In this work, the authors present a novel state estimate of the Amundsen and Bellingshausen Seas, along with the associated ice shelves and their cavities. Using techniques developed by the ECCO project, the authors iteratively optimize a numerical simulation of their target region using a suite of available observations. They show that these observational constraints increase the agreement between the numerical simulation and the observational suite, as evidenced by a decreasing cost function. The mean state of the model improves by this metric, although there is still room for improvement with respect to temporal variability on interannual and multi-year scales. Finally, the authors conduct a set of sensitivity experiments by turning off the optimization for selected aspects of the "control vector"; for example, they test the effect of optimizing wind by using the non-optimized wind, while leaving the other optimizations in place, and observing the resulting change in the cost function.

The paper addresses the highly relevant and challenging problem of simulating the Antarctic shelf seas and associated ice shelves. The Amundsen Sea region is especially relevant to the impacts of climate change and sea level rise, as it features especially high ice shelf melt rates; it is considered a vulnerable part of the ocean-ice system and is currently the focus of a UK-US Thwaites Glacier study, among other funded projects in the region. As such, the state estimate developed in this paper is highly relevant to a pressing scientific problem. The work described here represents a significant step in ocean-ice modelling, as there have been very few prior attempts to incorporate ice shelves into state estimates. In this context, this modeling effort has been especially successful and is encouraging to see.

The methods and assumptions in the paper are clearly stated and appropriate for the process of state estimation. For example, changing the viscosities and melt rate coefficients for better agreement is a standard approach in adjoint modelling and state estimation. The results support the conclusions of the paper. An especially interesting result is that the adjustment of wind has the strongest impact on the ocean state, while the adjustment of heat flux has the strongest impact on the sea ice state.

The work is reproducible to an extent, as the authors have uploaded their work via Zenodo and the NASA data portal. The adjoint-driven optimization procedure uses an expensive commercial license to derive the adjoint code from the forward code; the cost of this license prohibits this work from being fully reproducible in terms of repeating the optimization process, but the open-source tools available for algorithmic differentiation of MITgcm are not yet competitive with the commercial tool. It is still standard practice in state estimation to use the commercial tool. That being said, the authors have taken the appropriate steps to make their state estimation work reproducible to the extent possible. One can run the non-optimized and the optimized simulations using the setup that they have provided. (I have not tried to run it myself, but from what I can see, all the required files have been uploaded.)

Overall, this is a highly relevant and exciting piece of work. The paper is mostly clear, but it would be an even stronger paper with some additional editing. I have made some specific suggestions below, along with a few questions about their state estimation procedure.

-- Specific questions, comments, and suggestions –

Abstract

(Lines 1-5) Change “is impacting” to “impacts”. Add “the” before “global carbon”.

(Line 20) Change “that that” to “that”

(Around line 25) Add “the” before “global carbon cycle”. Add “the” before “ice shelf cavities”.

(Line 42) Change “descriptions” to “representations”

(Line 43) I’m used to seeing this described as the “4DVAR method”, as opposed to the “adjoint method”. I think either one is fine, but its’ good to be consistent and use the same term throughout the paper. Later in the paper this is called the “4DVAR method”.

(Line 47) I’ve heard the suggestion that we should call this “algorithmic differentiation” instead of “automatic differentiation”. Consider changing it if you agree.

(Between lines 60-65) Should “Green’s function” be capitalized or not? Be consistent.

(Line 69) I’m not sure why the reference to Figure 2 is here. At this point, the concepts of iteration, cost functions, and the individual terms have not yet been introduced. It may be a bit confusing here. Add a quick definition of what you mean by “iterations” here. Otherwise people may get confused between numerical forward integration iterations and optimization iterations.

(Line 73) What is the advantage of using LLC270 here? Can you tell us more about this choice, please.

(Line 84) Why aren’t these datasets used? Mentioning them raises the question as to why they weren’t used.

(Line 92) What is a bipolar grid? Expand on this please.

(Lines 97-98) Please say more about why the ice barrier was necessary, for readers not familiar with the Nakayama study.

(Line 100) Sorry, what is the “ongoing LLC270 optimization”? Is this a global state estimate? What’s the approximate resolution here?

(Line 101) Did you consider using the new ERA-5 product or perhaps MERRA? What was the reason for choosing ERA-Interim?

(Line 104) Here you call it the 4DVAR system for the first time. It's better to be consistent throughout the paper. Either call it the adjoint method or the 4DVAR method.

(Line 109) Did you use the M1QN3 that comes with MITgcm, or did you use Martin Losch's implementation of this approach?

(Line 111) Say more about this. What is the "pseudo-sea ice adjoint"?

(Line 117) How did you decide by how much to change the coefficients?

(Line 130) What is your point of comparison here? When you say "too cold and fresh", relative to what observations? Make it clear how you're doing this comparison.

(Lines 142-145) This is worded a bit confusingly. It's better to repeat the sentence instead of using the "respectively" construction. For example:

"For March in the AS, simulated sea ice in iteration 0 is larger than observations by X (Y% difference), and it iteration 20 it is larger by X (Y% difference). In the BS, ... (repeat)."

I would also include a percentage-wise comparison to give us a sense of the relative size of the difference.

(Line 150) Replace "2010-2014 simulation for unoptimized simulation" with "2010-2014 unoptimized simulation".

(Line 178) I would include these ice shelf melt rates as lines of reference on Figure 9. They could be straight horizontal lines, clearly labeled.

(Lines 190-205) It is a bit hard to tell from the text how the sensitivity experiments are done. Please be more explicit about how the tests were done and introduce the case names.

(224-226) So from what I understand, the sea ice cannot be adjusted directly in this state estimate, because there is only a pseudo-adjoint for the sea ice. Any changes in the simulated sea ice come from the adjustments of the surface forcing and ocean state. How do you think the solution would change if the sea ice could be included in the control vector? Since this is the discussion section, where perhaps a little informed speculation is allowed, can you speculate on how the addition of a sea ice adjoint would change the state estimate?

(Line 235) Change "is available" to "are available". Change "cost" to "the cost function" to be more explicit.

(Line 240) Add "the" in front of "seasonal".

(Line 242) The comparison with the mooring data in Figure 12 shows that the optimization procedure helped improve the agreement between the observations and the model state. I think it's worth emphasizing this positive improvement here.

(Line 246) You are using some moorings, correct? I assume this means that there are additional moorings that could be included in the future. Please could you clarify this?

-- Figures –

Figure 1. Add degree symbols to the latitudes and longitudes.

Figure 2 caption. Mention that the vertical scales are different in the different subplots. It's worth emphasizing.

Figure 3. This figure is not referenced in the paper, as far as I could tell. Add some text to the paper to describe this figure and to help the reader understand why you included it.

Figure 6 caption. Replace the last sentence with "In the iteration 20 simulation, potential temperature in these regions become warmer as mCDW intrusion into the ice shelf cavities in the AS are correctly represented."

Figure 8. This figure would be more impactful if you could plot the observed values in a subplot or two, instead of referring the reader to another figure. If this is possible, consider adding observational subplots for better comparison.

Figure 10. Add degree symbols to the lat/lon values. Also add a label to the colorbar; this will help readers quickly understand what they are looking at.

Figure 11. It is difficult to see the difference between these plots. Consider zooming in to the region where the changes are happening; at present there is a lot of blank space and space where the changes are very small. Please also consider plotting the differences instead; that may help clarify what the changes look like.

Figure 12. In panel (b), there are two black lines. It's not obvious to me why there are two. Please make them a bit different from each other and indicate where they come from.

-- Tables --

Table 4. Consider italicizing or otherwise visually indicating which basal melt rates that you changed, for quick visual reference. Or perhaps have these rows in bolt and remove the bold from the observational column and optimized column.

Very nice work!