays, for wind farms, a role analogous to that of the Betz limit for the actuator disk, while finite wind-farm corrections are analogous to finite-blade corrections, of which several exist and continue to be developed.

The arguments by vdLW are therefore only relevant to the question of what constitutes the best finite-farm correction methodology, but do not challenge the existence or validity of the underlying asymptotic limit itself.

In the paper describing the original model of Sørensen and Larsen (2021), later elaborated in Sørensen et al. (2024, 2026), a factor was introduced to correct for the finiteness of a real wind farm. This factor discriminates between edge wind turbines subject to free wind and turbines subject to wake effects within an infinite wind farm:

$$P_Y = (1-a)P_W + aP_S, \tag{2}$$

where P_Y is the yearly production of a representative wind turbine in a finite-size wind farm, P_W is the production of the same wind turbine in the fully developed wind farm boundary layer, and P_S is the production of an analogous solitary wind turbine (corresponding to a turbine located at the edge). The parameter $a = a(N_T, S, U)$ is a finiteness correction depending on the number of wind turbines, N_T , the distance between the turbines measured in rotor diameters, $S = L/D$, and wind speed, U, which takes values between 0 and 1, where $a = 0$ corresponds to an infinite wind farm, and $a = 1$ to a solitary wind turbine operating under free-stream conditions.

Hence, the production of a finite wind farm is assumed to consist of the sum of the production from aN_T freely operating wind turbines and from $(1-a)N_T$ wind turbines located in the fully developed wind farm boundary layer.

It is natural to combine this with Eq. (1), resulting in the following equation:

$$CF_{\text{finite}} = (1-a)\,CF_W + a\,CF_S, \tag{3}$$

stating that the maximum capacity factor of a finite wind farm is the weighted sum of the maximum capacity factor of solitary wind turbines and of turbines operating inside an infinite wind farm.

It is important to note that all three expressions define points located on the same curve. Hence, for a given wind farm size and location, solving Eq. (3) for ϕ_{WF} gives the maximum capacity factor of the wind farm.

To determine the value of a, different techniques have been explored. In the original paper (Sørensen and Larsen, 2021), the correction factor was assumed to depend on the number of edge turbines, hence

$$a = \frac{c}{\sqrt{N_T}},$$

with c being a constant. In the first version, the constant was set equal to 3, but after explicitly including losses in the model, it was later increased to a value of 5.

In later work Sørensen et al. (2024), the a-value was systematically determined using simulations of generic wind farms with the linear CFD solver Fuga, where a was found to vary between 4.1 and 5.9 depending on the specific wind farm. This was later parametrized into the following simple analytical expression:

$$a = a(S, N_T) = -0.045S + 1.611S^{1/4} - 0.966S^{1/8}N_T^{1/16} \tag{4}$$

In a recent publication by Sørensen et al. (2026), this expression was validated against annual production data from five different wind farms over periods ranging from 2004 to 2018. It was found that annual production using the expression could be predicted with a mean relative error of less than 4%. Examples of applying this model in combination with Eq. (3) are shown in Appendix C.

Despite the validation cases being very different with respect to topology, power density, and size, the comparisons between computed and measured capacity factors were found to be in excellent agreement. Hence, we have strong confidence in the proposed method for "stand-alone" wind farms, i.e., wind farms that are not exposed to shadow effects from neighboring wind farms.

However, when wind farms are installed close to existing wind farms or packed into large groups, the determination of the a-factor becomes more complicated.

Consequently, in the present work by SLS, the a-factor was computed by determining the number of edge turbines operating in free-stream conditions multiplied by a constant accounting for the penetration depth of free-flow equivalent-power (typically taken as 2.5 rows, except for smaller farms). This method makes it possible to account for the influence of nearby wind farms, local shadow effects, and the non-uniform distribution of wind turbines.

In many respects, this approach is conceptually similar to what vdLW attempted in their work. However, because the methodology adopted in SLS relied on manual inspection rather than full automatization, it avoided some of the implementation errors encountered in the vdLW approach. At the same time, the work by vdLW represents an important step toward the development of more systematic and automated finite wind-farm correction methodologies.

It may ultimately prove that the methodology proposed by vdLW performs better in many situations, particularly because of its automated character. We have taken considerable inspiration from the effort by vdLW and learned much while auditing their approach. In several cases, identifying errors and limitations in the vdLW implementation also helped us reflect upon and identify some of our own assumptions, limitations, and possible sources of bias.

Again, this discussion does not alter the main conclusions of SLS, but rather provides valuable direction for future refinement and improvement of finite wind-farm correction methodologies.

3 Reproducibility analysis and key results of the audit of the vdLW implementation

The audit of the vdLW implementation revealed two distinct categories of issues. The first category consists of analytical implementation errors and altered physical inputs directly affecting reproducibility of the published SLS results. The second category concerns geometry-processing and finite wind-farm pre-processing issues associated with the alternative N_{free} methodology proposed by vdLW. The following sections discuss these two categories separately.

3.1 Type 1 errors — Analytical implementation errors and altered physical inputs

In Figures 5 and 6(b) of the comment paper, vdLW present results obtained using the N_{free} values published by SLS and state: *"First, PyWake is used with N_{free} from Simão Ferreira et al. (2026) in an attempt to obtain the same results. Despite*

using the same input data as Simão Ferreira et al. (2026), different results are obtained possibly due to differences in model implementation."

This statement is incorrect because vdLW did not in fact use the same inputs as SLS. The discrepancies reported in Figures 5 and 6(b) of the vdLW work originate almost entirely from the use of altered physical inputs and operational assumptions, together with one coding error in the implementation of the capacity-factor calculation equation. The four identified errors consist of:

1. removal by vdLW of the hub-height wind-speed correction of the Weibull scale parameter, causing a different input of the local unperturbed wind speed (originally also present in vdLW's code),
2. change by vdLW of the turbine cut-in and cut-out assumptions compared to those used in SLS (originally also used in vdLW's code),
3. replacement of site-specific latitude values by a fixed latitude,
4. incorrect implementation of the wake-reduction equation used in the analytical capacity-factor integration.

Audit of the vdLW implementation further showed that the original SLS inputs and configurations were initially present in the vdLW code, but were later replaced by modified assumptions and alternative parameter choices. The fourth error is different: it concerns an incorrectly programmed equation in the analytical capacity-factor calculation. These points are documented in detail in the Appendix, including direct excerpts from the vdLW code.

Upon identifying the errors, we corrected the vdLW code and executed the calculations using the published SLS inputs. With the correct inputs and equations restored, the vdLW implementation reproduces the SLS published results exactly, with only negligible differences associated with numerical rounding.

Figures 1 and 2 present the progressive correction of the four identified Type 1 errors and their impact on the agreement between the vdLW implementation and the published SLS results. Figure 1 compares the original vdLW results with the corrected implementation after the restoration of the published SLS inputs, assumptions, and equations. Figure 2 shows the progressive reduction of the discrepancies as the four Type 1 errors are corrected sequentially.

Once these Type 1 errors are corrected, the vdLW implementation reproduces the published SLS results exactly when using the published SLS N_{free} values. This resolves the first reproducibility claim raised in the comment paper.

A second comparison can then be performed using the corrected vdLW analytical implementation and inputs, but replacing the published SLS N_{free} values by the alternative N_{free} values calculated using the vdLW methodology. This additional comparison isolates the impact of the alternative finite wind-farm correction methodology itself, without contamination from the Type 1 implementation errors identified above.

Importantly, once the Type 1 errors are corrected, the agreement between SLS and vdLW remains substantial even when using the vdLW N_{free} methodology. Figure3 shows the comparison between the SLS results and the vdLW results obtained using the corrected analytical implementation together with the vdLW N_{free} methodology. The comparison shows substantial overall agreement, with a linear regression slope close to unity and an R2 value of 0.79.

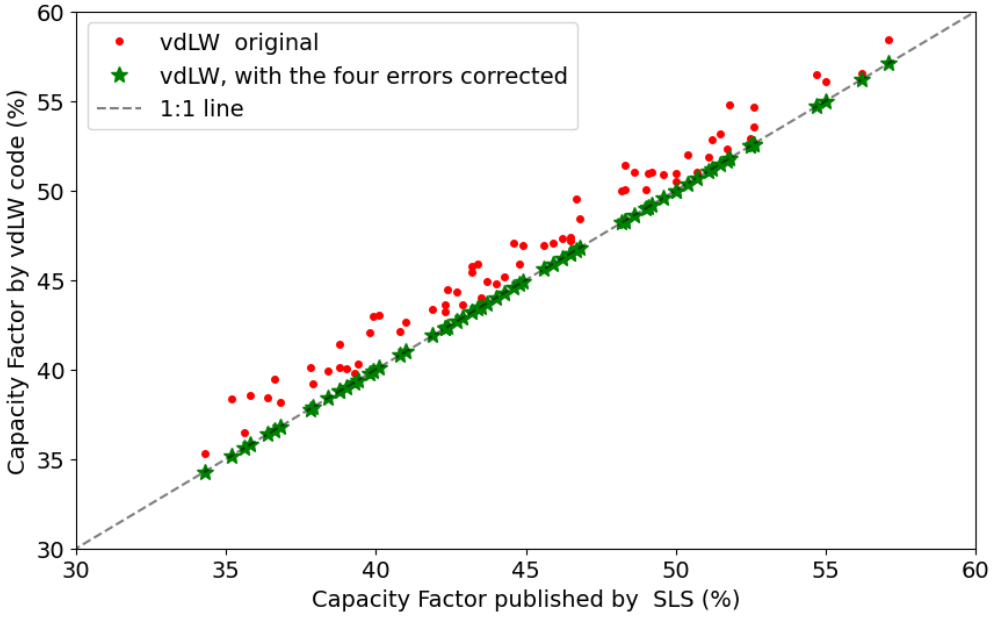

Figure 1. Relative difference between the capacity factors calculated with the vdLW code and the published SLS 2026 values, shown as a function of the published SLS 2026 capacity factor. The black circles show the original vdLW results, which exhibit discrepancies reaching up to approximately 9% of capacity factor. The red circles show the results after correcting Error 1 (removal of the hub-height wind-speed correction), demonstrating that this single implementation error accounts for most of the reported disagreement. The orange circles show the additional correction of Error 2 (incorrect use of cut-in and cut-out wind speeds), reducing the discrepancies to below approximately 1%. The yellow circles show the further correction of Error 3 (use of a fixed latitude instead of site-specific latitude), yielding near-complete agreement. Finally, the green stars show the results after correcting all four implementation errors, including Error 4 (incorrect implementation of the capacity-factor equation), resulting in essentially zero difference relative to the published SLS 2026 results.

This result is important because it demonstrates that the large discrepancies originally reported by vdLW were not primarily caused by the use of an alternative finite wind-farm correction methodology itself. Instead, a substantial fraction of the dis-
215 agreement originated from the Type 1 implementation errors discussed in the present section. Once these errors are corrected, the remaining differences associated with the vdLW N_{free} methodology become comparatively moderate.

Inspection of Figure 3 further shows that part of the remaining scatter is associated with a limited number of clear outlier wind farms. As discussed in the following section, these outliers are primarily linked to additional Type 2 geometry-processing and finite wind-farm methodology issues affecting the vdLW N_{free} estimation. The next section therefore evaluates the impact
220 of correcting selected Type 2 errors for a subset of the most strongly affected wind farms and demonstrates that the agreement improves further once these issues are addressed. The following section then evaluates the remaining differences associated with the alternative N_{free} methodology itself and the additional Type 2 geometry-processing and finite wind-farm methodology issues.

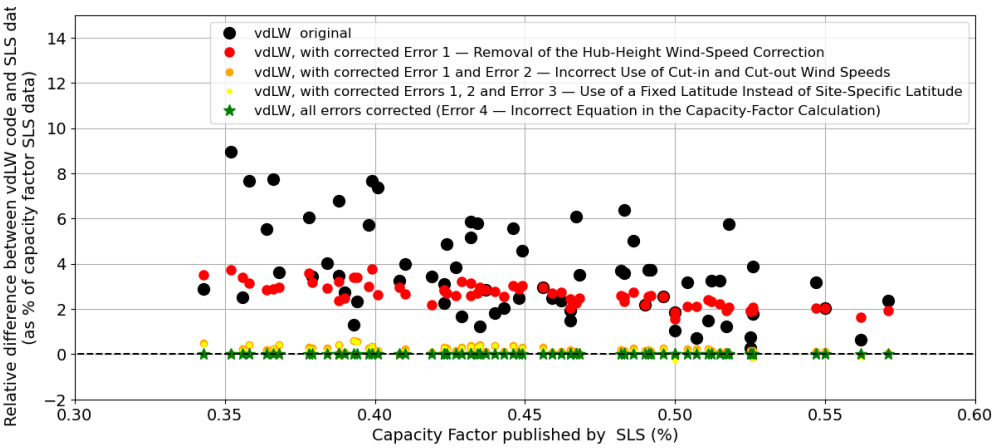

Figure 2. Comparison between the capacity factors published in SLS 2026 and the capacity factors reproduced using the vdLW code, before and after correcting the four implementation errors identified in this rebuttal. The red circles show the original vdLW results. The green stars show the results obtained after correcting the four errors discussed. After correction, the vdLW code reproduces the published SLS results with essentially perfect agreement along the 1:1 line.

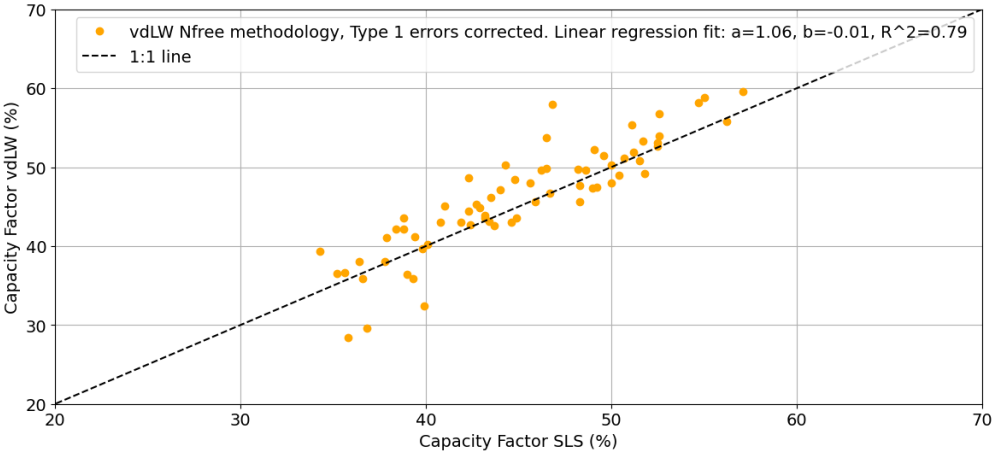

Figure 3. Comparison between the capacity factors published by SLS and the capacity factors obtained using the vdLW N_{free} methodology after correction of the four Type 1 analytical implementation errors identified in the present work. The vdLW calculations shown here use the corrected SLS analytical implementation, inputs, assumptions, and equations, while replacing the published SLS N_{free} values by the alternative N_{free} values calculated using the vdLW methodology. The comparison shows substantial agreement, with a linear regression slope close to unity (a=1.06), near-zero offset (b=-0.01), and (R^2=0.79). The remaining scatter is primarily associated with a limited number of outlier wind farms affected by additional Type 2 geometry-processing and finite wind-farm methodology issues discussed in the following section. The figure demonstrates that, once the Type 1 errors are corrected, the remaining differences associated with the vdLW N_{free} methodology become comparatively moderate.

3.2 Type 2 errors — Geometry-processing and finite wind-farm methodology errors

After correction of the Type 1 analytical implementation errors discussed in the previous section, the agreement between the SLS results and the vdLW implementation becomes exact when using the published SLS N_{free} values, and remains substantial even when using the alternative N_{free} methodology proposed by vdLW. The remaining differences therefore primarily concern the treatment of finite wind-farm geometry and the estimation of the finite wind-farm correction parameter N_{free}. Inspection of the vdLW preprocessing and geometry-processing methodology revealed several additional issues affecting:

- identification of wind-farm inlet-edge turbines,
- neighboring wind-farm interactions,
- wake-shadow regions,
- equivalent clean-inflow fractions,
- and finite wind-farm edge orientation relative to wind direction.

These issues are referred to here as Type 2 errors because they do not concern reproducibility of the analytical SLS framework itself, but rather methodological inconsistencies and physically unrealistic treatments in the vdLW finite wind-farm preprocessing and N_{free} estimation methodology. The identified Type 2 issues include:

1. omission of neighboring wind farms,
2. treatment of neighboring wind farms as infinite wake boundaries,
3. artificial connection of disjoint wind-farm regions,
4. creation of false neighboring regions suppressing existing wind farms,
5. incorrect identification of wind-farm inlet edges,
6. neglect of the temporal evolution of wind-farm clusters,
7. unrealistic estimates of turbines exposed to clean inflow,
8. neglect of the orientation of wind-farm edges relative to wind direction.

Several of these issues originate from the automated geometry-processing methodology used by vdLW, which can generate non-physical aerodynamic regions and unrealistic wake-shadow boundaries for complex wind-farm clusters.

Importantly, several of the Type 2 errors partly compensate each other. As a consequence, partial agreement with measured capacity factors can sometimes arise from cancellation of errors rather than from a physically consistent aerodynamic representation. The purpose of the present section is not to exhaustively correct all Type 2 issues for all wind farms, but rather to demonstrate that the remaining discrepancies identified in Figure 3 are strongly associated with a limited number of wind farms

affected by identifiable preprocessing and geometry-processing errors. Detailed discussion and documentation of the individual Type 2 issues are provided in the Appendix. More detail on the process used to improve the wind-farm input definitions, the individual wind farms affected, and the application of the selected Type 2 corrections is also provided in the Appendix. To illustrate this, selected Type 2 corrections are applied to a subset of the most strongly affected wind farms. Importantly, the vdLW methodology itself was not modified. Instead, corrections were introduced only through improved management of the wind-farm input definitions in order to reduce the impact of some of the most extreme preprocessing errors. Using corrected wind-farm input definitions, it was possible to reduce the impact of Type 2 Errors 1, 3, 4, 5, 6, and 7 for several of the most strongly affected wind farms. However, the Type 2 Errors denoted 2 and 8 are more fundamental to the vdLW methodology itself, and correcting them would require modification of the underlying vdLW finite wind-farm methodology and code rather than adjustment of the input preprocessing.

3.2.1 Example of Type 2 Geometry-Processing Inconsistencies

A representative example of the Type 2 geometry-processing inconsistencies is shown in Figure 4 for the Belwind–Nobelwind wind-farm cluster. The example illustrates how the vdLW automated geometry methodology can generate administratively defined and non-physical neighboring wind-farm regions that strongly distort the inferred aerodynamic environment.

In this case, disconnected turbine groups belonging to Nobelwind are artificially connected into a continuous neighboring-shadow region extending across most of the Belwind wind farm. Large parts of the generated shadow region contain no turbines at all, yet the vdLW methodology treats these regions aerodynamically as continuous upstream blockage. As a consequence, most of the Belwind wind farm becomes embedded inside a false neighboring wind-farm region, leading the vdLW methodology to conclude that Belwind experiences almost no clean inflow for this wind direction.

The example simultaneously illustrates another important Type 2 inconsistency: omission of neighboring wind farms from the geometry evaluation. In the original vdLW geometries published in their repository, the neighboring Borssele III-V wind farms are absent despite physically affecting the aerodynamic environment of the Belgian offshore wind-farm cluster.

The top panel of Figure 4 was generated by SLS using the vdLW code after correcting the missing-neighbor issue and explicitly plotting the neighboring turbines and generated shadow regions. The two bottom panels are reproduced directly from the original geometries published by vdLW in their Zenodo repository. Comparing these geometries immediately reveals both the absence of the Borssele III-V wind farms and the generation of the false neighboring-shadow region extending across Belwind.

Importantly, this behavior does not originate from the analytical SLS framework itself, but from the vdLW geometry-processing methodology used to estimate the finite wind-farm correction. The resulting aerodynamic behavior depends not only on the physical turbine layout, but also on administrative wind-farm grouping and labeling.

The example simultaneously illustrates several different Type 2 inconsistencies discussed in this rebuttal: creation of false neighboring wind-farm regions through artificial connection of disconnected turbine groups; suppression of entire wind farms under non-physical neighboring-shadow regions; omission of physically existing neighboring wind farms; unrealistic sup-

pression of clean inflow; dependence of the aerodynamic behavior on administrative wind-farm grouping; and inconsistent estimation of turbines exposed to freer inflow conditions.

The same methodology also generates self-blockage and inconsistent N_{free} estimates in other wind farms. In Northwind, turbines visibly exposed to incoming flow are suppressed by the generated geometry boundary, while in Nobelwind the methodology can generate unrealistically large N_{free} values because weakly exposed or artificially extended boundaries are treated as fully exposed inlet edges.

Importantly, vdLW themselves acknowledge that their methodology has limitations for wind farms embedded in interacting clusters. However, the impact of these limitations is not explicitly illustrated or discussed in the comment paper itself. The present example shows that many of the remaining differences between vdLW and SLS originate precisely from cases where the vdLW geometry methodology struggles to represent complex wind-farm clusters.

3.2.2 Comparison between the SLS results and the vdLW N_{free} methodology after mitigation of selected Type 2 errors.

Figure 5 shows the comparison between the SLS results and the vdLW N_{free} methodology after mitigation of selected Type 2 preprocessing issues for a subset of the most strongly affected wind farms. Importantly, the vdLW methodology itself was not modified. The improvements were obtained solely through improved wind-farm input definitions designed to reduce the impact of the most extreme preprocessing inconsistencies.

The comparison shows a clear improvement relative to Figure 3, with the regression slope moving closer to unity and the R^2 value increasing from 0.79 to 0.86. The remaining discrepancies are now largely concentrated in a smaller number of wind farms associated with the unresolved Type 2 methodology limitations discussed above.

The most outlier point ("Riffgat" wind farm) corresponds to a error Type 2 – *8 neglect of the orientation of wind-farm edges relative to wind direction*. Correcting it would require changing the vdWL methodology, and is outside of the scope of this rebuttal. These results further reinforce the conclusion that the remaining disagreement between SLS and vdLW is primarily associated with refinement-level sensitivity in the treatment of highly complex finite wind-farm geometries inside interacting wind-farm clusters, rather than irreproducibility of the SLS analytical framework.

3.3 Agreement of the vdLW methodology with the SLS theoretical limit framework

An important outcome of the previous sections is that, even with the limitations and preprocessing inconsistencies identified in the vdLW finite wind-farm methodology, the resulting wind-farm performance trends remain in substantial agreement with the SLS theoretical limit framework.

Figure 6 shows the operational offshore wind-farm data represented using the vdLW N_{free} methodology after correction of the Type 1 analytical implementation errors and mitigation of selected Type 2 preprocessing issues. The figure compares the resulting wind-farm performance against the theoretical limit curves proposed in SLS.

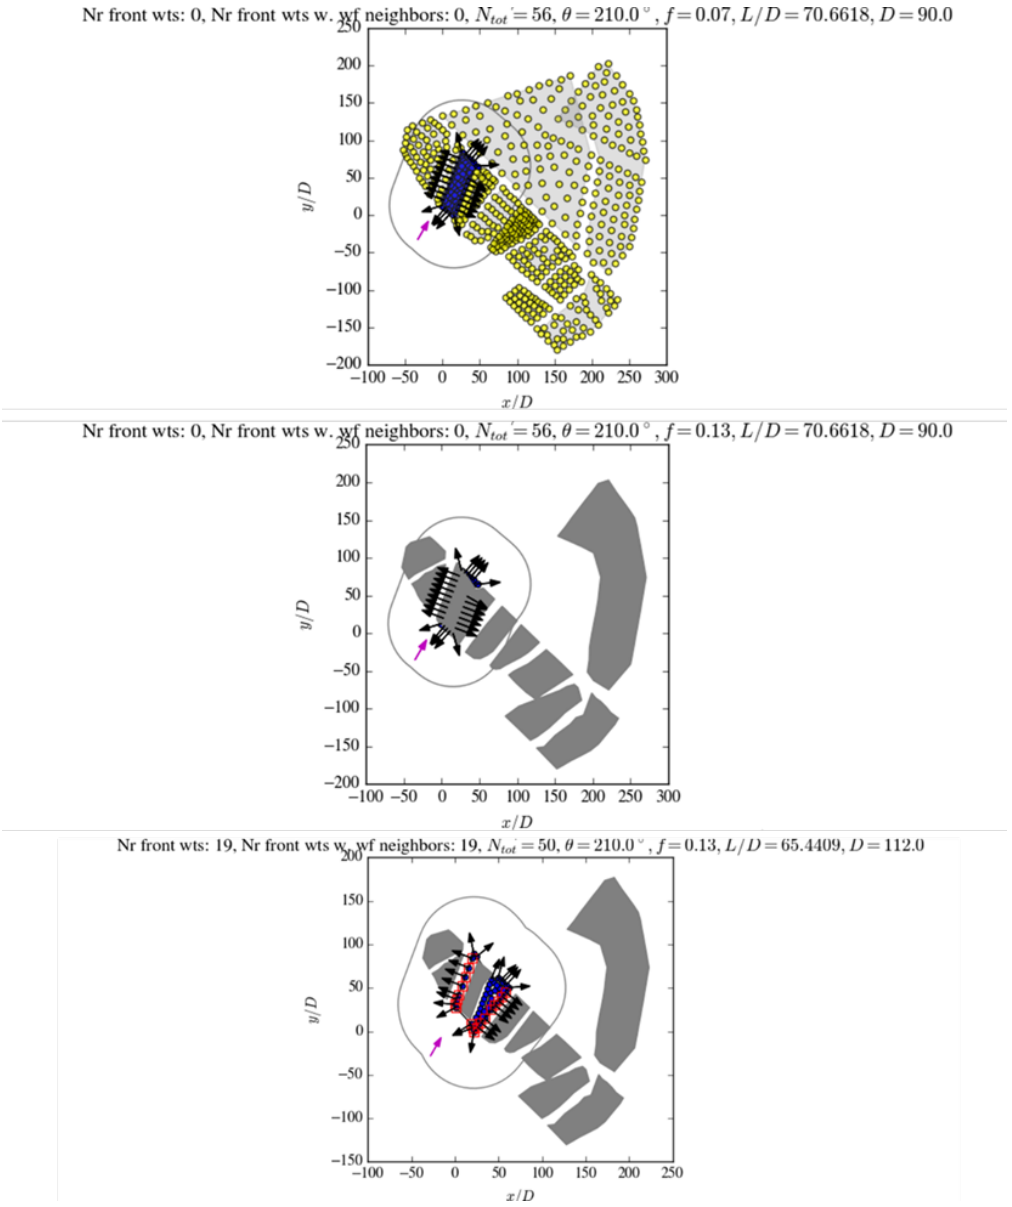

Figure 4. Dutch-Belgium cluster, with focus on the Belwind–Nobelwind example for a wind direction of 210°. The vdLW geometry methodology artificially connects disconnected Nobelwind turbine groups into a continuous neighboring-shadow region extending across most of the Belwind wind farm. The top panel shows the geometry generated by SLS using the vdLW code after correcting the missing-neighbor issue and explicitly plotting neighboring turbines and generated shadow regions. The two bottom panels correspond to the original geometries published by vdLW. The example simultaneously illustrates omission of neighboring wind farms (Borssele III-V), creation of false neighboring wind-farm regions, and artificial suppression of clean inflow caused by administratively generated aerodynamic blockage regions.

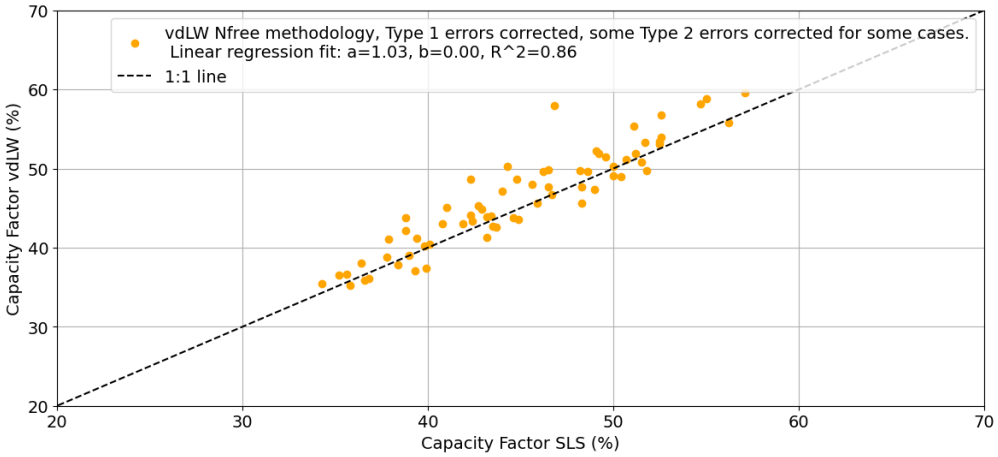

Figure 5. Comparison between the capacity factors published by SLS and the capacity factors obtained using the vdLW N_{free} methodology after correction of the Type 1 analytical implementation errors and mitigation of selected Type 2 preprocessing and post-processing errors for a subset of the most strongly affected wind farms. Importantly, the vdLW methodology itself was not modified. The improvements were obtained solely through improved management of the wind-farm input definitions in order to reduce the impact of several extreme preprocessing inconsistencies affecting the vdLW N_{free} estimation. The comparison shows further improvement relative to Figure 3, with a regression slope closer to unity $a = 1.03$, near-zero offset $b = 0.00$, and $R^2 = 0.86$. The figure demonstrates that mitigation of a limited number of identifiable Type 2 geometry preprocessing issues, edge identification, and depth of free turbines already leads to substantially improved agreement between the SLS and vdLW results, even without changing the vdLW code.

Despite the methodological differences between the SLS and vdLW finite wind-farm treatments, the operational wind-farm data remain in close agreement with the theoretical limit framework. In particular, the vdLW methodology still produces a clear clustering of the operational wind farms below the theoretical limit curve and close to the 90% theoretical limit curve.

This result is important because it demonstrates that the main conclusions of SLS regarding the aerodynamic limits of offshore wind-farm performance remain robust even when using the alternative finite wind-farm methodology proposed by vdLW.

A remaining systematic difference between the SLS and vdLW methodologies is associated with the treatment of inter-wind-farm wake interactions over distances larger than approximately one wind-farm scale. The vdLW methodology does not account for these longer-range inter-wind-farm wake effects, which leads to slightly higher predicted wind-farm performance and correspondingly lower estimated Wind Farm Wind Factor values.

This systematic effect is consistent with the small remaining difference in regression slope between the SLS and vdLW comparisons discussed in the previous sections. As a consequence, the vdLW methodology tends to place the operational wind farms at slightly lower Wind Farm Wind Factor values, resulting in data points that appear marginally further below the theoretical limit curve.

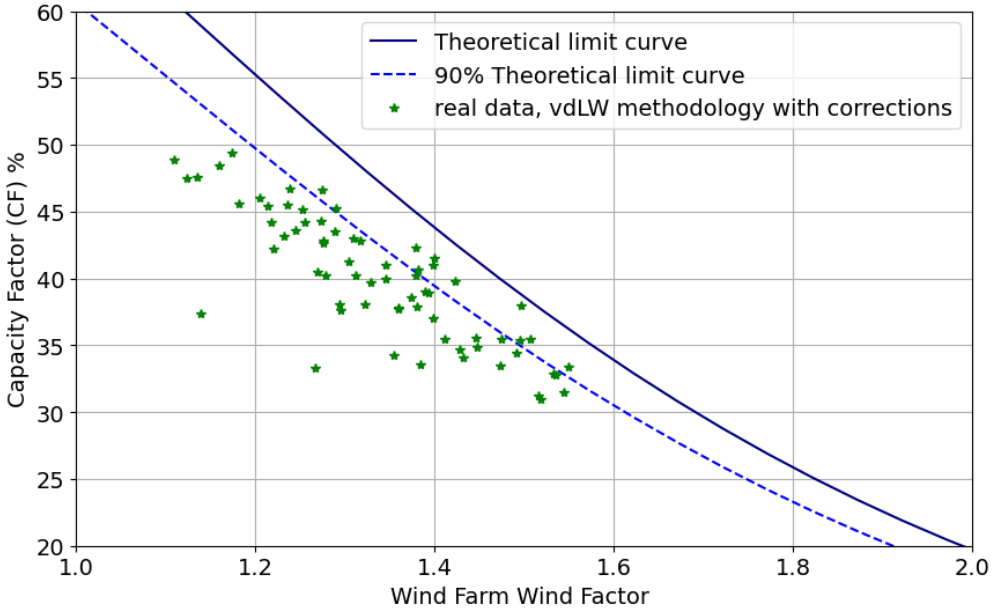

Figure 6. Operational offshore wind-farm data represented using the vdLW N_{free} methodology after correction of the Type 1 analytical implementation errors and mitigation of selected Type 2 issues, compared against the SLS theoretical limit curves. The solid blue line shows the SLS theoretical limit curve, while the dashed blue line shows the 90% theoretical limit curve. Despite the methodological differences between the SLS and vdLW finite wind-farm treatments, the operational wind-farm data remain in substantial agreement with the SLS theoretical framework, clustering below the theoretical limit curve and close to the 90% theoretical limit curve. A remaining systematic difference between the methodologies is associated with the incomplete treatment of long-range inter-wind-farm wake interactions in the vdLW methodology, which tends to produce slightly lower Wind Farm Wind Factor values and correspondingly places the wind farms marginally further below the theoretical limit curve. This result reinforces the robustness of the SLS theoretical framework and suggests that non-wake losses, including operational and curtailment losses, may play an even more important role in practical offshore wind-farm performance than currently assumed.

Importantly, this shift does not weaken the SLS theoretical framework. Instead, it reinforces the interpretation that non-wake losses, including operational losses, curtailment losses, and other non-aerodynamic effects, may play an even more important role in practical offshore wind-farm performance than currently assumed.

Overall, the comparison demonstrates that the vdLW methodology, despite its identified limitations, remains broadly consistent with the SLS theoretical limit framework and therefore reinforces rather than invalidates the central conclusions of SLS.

4 Conclusions

The present rebuttal demonstrates that the primary source of disagreement reported in the vdLW comment paper does not originate from irreproducibility of the SLS theoretical framework itself, but rather from implementation errors, altered physical inputs, and preprocessing inconsistencies in the vdLW reproduction of the SLS results and implementation of their proposed finite wind-farm correction methodology.

Audit of the vdLW implementation identified four Type 1 analytical implementation errors affecting the reproduced results. Once these errors are corrected and the published SLS inputs and assumptions are restored, the vdLW implementation reproduces the published SLS results essentially exactly, thereby resolving the main reproducibility claim raised in the comment paper.

A second analysis was then performed using the corrected analytical implementation together with the alternative finite wind-farm correction methodology proposed by vdLW. Even before correcting additional preprocessing and geometry-processing issues, agreement between the vdLW methodology and the SLS results remained strong, and improved further after mitigation of selected Type 2 preprocessing inconsistencies.

Most importantly, the present analysis shows that the discussion raised by vdLW concerns refinement of finite wind-farm correction methodologies rather than the validity of the underlying asymptotic theoretical limit itself. The infinite-farm equilibrium solution proposed by SLS remains a theoretical limit derived from atmospheric momentum-balance principles, while finite wind-farm corrections provide physically motivated methodologies for representing, through an equivalent effective flow state, the variation between freer inflow conditions near wind-farm edges and the fully developed flow conditions inside the wind-farm interior.

In this sense, the SLS theoretical limit plays, for wind farms, a role conceptually analogous to that of the Betz limit for the actuator disk. The asymptotic equilibrium solution defines the underlying theoretical limit for wind-farm performance, while finite wind-farm corrections represent finite-size effects in real wind farms. This is analogous to finite-blade corrections in rotor aerodynamics, where physically based corrections such as the Prandtl tip-loss correction refine the representation of real finite rotors without invalidating the underlying Betz-limit framework itself.

The comparison further demonstrates that the alternative methodology proposed by vdLW ultimately remains fully consistent with the SLS theoretical framework once implementation issues are corrected. In particular, operational offshore wind-farm data represented using the corrected vdLW methodology continue to cluster below the SLS theoretical limit curves and close to the 90% theoretical limit curve.

The present work further highlights that finite wind-farm correction methodologies remain an important and evolving research topic, particularly for highly complex wind-farm clusters and interacting neighboring wind farms. In this regard, the effort by vdLW toward more systematic and automated finite wind-farm correction methodologies is valuable and has provided useful inspiration for further refinement of the framework.

Overall, the present rebuttal reinforces the robustness, reproducibility, and physical consistency of the SLS theoretical limit framework, while clarifying that the remaining differences discussed in the vdLW comment are associated primarily with refinement-level treatment of finite wind-farm corrections rather than with invalidation of the underlying theory.

Because several of the discrepancies originally reported in the vdLW comment paper were shown here to originate from implementation errors, altered physical inputs, and geometry-processing limitations associated with interacting wind-farm clusters, interpretation of the original vdLW results should be made with caution unless the corrected implementation and the identified methodological limitations discussed in the present rebuttal are taken into account. In particular, the present analysis shows that a substantial fraction of the remaining discrepancies originates from cases where the current vdLW geometry methodology struggles to represent complex interacting wind-farm clusters, a limitation that is acknowledged by vdLW but whose impact on the reported disagreement with SLS was not explicitly quantified in the original comment paper. At the same time, we believe the effort by vdLW has been valuable and constructive, particularly in motivating further refinement and automation of finite wind-farm correction methodologies.

Appendix A: Responses to selected vdLW comments not addressed in the main text

The following appendix contains brief responses to selected comments by vdLW that are not already addressed in detail elsewhere in the rebuttal. The purpose of this appendix is therefore limited to clarifying a number of remaining technical and conceptual points raised in the vdLW comment paper.

---

vdLW statement in the Introduction

> "The wind industry and academia have developed a range of complex models to calculate such losses … which due to their complexity and computational cost, may not be accessible to policy makers."

Reply

We agree with this observation. One of the motivations of the SLS framework is precisely to provide a physically transparent and computationally inexpensive tool that can support first-order assessments of offshore wind-farm saturation effects and large-scale deployment scenarios.

High-fidelity mesoscale and CFD-based models remain essential tools for detailed studies. However, such models are computationally expensive, depend on multiple parameterizations and assumptions, and may themselves contain significant uncertainties associated with spatial resolution, atmospheric modeling, and wind-farm parameterizations.

For example, in a recent presentation at EGU 2026 ("Mapping the future offshore wind potential in Denmark: Assessment of 2050 wind farm scenarios"), A. Hahmann et al. state that:

> "The results carry uncertainty due to unvalidated aspects of the modeling."

Similarly, Balthazar Sengers et al., in a presentation at WESC 2025 ("Beyond standard Fitch: Validating WRF's wind farm parameterizations with offshore SCADA data"), state that WRF simulations suffer from a "Lack of validation of Wind Farm Parameterizations (WFP)" and report a "Yield underestimation of $\sim 3\%$."

These uncertainty levels are of the same order of magnitude as the mean relative error of less than 4% obtained using the much simpler SLS framework.

The purpose of the SLS methodology is therefore not to replace higher-fidelity models, but rather to provide a physically interpretable first-principles framework that can be used to estimate large-scale offshore wind-energy limits and identify potentially unrealistic deployment assumptions. At the same time, the validation results demonstrate that the framework also performs remarkably well even for comparatively small and geometrically irregular wind farms.

---

vdLW statement on Page 2

> "The plotted limit labelled as theoretical limit (solid line) in Fig. 4, 5 and 9 in Simão Ferreira et al. (2026) represents a normalized gross annual energy production (AEP), but this is not clear from the paper."

Reply

This statement is not correct. Only in the limiting case corresponding to a solitary wind turbine does the expression reduce to the normalized gross annual energy production of an isolated turbine. For all other cases, wake losses are explicitly included through the finite-wind-farm and infinite-wind-farm corrections, as described in the original paper (see Section 2 for the discussion of ϵ and the associated wake-loss formulation).

---

vdLW statement on Page 2

> "... the wind farm area ... is mathematically undefined for a wind farm layout with a concave shape."

Reply

This statement is not correct. A wind-farm area can be defined for arbitrary geometries, including concave geometries. Different geometric definitions may produce different numerical values, but this does not imply mathematical undefinedness.

---

vdLW statement on Page 3

> "... wind turbine model does not have a cut-in, $U_{\mathrm{cut-in}}$ and cut-out wind speed, $U_{\mathrm{cut-out}}$."

Reply

The validation calculations against the 72 offshore wind farms explicitly include cut-in and cut-out wind speeds. These are neglected only in the derivation of the idealized asymptotic theoretical limit.

As demonstrated in the Type 1 implementation-error analysis, the vdLW reproduction attempt used the idealized infinite operational range configuration instead of the physical validation configuration employed in the original SLS study. This modification directly changes the integrated capacity-factor calculation and was one of the dominant sources of the discrepancies reported by vdLW. Once the correct cut-in and cut-out wind speeds are restored, the agreement with the published SLS results improves substantially.

---

vdLW statement on Page 4

> "It should be noted that the infinite wind farm wake loss may never be reached for a large finite wind farm."

Reply

This is correct and entirely expected. The infinite-wind-farm solution is an asymptotic limit. Whatever the size of a finite wind farm, it will never exactly reach the asymptotic equilibrium state. However, even comparatively small wind farms already begin to tend toward the asymptotic solution. The infinite-wind-farm solution therefore remains a meaningful and physically relevant reference state within the theoretical framework.

---

vdLW statement on Pages 5–6

> "...we do not recommend our automated method to be used as a correction for finite wind farms situated in a wind farm cluster."

Reply

We find this statement difficult to reconcile with the central conclusions of the vdLW comment paper, because many of the largest discrepancies reported by vdLW originate precisely from wind farms embedded in clusters. As demonstrated in the Type 2 error analysis, the automated vdLW methodology can generate substantial geometry-processing inconsistencies and erroneous suppression of edge turbines under clustered conditions. The limitations acknowledged by vdLW therefore directly affect several of the central comparisons and conclusions presented in the comment paper.

---

vdLW statement on Page 8

> "ϵ_∞ can be approximated with a simple explicit expression, which was not shown in their work."

Reply

The fact that an implicitly solvable equation can also be approximated by an explicit fitted expression may be mathematically convenient in some contexts, but it does not materially change the underlying theoretical framework discussed here. The original formulation already provides an exact analytical relation that is computationally inexpensive, universally applicable within the model assumptions, and directly connected to the underlying physical interpretation. Replacing this with a fitted approximation valid only over a limited parameter range does not provide a clear advantage in the context of the present study.

vdLW figure 5 representation on Page 10

Reply

The long horizontal gray lines shown in the vdLW figure appear to suggest that a broad range of Wind Farm Wind Factor values may correspond to the same capacity factor. This interpretation is physically misleading and inconsistent with the formulation of the SLS framework.

For a fixed wind-farm geometry and operating condition, the SLS framework associates the wind farm with a specific physically consistent state. The theoretical relationship is therefore not intended to represent an arbitrary continuum of equally admissible operating conditions for a single wind farm.

As discussed in Section 2, not all mathematically possible values of ϵ correspond to physically realizable wind-farm states. The fact that a parameter may formally vary within an equation does not imply that all values are physically achievable for a given finite wind-farm geometry. The horizontal gray representation used by vdLW implicitly suggests that a finite wind farm could continuously vary between a state where all turbines operate as isolated turbines and a fully developed infinite-wind-farm equilibrium state. Such interpretations are not physically realizable for any finite wind-farm geometry.

The vdLW representation therefore gives a misleading impression that the model allows a broad continuum of admissible theoretical states for a single wind farm, whereas the SLS framework associates each wind farm with a unique physically consistent theoretical state.

Appendix B: Audit of the vdLW implementation and methodology: main errors and their origin

B1 Type 1 - Errors in the analytical capacity-factor implementation

B1.1 Error 1 — Removal of the Hub-Height Wind-Speed Correction

To calculate the capacity factor of a wind turbine, the correct local wind speed at the turbine hub height is required. Depending on the wind farm, these hub heights differ substantially.

Meteorological datasets, including the Global Wind Atlas dataset used in our work, provide wind-speed information at fixed reference heights. In the comment-paper code, the Global Wind Atlas data are accessed at $z_{\mathrm{Ref}} = 100\,\mathrm{m}$. This reference height does not necessarily match the actual turbine hub height H_t of each wind farm.

Therefore, the wind speed must first be correctly extrapolated/interpolated to the turbine hub height before it can be used in the analytical model. A standard and physically appropriate method is to apply the logarithmic atmospheric-boundary-layer profile (log-law correction). This correction was applied in the article by SLS.

Inspection of the code used by the authors of the comment paper shows this directly. Their own function `calc_windresource(...)` computes both the Weibull scale parameter at the 100m reference height, A_{Ref}, and the corrected hub-height value, A_H:

```
# Log interpolate A at hub height
AH = ARef * np.log(zH / z0) / np.log(zRef / z0)
```

However, this corrected value is not used in `calculate_minimalistic_model(...)`. In `run.py`, starting at line 173, their function instead reads `inputdata['lambda']`. The code contains the same required hub-height correction, but it is disabled:

```
# It seems that Sim o Ferreira et al. (2026) already corrected for this
# Awref = inputdata['lambda']
# Aws = Awref * np.log(Hs / z0) / np.log(Href / z0)
Aws = inputdata['lambda']
```

The commented lines show the appropriate correction from the reference height H_{ref} to the turbine hub height H_s using the logarithmic profile. The accompanying comment states: "It seems that Simão Ferreira et al. (2026) already corrected for this." However, the Excel input file supplied to the comment authors contains the reference height H_{ref}, always $100\,\mathrm{m}$, in a column labelled `Href`. The authors included this `Href` input in the prediction-function interface, but the quantity is not used inside the calculation itself. Furthermore, they use $z_{\mathrm{Ref}} = 100.0\,\mathrm{m}$ when retrieving the Global Wind Atlas data. Therefore, the correction should not have been removed. Instead, the code uses:

```
Aws = inputdata['lambda']
```

This means that the reference-height Weibull scale parameter is used directly as hub-height input to the analytical model, without correction to the actual turbine hub height.

As a consequence, the comment-paper calculations use incorrect wind speeds for the wind farms, corresponding to an incorrect atmospheric reference height.

This directly affects:

- the inferred geostrophic wind speed,
- the wake-reduction factor,
- the Wind Farm Wind Factor,
- and the calculated capacity factor.

This omission is one of the dominant causes of the discrepancies presented in the comment paper.

Importantly, the corrected hub-height Weibull scale parameter (A_H) is still calculated and exported in the output dataset generated by the preprocessing routines, despite the fact that the actual simulations in `calculate_minimalistic_model(...)` use the uncorrected reference-height value (`lambda`) directly.

As a result, the published output files contain hub-height-corrected wind-speed parameters that are not the quantities effectively used in the simulations themselves.

This substantially complicates independent verification and reproducibility of the calculations, because inspection of the exported data alone suggests that the physically appropriate hub-height correction had been consistently applied throughout the workflow.

The simulation workflow therefore becomes internally inconsistent, because the corrected hub-height quantities are generated and exported during preprocessing, while the subsequent simulations use the uncorrected reference-height values directly.

Importantly, once the hub-height correction is restored, the agreement between the reproduced calculations and the published results improves substantially.

A substantial fraction of the reported "non-reproducibility" in the comment paper is therefore a direct consequence of removing the hub-height correction used to generate the physically appropriate wind-speed input to the simulations.

B1.2 Error 2 — Incorrect Use of Cut-in and Cut-out Wind Speeds

SLS explicitly state in the discussion of Figure 4 of the SLS paper:

> "The infinite cut-out velocity and zero cut-in velocity are assumptions of the theoretical limit, while the wind farm simulations use a more realistic $3, \mathrm{m}, \mathrm{s}^{-1}$ cut-in and $25, \mathrm{m}, \mathrm{s}^{-1}$ cut-out velocity."

The validation calculations against real offshore wind farms therefore used:

- cut-in wind speed: $U_{\mathrm{in}} = 3, \mathrm{m}, \mathrm{s}^{-1}$,
- cut-out wind speed: $U_{\mathrm{out}} = 25, \mathrm{m}, \mathrm{s}^{-1}$.

These limits are required because real wind turbines do not produce power below cut-in wind speed and shut down above cut-out wind speed.

Inspection of the vdLW code shows that, in `run.py` lines 823–824, the authors configured their simulations using:

```
Uin = 0.0   # Cut-in wind speed [m/s]
Uout = 1e6  # Cut-out wind speed [m/s]
```

while the physically realistic turbine limits were commented out in lines 825–826:

```

# Uin = 3.0

# Uout = 25.0
```

As a consequence, the comment-paper simulations do not reproduce the validation configuration used in SLS, but instead use the idealized theoretical-limit configuration.

This directly alters the integrated capacity-factor calculation and creates an artificial discrepancy relative to the published validation results.

Importantly, once the correct turbine operating limits are restored, the agreement between the reproduced calculations and the published results improves substantially.

The correct validation inputs are already present in the vdLW implementation, but are commented out in the published calculations. Consequently, the reproduced calculations do not correspond to the published SLS validation configuration unless the original SLS inputs and assumptions are restored.

B1.3 Error 3 — Use of a Fixed Latitude Instead of Site-Specific Latitude

The geostrophic drag-law calculation depends on latitude through the Coriolis parameter:

$$f = 2\Omega \sin(\phi), \tag{B1}$$

here ϕ is the latitude of the wind farm.

Therefore, the latitude used in the model should correspond to the actual location of each wind farm. This is also consistent with the SLS methodology, where the geostrophic wind speed is inferred using site-specific conditions.

Inspection of the vdLW code shows that the latitude of each wind farm is read from the input data and stored in the output file. However, in `run.py` line 820, the actual model simulations are configured with a fixed latitude:

```
lat = 55  # Latitude [degree]
```

This fixed value is then passed to the model for all wind farms in line 907:

```
Cf_Free, Cf_Inf, Cf_Model, U_Free, U_Inf, U_Model,
G_Model, Phi_Model, Phi_Model_redo, phi_Model_a5p3,
Cf_Model_redo, Cf_Model_a5p3 =
calculate_minimalistic_model(
inputdata,
lat,
CP,
CT,
z0,
kw,
Uin,
Uout,
rho,
zRef,
Nwt_row_wf_neighbors
)
```

As a consequence, the Coriolis parameter is not calculated using the actual latitude of each wind farm, but using the same latitude for every wind farm.

This error is smaller than the hub-height wind-speed error and the cut-in/cut-out error, but it still modifies the geostrophic drag-law calculation and introduces discrepancies of approximately 1%.

Once the site-specific latitude is restored, the agreement with the published SLS results improves further.

B1.4 Error 4 — Incorrect Equation in the Capacity-Factor Calculation

The fourth implementation error occurs in the calculation of the wind-speed reduction used in the capacity-factor integration.

In `run.py` line 151, vdLW use:

```
eps2 = (1 + gam / delta) / (
1 + gam / kappa * np.sqrt(
Ctau * (Ur / Uh)**3.2 + (kappa / delta)**2
)
)
```

This differs from the equation used in SLS and in the original formulation of Sørensen and Larsen (2021). The equation should use the cut-out wind speed U_{out}, not U_h:

```
eps2 = (1 + gam / delta) / (
1 + gam / kappa * np.sqrt(
Ctau * (Ur / Uout)**3.2 + (kappa / delta)**2
)
)
```

Using U_h instead of U_{out} appears to result from an incorrect reading of the equation in Sørensen and Larsen (2021). It changes the wake-reduced wind-speed correction and therefore modifies the integrated capacity factor.

This error is smaller than the hub-height wind-speed error and the cut-in/cut-out error, but it still introduces discrepancies of approximately 1%.

Once this equation is corrected, the agreement with the published SLS results improves further.

Together with the previous corrections, this removes the remaining differences between the vdLW calculations and the published SLS results.

After correcting the four implementation errors discussed above, the vdLW code reproduces the published SLS results with negligible numerical differences for all wind farms. The discrepancies reported in the vdLW comment therefore do not originate from irreproducibility of the SLS methodology or results, but from identifiable implementation inconsistencies and coding errors in the vdLW reproduction attempt.

Importantly, the two dominant sources of discrepancy originate from the use of inputs and configurations different from those used in SLS, despite the correct implementations already being present in the vdLW code itself.

The results presented here demonstrate that the analytical framework and published results of SLS are fully reproducible when the correct equations, inputs, and physical assumptions are consistently applied.

B2 Type 2 — Errors in the Wind-Farm Geometry and Finite-Farm Methodology

The following examples demonstrate that the vdLW geometry methodology can substantially misrepresent the aerodynamic environment of offshore wind-farm clusters, particularly for irregular layouts, fragmented wind farms, and densely developed offshore regions.

B2.1 Error 5 — Missing Neighboring Wind Farms

The vdLW automated geometry algorithm only loads wind farms that are themselves targets of the capacity-factor assessment. As a consequence, neighboring wind farms that physically exist upstream are ignored if they are not themselves included as target wind farms in the evaluated dataset.

This creates an artificial overestimation of clean inflow directions and an underestimation of inter-farm interaction effects.

In reality, the atmosphere does not distinguish between wind farms that are targets of an evaluation and wind farms that are not. Any upstream wind farm can modify the inflow conditions experienced by downstream wind farms.

This issue becomes particularly important in densely developed offshore clusters, where neighboring wind farms strongly interact aerodynamically.

As a consequence, the vdLW methodology can incorrectly classify sectors as exposed to free inflow, even when substantial upstream wind-farm blockage physically exists.

This directly affects:

- the number of turbines exposed to clean inflow,
- the estimated finite wind-farm correction,
- the inferred wake losses,
- and the calculated capacity factor.

The resulting geometry therefore does not correspond to the real aerodynamic environment experienced by the wind farm.

This limitation is visible in `run.py` lines 865–866, where neighboring wind farms are identified only from the set of wind farms already loaded as targets of the capacity-factor assessment:

```
# Find wind farm neighbors

wf_neighbors, wf_neighbors_flag = get_wf_neighbor(
    wf_lats,
    wf_lons,
    inputdata['Name'],
    lat_dist=0.5,
    lon_dist=0.5
)
```

Because `wf_lats`, `wf_lons`, and `inputdata['Name']` only contain wind farms included in the evaluated dataset, any physically existing neighboring wind farm outside that list cannot be identified as a neighbor and is therefore absent from the finite wind-farm correction.

A clear example is the absence of the Borssele III-IV and Borssele V wind farms from the neighbor set. These wind farms physically affect the aerodynamic environment of Borssele I-II and of the Belgian offshore wind farms. If they are not loaded as neighboring wind farms, the vdLW algorithm overestimates the amount of clean inflow reaching these farms from those wind directions.

The code by vdLW generates figures for all evaluated wind farms including the detected neighboring wind farms and edge turbines. Furthermore, the source code shows that particular attention was given by vdLW to the Belgian wind farms and Borssele I-II through manual corrections of the edge-detection procedure in `run.py` lines 877–897.

The missing-neighbor issue therefore remained present in the published geometry evaluations despite the existence of explicit geometry-review procedures in the vdLW workflow.

Importantly, this issue could already have been largely corrected without modifying the vdLW code itself, simply by adding Borssele III-IV and Borssele V to the input list of evaluated wind farms. These wind farms are already present in the wind-farm database used by vdLW and therefore only needed to be included in the analyzed wind-farm list.

To demonstrate this issue, Figures B1 and B2 compare the geometry published by vdLW in their Zenodo repository for the evaluation of Borssele I-II at a wind direction of 300° with a corrected geometry generated by adding the missing neighboring wind farms to the input file of analyzed wind farms.

The impact of the missing neighboring wind farms on the detected inlet-edge turbines is immediately visible, reducing the number of turbines identified as exposed to clean inflow from 15 to 4. This demonstrates that the original vdLW geometry does not correctly represent the physical aerodynamic environment of the wind-farm cluster in some cases.

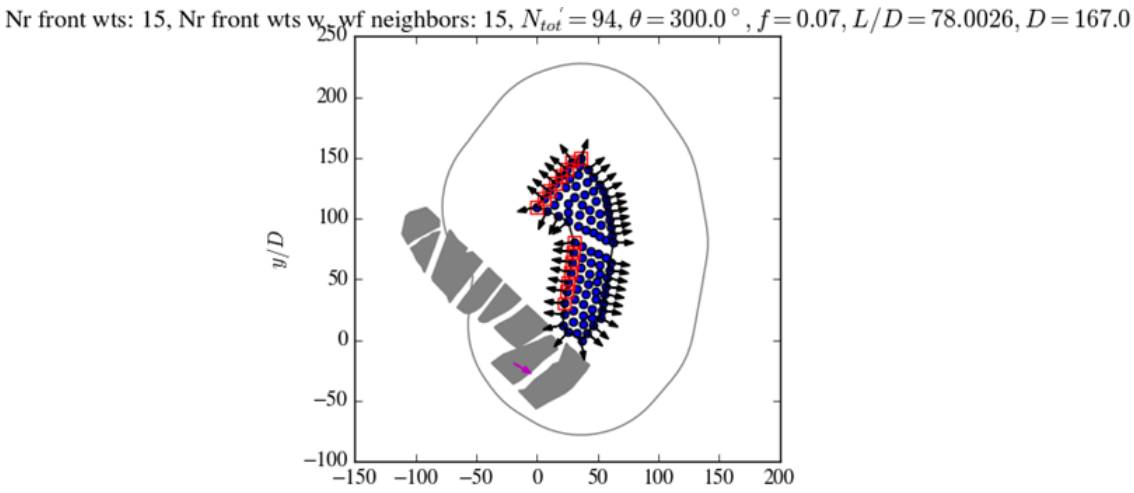

Figure B1. Original vdLW geometry for the evaluation of Borssele I-II at a wind direction of 300°. The neighboring wind farms Borssele III-IV and Borssele V are absent from the geometry evaluation, resulting in 15 turbines being identified as exposed to clean inflow. Pink arrows show the wind direction.

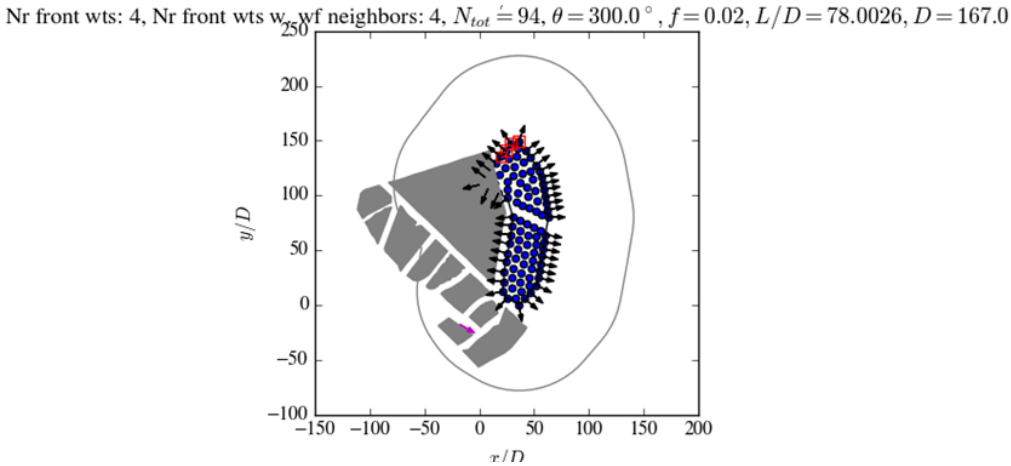

Figure B2. Corrected geometry obtained by adding the missing neighboring wind farms to the input dataset. The inclusion of Borssele III-IV and Borssele V strongly modifies the aerodynamic exposure of the wind farm, reducing the number of turbines identified as exposed to clean inflow from 15 to 4. Pink arrows show the wind direction.

B2.2 Error 6 — Treating a Neighboring Wind Farm as an Infinite Wake Boundary

The vdLW methodology treats an upstream neighboring wind farm as if it creates a fully blocking wake boundary for the target wind farm. If the target wind farm is located behind a neighboring wind farm for a given wind direction, then the target farm can be assigned zero inlet-edge turbines in that direction. In practice, this means that the downstream wind farm is treated

aerodynamically as if it were already operating fully inside an infinite-wind-farm equilibrium state, with no turbines exposed to freer inflow conditions. This assumption is not physically justified.

The problem is that the vdLW code applies this behavior even when the upstream neighboring wind farm contains only a single row of turbines. A single upstream turbine row does not physically cause the downstream wind farm to operate as if all turbines were embedded inside an infinite-wind-farm equilibrium layer. It only modifies part of the inflow.

A simple example illustrates the problem. Consider a 5×5 turbine array, treated as one wind farm. For a wind direction perpendicular to one side, the code would identify 5 inlet-edge turbines. With the vdLW factor of 2.5, this gives:

$$N_{\text{isolated}} = 2.5 \times 5 = 12.5$$

turbines equivalent out of 25 turbines, or 50% of the farm exposed to clean-flow conditions.

Now consider the same physical turbine layout, but suppose that the first upstream row of 5 turbines belongs administratively to a different wind farm, with a different project name or owner. Aerodynamically, the configuration remains essentially unchanged, and the aerodynamic response should therefore remain broadly similar. However, in the vdLW treatment, the upstream row is now treated as a neighboring wind farm. The downstream wind farm is then considered to be behind a neighbor, and the code can assign it zero inlet-edge turbines in that wind direction.

The result is that the same physical layout changes from an estimated 50% clean-flow contribution to approximately 20%, purely because the upstream row is assigned a different administrative wind-farm label. This is an artificial consequence of the geometry algorithm, not an aerodynamic effect. The atmosphere does not recognize ownership boundaries, concession labels, or project names. A row of turbines has the same aerodynamic effect whether it is called part of Wind Farm A or part of Wind Farm B.

This error directly affects:

- the detected inlet-edge turbines,
- the finite wind-farm correction,
- the inferred wake losses,
- and the calculated capacity factor.

The resulting geometry can therefore misrepresent the real aerodynamic environment of the wind-farm cluster.

To illustrate this issue, Figure B3 compares two physically identical turbine layouts for a wind farm array of 5×5 turbines. In the first case, all turbines belong to the same wind farm. In the second case, the upstream row is assigned administratively to a different wind farm. Although the aerodynamic configuration is essentially unchanged, the vdLW methodology produces a substantially different estimate of the number of turbines exposed to clean inflow.

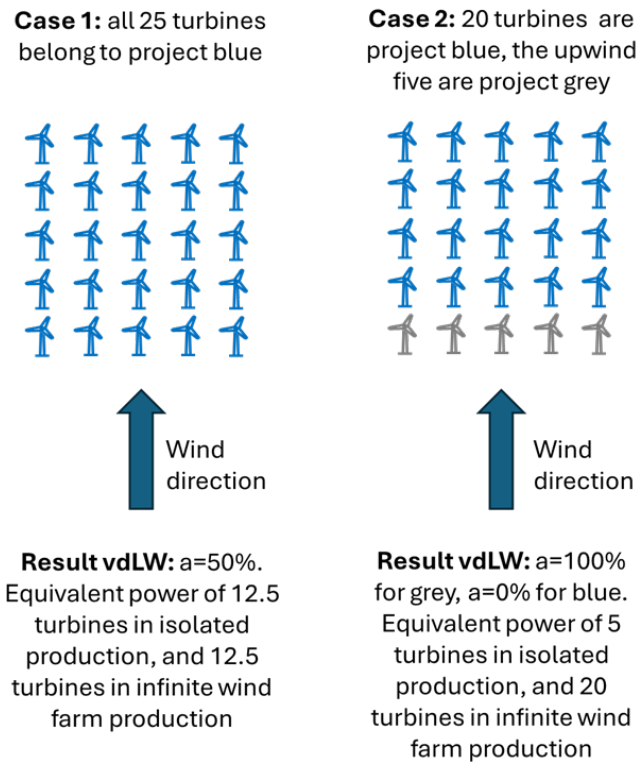

Figure B3. Comparison between aerodynamic and administrative interpretations of the same physical turbine layout. (a) A 5×5 turbine array treated as a single wind farm. For the incoming wind direction shown, the vdLW methodology identifies 5 inlet-edge turbines, resulting in an estimated clean-flow contribution equivalent to 12.5 turbines out of 25. (b) The same physical turbine layout, but with the first upstream row assigned administratively to a different wind farm. Although the aerodynamic configuration is essentially unchanged, the vdLW methodology now treats the upstream row as a neighboring wind farm and assigns zero inlet-edge turbines to the downstream wind farm, reducing the estimated clean-flow contribution to approximately 20% of the farm. The figure demonstrates that the vdLW geometry methodology can produce substantially different results for the same aerodynamic configuration purely because of administrative wind-farm labelling.

The same behavior is visible in figures generated and published by vdLW together with their code and Zenodo repository. Figure B4 shows the Greater Gabbard wind farm for a wind direction of $90°$. In this case, the vdLW methodology calculates $N_{\text{isolated}} = 0$, meaning that the wind farm is considered to experience no clean inflow.

The figure shows how a relatively narrow upstream turbine arrangement can be treated by the vdLW methodology as a complete aerodynamic blockage, effectively causing the downstream wind farm to behave as if it were operating fully inside an infinite-wind-farm equilibrium state.

The figure also points to the next geometry error discussed in this rebuttal: the automated method artificially connects separated turbine groups and geometries in order to create continuous upstream shadow boundaries.

The top panel corresponds to the original figure generated and published by vdLW together with their code and Zenodo repository. The bottom panel was generated using the same vdLW code, but additionally plotting the neighboring wind turbines explicitly, and changing the neighboring-wind-farm shadow region from opaque dark gray to semi-transparent gray.

Several important issues become immediately visible:

- First, in the most southern part of the geometry, the neighboring shadow region is generated by only a single row of wind turbines. Nevertheless, the vdLW methodology treats this as sufficient to produce a complete aerodynamic blockage.
- Second, the neighboring gray shadow region does not actually follow the neighboring wind-farm geometry. In several regions, the shadow boundary extends beyond the neighboring wind farm itself and even overlaps the target wind farm, Greater Gabbard.
- Third, an artificial wake-shadow region is created in areas where no neighboring wind turbines exist. This occurs because disconnected neighboring turbine groups are artificially connected by the vdLW geometry method. As a consequence, the empty space between separated neighboring wind-farm regions is effectively treated as producing complete aerodynamic blockage equivalent to infinite-wind-farm behavior for the downstream turbines.

These last two issues are highly relevant and are discussed further in the next geometry-error section.

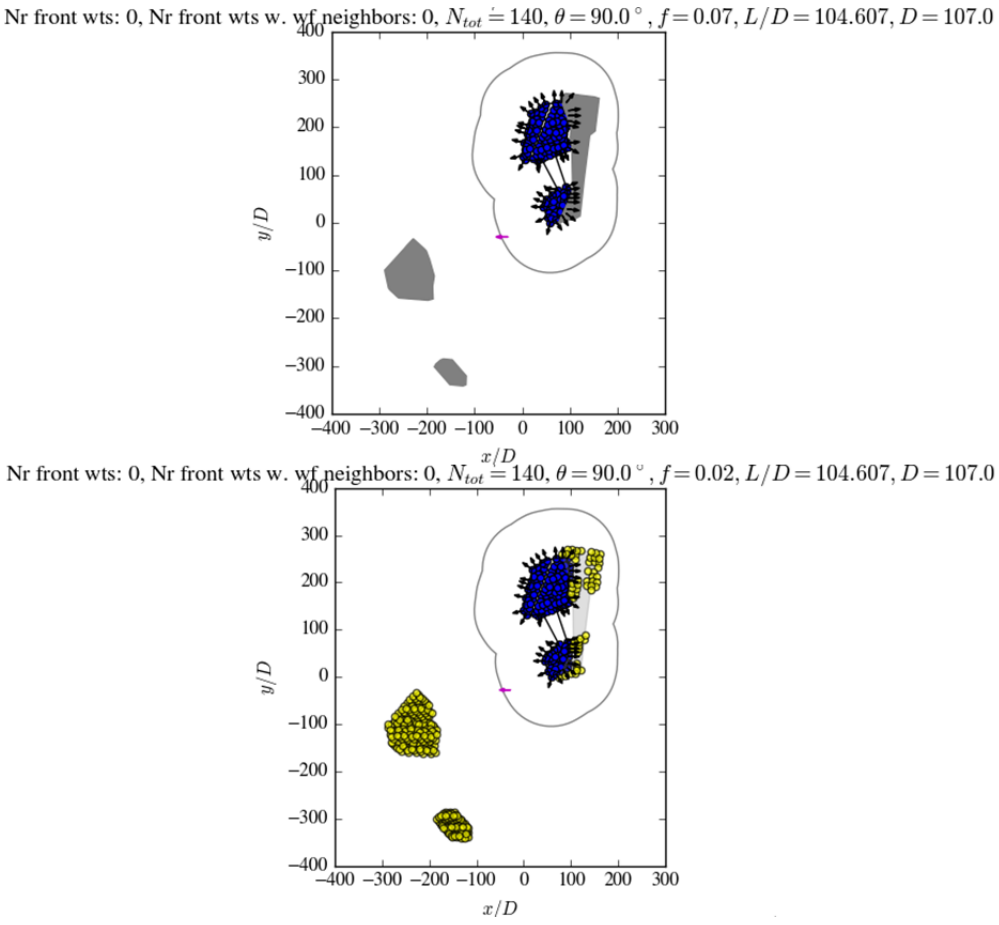

Figure B4. Greater Gabbard example for a wind direction of 90°. Top: original figure generated and published by vdLW. The methodology assigns $N_{\text{isolated}} = 0$, meaning that the wind farm is considered to experience no clean inflow. Bottom: same case generated using the vdLW code, but additionally plotting the neighboring wind turbines explicitly and rendering the neighboring shadow region semi-transparent. The figure reveals that parts of the neighboring shadow are generated by only a single upstream turbine row, that the shadow region extends beyond the actual neighboring wind-farm geometry, and that artificial shadow regions are created in empty areas without turbines due to the artificial connection of separated neighboring turbine groups. The figure demonstrates how the vdLW geometry methodology can generate non-physical aerodynamic blockage regions and unrealistically suppress clean inflow. Pink arrows show the wind direction.

B2.3 Error 7 — Creation of False Wind-Farm Regions by Artificial Connection of Disjoint Wind-Farm Regions

The vdLW automated geometry methodology can artificially connect separated turbine groups and disjoint wind-farm regions into a single continuous neighboring-shadow boundary.

As a consequence, empty regions without turbines can become treated as if they were occupied by a continuous upstream neighboring wind farm.

This creates artificial aerodynamic shadow regions.

The problem was already visible in the Greater Gabbard example discussed previously. In that case, disconnected neighboring turbine groups were artificially connected by the vdLW geometry method, generating a continuous neighboring-shadow region extending across areas where no turbines physically exist.

Once these artificial neighboring regions are created, they are then treated aerodynamically as continuous upstream blockage regions. The resulting artificial geometry can therefore suppress clean inflow both for the target wind farm itself and for neighboring wind farms.

This issue becomes particularly problematic when:

- neighboring wind farms are fragmented into separated turbine groups,
- disconnected sub-regions exist within the same administrative wind farm.

In these situations, the vdLW geometry methodology can effectively create aerodynamic barriers across empty ocean regions. This assumption is not physically justified.

A wake-shadow region should only exist where turbines physically extract momentum from the flow. Empty regions between disconnected turbine groups cannot behave as continuous aerodynamic blockage regions equivalent to infinite-wind-farm behavior.

This error directly affects:

- the detected inlet-edge turbines,
- the finite wind-farm correction,
- the inferred wake losses,
- and the calculated capacity factor.

The resulting geometry can therefore significantly misrepresent the real aerodynamic environment of offshore wind-farm clusters.

A direct example is shown in Figure B5 for the Gemini wind farm. In this case, the wind direction is 180°, corresponding to flow from the bottom of the figure toward the top.

The Gemini layout consists of two physically separated sub-farms. However, the vdLW geometry method constructs a continuous external boundary that artificially connects both regions.

As a consequence, the empty space between the two sub-farms becomes treated as part of a continuous aerodynamic boundary.

This artificial boundary then suppresses part of the clean inflow to turbines in the right-hand sub-farm. In particular, turbines that should physically experience direct inflow become partially shadowed because the vdLW method interprets the empty region between the two sub-farms as part of a continuous aerodynamic obstacle.

This behavior is aerodynamically inconsistent with the physical turbine layout. No turbines exist in the empty region between the two Gemini sub-farms, and therefore no aerodynamic wake extraction should occur there.

The example demonstrates how the vdLW geometry methodology can create artificial aerodynamic blockage regions purely as a consequence of the boundary-construction algorithm.

This can be tested directly by changing only the administrative labeling of the Gemini layout. When the two physically separated halves of Gemini are treated as two independent wind farms, Gemini 1 and Gemini 2 (simply by changing the input file of cases analyzed by the vdLW code), the artificial connecting boundary disappears. The aerodynamic layout and wind direction are unchanged; only the administrative grouping is changed.

With this change, the number of detected inlet-edge turbines increases from 21 to 31 for the same wind direction. This confirms that the difference is not caused by aerodynamics, but by the way the vdLW boundary algorithm connects disconnected sub-farms when they share the same wind-farm name.

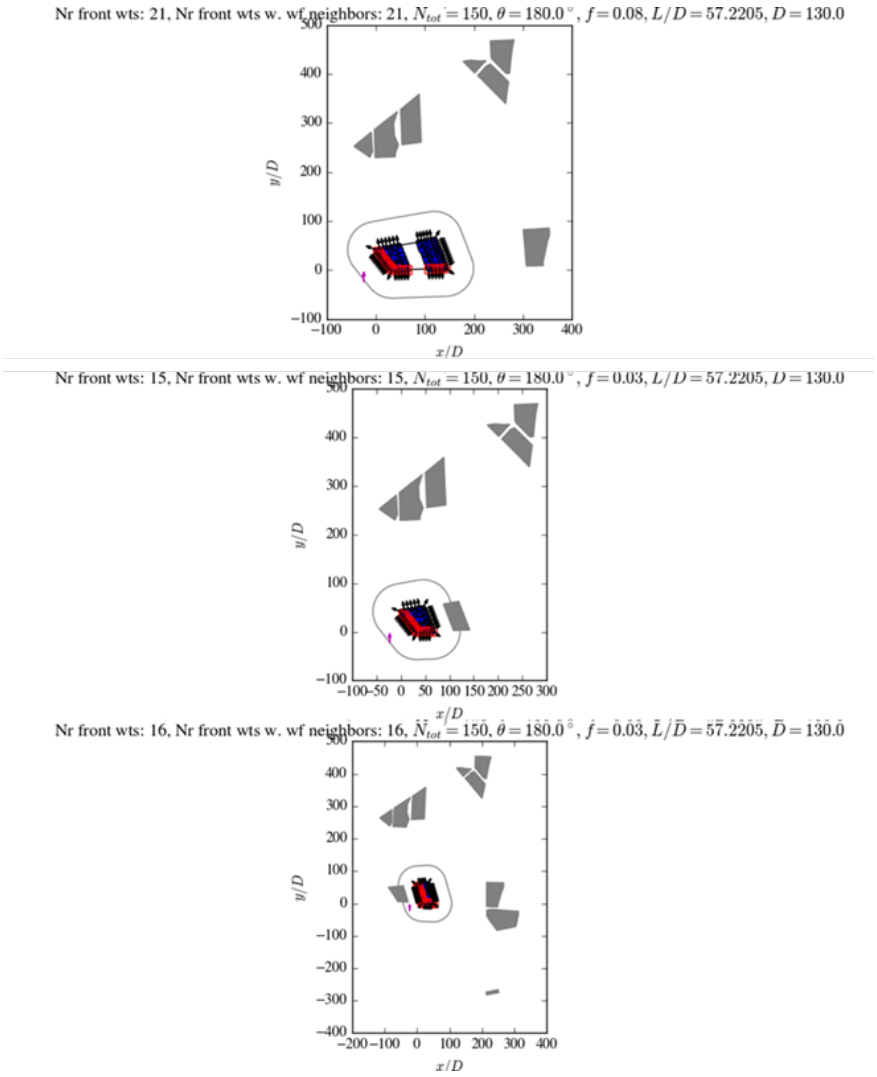

Figure B5. Gemini wind-farm example for a wind direction of 180°. Top: original vdLW treatment, where the two physically separated Gemini sub-farms are treated as a single wind farm. The vdLW boundary-generation method constructs a continuous external boundary connecting both regions, so the empty space between the sub-farms becomes part of an artificial aerodynamic obstacle. In this configuration, 21 inlet-edge turbines are detected. Middle and bottom: corrected administrative treatment, where the same physical layout is split into Gemini 1 and Gemini 2. The artificial connecting boundary disappears, and the number of detected inlet-edge turbines increases from 21 to 31 for the same wind direction. The figure demonstrates that the vdLW edge-detection methodology depends on administrative grouping rather than only on the physical aerodynamic layout. Pink arrows show the wind direction.

B2.4 Error 8 — Suppression of Wind Farms Under False Neighbor Wind-Farm Regions

The vdLW geometry methodology can generate neighboring wake-boundary regions that extend beyond the actual neighboring wind-farm geometry and overlap other wind farms.

As a consequence, wind farms can become artificially embedded inside neighboring-shadow regions, suppressing clean inflow and reducing the number of detected inlet-edge turbines. This issue was already visible in the Greater Gabbard example discussed previously.

In the vdLW methodology, neighboring wind farms are converted into neighboring-shadow regions through the construction of external aerodynamic boundaries. However, these generated neighboring-boundary regions do not necessarily follow the real turbine geometry. Instead, the generated regions can:

- extend beyond the physical neighboring wind-farm layout,
- overlap the target wind farm,
- and cover empty ocean regions where no turbines exist.

This creates a highly problematic situation. A target wind farm can become partially or fully embedded inside an artificial neighboring-shadow region that does not correspond to the real aerodynamic environment. The consequence is a suppression of detected clean inflow. This directly affects:

- the number of inlet-edge turbines,
- the finite wind-farm correction,
- the inferred wake losses,
- and the predicted capacity factor.

The problem is particularly severe because the artificial neighboring-shadow regions are treated aerodynamically as continuous upstream blockage regions, even in areas where no neighboring turbines physically exist.

As demonstrated previously for Greater Gabbard, the neighboring-shadow region generated by the vdLW methodology can overlap the target wind farm itself.

In some regions, the neighboring-shadow boundary extends significantly beyond the neighboring turbine layout, while in other regions disconnected neighboring turbine groups become artificially connected into continuous aerodynamic obstacles.

This behavior is aerodynamically inconsistent with the physical turbine layout.

Wake-shadow regions should only exist where turbines physically extract momentum from the atmospheric flow.

Artificially extending neighboring wind-farm boundaries into empty regions or across other wind farms introduces artificial aerodynamic suppression and distorts the real aerodynamic interaction between offshore wind-farm clusters.

A key example of this problem is the suppression of the Belwind (Phase 1 and 2) wind farm under the neighboring Nobelwind wind farm.

The Nobelwind wind farm was constructed several years after Belwind and bounds Belwind both to the east and to the west. Most of the Nobelwind turbines are located east of Belwind, while only a single turbine row is located to the west.

However, the vdLW methodology artificially connects the eastern and western sections of Nobelwind, creating a false neighboring wind-farm region that overlaps most of the Belwind wind farm.

As a consequence, the Belwind wind farm becomes almost completely embedded inside the false neighboring-shadow region, and the vdLW methodology concludes that Belwind experiences very limited direct clean inflow.

Figure B6 shows the Belwind case for a wind direction of 210° (south-southwest).

In the upper panel, generated by SLS using the vdLW code, the Belwind turbines are shown in blue and the neighboring turbines are shown in yellow. The gray region corresponds to the neighboring wind-farm boundary generated by the vdLW methodology.

The middle panel corresponds to the figure published by vdLW. The Belwind wind farm almost disappears beneath the false neighboring wind-farm region generated by the artificial connection between the Nobelwind turbine groups located west and east of Belwind.

To help the reader interpret the geometry, the lower panel shows the Nobelwind case itself for the same wind direction.

Comparing the upper panel with the lower two panels also reveals the absence of the Borssele III–V wind farms in the original vdLW treatment. This missing-neighbor issue was already discussed previously under Error 5.

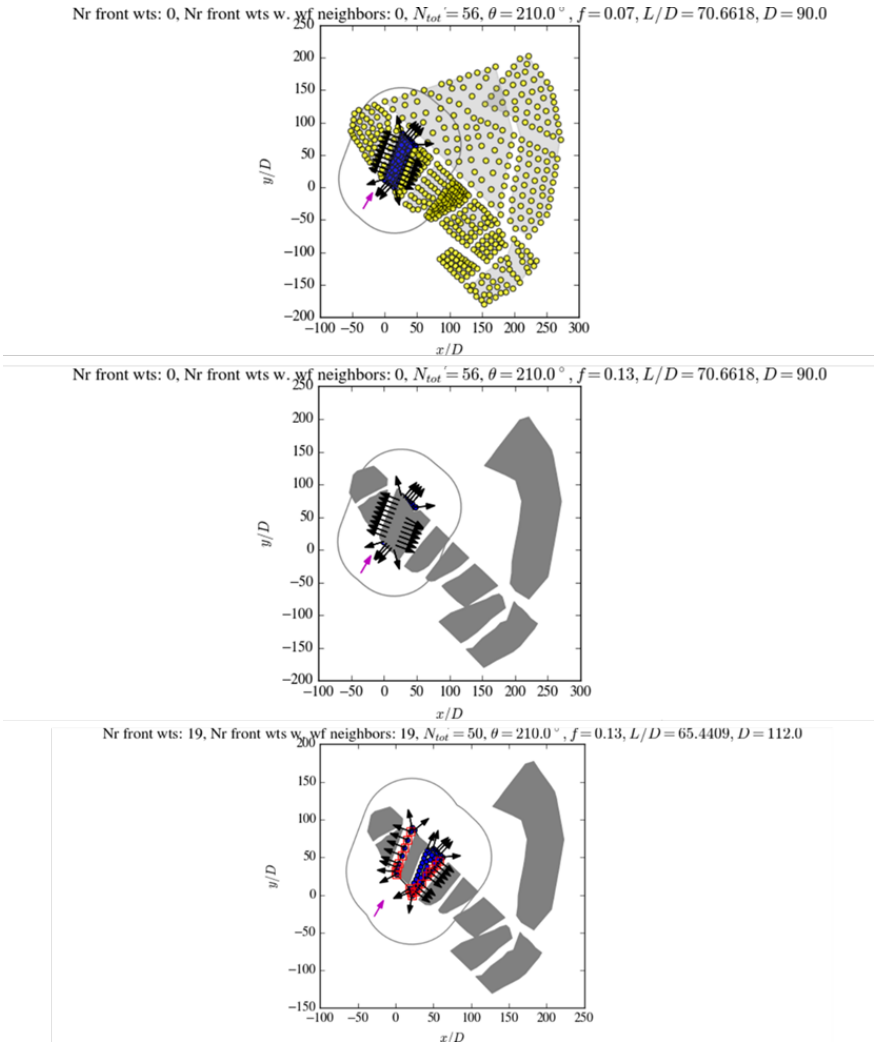

Figure B6. Belwind and Nobelwind example for a wind direction of 210°. Top: geometry generated by SLS using the vdLW code, showing the Belwind turbines in blue, neighboring turbines in yellow, and the neighboring wind-farm boundary region generated by the vdLW methodology in gray. The vdLW method artificially connects the Nobelwind turbine groups located west and east of Belwind, generating a false neighboring-shadow region that overlaps most of the Belwind wind farm. Middle: original figure published by vdLW for Belwind. The Belwind wind farm becomes almost entirely embedded inside the false neighboring wind-farm region, leading the methodology to conclude that Belwind experiences very limited clean inflow. Bottom: Nobelwind case for the same wind direction, helping visualize how the disconnected Nobelwind turbine groups become artificially connected into a continuous neighboring-shadow region. The figure demonstrates how the vdLW geometry methodology can suppress the aerodynamic exposure of entire wind farms through the generation of false neighboring wind-farm regions. Pink arrows show the wind direction.

B2.5 Error 9 — Failure to Correctly Identify the Wind-Farm Edge

The vdLW methodology uses an automated edge-detection procedure to identify the inlet-edge turbines of each wind farm. However, the detected wind-farm edge depends strongly on a non-unique geometric tuning parameter and can fail to correctly identify turbines directly exposed to incoming flow.

This step is critical because the estimated number of inlet-edge turbines directly determines the finite wind-farm correction and therefore strongly affects the predicted capacity factor.

However, the vdLW edge-detection methodology can fail to correctly identify turbines that are directly exposed to incoming flow.

A direct example is shown in Figure B7 for the Northwind wind farm, without neighboring wind farms, for a wind direction of $150°$. The figure was published by vdLW in the *Zenodo* repository that contains vdLW code and automatically generated figures.

In this case, the vdLW methodology detects only 9 inlet-edge turbines, while visual inspection shows that approximately 13–14 turbines are directly exposed to incoming flow. The problem is caused by the generated aerodynamic boundary used by the vdLW geometry methodology. Several turbines located on the aerodynamic edge of the wind farm fall slightly inside the generated boundary and are therefore incorrectly excluded from the inlet-edge set. As a consequence, turbines directly exposed to incoming atmospheric flow are incorrectly treated as interior turbines, directly affecting the finite wind-farm correction and the predicted capacity factor.

The vdLW code itself shows that the authors were aware that the boundary-generation method required ad-hoc tuning. In `run.py` lines 870–903, they define a default value for the concave-hull parameter and then manually change it for selected wind farms:

```
# A Concave hull method is used to determine the farm edge,
# which is not unique and depends on alpha.

# A larger alpha means more concave.
# A too large alpha can result in strange results.

alphas = [0.0001] * len(inputdata['Nt'])

alphas[inputdata.loc[inputdata["Name"] == 'Amrumbank West', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Kriegers Flak', :].index[0]] = 0.00025
alphas[inputdata.loc[inputdata["Name"] == 'Beatrice extension', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Belwind Phase 1', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Borssele I-II', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'DanTysk', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Dudgeon', :].index[0]] = 0.0005
```

```
alphas[inputdata.loc[inputdata["Name"] == 'Galloper', :].index[0]] = 0.00015
alphas[inputdata.loc[inputdata["Name"] == 'Gode 1 and 2', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Greater Gabbard', :].index[0]] = 0.00025
alphas[inputdata.loc[inputdata["Name"] == 'Gunfleet Sand', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Gwynt y M r', :].index[0]] = 0.001
alphas[inputdata.loc[inputdata["Name"] == 'Hornsea 1', :].index[0]] = 0.00025
alphas[inputdata.loc[inputdata["Name"] == 'Hornsea 2', :].index[0]] = 0.0002
alphas[inputdata.loc[inputdata["Name"] == 'Horns Rev 3', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Kaskasi', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Lillgrund', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Lincs', :].index[0]] = 0.001
alphas[inputdata.loc[inputdata["Name"] == 'Meerwind Sud/Ost', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Merkur', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Moray East', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Nobelwind', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Nordsee One', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Norther', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Northwester 2', :].index[0]] = 0.0005
alphas[inputdata.loc[inputdata["Name"] == 'Princess Amalia', :].index[0]] = 0.001
alphas[inputdata.loc[inputdata["Name"] == 'Race Bank', :].index[0]] = 0.001
alphas[inputdata.loc[inputdata["Name"] == 'Rampion', :].index[0]] = 0.00025
alphas[inputdata.loc[inputdata["Name"] == 'Veja Mate', :].index[0]] = 0.00025
alphas[inputdata.loc[inputdata["Name"] == 'Walney 2', :].index[0]] = 0.00025
alphas[inputdata.loc[inputdata["Name"] == 'Walney Extension', :].index[0]] = 0.00025
```

This confirms that the detected wind-farm edge depends on a non-unique geometric tuning parameter. The need for farm-specific tuning demonstrates that the edge-detection procedure is not uniquely defined and does not generalize robustly across wind-farm geometries.

Importantly, vdLW themselves acknowledge in the comment paper that the methodology can have difficulties identifying wind-farm edges in some cases, although the resulting impact on the finite wind-farm correction and predicted capacity factors is not discussed in detail.

For wind farms such as Northwind, this leads to an underestimation of turbines operating under isolated or clean-inflow conditions.

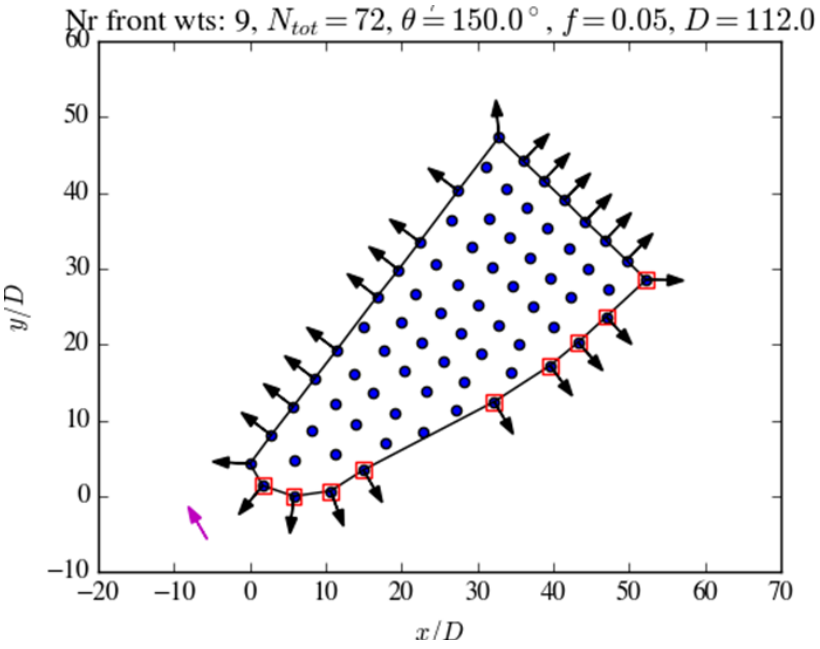

Figure B7. Northwind wind-farm example for a wind direction of 150°, without neighboring wind farms. The vdLW methodology detects only 9 inlet-edge turbines, while approximately 13–14 turbines are directly exposed to incoming flow. Due to the generated concave-hull boundary used by the vdLW geometry methodology, several turbines on the aerodynamic edge fall slightly inside the generated boundary and are therefore incorrectly excluded from the inlet-edge set. Pink arrows show the wind direction.

B2.6 Error 10 — Ignoring the Temporal Evolution of Wind-Farm Clusters

The vdLW evaluation assumes a fixed 2025 wind-farm configuration, despite the measured production data spanning years during which neighboring wind farms were progressively constructed and commissioned. The vdLW simulations therefore compare measured production from partially developed wind-farm clusters against aerodynamic environments corresponding to much later cluster-development stages.

As a consequence, the aerodynamic environment used in the vdLW simulations does not correspond to the real aerodynamic environment experienced by the wind farms during the years covered by the measured production data.

This creates a mismatch between:

- the real neighboring wind farms present during operation,
- and the neighboring wind farms included in the vdLW simulations.

The error is particularly important for rapidly evolving offshore wind-farm clusters, where neighboring wind farms were added progressively over time and the aerodynamic environment evolved continuously during the measurement period.

As a consequence, the vdLW methodology can:

- overestimate neighboring wake effects for earlier years,

- suppress clean inflow that physically existed during operation,
- and underestimate the number of turbines exposed to isolated or near-isolated inflow conditions.

This directly affects:

- the detected inlet-edge turbines,
- the finite wind-farm correction,
- the inferred wake losses,
- and the predicted capacity factor.

The issue is methodological. A reproducibility assessment using measured production data should reconstruct, as closely as possible, the real wind-farm environment corresponding to the years of operation used in the validation. Otherwise, part of the disagreement can originate from comparing historical production data against aerodynamic environments that did not yet physically exist during operation.

B2.7 Error 11 — Assuming Unrealistically Large Numbers of Turbines Exposed to Clean Inflow

The vdLW methodology estimates the number of turbines exposed to clean inflow using the relation:

$$N_{\mathrm{free}} = 2.5M,$$

where M is the number of detected inlet-edge turbines.

This factor of 2.5 is not a general physical constant. It comes from applying the finite-size correction with $a = 5$, corresponding to the characteristic finite-farm correction parameter used in SLS for a square wind-farm scaling approximation, and assuming that the exposed edge of a square wind farm corresponds to approximately $2\sqrt{N}$ turbines. This leads to the approximation that the number of turbines exposed to clean inflow is about 2.5 times the number of detected inlet-edge turbines.

However, applying this relation as a fixed multiplier to arbitrary wind-farm geometries is a misinterpretation of both SLS and the earlier work by Sørensen and Larsen.

The correction must depend on the total number of turbines in the wind farm and on the actual wind-farm geometry.

For many wind farms in the validation dataset, using a factor close to 2.5 may be a reasonable approximation. However, for smaller wind farms, or for geometries where the detected edge length is large relative to the total number of turbines, this approach can produce impossible results.

In particular, the method can estimate that more turbines operate in isolated-power mode than there are turbines in the wind farm.

This is physically inconsistent. A finite wind-farm correction cannot assign more clean-inflow turbines than the total number of turbines available.

This problem is visible for small wind farms such as Albatros and Northwester 2. In several wind directions, the vdLW method calculates N_{free} values larger than the total number of turbines in the wind farm.

For example, in the Albatros case shown in Figure B8, vdLW detects 7 inlet-edge turbines in a wind farm with only 16 turbines. Applying the fixed factor of 2.5 gives 17.5 turbines exposed to clean inflow, which is larger than the total number of turbines in the wind farm.

Similarly, for Northwester 2, vdLW detects 11 inlet-edge turbines in a wind farm with only 23 turbines. Applying the same factor gives 27.5 turbines exposed to clean inflow, again exceeding the total number of turbines.

These examples are particularly important because the vdLW method also fails to identify all edge turbines in several cases. In practice, different errors can partly cancel each other: under-detection of edge turbines can sometimes mask the overestimation caused by the fixed 2.5 multiplier. As a consequence, apparent agreement in some cases can partly result from compensating geometry and finite-farm-correction errors rather than from a physically consistent aerodynamic representation.

vdLW appear to be aware of this issue, as their code adds a limit check to ensure that the number of free turbines does not exceed the total number of turbines in `run.py` line 107:

```
1: Nfree = np.minimum(Nfree, Nturb)  # To make sure Nfree/Nturb is not larger than one, not
       yet implemented in PyWake
```

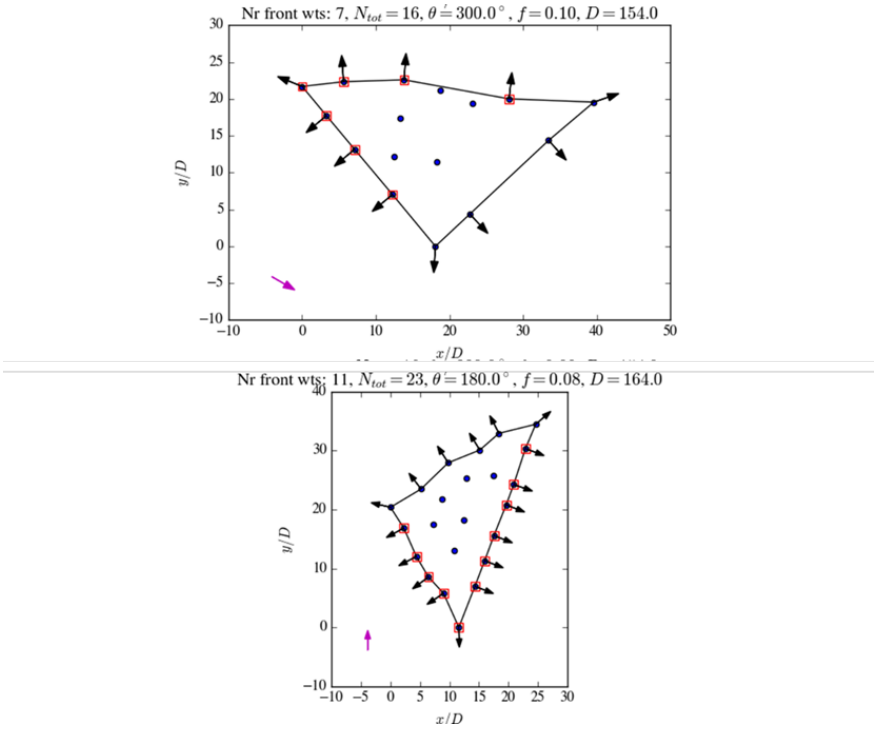

Figure B8. Examples of physically impossible clean-inflow estimates generated by the vdLW fixed-multiplier method. Top: Albatros wind farm, where vdLW detects 7 inlet-edge turbines out of 16 total turbines. Applying the fixed factor of 2.5 gives 17.5 turbines exposed to clean inflow, exceeding the total number of turbines. Bottom: Northwester 2 wind farm, where vdLW detects 11 inlet-edge turbines out of 23 total turbines. Applying the fixed factor gives 27.5 turbines exposed to clean inflow, again exceeding the total number of turbines. These examples show that the vdLW finite-farm correction can produce impossible values, especially for small wind farms. Pink arrows show the wind direction. The examples correspond to isolated-farm cases without neighboring wind farms. The presence of neighboring wind farms can reduce the detected number of edge turbines, partially masking the underlying overestimation introduced by the fixed multiplier.

B2.8 Error 12 — Ignoring the Impact of the Orientation of the Wind-Farm Edge Relative to Wind Direction

The vdLW methodology identifies inlet-edge turbines primarily from geometric boundary exposure.

However, not all geometric edge turbines experience the same degree of direct atmospheric inflow.

A turbine located behind an edge perpendicular to the incoming wind direction experiences substantially more direct inflow than a turbine located behind an edge nearly parallel to the flow.

In practice, the aerodynamic exposure depends on the angle between the incoming wind direction and the local edge orientation.

The vdLW methodology does not explicitly weight the aerodynamic exposure of edge turbines based on edge orientation relative to the incoming flow.

As a consequence, turbines located along weakly exposed or nearly tangential boundaries can be treated similarly to turbines located behind strongly exposed frontal boundaries. A boundary perpendicular to the flow behaves aerodynamically very differently from a boundary nearly aligned with the flow.

This can substantially distort the estimated number of turbines operating under isolated or near-isolated inflow conditions.

The problem becomes particularly important for:

- curved wind-farm boundaries,
- irregular geometries,
- elongated layouts,
- and interacting wind-farm clusters.

The issue directly affects:

- the detected inlet-edge turbines,
- the finite wind-farm correction,
- the inferred wake losses,
- and the predicted capacity factor.

A physically meaningful inlet-edge methodology should account not only for whether a turbine lies on the geometric boundary, but also for how strongly that boundary is exposed to the incoming atmospheric flow.

In many cases, this error partly compensates other errors discussed previously, such as the failure to identify all edge turbines and the unrealistic suppression caused by neighboring wind farms. For many wind farms, the aspect ratio is relatively small, or the long direction of the wind farm is approximately perpendicular to the dominant wind direction.

A particularly relevant exception is the Riffgat wind farm.

Riffgat has an aspect ratio of approximately 3.3, and its long direction is largely aligned with the dominant wind direction. It also does not have strong neighboring wind-farm interactions.

As shown in Figure B9, for a wind direction of $240°$, the vdLW methodology identifies a large number of inlet-edge turbines along the elongated side boundaries of the wind farm.

However, most of these turbines are located along edges that are nearly aligned with the incoming flow and therefore experience substantially weaker direct inflow than turbines located on a frontal boundary.

Despite this, the vdLW methodology treats these turbines as fully exposed edge turbines, and the resulting edge count is then multiplied by the fixed factor of 2.5.

This leads to a substantial overestimation of the number of turbines assumed to operate under isolated or near-isolated inflow conditions.

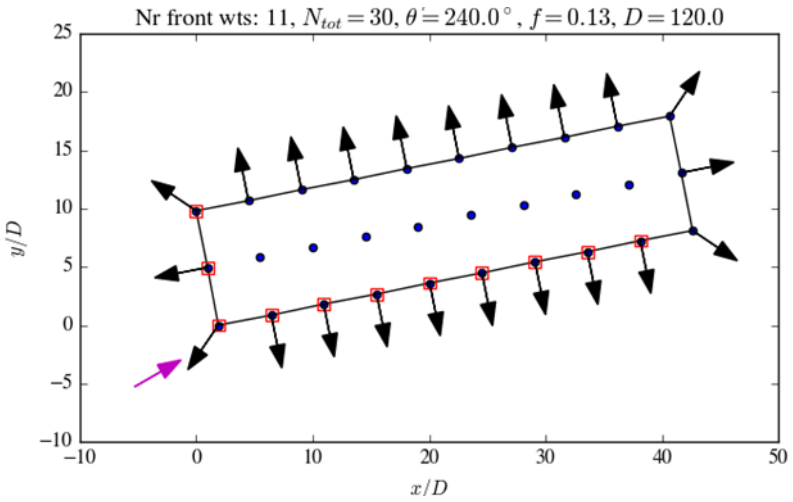

Figure B9. Riffgat wind farm for a wind direction of 240°. The elongated wind-farm layout is largely aligned with the incoming flow direction. The vdLW methodology identifies many turbines along the side boundaries as inlet-edge turbines, even though these boundaries are nearly parallel to the flow and therefore only weakly exposed to direct inflow. Combined with the fixed factor of 2.5, this leads to an overestimation of the number of turbines assumed to operate under isolated or near-isolated inflow conditions. Pink arrows show the wind direction.

Appendix C: Validation examples of the methodology adopted in the work of SLS

As an example, the methodology introduced in Section 2 is applied to determine the theoretical upper limit for offshore wind-energy extraction for three selected offshore wind farms. The wind farms were selected because they are largely free from neighboring-wind-farm shadow effects and because they represent substantially different topologies, sizes, and installed power densities. The selected cases are the Lillgrund wind farm, which is an unusually dense wind farm with a comparatively low capacity factor; the Rødsand II (Nysted 2) wind farm, which represents a medium-density configuration; and the Horns Rev 3 wind farm, which is a low-density wind farm with a very high capacity factor.

C1 The Lillgrund wind farm

The Lillgrund wind farm consists of 48 Siemens 2.3 MW wind turbines located in the Øresund strait between Copenhagen and Malmö, occupying an area of approximately 4.8 km^2. This corresponds to an installed power density of approximately 22 MW/km^2, which is among the highest values reported for an operational offshore wind farm.

The topology of the wind farm is shown in Fig. C1 (left), illustrating both the opening in the center of the wind farm and the irregular outer boundary. Figure C1 (right) shows the results obtained using the proposed technique for estimating the theoretical upper limit of energy extraction. The solid curve represents the theoretical maximum from eq. (1), while the dashed curve includes an additional 9% loss factor.

The blue point corresponds to the theoretical maximum capacity factor of a solitary turbine, whereas the red point represents the asymptotic infinite-wind-farm limit. Due to the extremely high installed power density, the infinite-wind-farm capacity factor is reduced to approximately 0.11. Applying the finite-wind-farm correction methodology yields an equivalent theoretical maximum capacity factor of approximately 0.39 for the finite-sized Lillgrund wind farm (black point), compared with an observed capacity factor of approximately 0.33 (magenta point).

Despite the unusually dense and irregular layout of the wind farm, the agreement between the predicted theoretical limit and the measured operational performance is very good.

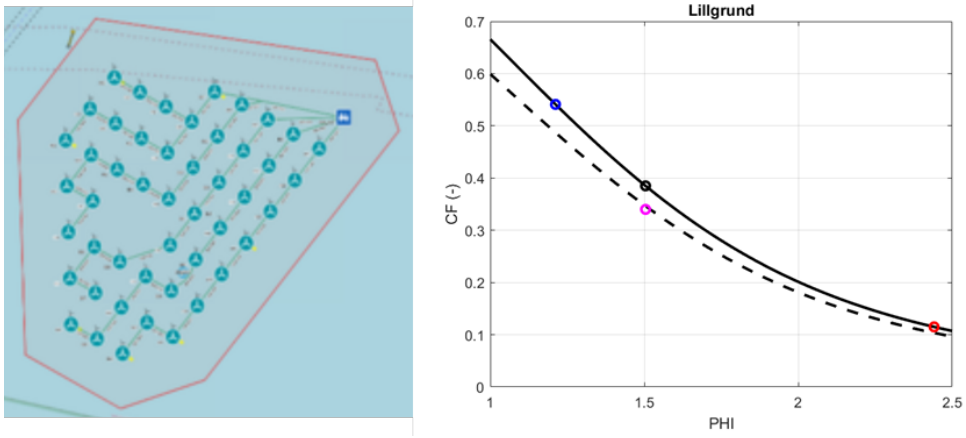

Figure C1. Lillgrund wind farm. Left: Wind-farm layout. Right: Capacity factor as a function of the Wind Farm Wind Factor (Φ). Solid line: theoretical maximum from eq. (1). Dashed line: theoretical maximum including 9% losses. Blue dot: solitary wind turbine. Red dot: infinite wind farm. Black dot: equivalent theoretical limit. Magenta dot: actual measured value.

C2 The Rødsand II (Nysted 2) wind farm

The Rødsand II wind farm consists of 90 Siemens 2.3 MW wind turbines located in the Baltic Sea south of the Danish island of Lolland, covering an area of approximately 34.8 km^2. This corresponds to an installed power density of approximately 5.9 MW/km^2, which is representative of many conventional offshore wind farms.

The topology of the wind farm is shown in Fig. C2 (left), illustrating the characteristic curved geometry of the wind-farm layout. Figure C2 (right) shows the resulting theoretical-limit analysis. Compared with the Lillgrund case, all capacity-factor values are shifted toward higher values due to the substantially lower installed power density.

Combining the solitary-turbine and infinite-wind-farm limits with the a-factor determined using eq. (3) together with the procedure of eq. (2) yields a theoretical maximum capacity factor of approximately 0.50, compared with an actual measured value of approximately 0.43.

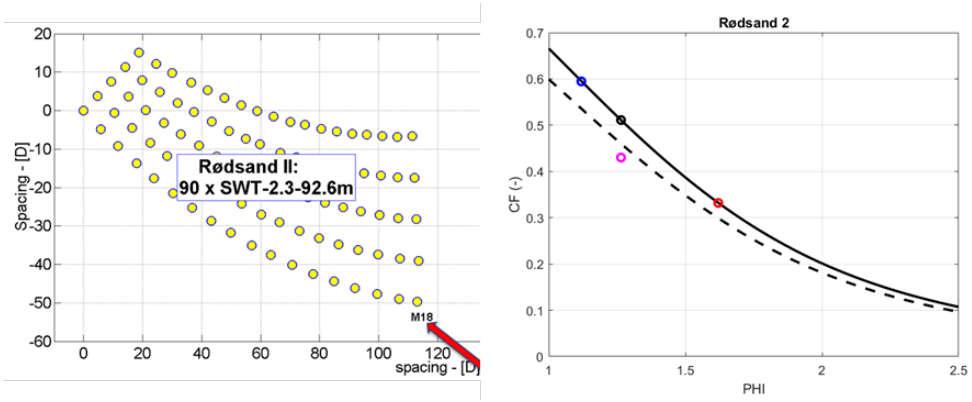

Figure C2. The Rødsand II wind farm. Left: Wind-farm layout. Right: Capacity factor as a function of the Wind Farm Wind Factor (Φ). The symbols are identical to those used in Fig. C1.

C3 The Horns Rev 3 wind farm

The Horns Rev 3 wind farm consists of 49 Vestas 8.3 MW wind turbines located in the North Sea approximately 25–40 km offshore from the west coast of Jutland, Denmark. The wind farm occupies an area of approximately 100 km^2, corresponding to an installed power density of approximately 4.1 MW/km^2, which is a comparatively low value.

The topology of the wind farm is shown in Fig. C3 (left), illustrating the irregular geometry of the layout. Figure C3 (right) presents the resulting theoretical-limit analysis. Due to the low installed power density, both the infinite-wind-farm limit and the corrected finite-wind-farm limit correspond to comparatively high capacity factors.

The equivalent theoretical maximum capacity factor is found to be slightly below 0.60, while the measured operational capacity factor is approximately 0.53.

Despite the three validation cases being substantially different with respect to topology, installed power density, and size, the comparisons between predicted and measured capacity factors show consistently good agreement. These examples therefore provide additional support for the methodology introduced in Section 2 when applied to stand-alone wind farms, i.e. wind farms not strongly affected by neighboring-wind-farm wake interactions.

However, when wind farms are located close to neighboring wind farms or become embedded within large wind-farm clusters, determination of the equivalent finite-wind-farm correction factor becomes substantially more complicated.

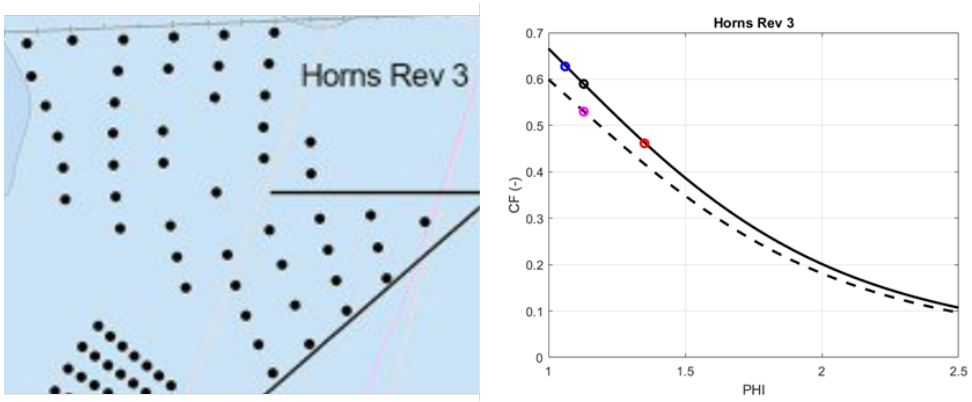

Figure C3. The Horns Rev 3 wind farm (with part of the neighboring Horns Rev 2 wind farm visible in the lower-left corner). Left: Wind-farm layout. Right: Capacity factor as a function of the Wind Farm Wind Factor (Φ). The symbols are identical to those used in Fig. C1.

References

Simão Ferreira, C., Larsen, G. C., and Sørensen, J. N.: A theoretical upper limit for offshore wind energy extraction, Cell Reports Sustainability, 3, https://doi.org/10.1016/j.crsus.2025.100573, 2026.

Sørensen, J., Larsen, G., and Pedersen, M.: Validation of revised minimalistic model for predicting energy production of offshore wind farms, in: accepted to Torque 2026 - The Science of making Torque from Wind conference, 2026.

Sørensen, J. N. and Larsen, G. C.: A Minimalistic Prediction Model to Determine Energy Production and Costs of Offshore Wind Farms, Energies, 14, 448, https://doi.org/10.3390/en14020448, number: 2 Publisher: Multidisciplinary Digital Publishing Institute, 2021.

Sørensen, J. N., Garcia, A. M. I., Larsen, G. C., Pedersen, M. M., and Fournely, D.: Extension and Validation of Minimalistic Prediction Model to Determine the Energy Production of Offshore Wind Farms, Journal of Physics: Conference Series, 2767, 092 022, https://doi.org/10.1088/1742-6596/2767/9/092022, 2024.