

Response to Reviewer #2

We greatly appreciate the time taken by the reviewer to read our manuscript. We have taken into consideration and addressed all comments, questions, and suggestions from the reviewer, and we feel that the revised manuscript is now substantially stronger as a result. Changes made to the text at the request of the reviewer have been highlighted in red in the revised manuscript. In the following, reviewer comments are repeated in blue italics and our responses are provided in the bulleted sections of text.

The authors present a hybrid parameter estimation technique which includes a gradient-free first optimization stage and a gradient-based second stage. The two-stage approach applied to a 1D coupled physical-biogeochemical ocean model appears to work well, but for a study focused on methodology, more details about it should be included.

General comments

The study makes a relatively straightforward case, the text in general and the description of the methods is mostly easy to follow, and the experiments are well motivated (maybe except for the TSE, see below).

However, the manuscript frames this study as one primarily focused on the demonstration of the parameter estimation method. And here, I think, more emphasis could be placed on the parameter estimation part.

(1) As a reader, it would be good to get a better idea of the computational cost in terms of runtime or number of model/function evaluations, in particular comparing the initial global search to the gradient-based second stage.

- We agree that it is important to clarify the computational expense of the proposed method, particularly since many of the decisions made were geared toward increasing computational efficiency. To address this point, we have now included a new Appendix B in the revised paper that outlines the cost of running each step in the methodology. In brief, the sampling computational time was 12,571, 12,247, and 45,000 cpu-hours for BATS, HOTS, and the multi-site calibration cases, respectively. The total computational time was 31,052, 21,279, and 81,951 cpu-hours, with 60%, 42%, and 45% of the compute time spent on the optimization portion of the two-step process.

(2) It seems sensible to use the two-stage (global+local) parameter estimation, but how does each stage contribute to the decrease in the cost function? Readers may ask if after 25000 simulations in stage 1, is there any need for stage 2? Or, how well does stage 2 do if started from the baseline?

- This is another great point and, in addition to exploring the computational cost of each stage in the optimization methodology, we have included more detail in Appendix B of the relative improvements in model agreement from the random sampling and the local optimization portions of the method. As shown in Figure B1, in all cases we can further reduce the error after the initial global step by doing local optimization. For BATS, HOTS, and the multi-site cases the error decreases by 31%, 91%, and 85% respectively. Note that most of the improvements are for HOTS which is initially just very bad.

(3) Studies like this one often include replicate experiments, why not include one here, to show that similar parameter values are recovered when running the estimation again.

- We agree that replicate experiments can be an important aspect of parameter estimation studies and text near the beginning of Section 4.2 noting this point. We have chosen not to do such experiments in the present study because of the random nature of the initial global search in the 51-dimension parameter space considered here. That is, with only 25,000 samples in the Latin hypercube step, we will likely start the second-phase gradient-based optimizations in the replicate experiments from a completely unique set of parameter values, resulting in different final parameter values. Based on our tests, however, the 25,000 samples in the initial global search are sufficient to ensure that each successive application of the overall method will result in better agreement with the observational data than the baseline values from Smith et al. – it is just the extent of the improvement that may differ.

(4) 25000 model evaluations are a lot, which basically prohibits the use of 3D models (BGC model coupled to 3-dimensional circulation models). What could be done to reduce runtime, would the authors consider the use of a more sophisticated gradient-free method (some are described in the introduction) instead of Latin hypercube sampling?

- We thank the reviewer for drawing attention to these important points. Regarding ways to reduce the runtime, there are two distinct ways to reduce the overall computational cost of performing the parameter estimation. The first is to cheapen the cost of running the model, while the second is to reduce the number of model evaluations performed. For the former, we cannot significantly reduce the cost of evaluating BFM17+POM1D, which is already effectively a physics-based surrogate for more detailed dynamical models. This is now discussed in the final paragraph of the Conclusion. To reduce the overall number of model evaluations, it is indeed possible that other gradient-free methods, such as genetic algorithms (GAs) could be used. However, for the present study, we instead chose Latin hypercube sampling (LHS) because many model evaluations can be easily performed in an embarrassingly parallel fashion. So, while the computational cost of running many LHS samples may be comparable to (or even slightly greater than) running a GA to convergence, the real time taken to perform the optimization is much less since there are limits to what can be parallelized when implementing a genetic algorithm. It is conceivable that truncated GAs could be used instead of the LHS, and we have mentioned this approach in the last paragraph of Section 2.2. To address the broader point from the reviewer, we have also added a new fourth paragraph to the Conclusions.

In the twin simulation experiment (TSE; section 4.1), it reads like all 17 state variables were used as synthetic observations. But not all 17 variables are part of the BATS or HOTS datasets, in fact, only 5 appear to be used. How do the results of the TSE change if only 5 state variables are observed? This would appear to be a much more important experiment than one using all the variables. Additionally, a little later in the section (l 312): "While these results may suggest that the least sensitive parameters could be excluded from the subsequent calibration studies, redoing the sensitivity study with our standard objective function reveals larger relative importance values for the full set of parameters.": This is a bit of an unsatisfactory result, could the discrepancy be due to including 5 compared to 17 observation types in the objective function?

- In response to this helpful comment, as well as a similar comment from the other reviewer, we have performed an additional twin simulation experiment (TSE) that more closely matches the annual, five-field optimization that is the focus of the parameter estimation at the BATS and HOTS locations. The results of this TSE are included as a new Figure 6 in the revised paper, with substantial additional text included at the end of Section 4.1 to discuss the results of this TSE. Briefly, we found in the annual TSE that more parameters had a relative importance greater than 0.01, as compared to the 30-day tests. This is indicative of a more complex optimization problem where many more parameters can affect the results. Consequently, although we recover the baseline parameter values across the range of relative importance values, there are still some parameters that we do not fully recover. However, most parameter values that do not reach the baseline value do at least approach the value. As we now discuss at greater length at the end of Section 4.1, the comparison of the 30-day and annual TSE results demonstrates the challenge of estimating many sensitive parameters in a complex objective function space, even when using a gradient-based approach. Concerning the last question from the reviewer, the discrepancy in sensitivities comes from having a different objective function. In one sense it is not a discrepancy and is instead the result of the problem being fundamentally different. However, in terms of making sure that the optimizer is decreasing the error as expected and demonstrating that DAKOTA and our model are interfacing correctly, we feel that the present TSEs are sufficient.

Specific comments

L 72: "in a high-dimensional BGC model across a range of ocean conditions, using a 1D model for vertical ocean mixing": Here it would be nice for the reader to explain a bit better what is meant by dimensionality: "high-dimensional BGC model" and "1D model for vertical ocean mixing" are used together and at this

point it is unclear if "high-dimensional" refers to 3 spatial dimensions (represented at various spatial locations as 1D vertical models) or if it is referring to the dimensions of the state space. More generally, it would be useful to describe what is to follow in a bit more detail early on. Even the formulation in line 5 "simultaneous parameter estimation at multiple ocean locations" is a bit ambiguous, and the kind of model setup (1D, 3D) that is being used in the study could be mentioned earlier.

- We appreciate this suggestion and have now attempted in the abstract, introduction, and several other places in the paper to be more specific when referring to "dimensionality".

L 98: I presume Π is a matrix of weights and $\Pi_{\{i,j\}}$ is a scalar weight?

- Yes, Π is a matrix of weights, while $\Pi_{\{i,j\}}$ are scalar weights for the corresponding field at the corresponding ocean site. In this work, we use an identity matrix for Π but include it in the generalized description of the methodology to emphasize that weights can be added. We point out that the weights can be formulated differently in the description of the method and in the conclusion, and the corresponding line noted by the reviewer has been updated to improve clarity.

L 99: "describes the misfit with observational (or other reference) data": For clarity, I suggest adding "of the model output" or similar.

- The corresponding line has been updated for clarity.

Eq 2: I would have expected that the $1/\sigma$ term (a weight) would be contained in Π . In fact, the subsection makes no further reference to Π , and it would be useful to explain here what it is used for.

- We thank the reviewer for bringing this potential point of confusion to our attention. We consider the $1/\sigma$ term to be a normalizing factor in the error function (Eq. 2) and different forms of this function, with different normalizations, are considered in Appendix A2. The weights Π are included here for generality and could be used to provide more emphasis on certain observational fields, for example those that are deemed to be of greater importance or for which uncertainties are lower. We have now added text in Section 2.1 after Eq. (2) to explain both these points.

Fig. 1: It is somewhat easy to trace the line from model run and comparison with data to "Calculate model error", feeding it into DAKOTA and obtaining new parameters as output with which to run the model again. but what do the dotted gray lines mean and why does the calculated model error bypass the interface? A minor complaint is that terms like BFM17 and $\partial A_j / \partial t$ have not been mentioned in the text when this figure is first referenced.

- To address the concerns, the caption for Figure 1 has been updated. The grey dotted lines are intended to show how the different components are implemented in practice – the error calculation output is formatted so that DAKOTA can read it, so it does not have to be interpreted by the interface.

L 151: "how quickly the values of J increase when the N_{random} simulations are sorted": That sounds a bit like J increases while the simulations are being sorted, to avoid confusion maybe use something like: "how quickly the values of J increase in the sorted N_{random} simulations".

- The proposed change was made for clarity.

L 164: "The QN algorithm reliably and efficiently converged to optimized solutions.": Is this a result from this study or a general observation? Maybe add "In our experiments ...".

- This change has been made.

L 165: "Similar to the ecosystem parameter estimation study by Matear (1995), we found that the conjugate gradient method failed to converge efficiently." It reads as if this is still meant as a justification for using QN over the conjugate gradient method. I suggest rephrasing it slightly: "In comparison, we found that the

conjugate gradient method failed to converge efficiently, a result similar to that in the ecosystem parameter estimation study by Matear (1995)."

- We agree with the reviewer that this point could be clearer and have made the suggested change.

L 192: "Each CFF is represented as a vector where each element is a constituent component concentration corresponding to a state variable.": Initially, I thought that this meant that the state variables are divided up into multiple vectors according to their type/function, which is not the case according to the following paragraph. This could be explained better without adding much more text, maybe "Each CFF is a vector representing a model variable divided into the elemental constituents represented by the model, e.g., the phytoplankton CFF is a 4-element vector containing the C, N, P and chlorophyll concentration of the phytoplankton variable (see below for more details)."

- The text has been edited to add clarity.

L 193: "BFM17 was simplified to be a general, but computationally cheaper, model that retains the essential BGC processes for modeling a phytoplankton spring bloom ...": Is this mean in comparison to BFM56?

- Yes, this statement was made for BFM56. The text has been edited to add clarity.

L 232: At this point, it is unclear from the text what parameter values are used. Interestingly, this information, plus comments about the manually adjusted parameters, is provided in the description of Fig. 4 in the next subsection, which is helpful to the reader. I would suggest moving the information that applies to both datasets (use of baseline model parameters, single-site model calibration and the multi-site calibration etc.) to the description of Fig. 3.

- A reference to the origin of the baseline parameter values is now made in the earlier subsection, with the later subsection being edited accordingly.

L 303: "nominal value": Does this refer to the optimized value?

- Nominal value refers to the baseline or target parameter values we are trying to recover over the course of the twin simulation experiment. The verbiage in the paper has been updated for clarity.

L 304: So some parameters were only perturbed down, because perturbing them up would hit the baseline + 25% boundary? What then makes up the second perturbation case, the optimized value? This could be explained better.

- The sensitivity analysis was performed by perturbing values by 5% except when this would exceed the standard bounds (i.e., those in Table C1), not the alternative +/- 25% bounds. The text has been updated to clarify this point.

L 305: "... defined as the maximum objective function evaluation between the two perturbation cases for each parameter." This sentence is difficult to understand, I suggest rephrasing it.

- Based on this comment as well as one from the other review, we have added Eq. (6) to make the calculation of the sensitivity factor clearer. We have also added corresponding text around this equation.

L 317: It would be useful to state what N_v is set to here. I presume it is 5, based on previous figures. But it is not stated directly in the text.

- Yes, we have used $N_v = 5$, corresponding to the five target observational fields. This is now explicitly addressed at the beginning of Section 4.2.

L 351: "improved by a factor of ... over 236 (for nitrate)": True, but this large improvement is more a function of the enormous misfit in the baseline experiment.

- This is indeed true, and we have added text at the beginning of Section 4.2.2 noting the enormous initial error at the HOTS location.

We thank the reviewer for these useful comments, and the paper has now been revised to address all the above points. Sincerely, the authors.