Author's response to RC1

Note: *the referee's comments are shown in blue, while the authors' responses are shown in black.*

This study evaluates a two step assimilation method to create snowpack simulations operationally. The topic is highly relevant to the journal. Congratulations to the authors for this strong piece of work. The paper is very interesting and well-written. I just have a few minor comments and questions listed below.

Thank you for reviewing our manuscript and offering constructive feedback. Addressing your comments will definitely enhance the quality of the study. Please find our detailed responses below.

L148: I don't understand how you upscale the precipitation multiplier from 25m resolution to 250m.

The precipitation multiplier at the model resolution is obtained by first computing it for each 25 m (sub-)grid cell of the digital elevation model (based on its respective slope $\beta_{25\,\mathrm{m},i}$), and then averaging these values over all 100 (sub-)grid cells contained within a 250 m grid cell:

$$\mathrm{pm}_{250\,\mathrm{m}} = \frac{1}{100} \sum_{i=1}^{100} \mathrm{pm}(\beta_{25\,\mathrm{m},i})$$

We will revise this sentence in the manuscript accordingly.

L205: I get that the choice of lower and upper bounds is already explained in another paper. However the rest of the section is very well explained and it feels weird not to get why the bounds are chosen. 20 cm in particular seems to be a small snow depth for an upper bound. I suggest you add a sentence to explain this range.

The bounds do not refer to the snow depth itself but to the observation uncertainty, or more precisely the observation error standard deviation σ_{obs}, which scales the difference between observed and simulated snow depth when computing the Gaussian weights (see Equation 5 in the manuscript). Setting the observation error standard deviation to 5 % of the observed snow depth balances the need for close agreement with the local observations against the representativeness error of point measurements. The bounds ensure a reasonable minimum observation error standard deviation of 5 cm for shallow snowpacks (snow depths below 1 m) and cap it at 20 cm for snow depths above 4 m, which are hardly ever observed at our stations. As you suggested, we will add this explanation to the revised manuscript.

L225-227: Either you already explained that or this sentence is unclear to me. Maybe precise which assimilation step you are talking about.

We will revise this paragraph to make clear that it refers to the second assimilation step, which adopts some design choices from the first assimilation step.

L235-237: Does this mean that the 36 particles of the prior simulations in the flat grid cells are the same? So after step 1, you "lose" the uncertainty of the posterior simulation and recreate no perturbations in all the flat grid cells of your domain?

Yes, that is correct. The first assimilation step deliberately collapses the probabilistic information onto a single set of forcing corrections for each station and assimilation window, which facilitates the spatial interpolation. The parameter perturbations of the second assimilation step are deliberately designed to default to the zero (i.e., unperturbed) member in flat terrain, so that they do not counteract the forcing corrections inferred at the flat-field station locations, keeping the two assimilation steps entirely independent of one another. We acknowledge that leaving flat grid cells unperturbed during the second assimilation step may be seen as a limitation of the current setup, even though only very few 250 m grid cells are entirely flat. We will state this explicitly in the revised manuscript and discuss possible extensions in the Discussion section (see also our response to the second referee, RC2, comment on Figure 3).

L240-250: I think it is very interesting for an operational point of view and rarely pointed at so I definitely think that you should keep this comment. But it could be misleading for some readers. With every new observation taken into account in the assimilation process of the PBS, the posterior simulation is updated over the whole period and propagates backward in time (and potentially increases in accuracy). My point is that at the end of June the simulated snow depth in March might not be the same as the simulation made in March with only the observation from January to March because you do not have the same list of $SCF^t_{obs} - SCF^t_{sim}/\sigma_{SCF}$. Therefore this sentence in particular could be misleading: "This equivalence implies that the presented results would be unchanged if the parameters were updated sequentially as new observations become available, which is how the method would be applied in an operational nowcasting context, as discussed in Section 4."

Thank you for pointing out this potential source of confusion. As you correctly stated, the equivalence between a particle batch smoother (PBS) and a particle filter (PF) without resampling only holds for the posterior at the end of a given assimilation window. Since the weights change with every new observation, the simulation for any given point in time produced with the observations available up to that point can generally differ from the corresponding state of the end-of-season posterior. What the equivalence does guarantee, though, is that sequentially updating the weights with each new observation, rather than waiting until the end of the season, yields the same final parameter estimates as presented here. We will rephrase the quoted sentence accordingly and make this distinction explicit in the discussion of the operational setup in Section 4.

L250-260: Add this step somehow in Figure 2, or add another figure in section 2.6 to show the whole method of the operational setup. I really believe the paper would benefit from showing the whole operational method at once.

Following your suggestion, we will add the interpolation/propagation step to Figure 2 to outline the full operational workflow.

L260-280: This question is linked to the one l 235-237 because I don't fully understand. With the first assimilation step you correct the meteorological forcings. However, I guess that you do not have snow depth stations in the highest elevations of the catchments and therefore the meteorological forcings at high elevations could be biased by the lower elevation parts. I do not get why the precipitation multiplier of the flat grid cells needs to be one? It could correct biases that were not corrected in the first assimilation step. I guess it is linked to the "goal of correcting only non-systematic, small-scale error" L368 but I do not get the rationale behind the goal.

Thank you for identifying this gap in our explanation. The rationale is as follows: in the first step, we correct the solid precipitation input at the flat-field station locations. Spatially explicit snow depth observations, on which the empirical relationship of the precipitation multiplier is based (Vögeli et al., 2016), show that snow depth on slopes is typically less than snow depth in flat terrain. Hence, the precipitation multiplier equals one in flat terrain and accounts only for the reduction of snow in sloped terrain, assuming the first step correctly adjusts the total amount of solid precipitation (as measured at flat locations). However, you are right that this design cannot correct residual biases in snow accumulation not captured in the first step, e.g., at higher elevations not covered by stations. Future work could therefore pursue a joint (single-step) assimilation of both observation types, using snow cover fraction (SCF) observations to also adjust the precipitation input for all grid cells, even though the limited information content of SCF observations remains a general constraint. We missed the opportunity to address this in the Discussion and will do so in the revised manuscript.

L332-335: If the LOO method also included a LOO snow depths data in the first assimilation step, could you have evaluated the results in terms of snow depths as well? I am not saying to redo the whole experiments but perhaps you could discuss that these evaluations only take into account the SCF while the end goal of such operational simulations often include access to the SWE, snow depth, etc.

Thank you for raising this point. We will address it in the revised manuscript in two ways: (i) by stating more explicitly wherever we evaluate against SCF observations, and (ii) by adding a section in which we evaluate the simulated snow depth for a small domain where we have photogrammetric snow depth observations. For more details, including the corresponding figure, please see our response to the second referee (RC2, comment on L321).

L367: You could add the correlations in supplementary or appendix.

Figure A shows the correlations between model parameter perturbations and various topographic variables for the 2022/23 season. As the correlations are generally weak, we feel that the figure adds little to the manuscript and would prefer not

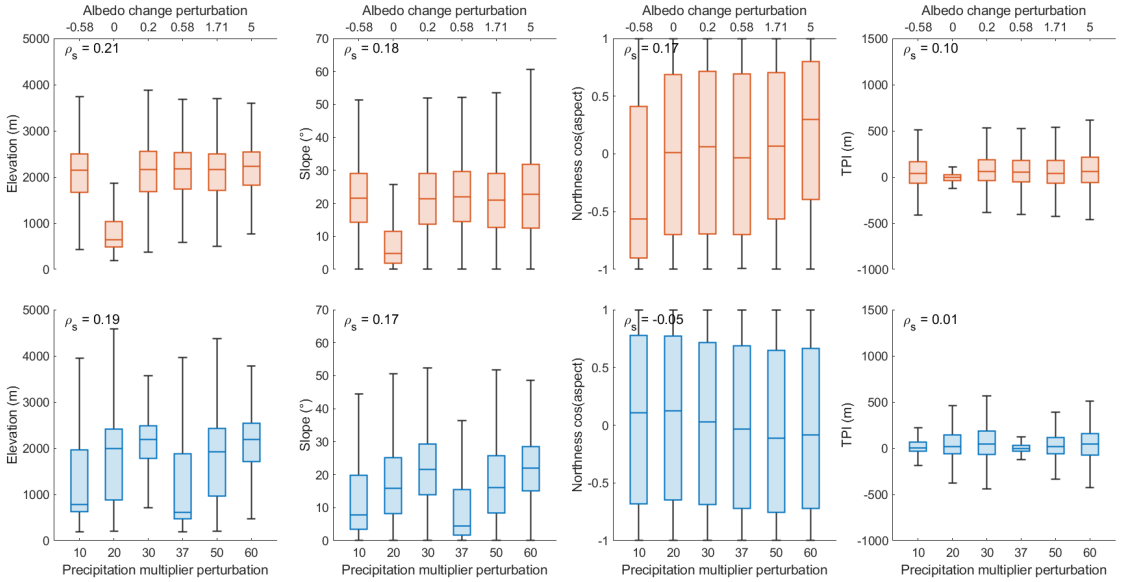


Figure A: Correlations between grid-cell-specific albedo decay perturbations (top row) and precipitation multiplier perturbations (bottom row) and elevation, slope, northness (i.e., cosine of the aspect) and topographic position index (TPI) for the 2022/23 season. The Spearman correlation coefficient ρ_S is given in the top right corner of each subpanel.

to include it. It would be interesting to further investigate topographic and/or climatological relationships at smaller scales to potentially improve the spatial interpolation, but this is beyond the scope of this study.

Fig 4 & 6: The forests are not taken into account in your analysis and I think that the green of the forest is too visible compared to the actual evaluation. Perhaps put the forest in white when you are showing the difference of SCF? In general, in the places where you didn't do the evaluation (glaciers, water etc.) I think that the grey (of what I guess is no data) is too close to the grey of the 0% difference in SCF which makes it hard to fully understand the figures.

We agree that the color schemes in Figures 4 and 6 are not yet optimal. We will follow your suggestion and mask the forest in the difference plots, and we will choose a "no data" color that is clearly distinguishable from the color of 0 % SCF difference (white is not suitable, as it is already part of the current colormaps).

L400-405: Have you tried other SCF parameterizations?

No, we have not tried other SCF parameterizations in the presented assimilation study. The SCF observations could be used to evaluate the SCF parameterization used operationally, but this is beyond the scope of this work. A recent study by

Sourp et al. (2026) compares different SCF parameterizations for snow reanalysis, which we will discuss in the revised Discussion section.

L500: I don't get why you end with this sentence when you said L450-455 that it is not a solution in your setup.

The respective sentence refers to a sequential PBS, in which the weights (and thereby the best parameter estimates) are updated with each new Sentinel-2 acquisition. In contrast, L450–455 discuss multiple fixed PBS assimilation windows within one season, such as one for the accumulation and one for the ablation period, which we do not expect to be beneficial. Thank you for pointing out this potential confusion. We will revise both passages to make the distinction clear.

References

Sourp, L., Pedinotti, V., Alonso-González, E., Jarlan, L., and Gascoin, S. (2026). Assessment of Snow Cover Fraction Parameterizations for High Resolution Snowpack Reanalyses. *Hydrological Processes*, 40(4).

Vögeli, C., Lehning, M., Wever, N., and Bavay, M. (2016). Scaling Precipitation Input to Spatially Distributed Hydrological Models by Measured Snow Distribution. *Frontiers in Earth Science*, 4:108.