Review for "Improved basal drag of the West Antarctic Ice Sheet from L-curve analysis of inverse models utilizing subglacial hydrology simulations" by Höyns et al.

In this manuscript, the authors analyse ensembles of basal drag inversions with a range of different values for their regularisation parameter, using six different variants of sliding laws; Weertman and Budd sliding with m=1 and m=3, with two different formulations of effective pressure within the Budd sliding law, N_{op} and N_{CUAS}. For each sliding law, an L-curve is produced and the authors identify a "best" value for their regularisation, along with maximum and minimum acceptable values. They explore various methods by which to improve the appearance of their L-curves, and carry out a sub-domain L-curve analysis which reveals differences in optimal regularisation values for areas of the domain with different glaciological settings.

Comparisons are made between the outputs produced by the "best" regularisation values in each case, and also between the maximum and minimum values in the authors' preferred case of non-linear Budd sliding using N_{CUAS}, for which differences in the structure of the output are discussed. Finally the authors attempt to validate their preferred basal friction field by drawing comparisons with previous studies of bed structures and potential subglacial lake positions.

General comments:

There are a lot of ideas presented in this manuscript, some of which are more convincing to me than others. The model setup and majority of the methodology is thoroughly explained, and entirely appropriate to the topics being addressed.

The main highlight of this manuscript for me was the idea of sub-domain L-curve analysis. The idea that different regularisation may be needed in areas which contain fundamentally different physical settings, and therefore major differences in the ice flow, is an important one. I would have been interested in seeing far more focus on this aspect of the work. Comparisons in those sub-domains of the outputs using the locally-defined best regularisation against using the globally-defined value would have been illuminating, and I think this very promising avenue should be explored more in the future.

Another highlight was the results displayed in Fig.15, showing very clearly why it is important to consider regularisation carefully, and the difference it can make in a model.

In general, I felt there was too much focus on the fine details of L-curve shapes without much explanation of why readers should care about this. The details of the L-curves and convergence are likely to be specific to ISSM, while other models have different inversion processes, and different regularisation parameters. Some models require choices for multiple parameters, which require multi-dimensional "L-surfaces", and would require entirely different analysis. To my mind, the more interesting and useful comparisons for a wider audience outside of ISSM modellers were made when it was shown what the effect of different regularisation values was on the actual outputs of the inversion (ie. basal traction fields). I wasn't convinced of the importance of how smooth the curves were from which the values came, or the statements being made regarding the comparative shapes of the curves.

I was also not particularly convinced by the comparison with proposed subglacial lake positions. Partly because it is not actually known whether lakes really do exist in these locations, but mostly that I do not see much meaningful correlation between the basal drag fields and the lake outlines in the figures.

This manuscript appears to rely a lot on reference to Wolovick et al. (2023), but I think the most important parts of methodology should be restated here. In particular, the procedure for picking λ_{best}, λ_{min} and λ_{max}, which are crucial to the results of this manuscript, is buried in an appendix in Wolovick et al. and should be stated clearly to help readers understand the L-curve figures.

Overall, I believe that this manuscript contains some very interesting and important work, but that it is somewhat hidden amongst a wide variety of ideas which could be organised in a better way, and some of which seem specific to the model being used. I think this manuscript could benefit from being edited down to a more concise form, and that some restructuring could benefit its readability.

I would recommend major revisions before this manuscript is accepted for publication.

<u>Specific comments:</u>

<u>Abstract</u>

Lines 10-12: It's a bit unclear what "best" and "worst" mean in terms of L-curve behaviour. Perhaps saying Pine Island produces the smoothest curves would be a better way of wording this?

Lines 16-17: I don't agree with using "more accurate" here. The choice of parameter, as the authors point out, is a balance between matching and regularising the observations. Accuracy could be interpreted as getting the closest fit to observations, which isn't the point being made. The regularisation parameter itself cannot be defined as being accurate or not, as there is no physical real-world comparison to make.

Line 18: “improved performance” in terms of what? Convergence, smoothness of L-curves? This should be specified.

Line 20: Should this be a comma rather than two separate sentences?

Section 1. Introduction

Lines 35-36: Could you expand on what exactly is meant by “the majority of high velocities are cause by sliding”, or provide a reference? Is this just saying that sliding causes higher velocities than flow driven by surface gradients, or is it claiming that changes to basal sliding have a larger influence than other processes such as mass balance?

Line 43: I’d say “represent” rather than “compute”. The parameters from inversions broadly represent a combination of several factors (as is pointed out in the next paragraph), rather than being a value which physically describes one thing.

Lines 87-89: This reads as if the Budd law doesn’t account effective pressure. Should be reworded for clarity.

Lines 98-100: Isn’t the inverse problem in glaciological models always ill-posed, regardless of sliding laws or domains? Could this sentence be explain further?

Section 2. Method

Some restructuring in this section could be beneficial, as there appears to be some repetition, and jumping between topics.

Lines 108-115: I don’t think this section is necessary. ISSM can be introduced at the start of 2.1 instead, and other parts stated in the relevant sections.

Lines 134-139: This could be moved into 2.3, with the rest of the inversion description.

Lines 166-176: The discussion of B fields seems to interrupt the description of the forward model. It could be better placed in the subglacial hydrology section as it is being discussed in relation to CUAS.

Lines 189-194: Why not introduce both effective pressure parameterisations in the same section (ie. move this to the next section).

Lines 233-236: These two explanatory sentences belong in the introduction (in fact, I think they restate something already in the introduction)

Line 243: Given how important the values of λ_{best}, λ_{min} and λ_{max} are for this work, I think the method should be shown here rather than just a reference.

Lines 276-280: Should the detail about the forward model solver be under the Forward Model section?

Section 3. Results

Lines 304-305: Why not just use a full range of [10^{-3} 10^{4}] for all L-curves?

Line 313: How are outliers identified? I assume these are points which lie over a certain distance from the smoothed tradeoff curve, in which case this should be specified. Or are they simply the inversions which struggled to reach convergence? In particular, for Fig. 8(d,g,h) it isn't clear why some of the outliers shown shouldn't be part of the curve.

Lines 318-329: Was the same smoothing of k_{init} and the same values of ε_{gttol} and Δ_{xmin} applied for all inversions for all sliding laws? Given that these can have an important impact of the results I assume a like-for-like comparison has been conducted, but it could read as if different smoothing was used in different cases.

Lines 335-342: As a more general point, is "best" the right word to use? The choice, as is always the case with L-curves, is quite subjective. I agree that the identified corner regions are a good guide for picking your regularisation parameter, and appreciate the more rigorous methodology behind picking a value rather than doing so by eye, but I would assume any pick within this region would be reasonable. Specifically, as noted here, in some cases the identified λ_{best} is very close to λ_{min}, and perhaps another argument could be made to say that a value halfway along the curve between λ_{max} and λ_{max} could be the best choice. Do you have any insight into what variability the choice within the corner region causes in the output? It would be interesting to see the differences between your λ_{best} and a value one might pick by eye, to give more of an idea of how much this matters.

Lines 343-345: Is it necessarily to be expected that λ values would be close to each other? What were the differences in the other inversions mentioned here which display larger variability, and why should that make them less trustworthy?

Lines 346-354 & Fig.7: The convergence could be specific to the particular inversion process of ISSM with the chosen optimisation algorithm of M1QN3. Appendix B1 of Barnes et al. (2021) shows improved performance (greater minimisation of the cost function) achieved using an interior point algorithm (Byrd et al., 1999). Do you have any thoughts on whether you might find similar differences in convergence between your cases using a different algorithm such as this?

Line 355: I think this section should be 3.2. It contains enough to stand alone. I like the idea of running inversions on subdomains to find whether different regularisation could be required under different physical conditions. I think the values of λ_{best} should be highlighted on Fig.8 as they are in Fig.6, to show the difference more clearly. I would personally also be very interested to see what difference would be seen in the inversion outputs going into a forward run, comparing a spatially-varying λ_{best} to the global value.

Fig 8: This figure, and some that follow, could be made clearer by highlighting the corner region in a different colour, rather than the thicker line currently used. I assume the circles with white in the middle are outliers, but this is not labelled on the figure. Some of the L-curves presented in this figure could benefit from an extended range of values, as the corners appear very close to the end of the intervals used.

Figs 9,10,12: The diamonds showing location of λ_{best} is not clear. The values of λ_{best} could also be shown.

Lines 464-471: I'm not sure I follow the reasoning for using the same λ_{best} in both cases. The whole point in the inversion process is surely to pick an optimal regularisation for each case, in order that the output is most suited to the individual problem. I would think it was more instructive to compare the outputs that come from each experiment's λ_{best} value. Also, the numbers given do not correspond to the λ_{best} values stated on Fig.6. If these are coming from different L-curves it should be made clearer.

Fig.11: Why not also show the Weertman output?

Lines 519-520: Could you explain how the chosen case is the "best estimate" of those on Fig.6?

Fig.15: This is a great figure for showing why regularisation choices are important, and the difference they can make!

<u>Section 4. Discussion</u>

Line 558: It may be desirable from a practical perspective, but for better results maybe picking a few different regularisation values based on ice speed or other obvious differences in the physical setting would be the way to go. I find this idea very interesting.

Figs 17,18: I don't personally find these to be particularly convincing. While there are a couple of places where the basal drag field lines up with the white lines, there are also places where the outlines cross areas of higher friction. What are the locations of these "possible" lakes based on? Is there a correlation with their positions in the effective pressure fields N_{op} and N_{CUAS}? Is similar correlation seen for the Weertman inversion output as the Budd ones?

<u>References</u>

Barnes, J.M., Dias dos Santos, T., Goldberg, D., Gudmundsson, G.H., Morlighem, M. and De Rydt, J., 2021. The transferability of adjoint inversion products between different ice flow models. *The Cryosphere*, *15*(4), pp.1975-2000.

Byrd, R.H., Hribar, M.E. and Nocedal, J., 1999. An interior point algorithm for large-scale nonlinear programming. *SIAM Journal on Optimization*, *9*(4), pp.877-900.

Wolovick, M., Humbert, A., Kleiner, T. and Rückamp, M., 2023. Regularization and L-curves in ice sheet inverse models: a case study in the Filchner–Ronne catchment. *The Cryosphere*, *17*(12), pp.5027-5060.