

Adaptation of Engineering Wake Models using Gaussian Process Regression and High-Fidelity Simulation Data

Andersson, Leif Erik; Doekemeijer, Bart; Van Der Hoek, Daan; Van Wingerden, Jan Willem; Imsland, Lars

DOI

[10.1088/1742-6596/1618/2/022043](https://doi.org/10.1088/1742-6596/1618/2/022043)

Publication date

2020

Document Version

Final published version

Published in

Journal of Physics: Conference Series

Citation (APA)

Andersson, L. E., Doekemeijer, B., Van Der Hoek, D., Van Wingerden, J. W., & Imsland, L. (2020). Adaptation of Engineering Wake Models using Gaussian Process Regression and High-Fidelity Simulation Data. *Journal of Physics: Conference Series*, 1618(2), Article 022043. <https://doi.org/10.1088/1742-6596/1618/2/022043>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

PAPER • OPEN ACCESS

Adaptation of Engineering Wake Models using Gaussian Process Regression and High-Fidelity Simulation Data

To cite this article: Leif Erik Andersson *et al* 2020 *J. Phys.: Conf. Ser.* **1618** 022043

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Adaptation of Engineering Wake Models using Gaussian Process Regression and High-Fidelity Simulation Data

Leif Erik Andersson¹, Bart Doekemeijer², Daan van der Hoek²,
Jan-Willem van Wingerden², Lars Imsland¹

¹Department of Engineering Cybernetics, Norwegian University of Science and Technology, 7491 Trondheim, Norway.

²Delft University of Technology, Delft Center for Systems and Control, 2628 CD Delft, The Netherlands.

E-mail: {leif.e.andersson, lars.imsland}@ntnu.no, {b.m.doekemeijer, d.c.vanderhoek, j.w.vanwingerden}@tudelft.nl

Abstract. This article investigates the optimization of yaw control inputs of a nine-turbine wind farm. The wind farm is simulated using the high-fidelity simulator SOWFA. The optimization is performed with a modifier adaptation scheme based on Gaussian processes. Modifier adaptation corrects for the mismatch between plant and model and helps to converge to the actual plan optimum. In the case study the modifier adaptation approach is compared with the Bayesian optimization approach. Moreover, the use of two different covariance functions in the Gaussian process regression is discussed. Practical recommendations concerning the data preparation and application of the approach are given. It is shown that both the modifier adaptation and the Bayesian optimization approach can improve the power production with overall smaller yaw misalignments in comparison to the Gaussian wake model.

1. Introduction

Wind energy is a core part of the effort to reach a carbon neutral society in Europe and worldwide. Nowadays about 5% of the global energy production is supplied by wind [1]. However, in the coming years the deployed capacity will increase rapidly in order to meet the high renewable energy targets. The European Commission plans to install 450 GW of offshore wind by 2050 [2]. Most of the wind turbines will be grouped together and installed near each other to reduce maintenance and deployment costs [3]. However, when a wind turbine extracts energy from the wind a *wake* behind the turbine develops. The wake is characterized by a reduced flow velocity and higher turbulence intensity. A turbine operating in a wake of another turbine extracts less energy and experiences higher load variations [4]. Clearly coordinated wind farm control strategies that dampen the interaction between turbines have the potential to reduce the levelized cost of energy [5] and make wind energy even more competitive.

Currently a promising approach to improve the power production with wind farm control is wake steering. A yaw offset is applied to the upstream turbine misaligning the rotor plane with the inflow and deflecting the wake away from the downwind turbine. Significant power gains can be achieved downstream resulting in overall power gains in the wind farm. Moreover, field



experiments to prove the effect of wake steering have shown promising results [6–8].

Model-based wind farm control usually relies on simplified surrogate models to estimate the wake interactions. These models often calculate the steady-state situation and can represent the general behavior of wakes [9–11]. Nonetheless, the complex dynamics of the air flow and the large range of spatial and temporal scales make the development of an accurate surrogate model challenging [12]. Improving and developing new surrogate models is a very active research field [13–16]. Model-free optimization methods were investigated to drive the wind farm iteratively to a plant optimum without the requirement of a possible inaccurate model [17, 18]. However, these methods suffer from long convergence times.

Modifier-adaptation (MA) is a real-time optimization (RTO) strategy that uses measurement data to correct the cost and constraint functions of the optimization problem directly, and reaches, under suitable assumptions, true plant optimality upon convergence [19]. Consequently, MA can correct the inaccurate surrogate model using plant measurements and provides a viable closed-loop model-based wind farm control solution.

Recently, a method was proposed that combines MA with Gaussian process (GP) regression for wind farm optimization [20]. The GP regression is used to correct for the plant-model mismatch. In [21] the approach was improved. Instead of identifying the plant-model mismatch of the power production of the wind farm, the plant-model mismatch of the power production of each single turbine was identified with multiple GP regression models. The multiple-input and multiple-output approach needs less data to achieve the same performance and scales better for large wind farms compared to the multiple-input and single-output approach.

In this article, MA and GP regression (MA-GP) are used to optimize the power production in a wind farm using data generated with the high-fidelity simulator SOWFA [22]. It is shown that the approach can handle noisy and dynamic data. Moreover, the approach achieves satisfactory result with a small data set.

The article is organized as follows: In section 2 the methodologies used in the article are introduced. It follows a detailed description of the Case study in section 3. The article ends with a conclusion.

2. Methodology

This section gives a brief description of the MA-GP approach. A more detailed description of the MA-GP approach for wind farms can be found in [21]. First, the MA approach and objective are introduced in Section 2.1, then the GP regression is described in 2.2 and finally the MA-GP approach is presented in Section 2.3.

2.1. Modifier Adaptation

The objective of the modifier adaptation approach is to find the optimal steady-state operation point of the plant. The challenge in real-time optimization is that usually an exact input-output map of the plant is unknown and instead a model approximation is exploited in the optimization. Inevitably the plant-model mismatch introduced by the model can lead to sub-optimal performance of the real-time optimization. MA improves the optimization by correcting the objective and constraint functions of a real-time optimization problem using plant measurements.

In context of the wind farm optimization, the model consists of a wake and turbine model. The objective is to maximize the power production and the optimization variables are the turbine yaw angles.

A computational bottleneck of the MA approach is the calculation of the plant gradients, which are required in the correction of the objective and constrained functions. The use of GPs was proposed by [23] to overcome this bottleneck.

2.2. Gaussian process regression

GPs are based on kernel methods [24]. They aim to describe an unknown function $f : \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ from data. It is assumed that the noisy observation of $f(\cdot)$ are given by

$$y = f(\mathbf{u}) + \nu, \quad (1)$$

where the noise ν is Gaussian with zeros mean and variance σ_ν^2 and \mathbf{u} is the input, which is assumed to follow a multivariate Gaussian distribution. A covariance function and mean function determine the smoothness and continuity properties of the underlying function [25]. In this article a zero mean function and the automatic relevance (ARD) squared-exponential (SE) covariance function are chosen as a default choice:

$$k(\mathbf{u}_i, \mathbf{u}_j) = \sigma_f^2 \exp\left(-\frac{1}{2} \boldsymbol{\tau}_{ij}^T \boldsymbol{\Lambda}^{-1} \boldsymbol{\tau}_{ij}\right), \quad (2)$$

where σ_f^2 is the covariance magnitude, $\boldsymbol{\Lambda} = \text{diag}(\lambda_1^2, \dots, \lambda_{n_u}^2)$ is a scaling matrix and $\boldsymbol{\tau}_{ij} = \mathbf{u}_i - \mathbf{u}_j$. Alternatively, the Matérn $5/2$ covariance function (M $^{5/2}$) is chosen,

$$k(\mathbf{u}_i, \mathbf{u}_j) = \sigma_f^2 \left(1 + \sqrt{5 \boldsymbol{\tau}_{ij}^T \boldsymbol{\Lambda}^{-1} \boldsymbol{\tau}_{ij}} + \frac{5}{3} \boldsymbol{\tau}_{ij}^T \boldsymbol{\Lambda}^{-1} \boldsymbol{\tau}_{ij}\right) \exp\left(-\sqrt{5 \boldsymbol{\tau}_{ij}^T \boldsymbol{\Lambda}^{-1} \boldsymbol{\tau}_{ij}}\right), \quad (3)$$

which is less smooth and better suited for practical optimization problems [26].

The hyperparameters $\boldsymbol{\psi} := [\sigma_f, \sigma_\nu, \lambda_1, \dots, \lambda_{n_u}]^T$ are inferred from data maximizing the log marginal likelihood. The output of a prediction with a GP is a mean and variance value. Consequently, it is possible to evaluate the uncertainty of an operating point and even include this information in the optimization, which is a desired property.

2.3. Modifier Adaptation with Gaussian processes

The idea is to identify the difference between plant and model with a GP using plant measurements.

In context of this article the GP regression model creates an input-output map of the control inputs to the plant-model mismatch of the power production for each wind turbine. In fact, GP regression models are used to correct the power outputs of the approximate model. The resulting MA-GP model is used in the optimization to compute optimal yaw inputs for each wind turbine. The inputs and the difference between the measured and estimated outputs of plant and model, respectively, are used to update the data set and the hyperparameters of each GP regression model. The power measurements of the plant are filtered with a moving average filter (Fig. 1). Filtering is necessary since the steady-state behavior of the plant is represented by the approximate model.

2.4. Bayesian Optimization

The MA-GP approach is compared with a GP regression approach that identifies the plant without additional model (Bayesian optimisation (BO) approach). The only difference between the MA-GP and the BO approaches is the inclusion of the *Wind farm model* in the MA-GP approach. Consequently, the BO approach identifies the power production of each turbine directly instead of the plant-model mismatch.

¹ The wind farm picture is by Erik Wilde from Berkeley, CA, USA <https://www.flickr.com/photos/dret/24110028330/>, *Wind turbines in southern California 2016*, <https://creativecommons.org/licenses/by-sa/2.0/legalcode>

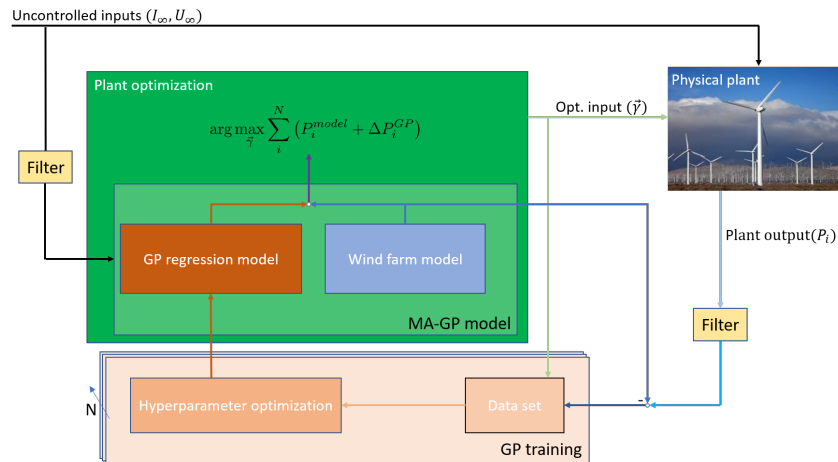


Figure 1: The basic idea of the MA-GP scheme for a wind farm ¹.

3. Case study

The proposed MA-GP approach is tested in a high-fidelity simulation for a wind farm with nine DTU 10 MW turbines (Fig. 2). It is the same setup as presented in [12]. The simulation settings are displayed in Tab. 1.




Figure 2: Top view of a snap shot on the wind farm simulated with SOWFA [12]. The dotted boxes indicate the three subsets in the wind farm.

Table 1: Simulation settings for SOWFA.

| Variable | Value |
|---|-------------------------------------|
| Timestep | 0.2 s |
| Cell size (near rotor) | 2.5 m × 2.5 m × 2.5 m |
| Cell size (outer region) | 10.0 m × 10.0 m × 10.0 m |
| Blade epsilon, ϵ | 5.0 m |
| Turbine layout | 3 × 3 turbines at 5 D × 3 D spacing |
| Freestream wind speed, U_∞ | 8.0 m/s |
| Freestream turbulence intensity, I_∞ | 5.0 % |
| Freestream wind direction, ϕ | 0.0 rad |

3.1. Data preparation

The generator power at each turbine is measured. A moving average filter with a time horizon of five minutes is used to filter the measurements. Each operating point (OP) is simulated for 10 min. The transient of the step response is excluded by removing the first five minutes after a change of an OP from the data. Consequently, for each OP one power measurement was included into the training sets of the GPs. The initial training set consisted of 20 OPs. However, several OPs were similar or even the same such that effectively only 15 different OPs were available. Nonetheless, equal and similar operating points can be used to estimate the noise of the process, so they contain some information. The lower and upper boundaries of the yaw misalignment of the three turbine rows are $\gamma_l = [-25^\circ, -20^\circ, -7^\circ]^T$ and $\gamma_u = [18^\circ, 10^\circ, 5^\circ]^T$, respectively. A larger variation for the upstream turbines was allowed since it is expected that the optimal OP requires larger yaw misalignment of the upwind turbines.

3.2. Surrogate model

The wake is approximated by the Gaussian wake model [11] and the wind turbines by the actuator disc theory [27]. The model adaptation method proposed by [12] is applied to improve the power estimates of the surrogate model. The ambient wind velocity and turbulence intensity are estimated by minimizing the difference between the time averaged power measurements of SOWFA and the surrogate model.

3.3. Results - Upstream turbines

It is difficult to visualize the objective function of a plant with nine control inputs. Therefore, the model identification of the upstream turbines, using MA-GP and BO, is investigated and compared (Fig. 3). The power production of turbine 1 with the MA-GP approach is well captured (Fig. 3a). It aligns with the expected values from SOWFA [12]. The BO approach (Fig. 3c) gives similar predictions in the vicinity of the measurement points. At the boundaries (large absolute yaw angles) the mean values of the power production flatten. For even larger yaw angles the power production even increases again. The method converges to about the mean value of the data points and the variance increases considerably. For the MA-GP approach this behavior cannot be observed. For the large yaw angles a constant bias is corrected and the variance converges quickly to a constant value.

For turbine 3 the uncertainty in the prediction of the power production is larger. For the MA-GP approach (Fig. 3e) an almost constant bias correction independent of the yaw angle is applied to the turbine. For the BO approach (Fig. 3g) the power production is almost independent of the yaw angle and rather flat. It follows the mean of the data set. The reason is the large noise of the measured power production. The power production of the yawed case is similar and sometimes even larger than for the non-yawed case. Consequently, the GP regression identifies a large measurement noise. The variation in the data is then mainly explained by the measurement noise without strong variation in the estimate of the mean value.

The figures in the right column (Fig. 3b - 3h) show the results with some more data points. The turbines are operated with a yaw angle between 20° to 30° . The MA-GP approach extrapolated quite well. Only the power production at turbine 1 is slightly corrected downward. The variance estimate at turbine 3 is corrected. The model uncertainty is increased while the measurement noise estimate is reduced. The previously extrapolated results of the BO approach for large yaw angles is corrected significantly.

3.4. Practical recommendations

Overfitting can be an issue if only a few data points are used in the identification, cf. Fig. 4. Here, the measurement noise σ_v and the length scales $\lambda_1, \dots, \lambda_{n_u}$ are estimated to be very small. Therefore, the data points are fitted almost perfectly by the GP giving the mean estimate a not

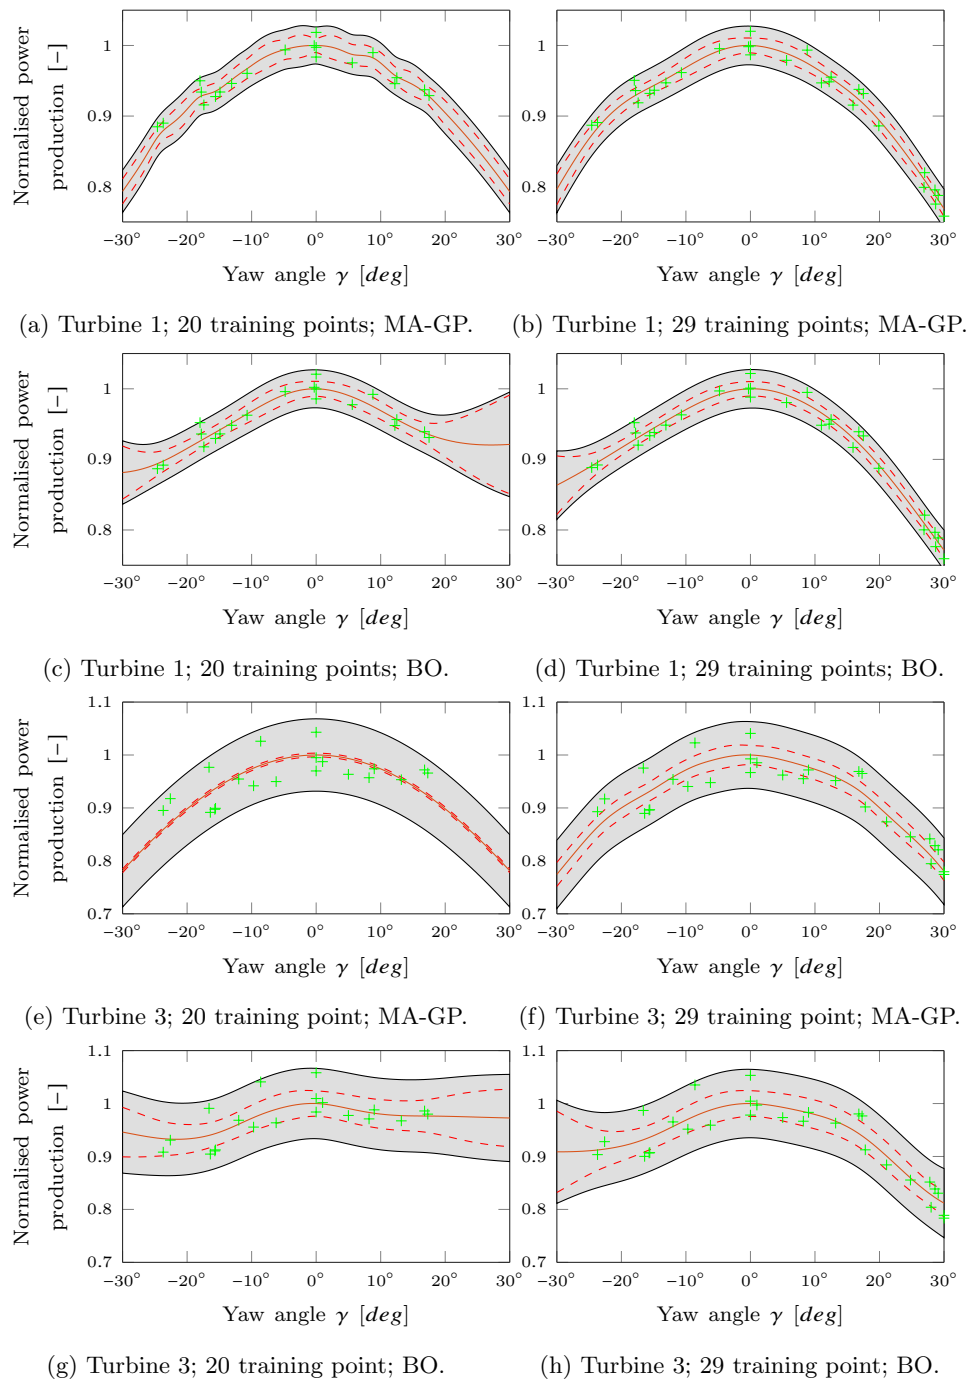


Figure 3: The normalized power production of the turbine 1 and turbine 3 identified with the MA-GP and BO approaches. The orange line represents the mean value, the red dashed line indicates the 95% certainty range of the model and the grey region indicates the 95% certainty range of model plus measurement noise. The green crosses represent (filtered) measurement points.

smooth behavior. Including some similar operating points (with different outputs), which help to estimate the measurement noise, can prevent this phenomenon. Another solution is to include regularisation in the hyperparameter estimate preventing large hyperparameter values.

As expected, the $M^{5/2}$ covariance function is less smooth and certain about the estimates than




Figure 4: GP regression with overfitting.

the SE covariance function. With few data points the M^{5/2} covariance function is more prone to overfitting. Otherwise, both covariance functions show similar results in estimating the power production of the first turbine row.

3.5. Results - Optimization of plant power production

The results of the optimization over several iterations will be discussed. The algorithms are initialized with 20 data points. The optimal operating points are calculated using the MA-GP and BO approach. Subsequently, the optimal operating point is applied in the SOWFA simulation to get a new power measurement.

Comparing the two covariances and the MA-GP and BO algorithms would have led to many time-consuming SOWFA simulations. Therefore, only for two methods the optimized yaw angles were sent to SOWFA. For the other approaches a *quasi-SOWFA* plant model was created. A three-turbine wind farm was simulated with a GP model including all available data. This was achieved by scaling the power production of each SOWFA run with the power production of the greedy setup and the same wind input. It removed the variance in the power production. Afterwards the turbine rows can be separated, and the power production of a three-turbine wind farm can be optimized or predicted using a larger dataset. With this trick we were able to represent *quasi-SOWFA* predictions. The power predictions in the resulting three-turbine plant is very accurate for each turbine row of the nine-turbine wind farm.

The results of ten iterations using the different algorithms and the quasi-SOWFA predictions as a plant model are shown in Fig. 5. After ten iterations the power production increase ranges between 21.8% to 22.5%.

The optimal yaw misalignment is about $\vec{\gamma} = [30^\circ, 12^\circ, 0^\circ]^T$ for each of the rows and a power gain of about 24% (in actual SOWFA runs) can be achieved. This optimum was found with the scaled three turbine plant. The yaw misalignment of the second turbine varies between 11° to 26° for both the BO and MA-GP approach, which indicates that the total power production of the plant might be less sensitive to the yaw misalignment of the first downwind turbine. The power production at the optimum is more distributed over all turbines. The second and third turbine produce almost the same power.

At the initialization the MA-GP approach is superior to the BO approach. After ten iterations the differences between the three best methods is within the expected variance of the power production estimate. In addition, the absolute yaw misalignment is similar for both approaches. Including an exploration term in the objective function may help to converge closer to the optimum. In comparison to the optimization using solely the wind plant model the BO and MA-GP approaches increase the power production about 2% to 3% with overall smaller yaw misalignments. Even at the initialization the power production is about 1.5% to 2.5% higher.

Even though both approaches have a similar performance it is recommended to use the MA-GP approach. Without constraints, the BO approach tends to very large yaw angles in the first

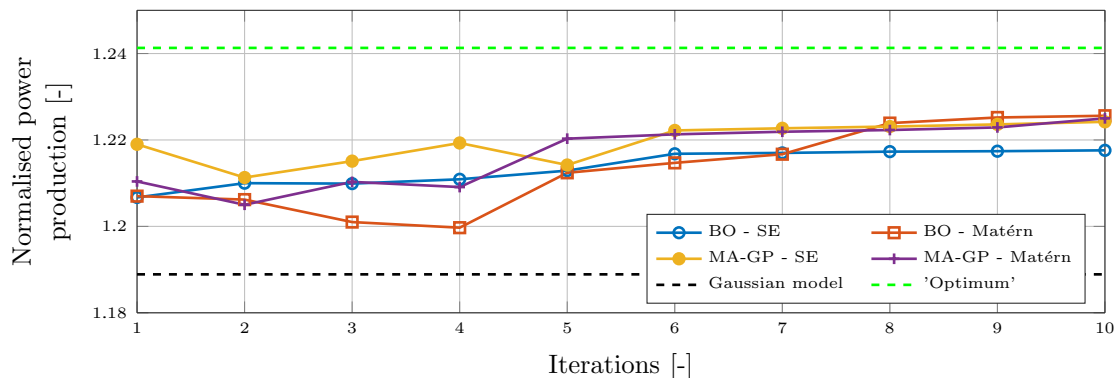


Figure 5: Operating point normalized by the greedy power production over the ten iterations. The lower dashed line shows the optimal power production of the Gaussian wake model. The upper dashed line shows the optimal power production of the scaled three-turbine plant. It is the best power production reached in the SOWFA simulations.

turbine row because of the behavior described in Sec. 3.3. Moreover, a similar behavior can be observed for the last turbine row, which tend also sometimes to large yaw angles. These issues cannot be observed in the MA-GP approach. Solutions to prevent this problem are:

- Filtering of the optimization results:

$$\vec{\gamma}_{k+1} = \vec{\gamma}_k + \kappa(\vec{\gamma}_{k+1}^* - \vec{\gamma}_k), \quad (4)$$

where $\vec{\gamma}^*$ are the optimal yaw angles and κ a filter variable.

- Constrain the optimization variables so extrapolation is avoided,
- Include the variance in the objective function

$$\vec{\gamma}_{k+1}^* = \arg \max_{\vec{\gamma}} (\boldsymbol{\mu}^2 - \tau \boldsymbol{\Sigma}), \quad (5)$$

where $\boldsymbol{\mu}$ represents the power production of the wind farm, $\boldsymbol{\Sigma}$ the variance in the power production estimate and τ a weighting factor.

- Use the trust-region approach proposed by [28]

If the applied constraints allow only a small feasible region they should be relaxed if the algorithm converged to a boundary. Furthermore, the penalization of the variance should also be relaxed once the data set grows to allow exploration. Otherwise, it may happen that the algorithm converges to the most certain point. On the other hand, the variance only increases strongly for the BO approach if the approach extrapolates. Therefore, it can be expected that the variance term has not a strong influence on the objective function if weighted well.

3.5.1. Practical recommendations regarding the filter horizon Previously a five-minute horizon was used to filter the power outputs of the turbines. A shorter filter horizon increases the variance of the data. We like to present some preliminary observation if a filter horizon of 2.5 min is used, but the data in the GP doubled. We observe a tendency to overfitting in the GP regressions. Moreover, in the BO approach the power production becomes less sensitive to yaw angle changes (Fig. 6b).

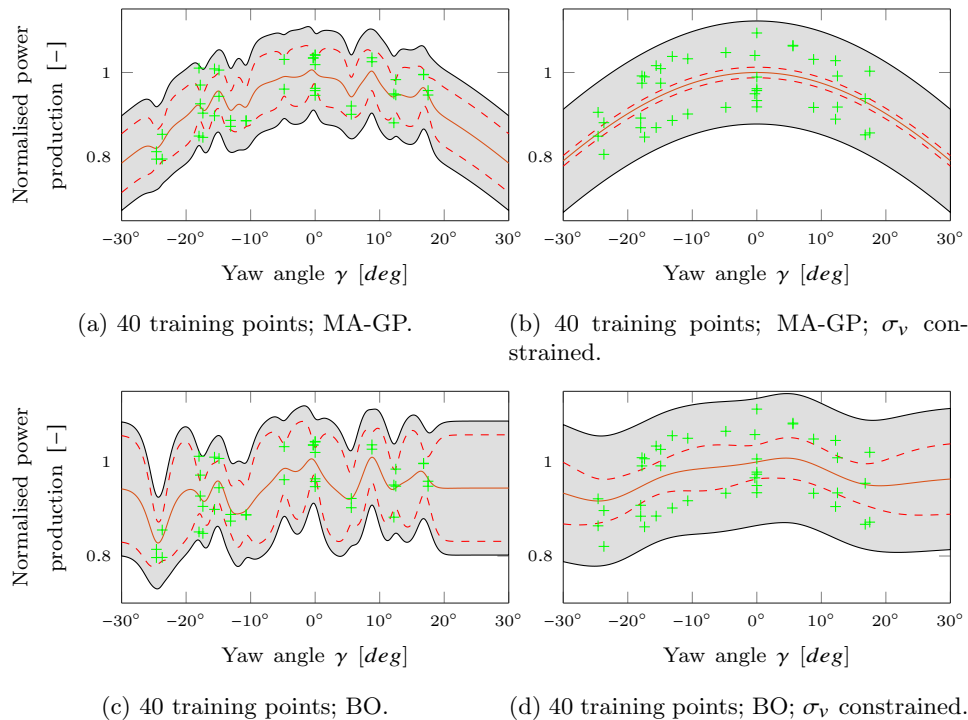


Figure 6: The normalized power production of turbine 1 and turbine 3 identified with the MA-GP and BO approaches using a filter horizon of 2.5 min. The left column shows the results with unconstrained hyperparameters. The right column shows the result with constrained hyperparameters.

Again, the $M^{5/2}$ covariance function is more prone to overfitting than the SE covariance. Overfitting can be reduced constraining the hyperparameters or simply by a larger dataset with different operating points. The optimization of the power production using the initial data set of 40 data points is negatively affected by the shorter filter horizon for both the BO and MA-GP approaches. Both approaches show similar performance. Interestingly, the SE covariance behaves superior to the $M^{5/2}$ covariance function. It is probably a direct consequence of the smoother property of the SE covariance. If the noise term σ_v is lower bounded the performance of the MA-GP approach increases. For the BO approach the results are inconclusive. Both tendencies were observed. It can be observed that smaller yaw misalignments are used, which indicates that constraining the σ_v might lead to a more robust approach in case of very noisy data. A more in-depth study using Monte-Carlo simulations must be performed to draw more conclusive results.

4. Discussion & Conclusion

In this article the MA-GP approach proposed by [20] is applied to a nine-turbine wind farm simulated with SOWFA. Alternatively, to the MA-GP approach the BO approach is tested, which uses no a priori plant model. For both approaches the multiple-input multiple output identification scheme proposed by [21] is applied. Even though a relatively small data set is used as initial training set, which contained a large amount of sub-optimal counterclockwise yaw settings, the algorithms can converge quickly to a near optimal solution. For the reported power gains, it must be kept in mind that the inflow in SOWFA is quasi-static, the free-stream turbulence intensity is relatively small and the plant alignment with the inflow is favorable for wake steering.

The MA-GP and BO approaches can improve the power production compared to the Gaussian wake model. Moreover, the expected better performance for clockwise yaw misalignment is identified and the surrogate model improved. Based on the current results a clear recommendation for either the MA-GP or the BO approach cannot be given. Both approaches showed similar performance. The MA-GP approach might be better suited for small data sets and large wind farms because of its better extrapolation property. Even though both covariance functions show similar performance the smoothness of the SE covariance seems to be a favorable feature making it more user-friendly.

In future work the recommendations and observations pointed out in the article should be verified in Monto-Carlo simulations. In such a setup the most accurate GP model will be used as the plant model to avoid time-consuming SOWFA simulations.

References

- [1] Veers P, Dykes K, Lantz E, Barth S, Bottasso C L, Carlson O, Clifton A, Green J, Green P, Holttinen H, Laird D and et al 2019 *Science* **366**
- [2] Freeman K, Frost C, Hundleby G, Roberts A, Valpy B, Holttinen H, Ramírez L and Pineda I 2019 *Our Energy, our Future - How Offshore Wind will help Europe go Carbon-Neutral* (Brussels, Belgium: Wind Europe)
- [3] Fleming P A, Ning A, Gebraad P M and Dykes K 2016 *Wind Energy* **19** 329–344
- [4] Steinbuch M, De Boer W, Bosgra O, Peeters S and Ploeg J 1988 *Journal of Wind Engineering and Industrial Aerodynamics* **27** 237–246
- [5] Boersma S, Doekemeijer B, Siniscalchi-Minna S and van Wingerden J 2019 *Renewable Energy* **134** 639–652
- [6] Fleming P, Annoni J, Shah J J, Wang L, Ananthan S, Zhang Z, Hutchings K, Wang P, Chen W and Chen L 2017 *Wind Energy Science* **2** 229–239
- [7] Fleming P, King J, Dykes K, Simley E, Roadman J, Scholbrock A, Murphy P, Lundquist J K, Moriarty P, Fleming K and et al 2019 *Wind Energy Science* **4** 273–285
- [8] Howland M F, Lele S K and Dabiri J O 2019 *Proc. of the National Academy of Sciences* **116** 14495–14500
- [9] Barthelmie R J, Hansen K S and Pryor S C 2013 *Proc. of the IEEE* **101** 1010–1019
- [10] Annoni J, Seiler P, Johnson K, Fleming P and Gebraad P 2014 *2014 American Control Conf.* (Portland, OR) pp 2517–2523
- [11] Bastankhah M and Porté-Agel F 2016 *Journal of Fluid Mechanics* **806** 506–541
- [12] Doekemeijer B M, van der Hoek D C and van Wingerden J W 2019 *2019 IEEE Conf. on Control Technology and Applications* (Hong Kong, China) pp 284–289
- [13] Martínez-Tossas L A, Annoni J, Fleming P A and Churchfield M J 2019 *Wind Energy Science* **4** 127–138
- [14] Sun H and Yang H 2018 *Applied Energy* **226** 483–493
- [15] Kabir I F S A, Safiullah F, Ng E and Tam V W 2020 *Energy* **193** 116761
- [16] Ge M, Wu Y, Liu Y and Yang X I 2019 *Renewable Energy* **141** 46–56
- [17] Marden J R, Ruben S D and Pao L Y 2013 *IEEE Transactions on Control Systems Technology* **21** 1207–1214
- [18] Ciri U, Rotea M A and Leonardi S 2017 *2017 American Control Conf.* (Seattle, WA) pp 25–30
- [19] Marchetti A, Chachuat B and Bonvin D 2009 *Industrial & Engineering Chemistry Research* **48** 6022–6033
- [20] Andersson L E and Imsland L 2020 *Wind Energy Science* **5** 885–896
- [21] Andersson L E, Bradford E C and Imsland L 2020 *2020 American Control Conference [accepted]* (Denver, CO)
- [22] Churchfield M J, Lee S, Michalakes J and Moriarty P J 2012 *Journal of Turbulence* N14
- [23] de Avila Ferreira T, Shukla H A, Faulwasser T, Jones C N and Bonvin D 2018 *2018 European Control Conference* (Limassol, Cyprus) pp 465–470
- [24] Rasmussen C E and Williams C K 2006 *Gaussian Processes for Machine Learning* vol 2 (Cambridge, MA: MIT press)
- [25] Snelson E and Ghahramani Z 2006 *Advances in Neural Information Processing Systems* ed Weiss Y, Schölkopf B and Platt J C (Cambridge, MA: MIT Press) pp 1257–1264
- [26] Snoek J, Larochelle H and Adams R P 2012 *Advances in Neural Information Processing Systems* (Cambridge, MA: MIT Press) pp 2951–2959
- [27] Burton T, Jenkins N, Sharpe D and Bossanyi E 2011 *Wind Energy Handbook* (Chichester, UK: John Wiley & Sons)
- [28] del Rio-Chanona E A, Graciano J E A, Bradford E and Chachuat B 2019 *IFAC-PapersOnLine* **52** 52–57