Interactive comment on “SICOPOLIS-AD v1: an open-source adjoint modeling framework for ice sheet simulation enabled by the algorithmic differentiation tool OpenAD” by Liz C. Logan et al.
Reviewer 2: Lizz Ultee (Referee)

We would like to thank the reviewer for her careful read of the manuscript and constructive comments. In the following, we respond comment-by-comment (our replies are in bold-face).

General comments: The manuscript by Logan et al. describes the generation of an adjoint code by algorithmic differentiation (AD) for the ice sheet model SICOPOLIS. The authors give a clear explanation of the motivation for adjoint modelling and for the use of AD in generating adjoint code, and they produce and interpret simple example applications for both Antarctica and Greenland.

I thank the authors for producing a well written and easy to follow manuscript. As a C1 non-expert in adjoint modelling, I found P3 an excellent description of its context and capabilities. I found the demonstration sensitivity analysis in Figures 2-3 interesting as well.

I am not the best person to comment on validation of the adjoint vs. finite-difference code, but I found the authors’ explanations generally supportive of their conclusions. Perhaps a bit more data (e.g. absolute value of the QoI) could be given to support the comment about numerical noise producing high misfit in Table 1, Columns 5 and 6.

I agree with the authors that understanding uncertain input variables and models’ sensitivity to them is important for contextualizing ice sheet/sea level projections. However, I disagree with the framing of the first paragraph of the introduction, namely that (A) effective adaptation to/mitigation of sea level rise relies on reducing uncertainty in projections and (B) development of more sophisticated ice sheet models will help reduce uncertainty.

Regarding (A): The social-science literature of climate adaptation discusses assorted factors that affect adaptive capacity, many of which have little to do with the state of the science. If the authors are interested, they could refer to e.g. Lemos and Rood 2010 (WIREs Climate Change) for a discussion of the "uncertainty fallacy" in climate science.

Regarding (B): It is intuitive that improved understanding of ice sheet dynamics will help us produce models that give more physically-consistent ("predictable") results. But physical consistency does not always translate to less uncertainty. For example, models that include "tipping point" dynamics (or hysteresis and multiple steady states) are arguably more sophisticated than those that do not, yet future projections over a range of climate scenarios may show a wider, not narrower, range when tipping point dynamics are included. Initial efforts to improve model sophistication by including newly-understood or newly-proposed processes can also increase uncertainty in terms of inter-model or inter-scenario spread. A notable example is the widening of 21st-22nd C2 century sea level projections shown in DeConto Pollard 2016 when the dynamics of marine ice cliff collapse and hydrofracture were included.

In a revised version, I suggest the authors strengthen the framing of the first paragraph to focus on the need for context and improved understanding, rather than leaning on the uncertainty angle.

These are really interesting points and We are glad you brought them up.
Regarding (A): Ok. Have removed the first sentence.

Regarding (B): you are correct, this work does not necessarily contribute to *reducing* uncertainties, but instead to provide a more complete characterization of uncertainties through calculation of comprehensive model sensitivities (which goes toward quantifying parametric and initial condition uncertainties). We have clarified this point in the introduction, and now omit the previously stated goal of "reducing uncertainties".

We note that the purpose of the introduction is to set the stage for the following discussion on the need to understand how *a particular* model responds to perturbations. Equipped with an adjoint model, their behaviors can be interrogated more rigorously, leading to the improved context and understanding needed to appreciate their projections. We also note that this work does not actually perform any predictions (or projections).

The revisions I suggest to the introduction and in the specific comments below are relatively minor and should not impede publication. I imagine that many ice-sheet modellers will be interested in what the authors have shown here, and I look forward to reading follow-up studies using SICOPOLIS-AD.

Specific comments:

Figures 2 and 3 - both figure captions state that the [B] subplots illustrate sensitivities to July temperature. The text on P13,L3, Table 1, and later discussion refers to January or "summer" precipitation. Is there a mistake in the figure captions, or do the figures depict something not discussed in the text?

Good catch. The captions said "July" in error. We corrected it to "January". All the tests were for summer (January for Antarctica, July for Greenland).

P4,L11 - "A model that can..." i.e. a forward model that can achieve the state deemed optimal by the Lagrange multiplier method? Does "reproduce the optimal behavior" refer to a model-vs-model or a model-vs-observation comparison?

Clarified: we meant model-vs-observation.

P5,L10 - Is the adjoint code acceptable if the finite-difference-derived sensitivities approximate the adjoint-derived sensitivities, or is it the other way round? Intuitively I would expect that we accept the adjoint code if the sensitivities it produces approximate those derived by finite difference; that is, the finite-difference sensitivities are the "standard" against which the adjoint code is judged.

In theory, it is the other way round. AD applied to discretized code produces the derivative of this implementation to very high precision, unlike finite-differencing, which depends on the order of the finite-differencing scheme and the epsilon chosen (the monograph by Griewank and Walther, 2008, which we cite, discuss this point in great detail). In practice, and for the purpose of this work, we wish to test and ascertain, that the AD tool has produced "correct" adjoint code, for which we use finite-differencing as a reference, but acknowledge that accuracy may be lower.

P5, Eq 3 - Given the explicit mention of "tolerance" in line 10, I might write the righthand side of this equation with = $\partial J \; \partial x + \delta$, where δ is the accepted tolerance.

There are two issues with writing it this way: (1) the tolerance so defined is unit-dependent (and thus may change for different physical variables of the extended control vector); (2) for very small sensitivities, the F.D. may pick up numerical noise, leading to large relative differences to the adjoint-generated derivative. Since our goal here is not to implement high-order F.D. schemes, we prefer to leave the discussion at this elementary level. Ultimately, judicious (or cautious) application of adjoint sensitivities in detailed studies should re-affirm the adjoint to be accurate (in the oceanographic context, please see e.g., Pillar et al. 2016; Smith and Heimbach 2019).

P5,L18 - This is a dense list of references without much discussion. Given that this manuscript focuses on an ice-sheet application, the study might be well-served by adding another paragraph to discuss specific distinctions among these past efforts in adjoint modelling of ice sheets and any notable contrasts with the present study.

Following your suggestion, we separated the reference, based on different applications (we still kept it brief, as the purpose of this paper is not to provide a review of the subject).

P10,L2-5 - What is the initial geometry? It is a bit unclear whether the experiment is 100 yr near equilibrium, a 100-yr spin-up, or something else. I don't know that it matters for the adjoint process, but it would be nice to have some more clarity.

Good questions. The revised version clarified this. We use the Antarctic ice sheet geometry of Fretwell et al. (2013), now clarified near the end of the paragraph.

P13, Fig 3 - I understand that the point of showing the logarithm of the absolute value of the sensitivities is so that the reader can compare their order of magnitude, both across parameters and within-parameter spatial variability. Is there a reason that each subplot uses a slightly different colormap? Could one colormap be applied to all subplots to facilitate intercomparison?

The subplot intercomparison is indeed made more difficult when the color bounds are not uniform across the panels. Unfortunately, given the large value differences between the different control runs, pattern would simply be lost. As for the blue-red vs blue-yellow color maps, these were selected with purpose: blue-red is a natural fit for positive-negative, where the absence of color (white) matches with zero. Because the other plots are meant to display overall pattern, we wished to reserve the white, zero-centered scale for the positive-negative patterns.

P17,L4-6 - Is there a use for adjoint modelling in distinguishing between different possible parameterizations of these processes?

Yes! Absolutely a use for adjoint models there, and we include this now.

Technical comments:

Done: P2,L23 - For balance, an example of a quantity that "parameterize[s] subgrid-scale processes or empirical constitutive laws" would be helpful. Perhaps iceberg calving or the routing of surface/basal meltwater would be an appropriate example to include?

Done: P2,L30 - "key quantities of interest that represent integrated quantities of an ice sheet" -> "ice-sheet-integrated quantities of interest"?

Done: P9,L22 - "...inherently non-differentiable, sometimes required..." is a comma splice. Replace comma by "and"?

Done: P16,L11 - "Thus ... within the main trunk of SICOPOLIS" is a very long sentence and I had to read it multiple times to understand. Consider streamlining.

Done: [re-phrased] P16,L18 - The phrase in the parentheses is hard to parse and might be a run-on. The authors might consider replacing the remarks in parentheses with another full sentence or two to flesh out the thought.

Done: [good observation] P20,L1 - "Precipitation" or the "sensitivity to precipitation" is almost entirely positive?

Done: P20, L3 - Typo "qsimulations"

References cited:

--

Pillar, H. R., Heimbach, P., Johnson, H. L., & Marshall, D. P. (2016). Dynamical Attribution of Recent Variability in Atlantic Overturning. Journal of Climate, 29(9), 3339–3352. http://doi.org/10.1175/JCLI-D-15-0727.1

Smith, T., & Heimbach, P. (2019). Atmospheric Origins of Variability in the South Atlantic Meridional Overturning Circulation. Journal of Climate, 32(5), 1483–1500. http://doi.org/10.1175/JCLI-D-18-0311.1