Thank you for taking the time to review our manuscript and for your valuable comments, which have helped us enhance the quality of the paper. Below, we include your comments in **black**, followed by our responses in **blue**.

Paper Title:

Wind turbine wake detection and characterisation utilising blade loads and SCADA data: a generalised approach

General summary:

The manuscript presents a 3-stage methodology for turbine-based wind sensing, wake detection and wake characterization. Core of the methodology is a deliberate combination dimensionality reduction and machine learning methods, linking turbine data and wind field information. Both training and testing use aeroelastic simulations coupled with the dynamic wake meandering model and Mann turbulence. The wind sensing shows convincing results, both qualitative and quantitative. Wake detection and characterisation are mainly assessed qualitatively and show mostly good performance. The shortcomings of the methodology and possible improvements are discussed.

General comments:

1. The abstract is good and expressive!

Thank you for your feedback. We are happy to hear you liked it.

2. Regarding the background chapter: The paper should be concise and focus on its main topic. A generic literature review of >6 pages is not appropriate in this context, especially since these topics are not picked up in the discussion section. The target audience can be expected to have a wind energy background. The interested reader will not choose this paper to learn about ABL, momentum theory or wake physics. Bottom line: It's suggested to remove section 2 "Background" and include the literature review with relevance to the paper topic (mainly contained in subsections 2.4 - 2.6) in the introduction section. (Accordingly, no specific comments were done for section 2 at this point).

We agree that the literature review could be more concise. We have removed the Sections 2.1 – 2.3 (original manuscript numeration) entirely, your argument that these concepts are not explicitly used in the rest of the paper is valid. Section describing the numerical wake models (2.4 in the original manuscript) was converted into a short description of the which can be now found in lines 132-135. As per suggestion, Sections

2.5 and 2.6 (original manuscript numeration) are condensed and integrated into the introduction.

3. The manuscript mixes past and present tense (e.g. in sections 3.2, 3.4, 3.5.2, 5.2). Please formulate in present tense where possible.

Agreed, the tense should be consistent across the paper. We have modified for the revised manuscript, all paper is now in present tense.

4. The language could be more concise. There are many instances of "As explained earlier, …", "First of all, …", "With that in mind, …", etc.

We believe that in some places the linking words like the ones you mentioned are helpful, as they allow the reader to connect the points in the narrative. However, considering the length of the paper, we agree that it could be more concise. For the revised manuscript, we have simplified and made it less wordy wherever possible.

5. At many instances in the paper, simulation parameters and numbers of runs are mentioned (training/sensing tests/performance tests). Gathering all that information in one concise table would be helpful.

In the original manuscript, for the sake of brevity, we aimed to summarise this information in the text wherever possible. However, we do agree that this information would be easier to digest in the table format. We believe that adding two tables – one showing the training simulation runs, and one showing testing simulation runs, could be better than one unified one. This is a justified by a different distribution of ambient U and I values in these two cases (e.g. 6 U values for testing, 100 U values for training), and additional column for training (describing the wake impingement case). We also think that the reader shouldn't have to jump several pages to see the summary of the information from specific part of the paper. In the revised manuscript, new tables are numbered 1 and 4, for training and testing, respectively.

6. The manuscript has individual "Results" and "Discussion" sections, which is good. Yet, the results section already contains aspects that would belong in the discussion section (e.g. Line 488-496, 520-521, 541-542, 569-570). On the other hand, the discussion section is very brief and further lacks a comparison to existing methods. Please revise and make sure to have a clear distinction, possibly ending up with a shorter but more concise results section and a more in-depth discussion section.

We acknowledge your feedback here and agree that the in-depth analysis should be rather kept in the Discussion section. We have revised the two sections: in the revised manuscript, Results is now shorter and containing only presentation of results, and Discussion has following subsections: Evaluation methods, Framework performance, Applicability and Current limitations. With regards to your comment about lacking comparison to existing methods, we have answered this in the response to the other reviewer, which we cite below:

“Although we do agree that this additional metric would be informative, we believe that adding a discussion subsection that would satisfy your comment is outside of the paper scope. The main novelty of this work is the introduction of a modular approach to the wake estimation problem – a generalised framework being a combination of several models. The models that we implemented can be easily swapped for something more accurate, and we want to encourage the community to use their methodologies this way.

Furthermore, a direct comparison with other detection/ characterisation studies would be difficult due to fundamentally different assumptions. Other works that consider the wake detection aspect, such as (Onnen et al., 2022) or (Bottasso et al., 2018), do not test their solutions under a full range of wind conditions as we did. Existing simulation-based wake characterisation studies such as (Onnen et al., 2022) use different approaches for establishing the reference.”

7. In section 4.2 and 4.3 the wake detection is mainly assessed qualitatively and visually. The detection ratio (in Fig. 14&15) is the proportion of detected wakes out of all sample slices, but not with respect to a reference. Using the DWM model, the wake positions of the simulations should be available. It is suggested to use this information and show a quantitative performance metric of wake detection. Furthermore, it’s suggested to use the RMSE of the estimated wake position as a metric of the characterization. Additionally, this information could help to unravel the unexpected behaviour of the detection at 5 m/s in Fig. 14.

Thank you for your comment, we agree that quantitative analysis would make this study better. We have spent significant time looking into the version of the DWM model that we’re using, attempting to apply your suggestion, here’s our conclusions:

a) wake detection: when computing the wind field at the receiving turbine, multiple meandering wake deficit profiles are in general imposed on an otherwise clean "Mann box", including additional added turbulence determined by the shape of the deficit profiles. To combine multiple overlapping wake deficit profiles, some rule is applied pointwise, that is, independency at each relevant grid position in the wind field. The rule we used was "pick the maximum decifit at the point", but other rules are possible. As a result, the wake deficit profile that is applied to the Mann box does not necessarily have a simple shape with a well-defined location of maximal deficit. For majority of wind directions where the flow is coming from the inside of the wind farm, there are always some influencing turbines registered by the code – even if they are too far away to have a clear effect on the ‘receiver’ device. As a result of the above, a reference such as ‘waked/not-waked for a given wind direction’ is impossible to establish directly from the simulations. This is partially dictated by the nature of our study - we have tested the detection performance for all wind directions, where the wake shedding turbines are at various upstream distances, hence the rate of wake breakdown/lateral displacement is also varied. Furthermore, to the best of author's knowledge, there isn't a widely-recognised definition of 'wake impingement' (e.g. by means of reduced power output) we could use here. As a work-around to this issue, we have provided a reference for the quantitative analysis in a 'synthetic' way. We've trained a new classifier analogically to the process described in Methodology, with the only difference being that the training dataset is derived from simulated wind fields, not the estimated ones. Without the bias from the wind field reconstruction, this classifier achieves approx. 99% accuracy under integrated testing, and its classifications are thus assumed to be precise enough to become 'ground truth' reference. Such a reference allows to consider the effects of varied wake dispersion under different ambient conditions, and arguably fits this analysis better than a general impingement definition based on inflow angle. All in all, the goal of our wake detector is to identify "clear wake impingement from a nearby device" as we define it in the article. The reference as seen in Fig. 12 (revised manuscript) are sensible, as the impingement ratio decreases with increasing wind speed – just what would happen in real life, as operating in the above rated region makes wake less pronounced. We consider the limitations and bias from this approach in the Discussion section.

b) wake characterisation: we have investigated how to best provide the reference for the estimated wake properties, and we've decided to establish it by fitting a 2D Gaussian function to YZ slices of the raw, simulated wind field. This approach is preferred over using the meandering wake centres applied internally in the DWM model, as turbulent fluctuations in the synthetic wind field — along with additional imposed turbulence – can cause the actual wake experienced by the turbine to deviate from the calculated position. Moreover, this method naturally accommodates the interaction of multiple, combined wakes. A similar approach for getting reference by fitting a Gaussian is used in other studies, like (Lejeune et al., 2022). We consider the limitations and bias from this approach in the Discussion section.

Specific Comments:

Abstract

1. Line 9: "virtual wind farm" – Please state the test environment here. It should be clear from the beginning how the method is tested. Especially since the title does not tell whether it's in field/simulations/wind tunnel.

Good point – we have revised the manuscript and added an appropriate sentence accordingly: "The framework is tested in a simulation environment incorporating the Mann turbulence model, DWM model for generating wakes and BHawC aeroelastic code."

2. Line 15: Partial wake conditions are not necessarily harder to track. The high load-imbalance along the rotor can even make it easier to estimate in comparison to a full wake. Your results, e.g. the findings of Fig. 11, do not seem to support this statement. It's suggested to leave out that sentence.

We agree, line 15 did contradict the results. We have cut out this statement from the sentence, replacing it with 'more turbulent conditions'.

Introduction

3. Line 19: "at the dawn of 2023" – please consider replace by "by end of 2023"

Sorted – we changed it to 'by the end of 2022', which is what the source used explicitly says.

4. Line 35/36: Cyclic pitch control is just one type of "dynamic induction control", which should be mentioned here as the general discipline. It splits up into Pulse and Helix. Frederik et al. focus on Helix. Please adjust and add a more general source.

The authors of a recent comprehensive review article on the subject (Meyers et al., 2022) describe the Pulse as dynamic induction control and Helix as dynamic individual pitch control. These two terms are also differentiated by the authors of the Helix method (Frederik et al., 2020). We believe it is valid to use different terms here: Pulse uses sinusoidally varying thrust coefficient (Munters and Meyers, 2018), which is why dynamic induction control fits perfectly; while Helix varies the tilt and yaw moment at the rotor, hence the induction factor is not directly modified. With that in mind, we have adapted the sentence towards: "(...) some examples include wake steering by introducing yaw offsets (Howland2020, Siemens2019), or dynamic individual pitch control (Frederik2020) and dynamic induction control (Munters2018) to induce enhanced mixing in the wake."

5. Wind farm flow control techniques are mentioned in the introduction, but the connection to the tracking task within the framework of closed-loop control is lacking. The outlook and final role of the wake tracking and characterization should be mentioned.

We do agree that the paragraph could be modified to better emphasise on the key role of wake 'sensing'. We have improved the narrative and added relevant source: "(...) To use the above mentioned approaches in a closed-loop control scheme, dynamic information on whether the impinging wake is being successfully redirected or dispersed is required (Raach2016). Moreover, before starting the control action, there first needs to be a confirmation that a turbine is indeed wake-affected to facilitate an intervention to its normal operating cycle. For the sake of the discussion in this work, we will term these two flow control prerequisites wake detection and wake characterisation. By the former we refer to (...)"

6. Line 49-53: This is not entirely true. The approach of (Bottasso et al., 2018) is used for impingement detection and EKF-based approach in (Onnen et al., 2022) further links wake-presence to the observability. Yet, the 3-stage approach of this manuscript is a novelty and the justification for initial unbiased wind field reconstruction, as mentioned in line 57, is there. Please elaborate in the paragraph and differentiate between the approaches.

We do agree that our statement of research gap lacked accuracy in its original form. The sources you mention consider wake detection, which is why we shouldn't disregard their contribution to the field. Our main novelty is the generalised character of the framework, which combines relatively simple models to achieve unbiased wake estimation performance. These models could be easily replaced with other (more advanced and potentially better performing) solutions, if someone wanted to repeat our methodology. With that in mind, we have modified the entire introduction: a) to include the relevant literature review, which was initially in Background; b) to state the research gap more accurately.

7. Line 55-56: "For this reason, the vast majority of methods developed so far are not yet applicable to real world operations." This is a too strong statement, considering that there are field validations for these other methods, see (Schreiber et al., 2020), (Lio et al., 2021). As said in the previous comment: The research gap exists, but it is not accurately described in this introduction.

As mentioned above, we have significantly modified the Introduction section to address your feedback on the research gap statement.

8. Line 65-67: "It is highlighted that focus of the current work is that of developing a solution which is able to confidently assert when a turbine is impinged by a wake from a nearby turbine, as this information is critical to farm level wake steering control."– Please rephrase this sentence or make it two sentences.

The sentence was a bit convoluted, we have reworded it towards: "This work focuses on developing a novel solution that can confidently assert when a turbine is impinged by a wake from a nearby machine – this being a key factor for farm-wide wake steering control."

9. Line 64-65: "A full end-to-end methodology is presented, which aims to provide both a demonstrator and performance benchmark for generalised wake detection and

characterisation methods of this type." – This would be a good place to mention the test environment (aeroelastic simulations, turbine type, DWM model, ...).

Agreed, added an appropriate sentence describing the overall training/testing setup: "The models are trained and tested for a full range of wind directions within a virtual offshore wind farm. The simulation environment incorporates Mann turbulence wind boxes, Dynamic Wake Meandering (DWM) model for generating wake interactions, and aeroelastic code to compute the turbine response."

Methodology

10. Line 295: why wind farm simulations, when only two turbines are used? Relating to intro: "generalized approach"

Indeed, at the current stage, only two turbines are used during training of the framework. Our focus was on isolating clear wake impingement conditions, so the capabilities of the image-recognition-based wake detection can be assessed. The fact that the models can be trained using only two devices (with data extracted only from the downstream turbine), should be considered an advantage when considering the field application. For example, if the 'true' wind fields for training would be acquired with LIDAR, the additional hardware would only need to be installed on one turbine. The entire wind farm layout with several turbines at once, and with wind coming from all directions, was used during the testing stage. We believe that this criterium justifies the use of wording 'generalised approach'.

11. Line 290-295: Please state the turbine type, diameter and the spacing between "emitting" and "receiving" turbine.

For confidentiality reasons, Siemens Gamesa prohibits us from publishing the details of the virtual wind farm/turbine we are using. As a result, we can't include the information on turbine type or diameter. We could however add the spacing between the turbines, expressed as a multiplication of diameter D. We included that in the revised manuscript in line 118.

12. It is unclear, whether the simulation environment includes just these two turbines or the whole wind farm.

The simulation environment includes the entire wind farm layout. Although we heavily rely on that aspect during testing, for training we purposefully choose turbines from first two rows and select specific wind directions to limit the interactions to just two devices (emitter and receiver). We've done it to clearly differentiate between the effects of wake impingement and standalone atmospheric turbulence; we didn't want to use wakes originating from the inside of the wind farm, as their shape could be distorted due to superposition effects. We agree that this could be potentially confusing to the reader if not explained properly; we have added a necessary explanation to Section 2.2. of the revised manuscript, it can be found in lines 115-118.

13. Line 309: "Figure 6 illustrates the process for training the wind sensing model and demonstrates its post-training performance in producing wind field estimations." The application scheme of the model is shown here, but not post-training performance. Please adjust the sentence.

Agreed, we have revised the sentence as suggested.

14. Eq(1) & Eq(2) (and possibly others): don't use italic font for sine and cosine and subscript text (except variables).

Thank you for spotting this. We have corrected this in the revised manuscript.

15. Line 344ff: This is the first time that higher harmonics are mentioned. The Coleman transform was only described for the 0P / 1P harmonics. Also: Fig. 7 names an "original load", which suggests a (non-transformed) blade load. Meanwhile, it's said in line 344 that the rotor loads are decomposed into their frequency components. Is it correct, that you e.g. calculate the 3P share of a yaw moment? Or do you calculate the 3P of a blade load? A flow chart of the pre-processing steps would help.

Agreed, the original manuscript lacked clarity when describing the turbine response preprocessing steps. We have added a new block diagram (Fig. 5 in revised manuscript) that visualises the process. Also, we have modified the figure describing the frequency decomposition (Fig. 6 in revised manuscript), replacing the confusing 'Original load' with proper variables describing the rotor loads (yaw/ tilt moment etc.).

16. Line 351ff: Please elaborate on the temporal dependency and time lag. Which temporal scales of the turbine dynamics are you addressing here?

Our implementation of lagging is a simplified 'memory' functionality. With a simple implementation, it makes wind sensing less instantaneous and helps to capture the turbine response across several seconds. We have expanded on the paragraph.:

"In order to capture the short-term temporal dependencies and patterns in time series of all wind sensing inputs, the features are embedded with their lagged values. For each sample in a time series, two additional features expressing the past value of the curve are added. These lagged features are obtained by shifting the time series by 4 and 8 seconds from the current time stamp. This effectively makes the estimation more stable and noise-resistant, as the wind slice is reconstructed with turbine response across several seconds. Specific lag values used are determined by testing the framework's performance with a few different configurations and choosing the one that gives the best overall results."

17. Subsection 3.3.3: The DTC is a nice choice and the dimensionality reduction shown in Fig.8 looks appropriate. My only point regarding the wind field parametrization is: The here shown YZ-wind field slices are rectangular, while the rotor swept area is circular. The corners of the wind field thus include non-observable features, but could still influence the lower-order share of the DTC outputs. Were the plain rectangular slices used for the training? Or was any weighting or masking applied? Please comment on whether you expect an impact here.

We've used rectangular slices during the wind sensing training. At the current stage, there was no masking applied. You are absolutely correct that there are consequences of doing so – we have analysed the YZ-wise RMSE of wind sensing in the Results section (see Fig. 11 in revised manuscript). It is clear that the corners have the highest error due to blades being basically unable to clearly sense the flow fluctuations in these regions. For the revised manuscript, we included a discussion on that aspect in the Discussion section (Framework performance – Wind sensing subsection). We will consider applying masks in future work.

18. Section 3.3.4: Please add some more details and a literature source to the used regression approach.

To address your comment, we have expanded the description of our linear regression implementation (paragraph starting in line 215 of revised manuscript). Moreover, we have added a source describing the localised linear regression in more detail (Cleveland, 1988).

19. Line 392ff: Is the distinction of the four classes based on the constellation of the simulation run or based on the instantaneous wake position (which could differ due to the employed DWM model). Also: Please define the overlap margins, from which you categorize full/partial/no wake impingement.

The classes are defined based on the simulation setup – specific wind direction refers to specific class. The wind direction differs by 5 degrees between the fully and partially impinged cases, we have added this information to the revised manuscript. When the wake deficit falls between the CNN's understanding of partial and full impingement, the classification is less certain (e.g. 60% full, 40% partial). It is however not an issue for the overall performance, as shown in the specific wind fields we have analysed sample by sample. When the wake meanders from the centre to the side, the CNN changes its output appropriately (from full to partial impingement), and the two-dimensional Gaussian gets fitted nonetheless.

We are not entirely sure what you mean by 'overlap margins'. If you refer to quantifying the degree of misclassifications, we have added an appropriate confusion chart (Fig. 9 in revised manuscript) that shows which classes are the main source of error.

20. Section 3.5.1: The fitting function does not fully reflect the fit that was probably implemented. To fit a wind field with wake deficit, it should be U = u_ambient – f_G, here considering that parameter A is negative.

Agreed, the function from the original manuscript did not fully represent the fitting procedure. We have added an equation (Eq. 5 in revised manuscript) for the reversed wake deficit, which the bivariate Gaussian is actually fitted on.

21. Section 3.5.2: How does the low-pass filter deal with falsely identified no-impingement instances?

In current implementation, the moving average filtering is performed before the removal of samples identified as 'no wake impingement' from the characterised wake time series. This removal is implemented as assignment of NaN values to specific time stamps. This solution isn't perfect, as there are several gaps in time stamps where the wake can be assumed to be present. We have potential ideas on how this problem could be solved (Kalman filtering), which we discuss in the Discussion chapter. To make the reader aware of how the current implementation works, we have added a subsection 'Treatment of non-impinged samples' (indexed as 2.5.3 in revised manuscript):

To satisfy the assumption that wake characterisation should not be performed for wind samples without clear impingement, the filtered wake properties are treated for a given time stamp i as follows:

$$y_c^{\text{filt}}(i), \sigma_y^{\text{filt}}(i) = \begin{cases} y_c^{\text{filt}}(i), \sigma_y^{\text{filt}}(i), & \text{if class}(i) = \text{fully/partially impinged} \\ \text{NaN}, & \text{if class}(i) = \text{no detectable impingement} \end{cases} \quad (9)$$

22. Line 456: Please rephrase this sentence.

We have rephrased the paragraph accordingly.

23. Fig. 11: Please add more details to the figure caption.

We have expanded the description: "Typical YZ distribution of RMSE in wind sensing (normalised by the corresponding Uamb value). Rotor outline marked with dotted white line."

24. Section 4.1: a diff-plot between estimated and simulated wind field would help to analyse, whether systematic or just random differences exist (e.g. the central deficit mentioned in line 500)

Fig. 11 in the revised the manuscript provides the information on wind sensing accuracy across the YZ plane. We believe that a diff-plot would be in this case redundant.

25. Line 501-503: "A slight disagreement would not be normally problematic; however, in this case where the flow has little overall turbulence, a subtle deficit like this can't 'hide' among other eddies, which could potentially result in classifying the

mentioned samples as being wake impinged.” This sounds rather complicated and nested. Please rephrase the sentence.

We have reworded the entire paragraph due to modifying the Results section.

26. Fig. 14 shows that simulations for all wind directions are on hand. Please report this more explicitly in section 3.6.

We have reworded the section appropriately, now the testing setup is described more explicitly.

27. Fig. 14 @ 5 m/s: why is detection ratio different here? In partial load range, the non-dimensionalized wake should be similar, thus limited impact on the sensing is expected. Also: It would be good to know the turbine’s cut-in wind speed. At 5 m/s undisturbed ambient wind speed, a wake-exposed turbine might experience a rotor-effective wind speed below cut-in.

There is a wind sensing anomaly that introduced a ‘fake’ wake to 5 m/s cases. We have now added additional metrics that quantitatively measure how often this anomaly occurs (confusion matrix in Fig. 9 of revised manuscript). The new Discussion section addresses this issue in more detail.

Unfortunately, due to SGRE’s data protection, we are unable to include turbine details such as cut-in wind speed. We have nonetheless addressed this aspect in the Discussion/Applicability subsection, where we comment on the implications of poor performance in 5 m/s region.

Discussion

29. Line 591ff: “. Implementing other approaches such as Long short-term memory (LSTM) networks could potentially allow for the forecast to be extended to predict wake locations and meandering behaviour a few minutes ahead. Further research needs to be conducted to investigate these leads.” Please provide a source here and a stronger supporting argument for the claim that the (stochastic)meandering wake location can be predicted by the receiving turbine. Otherwise, please consider softening this statement.

We may have indeed exaggerated the potential of using LSTMs here. We were unable to find proper sources supporting our statement of ‘a few minutes ahead’; the literature generally suggests shorter time scales of forecast. With that in mind, we’ve softened the discussion and added sources:

“ Relevant literature (Luo2024, Zhou2023) shows that Long Short-Term Memory (LSTM) networks could potentially provide a short-term forecast of the wake dynamics, thus providing an alternative solution. Further research needs to be conducted to investigate these leads.”

References

30. Line 700: incomplete reference (journal, DOI)

Sorted.

References

Bottasso, C. L., Cacciola, S., & Schreiber, J. (2018). Local wind speed estimation, with application to wake impingement detection. Renewable Energy, 116, 155–168. https://doi.org/10.1016/j.renene.2017.09.044

Lio, W. H., Larsen, G. C., & Thorsen, G. R. (2021). Dynamic wake tracking using a cost-effective LiDAR and Kalman filtering: Design, simulation and full-scale validation. Renewable Energy,172, 1073–1086. https://doi.org/10.1016/j.renene.2021.03.081

Onnen, D., Larsen, G. C., Lio, W. H., Liew, J. Y., Kühn, M., & Petrović, V. (2022). Dynamic wake tracking based on wind turbine rotor loads and Kalman filtering. Journal of Physics: Conference Series, 2265(2), 022024. https://doi.org/10.1088/1742-6596/2265/2/022024

Petrovic, V., Jelavic, M., & Baotic, M. (2014). Reduction of wind turbine structural loads caused by rotor asymmetries. 2014 European Control Conference, ECC 2014, 1951–1956. https://doi.org/10.1109/ECC.2014.6862484

Schreiber, J., Bottasso, C. L., & Bertelè, M. (2020). Field testing of a local wind inflow estimator and wake detector. Wind Energy Science, 5(3), 867–884. https://doi.org/10.5194/wes-5-867-2020

Meyers, J., Bottasso, C., Dykes, K., Fleming, P., Gebraad, P., Giebel, G., Gocmen, T., and van Wingerden, J.-W. (2022). Wind farm flow control: prospects and challenges. Wind Energy Science, 7, 2271–2306. https://doi.org/10.5194/wes-7-2271-2022

Munters, W. and Meyers, J. (2018). Towards practical dynamic induction control of wind farms: analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines (2018). Wind Energy Science, 3, 409–425. https://doi.org/10.5194/wes-3-409-2018

Frederik, J. A., Doekemeijer, B. M., Mulders, S. P., and van Wingerden, J.-W. (2020). The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms. Wind Energy, 23, 1739–1751. https://doi.org/10.1002/we.2513

Cleveland, W. S., Devlin, S. J., and Grosse, E. (1988). Regression by local fitting. Journal of Econometrics, 37, 87–114. https://doi.org/10.1016/0304-4076(88)90077-2

Luo, Z., Wang, L., Fu, Y., Xu, J., Yuan, J., and Tan, A. C. (2024). Wind turbine dynamic wake flow estimation (DWFE) from sparse data via reduced-order modeling-based machine learning approach. Renewable Energy, 237, 121 552, https://doi.org/10.1016/j.renene.2024.121552

Zhou, L., Wen, J., Wang, Z., Deng, P., and Zhang, H. (2023). High-fidelity wind turbine wake velocity prediction by surrogate model based on d-POD and LSTM. Energy, 275, 127 525, https://doi.org/10.1016/j.energy.2023.127525