Review of Frank et al. (2023): Reconciling ice dynamics and bed topography with a versatile and fast ice thickness inversion

Summary

This paper presents a new method for inverting for ice thickness and, optionally, basal friction at a glacier by iteratively minimising misfit between observed and modelled dh/dt through adjusting an initial guess at the bed (to which the method is insensitive), and by adopting the same approach with observed and modelled surface velocity for the basal friction field. The authors test their method on a synthetic ice cap in order to quantify its performance and sensitivity, and then apply it to Kronebreen, a tidewater glacier in Svalbard, finding that the method generally reproduces the bed well, and is fairly robust when subjected to a wide range of model parameters and input data perturbations. They also show that their method allows the creation of a self-consistent model glacier that can then be used in prognostic simulations without any relaxation.

I think this is a good paper that makes a useful contribution to the field of glacier thickness inversion, but I have a few major concerns around the general framing and structure of the paper that I feel need to be addressed before I can recommend it for publication. I do not think any of these are particularly difficult to solve, however, and I think the authors should be congratulated on coming up with an ingenious new solution to a difficult conundrum!

Samuel

Major points

- The title and abstract of the paper major on this being a fast method for inverting for thickness, but I think the paper doesn't necessarily back that up.
    - An easily fixable issue here is that the paper never actually says how fast the method is, beyond saying it's 'fast'. The authors present results in terms of the number of iterations they took to achieve, but never say how long it takes to run an iteration. So, somewhere, I think the paper has to provide an actual time per iteration or per experiment to allow the reader to judge how fast the method is
    - I feel it's also difficult to say the method is fast when its speed depends on the user's choice of ice-flow model. Presumably, if you were using a full-Stokes model, the method would be considerably slower?
    - I also think that the speed of the method is not really the important thing about it. I would say the real advantage of the method is that you end up with a self-consistent model glacier that you can stick into a prognostic simulation without having to do any relaxation. After all, you can invert for thickness very quickly if you just assume perfect plasticity and have a DEM. The resulting glacier will be a horrible mess, but it will be *fast*. What we don't have so much yet is a method that shortcuts all that relaxation running time, and that's where this method really innovates, I think.
    - So, overall, I would strongly encourage the authors to reframe the paper towards the 'self-consistent model glacier' aspect and away from the 'fast' aspect. I think changing the title and the abstract would go a long way towards doing that – the discussion already points this way, so I really think this is more-or-less a cosmetic front-end change than anything else. Though I think the authors really do still need to give somewhere an idea of how long the method takes to run in actual time, rather than just giving a number of iterations, as that's a useful comparator with other methods, regardless of any claims the authors do or don't make about how fast the method is.
- Section 3.2 (on model sensitivity) needs some rewriting. I'm very glad the authors tested the sensitivity of the model to both input errors and model parameters – that's a really useful thing to do – but I think this section needs quite a bit more work to make it really helpful for the reader. As it stands, it's much too vague and imprecise to provide a clear picture of model sensitivity – I came away from it feeling I didn't really have much more of a clue as to how sensitive the model really was compared to when I started reading it. I know the authors give a lot of the details in the

Appendix, but I think there are some headline statistics the authors could add to the main body of the text to make this section useful without having to make the reader dig down into the appendix.

  - In particular, in Section 3.2.1, it would be helpful if the authors could provide some quantification of how much the perturbations to the input data impacted on the final bed/friction field, beyond just saying 'large bed responses', which is a bit vague – 10 m, 50 m, 100 m?. The authors state that the ice volume is pretty insensitive to all the perturbations (though I would like some idea of what the percentage change in volume across all the simulations is), and helpfully explain all the mechanisms by which the errors feed through into glaciological changes, but I didn't really get any sense from the section as written what the magnitude of those glaciological changes were. As it stands, the section leaves open the possibility that it could well be that the perturbations don't affect the overall ice volume, because they lead to the bed being 500 m too high in half the domain and 500 m too low in the other half. I would also want to know how much the authors had to increase theta to mitigate the issues – 1.5x, 3x, 5x, 10x? - and I think there needs to be a bit more detail around 'more dependent on initial conditions'. The best solution is perhaps a table showing the mean bed and dh/dt mismatch for each simulation (or some other metric of the model's performance), and then maybe a figure showing some of the final beds to show what 'more dependent on initial conditions' actually means.
  - For Section 3.2.2, I would similarly like some idea of the range of values of beta and theta the authors tested (the authors give the range of values for T, for example, in Section 3.2.1; the same should be done here). The authors repeat their previous observation about larger theta leading to greater dependence on initial conditions here, with a little more detail as to why, but I'd still like to get some sort of sense of what that looks like or means in practice without having to dig through the appendix. If the authors do put together some sort of figure to show that, I think it fits better here than in 3.2.1, as this is the section about explicitly testing theta (in which case, the authors could just stick a 'see below' in Section 3.2.1 when they mention that larger theta leads to greater dependency).

- As a more general point, I'm wondering what would happen if the authors updated friction first, then the bed. Would the method converge towards the same friction/bed combination, or would you get something different? Fundamentally, how do the authors deal with the problem of equifinality (as in, there are lots of possible friction/bed combinations that would reproduce the observed variables – how do the authors know theirs is a sensible one?)? Some discussion of this would be a useful addition, either in Section 5, or elsewhere where the paper discusses the robustness of the model. One additional simulation changing the order of the friction and bed updates would, I think, be very instructive and I would encourage the authors to see if it makes a difference.

Minor points

- p.1, l.14: 'fieldwork' not 'field work'
- p.2, l.24: 'allow the derivation of' not 'allow to derive'
- p.2, l.30: 'former' not 'previous'
- p.2, l.47: Some examples of where the conditions/assumptions in the SIA aren't met (sliding glaciers…) would be good here - here the text assumes that people will know and don't explain, unless the reader goes back a few lines and works out that 'considering only internal shear' means 'not valid where there's something else going on'. It just makes the reader's life slightly easier.
- p.2, l.51: The text doesn't quite explain explicitly why using beds derived using standard thickness-inversion processes will require you to do model relaxation if you use them as inputs in a higher-order model. The text states previously that you'll get errors in the ice thickness using one of the standard approaches, but why do those errors then require model relaxation (the velocity and thickness fields won't match so there'll be artefacts so...)? It might also be worth explaining what model relaxation actually is. As with the previous point, the text is perhaps assuming slightly too much knowledge on the part of the reader and making them work harder than they perhaps should in the introduction!
- p.2, l.55: 'the present day'
- p.3, l.57: 'allows the inclusion of' not 'allows to include'
- p.3, l.80: 'way' not 'part'
- p.3, l.82-3: How well do we actually know the mass balance at most glaciers? This approach assumes the surface variables are well-known (usually true for the surface topography) and that the

errors in the model are therefore entirely due to bed errors, but I'm not convinced that assumption is valid at most glaciers. I realise this is dealt with later on, but some acknowledgement here that it is an assumption that isn't always going to be true would be good.

- p.4, l.89: 'point out' Habermann et al is plural.
- p.5, l.126: 'that gently induce'
- p.5, l.127: 'that it cannot otherwise reproduce'
- p.5, l.132: 'that could lead' I think I'll stop pointing these out now, but the general rule is that 'which' is used when introducing a new clause, otherwise use 'that'. Pretty much, if you've stuck a comma in, use 'which', if there's no comma, use 'that'.
- p.9, l.219: 'error' not 'errors'
- p.9, l.220: 'lead to similar outcomes' not 'induce similar mechanisms'
- p.10, l.252: 'the subject'
- p.13, l.332: 'it can be higher' – you're talking about the mismatch (singular). Linked to this, how much higher is 'higher' here?
- Section 4.5: I'm not entirely convinced by the post-processing step here, as it seems a little counter-intuitive to me to use an equation based on the SIA to correct errors caused by using the SIA. I don't deny that it works, I just don't like it conceptually. To be fair, I'm not sure what else could be done that would be reasonable, but I suspect readers might react in the same way as me, so if the authors can just add a sentence acknowledging that using the thing that caused the problems to fix the problems is what's going on and that the reader's possible confusion is understandable, it might make it a slightly easier pill to swallow.
- p.17, l.418: 'such as' not 'such that'
- Section 5: I think the approach presented in this paper could be usefully applied to the ITMIX dataset to benchmark it against other widely-used inversion methods. Are the authors considering doing this? Either way, a statement on yes/no and why/why not towards the end of the section would be good, because it's otherwise a bit of an elephant in the room