Effectively using multifidelity optimization for wind turbine design
Wind Energy Science

John Jasa, Pietro Bortolotti, Daniel Zalkind, Garrett Barter

Notes:

- Our direct responses to reviewers' comments are written in blue in this response.
- Actions taken to address the reviewers' comments are written in red in this response.
- Reviewer 1 modifications are highlighted in red in the manuscript.
- Reviewer 2 modifications are highlighted in yellow in the manuscript.
- Other modifications are highlighted in blue in the manuscript.

Reviewer 1 Comments and Response

Comment 0: My initial impressions is that this is a good article and that it should be accepted to this journal with some revisions.

The authors have looked at a multifidelity algorithm that uses additive surrogates based correction models. Furthermore, the approach does not require gradients from the high-fidelity model. The author made these choices so that the method is general purpose. The authors then applied this method to a blade design problem, controller tuning problem and finally a wind farm layout optimization problem. In all cases the authors were able to show that the multifidelity method was able to obtain a better solution (according to the high-fidelity analysis) than single discipline low fidelity optimization. While at the same time solving these problems faster than the single fidelity high fidelity optimization. The article is interesting and relevant work to the community and has been well written. However, there are some areas for improvement and I feel that some revisions can help make a stronger manuscript. Those are discussed as follows:

Response: Thank you for reviewing our work, we appreciate your insights and suggestions.

Comment 1: First, I have concerns about the multifidelity algorithm that I think the authors should comment on in greater detail. The first point of concern is that none of the multifidelity optimization solutions matched the high-fidelity solution. This suggests that at the limit the corrected low fidelity problem is not equivalent to the high-fidelity solution. Some corners of the multifidelity literature try to show that the corrected multifidelity algorithm is equivalent or not. It would be good for the authors to comment on whether their algorithms tries to achieve this or not. This lack of equivalence is obvious in the design solution for the blade optimization. For readers and engineers that are more concerned with the final design configuration than the objective, it would be helpful if the author could comment on that and explain to what extent the surrogate correction played a role in this discrepancy.

Response: Thank you, we've added two additional paragraphs in Sec. 2.3 and 3.1.3 to address your concerns in both the general sense and highlighting why the blade optimization results differ. In short, the approximation is interpolative, meaning the corrected low-fidelity model matches the high-fidelity model at the high-fidelity sampling point, but not elsewhere in the design space.

Comment 2: Also I think that this shortcoming is related to convergence of the surrogate to the high-fidelity model. The authors show some plots that qualitatively show convergence. However, it is small differences in the gradients at the limit that ultimately define an optimal solution. I think that it would be good for the authors to show quantitatively the convergence rate of the surrogate model at the optimal high fidelity optimization point for both the solution and the gradients.

Response: Thank you for this great suggestion. We've added a figure showing the function and gradient values for the multifidelity approximation, showing how they converge to the high-fidelity results as the trust-region method iterates. This is in Sec. 2.2.

Comment 3: Furthermore, surrogate based methods suffer greatly from the curse of dimensionality. Despite this, I was quite pleased that the authors were able to obtain good results despite this typical weakness. To overcome this weakness the authors explain how they used a specific type of surrogate that internally performs model reduction to overcome this dimensionality. I personally would like to know more details about this. For example, do many dimensions cause the surrogate longer to start converging? How many high-fidelity evaluations were needed before good convergence was achieved? Maybe more details on the internal model reduction methods used in their surrogate would be nice.

Response: Thank you, we have added an additional paragraph highlighted in red in Sec. 2.3 to address these concerns. This paragraph offers a short summary of the advantages of the Kriging partial least squares (KPLS) method, but also encourages the reader to seek the original work by Bouhlel et al. [2] for additional detail.

Comment 4: At the end of the article in figure 10, the authors summarize the different cases purely in terms of computational savings. However, I think different readers are more interested in a range of metrics. Since none of the multifidelity solution achieved the same performance of the pure high-fidelity, I think that you also need to include this in the summary. Furthermore, sometimes it's not the objective value that matters, but the final design configuration. Thus, I think that it's also important that you discuss limitation. Finally many readers would consider the distance in the design space from the high fidelity solution an important metric. So I would also show this in the summary. Most would see the objective and final design configuration more important...

Response: Thank you, we understand your point about providing more detail on the differences between the optimal designs for high fidelity and multifidelity methods. We have reworked this section to more clearly show the performance along with the computational cost. Specifically, we've removed Fig. 10 and replaced it with a table summarizing the three problems' results. For the design-space differences between the multifidelity and high-fidelity results, we note that this discussed in each of the case studies. We have also expanded on the limitations of this method both in individual case studies' sections and in the conclusion paragraph.

Comment 5: Concerning the cited literature, you have many citations of the work within wind energy. However, I think you could give a better summary of the different contributions in surrogate based optimization. This is an old topic in the aerospace community. There are a mix of different methods that have a range of speed improvements and accuracy in terms of reproducing the high fidelity solution. It would be good if you could present these different methods and explain how your work fits with these.

Response: Thank you, we added an additional paragraph highlighted in red in the introduction to further place the present work within the context of other surrogate-based optimization efforts.

Comment 6: Another comment is the choice of problems. The IEA Task 37 has created various optimization test cases for this purpose. It would be better for the larger community if your multifidelity was applied to these test cases so that the community can see how your results compare with the larger community results. This is of course a big task ... but maybe this is something to consider for future articles.

Response: Thank you for this suggestion. We will prioritize including IEA Task 37 test cases in continuations of this work. We also changed text in the paper to reflect the open-source nature of each application case so that readers can run, adapt, and use the problems with their own methods for comparison.

Comment 7: In your future work discussion, I would elaborate a bit more on how this method could be improved. The multifidelity literature in aerospace shows faster optimization with better correspondence with the pure high fidelity solution. You made choices to improve generality and applicability, so this not a criticism, but I think in light of other methods it would be good to comment on what would be needed to get closer the high fidelity solution faster.

Response: Thank you, we have added to the future work section to address possible changes to this method to improve convergence. Specifically, we have called out the use of high-fidelity gradients and expected improvement algorithms as mechanisms to improve on the current methodology.

Reviewer 2 Comments and Response

Comment 0: This manuscript talks about multifidelity optimization in wind energy. Now "multifidelity optimization" is a large umbrella under which many different approaches fall - things that people have been doing for years in order to speed up their optimizations, but also more recent ideas how to do this in a more structured way. The manuscripts starts with an informative review of some of the literature and recent domain-specific works that explicitly considers the multifidelity optimization idea, and this contains useful references and is valuable. The authors then quickly describe their method of choice, which is the use of trust-region optimization with a "Kriging partial least squares" surrogate model. Unfortunately no details at all are given about their particular algorithm used and what parameters or other choices have been made. This is a major weakness of this paper, especially so since trust-region optimization is not something that most people will be familiar with (even though it is extensively used with highly nonlinear problems). The used surrogate model is a modern variant of the Gaussian process model, which has received much attention recently (especially in wind energy applications), due to its intrinsic ability to estimate the uncertainty in its predictions. I would assume that this feature will be used when performing trust-region optimization with it, so this should be explained and discussed.

Response: Thank you for your detailed and constructive summary. We have added much more detail about our multifidelity trust region method, including an algorithmic description, additional details for the surrogate model creation, and links to repositories containing both the framework and specific cases studied in this paper. Additionally, we go into more detail discussing why we chose the methods we use for the surrogate model, highlighting the benefits and drawbacks of this method, especially regarding dimensionality. Specific changes to the paper are discussed in response to your other comments below.

Comment 1: There are many interesting options and ideas here that would interest readers, and generally these methodological details are more interesting than the shown example studies and their results. In fact, as long as the method is not fully clear, the results are not that interesting - to a scientifically-minded reader. One might be inclined to forgive the authors this lack of details here given the typical space requirements of journals (and the cost of publication), but without this information the article reads more like an advertisement for the authors' method than what was probably intended. At the very least a reference to supplementary material (e.g. in the form of a non-peer reviewed technical report - this is a prime example of why we still need and should value this type of publication!) should be given where readers can find all implementation details.

Response: We appreciate your findings that the paper is lacking in reproducibility and detail on how readers could use this approach. We have expanded the paper and included much more detail on our method and how users can apply this to their research, especially in the methodology and conclusion sections. Additionally, we have focused more on providing details so that readers can adapt and use these methods in their own work.

Comment 2: The manuscript continues to discuss three example studies where this approach - correcting a low fidelity simulation with a meta-model based on few evaluations of a high-fidelity simulation - seems to perform somewhat better than either optimization based on the low-fidelity model alone (being more accurate than this) or when using the high-fidelity model alone (being faster than this). Again, it is unfortunately completely unclear how the optimization with these models (only the low- or high-fidelity ones) is performed!

Response: The article aims to demonstrate that a multifidelity optimization approach works better than single-fidelity optimizations. Two of three case studies clearly support this conclusion. The third case study, in Section 5.3, instead shows that when the number of design variables is too high, the multifidelity optimization approach struggles to match the high-fidelity optimum (although it still guarantees lower computational costs). The single-fidelity design studies adopt the same identical objective, design variables, and constraints as the multifidelity studies. The only difference is that in the single-fidelity studies the optimizer is fed only with the results from the low- or high-fidelity levels. As discussed in Section 2, the multifidelity optimizer is exposed to both and we have expanded the algorithmic description of the specific trust-region approach. Hopefully this enhanced description gives greater insight and understanding into the results as well.

Comment 3: Somewhat surprisingly, the high fidelity models used in the case studies have only slightly higher complexity than the low fidelity models. For example, in the blade design study, both models are based on blade element momentum (BEM) theory. The low fidelity model is using the original steady state BEM approach, whereas the high fidelity model is using its unsteady generalization. For this pedagogic demonstration this is sufficient, but in real applications one would expect to use some kind of CFD simulations as the high fidelity model, I assume?

Response: We fully agree with the reviewer, the two rotor aerodynamic models are quite close in terms of fidelity. In addition to the steady and unsteady nature of the low- and high-fidelity models respectively, the high-fidelity model had more degrees of freedom enabled during the analysis, including a torque-pitch controller. This further added realistic differences between fidelity levels.

During the study, we did try to elevate the fidelity of the high-fidelity model, moving from BEM to a free-wake vortex model. Unfortunately, the latter model proved somewhat inadequate for optimization, with results leading to several unanswered questions. These questions are currently being addressed by a dedicated analysis effort. Once a more reliable higher fidelity model is made available, single- and multifidelity studies will be run. As the reviewer notes, for the pedagogic objective of this work, we deemed it appropriate to stick to more reliable and conventional rotor aerodynamic models. We added a sentence to the end of Section 6 highlighting that a larger offset in the fidelity of the rotor aerodynamic models will be investigated in the future.

Comment 4: All in all, the topic is relevant, and the authors seem to have some success with their approach. If now they could only describe the approach with enough details that readers could reproduce it (or learn from this paper how to do something similar with their own wind energy optimization problems), then this would be a very nice contribution to the literature.

Response: Thank you for your constructive feedback. We have fully detailed the trust region methodology in the paper and referenced the open-source code repository in the conclusion to reinforce that readers can directly apply the multifidelity method to their own optimization problems.

Comment 5: The authors need to explain the details of the trust-region optimization and KPLS surrogate model used, in sufficient detail that readers can (potentially) reproduce the results in the paper.

Response: Thank you for the suggestion to improve the details provided for the trust-region method and the KPLS surrogate model. We have clarified how the trust-region method operates by including a step-by-step algorithm, also noting that we adapted the method from March and Willcox [3]. Additionally, we have provided the URL for the open-source Surrogate Modeling Toolbox that supplies the KPLS method.

Comment 6: page 4: "This approximation is devised such that it is equal to the high-fidelity model at the points where we have high-fidelity data." - So it is what we would call an interpolating approximation. Note that this property is not strictly necessary (although it usually makes sense).

Response: We have added a sentence in Sec. 2.3 highlighting that we use an interpolative approximation function but that it is not strictly necessary.

Comment 7: Fig. 2: It is unclear how the trust region is changing between iterations.

Response: We have clarified how the trust region changes between iterations by adding in Alg. 1 and adding text in Sec. 2.2 explaining this process.

Comment 8: Fig. 3: There seem to be samples missing in the figure (e.g. in Fig. 3a, the function depicted is not piecewise-linear between the sample on the far left and the cluster of samples in the right third). Also, are all three figures based on the same number and location of samples from the high-fidelity model?

Response: All three figures are indeed based on the same number and location of samples. Additionally, the previous versions of the figures showed three quantities: the low-fidelity model, high-fidelity model, and the corrected low-fidelity model defined as $f_{\text{corrected}} = f_{\text{low-fidelity}} + f_{\text{surrogate}}$. We have modified the figures to also show the actual corrective function (surrogate) values as well. This better highlights the behavior of the corrective function between the models. Additionally we've added text in Sec. 2.3 clarifying what the plots show.

Comment 9: In the blade design study, the low-fidelity optimized blade must be cheaper (less materials used) than the high-fidelity or multifidelity optimized blade? By how much? The comparison (Table 2) would be more informative and convincing if the optimization was using the same route (e.g. by enforcing only increases in twist and chord?) instead of converging on completely different paths (toward completely different local optima).

Response: The blade design study is an aerodynamic optimization study whose objective is the maximization of rotor aerodynamic power. To keep things as simple as possible and focus on the optimization approach, rotor costs are not included in the analysis and, more importantly, are not exposed to the optimization solver. The solution space of a rotor aerodynamic design process is often flat, see for example the case study developed within the IEA Wind Task 37 on systems engineering described in McWilliam et al. [4]. In a real wind turbine design process, multiple design objectives and constraints guide the design process of a real wind turbine rotor as detailed in many publications, such as Bortolotti et al. [1]. This said, we find that the study presented here is more informative and convincing when the low- and high-fidelity models push the optimizer in opposite directions, with the optimizer still succeeding to converge to an optimum that is only marginally lower than the high-fidelity optimum. We have added a paragraph about these considerations at the end of Section 3.1.3.

Comment 10: How the method works is intuitively clear in 1D (second case study), but what happens in higher-dimensional cases? The first case study is 7 dimensional, the third one is 14 dimensional. How does the approach work with such high-dimensional data, what are its limits or challenges with it? For example, how are the surrogate models initialized, and does this effort not become too computationally expensive in higher dimensions?

Response: This is a good question, thank you. Our results support your comment in that the case study (#3) with the most design variables showed the mediocre improvement for the multifidelity method. We've added a brief discussion on this to the beginning of Sec. 3 for the case studies, also highlighting where else in the paper we discuss this. We've also expanded on the dimensionality scaling of KPLS in the surrogate description section, including initialization and training cost trends.

Comment 11: Fig. 6: Abbreviations of what simulation outputs are shown should be explained. Not every reader is so familiar with OpenFAST software that they will immediately understand what they see here.

Response: Thank you, that is a good point. We have updated Fig. 7 to show simulation output names and units instead of their OpenFAST channels.

Comment 12: page 6: "By not using simple additive or multiplicative factors ..." - This seems to suggest that only the last example uses a surrogate model as correction function, and the first two case studies use something simpler? Or is it just this statement that is somewhat misleading?

Response: Thank you for noticing this. It was intended to compare our approach to other methods in the literature, not the previous case studies. We have removed the phrase "By not using simple additive or multiplicative factors" to avoid confusion

Comment 13: Where is the data availability statement (required by the journal)? Where can readers find the code and the data used for the results in this paper?

Response: We have added the code and data availability statement and provided a link in the manuscript to the doi: `https://doi.org/10.5281/zenodo.6109699`.

Thank you for your diligent comments and suggestions.

References

[1] Bortolotti, P., Dixon, K., Gaertner, E., Rotondo, M., and Barter, G.: An efficient approach to explore the solution space of a wind turbine rotor design process, Journal of Physics: Conference Series, 1618, 042 016, doi:10.1088/1742-6596/1618/4/042016, 2020.

[2] Bouhlel, M. A., Bartoli, N., Otsmane, A., and Morlier, J.: Improving Kriging Surrogates of High-Dimensional Design Models by Partial Least Squares Dimension Reduction, Structural and Multidisciplinary Optimization, 53, 935–952, doi:10.1007/s00158-015-1395-9, 2016.

[3] March, A. and Willcox, K.: Provably convergent multifidelity optimization algorithm not requiring high-fidelity derivatives, AIAA journal, 50, 1079–1089, 2012.

[4] McWilliam, M. K., Zahle, F., Dykes, K., Bortolotti, P., Ning, A., Gaertner, E., Macquart, T., Merz, K., and Ruiz, A. I.: IEA Wind Energy Task 37-System Engineering-Aerodynamic Optimization Case Study, in: AIAA Scitech 2021 Forum, p. 1412, doi:10.2514/6.2021-1412, 2021.