Review of Wagner et al. (2026) 10.5194/egusphere-2026-4575

General comments

This manuscript by Wagner et al. presents openAMUNDSEN-DA v0.9, a particle filter (PF) workflow coupled with the open source snow model openAMUNDSEN, supporting the assimilation of station snow depth and SWE, as well as aggregated satellite-based fSCA, WSF, and WSLA, as illustrated in a case study from the Rofental catchment.

There are many strong aspects of this contribution. The workflow is open source, containerized, archived on Zenodo, and configured through transparent files with keyed random-number streams, making the experiments genuinely reproducible. The figures are generally very clear, and the assimilation of Sentinel-1 wet snow information is an exciting and still largely untapped direction for snow DA. Overall, coupling a PF to an advanced snow model like openAMUNDSEN is a promising contribution to GMD. That said, we have three general concerns.

1. The framing of the contribution. The manuscript seeks to present a new ensemble-based snow DA framework that complements existing ones, while what is really presented is a PF workflow for one snow model, implemented with one particular and quite restrictive variant of the PF. The umbrella term *"ensemble-based DA framework"*, used in the title, typically denotes a software suite (e.g. PDAF; Chen et al., 2025) that offers several ensemble methods from which the user can select (e.g., ensemble Kalman filters and smoothers as well as particle methods), whereas openAMUNDSEN-DA implements only one. Moreover, the limitations attributed to existing snow DA frameworks apply equally to openAMUNDSEN-DA. We suggest that the title, abstract, and discussion describe the contribution as it is, namely a PF workflow, and focus on what this workflow enables rather than on what other frameworks are claimed to lack. Related to this, redrawing the forcing perturbations from the prior at each window means that the scheme does not actually filter them in the forcing formulation (see specific comment on L139).

2. Use of spatially constant parameters and aggregated observations. Updating domain-wide perturbation parameters using domain-aggregated satellite retrievals is a major simplifying assumption that discards most of the spatial information in the observations, and it should be stated upfront. Your own results illustrate the consequences: the persistent and growing offset between the station residuals in Fig. 5, and the mismatched local wet snow patterns in Fig. 6. Spatially distributed or semi-distributed parameters would retain much more of this information.

3. The evaluation of the case study. Computing skill scores relative to a deterministic open loop with degraded forcing makes the improvement appear larger than it is, whereas scores relative to the prior would isolate the effect of each update. The case study also rests on a single season, eight hand-picked updates, and a single random seed. We understand that the aim is demonstration rather than validation, but the choice of updates should be justified and some robustness shown.

In summary, this manuscript presents a promising workflow that is clearly suited for eventual publication in GMD. We recommend major revisions to address the general comments above and the specific comments below.

Specific comments

1. Abstract: It should be made clear that your workflow updates domain-wide, rather than spatially distributed, parameters using highly aggregated (rather than pixel or grid cell level) satellite retrievals.

2. L1: Change *"into physically based snow modeling"* to *"into physically-based snow models"*.

3. Title & L3: We encourage you to be more specific. Your contribution is not a general ensemble-based snow data assimilation framework, but rather a particle filtering workflow for openAMUNDSEN. The current framing overstates the contribution given that you only use one particular scheme, a particle filter, instantiated in a very particular way.

4. L25: Consider providing some examples of satellite-based snow products.

5. L26: Not sure how *"However,"* fits here? Consider *"To get the best of both worlds,"* or similar. You also do not make it clear why this wealth of data is not all we need in snow hydrology (it usually isn't).

6. L32: Change *"Kalman filter"* to *"ensemble Kalman filter"* since very few, if any, snow data assimilation studies use the classical Kalman filter which requires linear models.

7. L40: The book by Chopin and Papaspiliopoulos (2020) is a great reference for particle methods, but it is not about snow so it does not fit in this sentence.

8. L41: Here you should add van Leeuwen et al. (2019), which is specifically about high-dimensional geophysical systems, while changing *"Van"* to *"van"* in your reference to van Leeuwen (2009) for consistency. Note also that resampling only addresses weight degeneracy at the cost of path degeneracy (sample impoverishment) (Murphy, 2023), and does not prevent weight collapse in high dimensions (van Leeuwen et al., 2019).

9. L48: *"organized around certain snow models"* while this is often true, this limitation applies equally to your contribution that is built around openAMUNDSEN. Moreover, the MuSA framework has recently been extended beyond the widely used FSM2 model to also include simpler temperature index snow models (Aalstad et al., 2026), with more complex models in the pipeline. Once more, there is plenty of room for your contribution, but one should be careful trying to criticize existing frameworks while not making it clear that those limitations apply equally to your own workflow.

10. L49: The limitations have not been reviewed, unless being designed for particular applications is supposed to be a limitation. There are certainly limitations that are worth reviewing, so we suggest that you either provide an overview of the limitations of existing frameworks or focus on your contribution.

11. L51: Here too, you make it seem as if existing schemes like CrocO, FSM2OSHD, and MuSA fail to provide such a workflow, but then how are these used effectively either operationally or for research purposes?

12. L58: Please revise this sentence. You are implicitly perpetuating the common misconception that physically consistent data assimilation in (presumably multilayer) complex snow models requires the use of particle methods and precludes the use of ensemble Kalman methods. This is simply not true, as demonstrated in many applications of ensemble Kalman methods to multilayer models like FSM2 (e.g. Mazzolini et al., 2025; Alonso-González et al., 2026). What actually matters for consistency is not whether or not you use ensemble Kalman or particle methods, but whether you use the forcing or model-state formulation of the data assimilation problem (see Evensen et al., 2022) as pointed out in Alonso-González et al. (2022). There are certainly good reasons for using particle methods, such as the lack of a linearity assumption and the use of non Gaussian priors or likelihoods, but consistency is not a good reason.

13. L60: *"sequential PF"*, here and throughout the paper you do not need to add sequential in front of PF. The filter (F) is already doing that work, so the sequential label is redundant. Where you introduce the theory you could mention that the particle filter is a sequential data assimilation scheme, otherwise just say *"PF"*.

14. L78: The fact that you are assimilating domain-scale (rather than pixel or grid cell level) satellite retrievals to update spatially constant forcing perturbation parameters should come much earlier than this. While this choice is partly understandable from the point of view of particle degeneracy, it is an unusual choice as well as a major simplifying assumption and source of information loss in your current workflow that should be stated up front earlier, ideally both in the abstract and at the end of the introduction.

15. L92: It is still not clear how openAMUNDSEN-DA complements or adds additional approaches to the existing frameworks you mention, except for (at least algorithmically) being a simplification of these applied to a different model. The fact that you built a particle filter workflow for openAMUNDSEN is great and a worthwhile contribution on its own merit. Why do you not focus on that instead of how you add to existing frameworks, which it is not really clear that you do at present.

16. L136: The way you use the term *"event"* (and derived terms such as 'event weights' and 'event analysis') is non-standard in data assimilation and conflicts with the established meaning of 'event' in probability theory. Since the manuscript already uses 'step' for the same concept in the project hierarchy (Fig. 2), consider using 'assimilation step', 'update' or 'analysis' consistently throughout.

17. L137: As with *"assimilation event"*, here (and elsewhere) you use a non-standard term *"materialized ensemble"*. Change this to *"resampled ensemble"*, which is the standard term for exactly what you are doing here.

18. L139: This is not rejuvenation of the posterior approximation but a complete redrawing of the forcing perturbations from an unconstrained prior. In SMC, particle diversity after resampling is *"rejuvenated"* by moving the resampled particles while preserving the target distribution, e.g. via MCMC moves as in the resample-move algorithm (Gilks and Berzuini, 2001), or approximately by jittering (Farchi and Bocquet, 2018). Although the snowpack state is carried forward, completely redrawing the perturbations from an unconstrained prior discards everything learned about the forcing errors from previous observations, and resets the prior on the perturbation parameters at every window. This is a missed opportunity to perform proper filtering in the forcing formulation, where the posterior perturbation parameters would inform the prior for the next window. It also makes the assumed forcing error model depend on the timing of the observations, which is hard to justify physically. A more natural approach would be to follow Magnusson et al. (2017) and give the perturbation parameters explicit dynamics in time (e.g. via jittering), which would provide both actual rejuvenation and filtering of the perturbation parameters across, rather than independently within, windows.

19. L164: While point observations such as station snow depth are representative of a local point, they are generally not representative of the grid cell scale of a model (100 m in your example). Is this within grid cell representativeness error mentioned explicitly anywhere? Moreover, the phrasing *"how far they inform"* grid cells could mislead the reader, as it suggests that distant grid cells are less affected, whereas the domain-scale scheme does not allow this kind of localization. We suggest *"how much they inform"*.

20. L166: *"they contribute uncertainty weighted to the comparison with the model simulation"*, this does not make sense as it stands, consider rewording. If anything the observations provide information (in the sense of reducing uncertainty measured via entropy) weighted by their precision (inverse error variance).

21. L189: Here you could reformulate to *"The sequential assimilation of snow observations in openAMUNDSEN-DA is performed via a PF at the domain scale"*, then drop the redundant *"sequential"* elsewhere.

22. L192: Once more this misconception appears more explicitly, please remove this sentence on ensemble Kalman methods which is not supported by recent literature (see comment on L58). This is about the forcing versus model-state formulation, not ensemble Kalman versus particle method choice. Incidentally, averaging across incompatible snow states and layer structures would be just as much of a problem with a particle method if you wanted to compute say layer-wise posterior mean quantities.

23. L195: Consider rewording *"meteorological forcing perturbations sample uncertainty"* to *"uncertainty is encoded via meteorological forcing perturbations"* or similar. Furthermore, it seems like here you are citing Evensen et al. (2022) as an appeal to the aforementioned forcing formulation. It is worth noting though that the forcing formulation does not require a deterministic model (i.e. no model error terms), that is actually a different additional assumption called the strong constraint (perfect model) assumption. So your approach corresponds to a particular instantiation of the fairly standard strong constraint forcing formulation of the snow data assimilation problem as outlined in Alonso-González et al. (2022), but used implicitly at least as far back as Margulis et al. (2015) and probably earlier.

24. Equation 1: In your special case of the snow data assimilation problem it makes sense that the observations y_k are scalars, but the model state should surely be denoted as a vector $x_k \to \mathbf{x}_k$. Moreover, this equation is a slight simplification of the problem you are actually implicitly solving, since it is the forcing perturbations that are being updated (weighted) directly and then lead to an indirect update of the (conditionally deterministic given the forcing perturbations) model state variables like SWE and snow depth. Keeping this implicit has become kind of standard so it is fine, but it hides the fact that you are using a forcing formulation. Making the forcing formulation more explicit in your presentation would also have made it clearer that you adopt an extreme memoryless approach that resets the perturbations to an uninformed prior after each observation, as opposed to say gradually injecting noise on the current estimate, where there is considerable room for improvement.

25. L197: Change *"assimilation date k"* to *"assimilation time t_k indexed by k"* since an index is not a date.

26. L199: *"x_k denotes the model trajectory"*, if k indexes time then this would just be the state vector at the observation time, whereas the state trajectory would be $\mathbf{x}_{1:k}$?

27. L203: Change *"observation standard deviation"* to *"observation error standard deviation"* and remove the *"as Gaussian likelihood"* since this is repetition (twice in the same sentence).

28. L208: Nice that you are doing calculations in log-space and using log-sum-exp. Even if your equation does not show it, presumably this is evaluated using the standard log-sum-exp trick to avoid numerical underflow and overflow (Murphy, 2022; Chopin and Papaspiliopoulos, 2020), see Aalstad et al. (2026) for details in a snow DA context. Please state this more explicitly. Although this is unlikely to matter with few observations per update, it becomes essential when assimilating many observations jointly. This occurs both in spatio-temporal assimilation of distributed observations and for longer assimilation windows in batch smoothing, both of which might be ways to extend your framework in the future.

29. L214: What is the problem with perturbing zero nighttime shortwave radiation multiplicatively, it would remain zero? This seems like an unnecessary complication and confusing detail.

30. L219: Clarify if the Magnus equations are applied to perturbed or unperturbed air temperature, we assume the former which is more consistent? As for the equation, we would recommend citing Alduchov and Eskridge (1996) which is a standard and more specific reference for the Magnus equations, then you can also clarify more explicitly which form you are using (see their Table 1) where your constants suggest SA90.

31. L225: It is true that priors must be carefully chosen, how were they chosen for your example?

32. L230: Discussion of future work (i.e., from *"Future work will. . . "* to the end of this paragraph) does not belong in the methods but rather an outlook or similar subsection in the discussion section.

33. L235: One could argue that all forcing formulation schemes conserve the members' mass and energy consistency, especially when the strong constraint assumption (as you do here) is made.

34. L240: Please provide a reference for the ESS, e.g. Elvira et al. (2022) provides a thorough treatment including limitations, otherwise it looks like you are coining this concept here.

35. L243: Change *"low values indicate that the weights are concentrated on few members, while values near N_e indicate flatter weights"* to something like *"low values indicate degeneracy where the weights are concentrated on a few members, while values near N_e indicate a diverse ensemble with more uniform weights"* which also explains what the ESS is actually used to diagnose.

36. L244: Change *"analysis is"* to *"expectations are"*, strictly the analysis (update, inference, assimilation) step with a particle method is the weight calculation in your Eq. (4) whereas computing posterior expectations (means, variances etc...) is a post processing step. Great that you use the weights to do this directly and not the resampled ensemble, as you avoid introducing additional sampling error.

37. L247: Similarly, change *"weighted analysis"* to *"weighted ensemble"*.

38. L248: Right, this corresponds to the well known (and often quite degenerate) sequential importance sampling (Chopin and Papaspiliopoulos, 2020) algorithm that the particle batch smoother (PBS Margulis et al., 2015) is built on. Nice that you point it out, but it is not exactly a secret in the snow data assimilation literature as it has been emphasized by e.g. Alonso-González et al. (2022) and used recently by both Aalstad et al. (2026) and Oberrauch et al. (2026) to build more advanced particle-based snow data assimilation workflows. So we would suggest that you mention this once in passing, but there is no need to mention it again on L405. Your point here about SIS also raises an intriguing possibility, why did you not try to run the PBS in your setup? If you simply turned off the prior redrawing then your current implementation would be a PBS that would likely also allow you to effectively assimilate satellite data at the grid cell (rather than domain) level as done in snow reanalysis (Margulis et al., 2015).

39. L251: As mentioned, *"rejuvenation"* is already a loaded term in particle methods that means something else, something like *"redrawing"* or *"resetting"* would be a better name for what you are doing when you sporadically (at the start of each window) redraw from an unconstrained prior. A more standard state space modeling

approach and likely better way to avoid impoverishment (path degeneracy) would have been to impose continuous (i.e., every time step) dynamics on your perturbation parameters via a possibly auto-correlated stochastic noise term as in Magnusson et al. (2017).

40. L255: Resampling and process-noise are only basic necessary but by no means sufficient conditions for non degenerate particle methods in higher dimensional geophysical problems (van Leeuwen et al., 2019). Evidence for this can be seen from the fact that particle methods are still not widely used in such settings, including (with a few exceptions like the LETPF) numerical weather prediction, where ensemble Kalman methods are more ubiquitous. Avoiding degeneracy in more challenging (e.g. spatio-temporal) snow data assimilation problems will require more sophisticated particle methods including better localization (Farchi and Bocquet, 2018), tempering (Murphy, 2023), and proposal adaptation methods (Aalstad et al., 2026, and references therein).

41. L261: *"scores comparing the posterior ensemble distribution"*, strictly speaking this only holds for the CRPS-based score, the other two are based on point estimates from the ensemble.

42. Eq. (6): Using the open loop, which appears to be a deterministic reference run with purposely degraded forcing, as the reference in these metrics artificially makes your DA look better. For the CRPS this is particularly problematic, as the open loop is a degenerate distribution, so its CRPS reduces to the absolute error, which will generally be higher than that of the prior ensemble even before any assimilation. It would be more informative to express the skill scores relative to the prior within each window, e.g. $\mathrm{CRPSS} = 1 - \mathrm{CRPS}_{\mathrm{post}}/\mathrm{CRPS}_{\mathrm{prior}}$, which isolates the effect of each update without conflating it with the degraded reference run. It is also not clear why you need scores other than the CRPS-based score: what do the two deterministic scores add to this probabilistic score?

43. L265: Please provide a reference for the CRPS, such as Hersbach (2000).

44. L279: As mentioned, there is still a remaining implicit representation error between the point scale observations and the 100 m grid scale. Presumably this is absorbed into the observation error standard deviation.

45. L281: Change *"uncertainty-based weighting"* to *"precision weighting"* since the inverse error variance is known as the precision, which is a more precise term than the more general term uncertainty which can refer to other measures of spread such as entropy.

46. L282: Could you please clarify the point you are making here, other than being computationally lightweight, is your approach somehow better or more justified than a gridded (spatio-temporal) update as done by Cluzet et al. (2022) that you cite? To us it seems like too extreme a simplification that leads to a loss of information and there is not really any (spatial) propagation going on when the perturbation parameters are constant in space. The snow DA field has been advancing towards increasingly sophisticated spatio-temporal approaches to better represent the complexity of forcing error distributions, whereas this spatially constant approach seems to run counter to recent developments in the field.

47. L284: Consider explicitly describing how the observation error depends on snow depth: Eq. (7) combines a constant absolute error and a relative error in quadrature, so the error standard deviation grows with snow depth.

48. L285: The choice of p_s is unfortunate as this is often used for surface pressure, consider some variation like ϱ_s.

49. L304: Consider changing sigma to observation error standard deviation, or at least defining that somewhere.

50. Eq. (9): Great to see this functional form for the fSCA observation error standard deviation, with the largest error standard deviation for mixed pixels. A very similar form has been used for sea-ice concentration in TOPAZ (Sakov et al., 2012) and was found to be empirically well supported for fSCA in Aalstad et al. (2020), which could help justify your choice.

51. L328: Change *"scalars"* to *"spatially-aggregated scalars"*. Moreover, this sentence only mentions the computational tractability upside of this choice but not the information loss that comes with aggregating and not updating spatially distributed parameters.

52. L346: Other than presumably being readily available, why did you choose to adopt interpolated station data as opposed to more widely used forecast or reanalysis data as atmospheric forcing?

53. L380: Add *"Evaluating the"* before *"ESS ..."*.

54. Figure 3: Great that you specify the source of the land cover map, what about the DEM?

55. L385: Consider instead using the notation σ_a and σ_d for the air temperature (T_a) and dew point temperature (T_d) perturbation standard deviation. Using σ_{RH} is misleading when you perturb dew point temperature.

56. L392: With a 500% relative uncertainty at Latschbloder in Eq. (7) with a floor of 0.3 m you are effectively ignoring this data since it will have a completely negligible effect on the ensemble. If these data are not trusted, it may be clearer to simply exclude them.

57. L395: Is the proportion for the 0.4% threshold computed on the whole snowpack or in the upper layer only?

58. Figure 4: Why is the prior ensemble not shown for each assimilation window in the time series (a)-(d), is it because it is too similar to the posterior? Please also specify if the shaded area is for the total ensemble spread or for some percentile range.

59. Table 3: This table is hard to follow as there are no headers in each column which makes it seem inconsistent across rows. Maybe it is just more of a list than a table. Some entries are also unclear or raise questions. As such, we suggest you change *"Prior ensemble"* to just *"Ensemble"* (this is not just about the prior?), *"Rejuvenation"* to *"Prior"* (as discussed, you are not rejuvenating but just resetting to unconstrained forcing perturbation parameters by drawing from the unconstrained prior), add a mean for the air temperature perturbation (presumably $\mu_a = 0$) for consistency and transparency, Change *"lognormal factor"* to *"Multiplicative"* (or Multiplicative lognormal) for consistency (that is what log-normal priors are used for, namely multiplicative perturbation parameters), change *"Relative humidity"* to *"Dew point temperature"* which is what you are actually perturbing and σ_{RH} to σ_d while also specifying μ_d, change *"Direct model equivalent"* to *"Assumed model equivalent"* (due to residual representativeness error).

60. Figure 5: Panels (a)-(e) (DA1-DA5) show that there is a consistent offset in the residuals between the two sites. On the one hand it makes sense that Latschbloder does not match as well since the data there are effectively ignored. On the other hand, it exposes the problem with your spatially constant perturbation parameter model: while one parameter set can fit snow depth at Proviantdepot reasonably well, those same parameters do not work well at Latschbloder. This raises the question of why you do not allow for spatially distributed (either completely distributed or semi-distributed) parameters. If the counter argument is that you really do not trust the snow depth data at Latschbloder at all, it would be clearer to exclude it.

61. Figure 6: Nice Figure! Some small suggestions: Calling the variable in panels (a)-(d) the wet snow fraction (WSF) is misleading because the WSF that you assimilate is what you get when you average this (binary) quantity over the valid part of the domain. This should instead be called the wet snow status or similar, where panels (b)-(c) show the ensemble average of this binary variable (what you call 'probability' elsewhere) but distributed in space. Incidentally, this figure also shows all the information you are throwing away by assimilating domain-wide WSF: with spatially constant parameters it is possible to match the domain-wide WSF by having the wrong local patterns in the wet snow map, whereas distributed parameters could help match these spatial patterns. Moreover, if possible, dry and no-snow pixels in panels (a)-(d) could have a different color to ease interpretation.

62. Figure 7: Panels (a)–(d) are potentially misleading since they do not show the same quantity. The model snow cover probability expresses how confident the ensemble is that a grid cell is snow covered in a binary sense, whereas the satellite product gives the fraction of each pixel that is snow covered, i.e. the probability that a randomly chosen point within the pixel is snow covered. Plotting both on the same 0–100% scale invites a direct comparison that is not meaningful. Please compare like with like, either by binarizing the satellite retrievals or by using a subgrid parameterization to obtain ensemble-mean fSCA at the grid scale. The former is simplest, since binary snow cover is what your ensemble-based snow cover probability predicts. This would also allow a quantitative evaluation using e.g. the Brier score (the binary cousin of the CRPS).

63. L430: It is still not entirely clear how this example currently provides the foundation for developing new DA methods. While the use of a particle filter with the openAMUNDSEN model is novel, and observations like WSF remain mostly untapped in snow DA, it is not evident that you have demonstrated a path towards developing new snow DA methods here. Again, we are not questioning the value of this contribution in terms of developing a snow DA workflow with the particle filter for an advanced snow model, but rather how you are framing your contribution. We ask that you tone down the language to align with the contribution itself in light of developments in the snow DA field as a whole. Just to be clear, we are not only referring to our own work with MuSA but other recent developments such as with the more operational FSM2OSHD. Furthermore, we do see how this workflow *could* in relatively short order be extended to help develop new snow DA methods, just not in its current state where the setup is quite restrictive in terms of how the particle filter is implemented with scalar domain-wide data assimilation and spatially constant parameters that are reset after each observation. This is more about the decisions you have made while developing the framework that only implements snow DA in a relatively narrow sense while not exploiting developments in spatio-temporal snow data assimilation related to e.g. localization as well as new particle methods.

64. L433: *"PF workflow"* seems like a much better name for what you have developed here rather than *"ensemble-based snow data assimilation framework"*. Your workflow nonetheless has the potential to be developed into a full blown snow DA framework in the future.

65. L441: Playing devil's advocate here, you could have adopted much of the same snow DA methods as used in CrocO, FSM2OSHD, or MuSA, the only thing that would change is the snow model itself. Instead, you chose to develop a new PF workflow. On the one hand, this has the advantage of making it easier to tailor this to your model which makes sense. On the other hand, the spatially constant parameter PF variant you have implemented is a strong simplification of what is possible in the aforementioned existing snow DA workflows. So arguably you do occupy the same software niche, and the possibility of developing a more synergistic workflow using existing tools is partly a missed opportunity.

66. L450: To be easily transferable to other locations, your workflow should support globally (or at least regionally) available forcing data such as from atmospheric reanalysis or numerical weather prediction as well as potentially downscaling thereof. Perhaps this is already implemented in openAMUNDSEN, but it is not mentioned explicitly in this study. If not, the claim to be easily transferable seems like something of a stretch. Again, this is mostly a language issue where if you would have written *"can be transferred"* (dropping *"easily"*) this would be easier to grasp.

67. L460: It makes more sense to refer to these as *"hyperparameters"* rather than *"parameters"* to avoid confusion with the perturbation parameters that you are (implicitly) inferring in your forcing formulation.

68. L462: Great that you acknowledge this information loss with WSF, but it also applies to fSCA.

69. L468: More robust than what? Spatio-temporal snow data assimilation such as with a localized particle filter? This is a strong claim and you have not shown this. Just because something is easier to apply it does not follow that it is more robust. It is also unclear what you mean by robust which is one of those words that can mean many things depending on the context.

70. L470: A natural extension of the subdomain approach that you suggest here, which would presumably be fairly straightforward to implement if your subdomains do not need to be spatially contiguous, would be to adopt a semi-distributed snow DA approach (Fiddes et al., 2019; Cluzet et al., 2022). For example, you could apply the clustering approach suggested in Fiddes et al. (2019) which would be straightforward to adapt from the PBS to the PF workflow you use here. In this way, you would keep more of the information in your satellite retrievals and be able to infer (semi) spatially distributed forcing perturbation parameters. We are not suggesting you go semi-distributed now, but this is a relatively straightforward extension worth considering in the future.

We hope these comments are helpful. Several suggestions would be relatively straightforward to implement given the solid foundation you have built and could make openAMUNDSEN-DA an even more valuable tool for the snow DA community, either now or in future versions. We look forward to the revised version of this promising manuscript.

Kind regards,
Kristoffer Aalstad and Marco Mazzolini

References

Aalstad, K., Westermann, S., and Bertino, L.: Evaluating satellite retrieved fractional snow-covered area at a high-Arctic site using terrestrial photography, Remote Sensing of Environment, 239, 111 618, https://doi.org/10.1016/j.rse.2019.111618, 2020.

Aalstad, K., Alonso-González, E., Pirk, N., Westermann, S., Willmes, C., and Yang, R.: Evolving beyond collapse: an adaptive particle batch smoother for cryospheric data assimilation, Geoscientific Model Development, 19, 8565–8596, https://doi.org/10.5194/gmd-19-8565-2026, 2026.

Alduchov, O. A. and Eskridge, R. E.: Improved Magnus Form Approximation of Saturation Vapor Pressure, Journal of Applied Meteorology and Climatology, 35, 601 – 609, https://doi.org/10.1175/1520-0450(1996)035<0601:IMFAOS>2.0.CO;2, 1996.

Alonso-González, E., Aalstad, K., Baba, M. W., Revuelto, J., López-Moreno, J. I., Fiddes, J., Essery, R., and Gascoin, S.: The Multiple Snow Data Assimilation System (MuSA v1.0), Geosci. Model Dev., https://doi.org/10.5194/gmd-15-9127-2022, 2022.

Alonso-González, E., Harpold, A., Lundquist, J. D., Piske, C., Sourp, L., Aalstad, K., and Gascoin, S.: Ensemble-based data assimilation improves hyperresolution snowpack simulations in forests, The Cryosphere, 20, 209–225, https://doi.org/10.5194/tc-20-209-2026, 2026.

Chen, Y., Nerger, L., and Lawless, A. S.: A Python interface to the Fortran-based Parallel Data Assimilation Framework: pyPDAF v1.0.2, Geoscientific Model Development, 18, 8235–8252, https://doi.org/10.5194/gmd-18-8235-2025, 2025.

Chopin, N. and Papaspiliopoulos, O.: An Introduction to Sequential Monte Carlo, Springer, https://doi.org/10.1007/978-3-030-47845-2, 2020.

Cluzet, B., Lafaysse, M., Deschamps-Berger, C., Vernay, M., and Dumont, M.: Propagating information from snow observations with CrocO ensemble data assimilation system: a 10-years case study over a snow depth observation network, The Cryosphere, 16, 1281–1298, https://doi.org/10.5194/tc-16-1281-2022, 2022.

Elvira, V., Martino, L., and Robert, C.: Rethinking the Effective Sample Size, International Statistical Review, 90, 525–550, https://doi.org/10.1111/insr.12500, 2022.

Evensen, G., Vossepoel, F. C., and van Leeuwen, P. J.: Data Assimilation Fundamentals, Springer, https://doi.org/10.1007/978-3-030-96709-3, 2022.

Farchi, A. and Bocquet, M.: Review article: Comparison of local particle filters and new implementations, Nonlinear Processes in Geophysics, 25, 765–807, https://doi.org/10.5194/npg-25-765-2018, 2018.

Fiddes, J., Aalstad, K., and Westermann, S.: Hyper-resolution ensemble-based snow reanalysis in mountain regions using clustering, Hydrology and Earth System Sciences, 23, 4717–4736, https://doi.org/10.5194/hess-23-4717-2019, 2019.

Gilks, W. R. and Berzuini, C.: Following a moving target-Monte Carlo inference for dynamic Bayesian models, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 63, 127–146, https://doi.org/10.1111/1467-9868.00280, 2001.

Hersbach, H.: Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems, Weather and Forecasting, 15, 559 – 570, https://doi.org/10.1175/1520-0434(2000)015<0559:DOTCRP>2.0.CO;2, 2000.

Magnusson, J., Winstral, A., Stordal, A. S., Essery, R., and Jonas, T.: Improving physically based snow simulations by assimilating snow depths using the particle filter, Water Resources Research, 53, 1125–1143, https://doi.org/10.1002/2016WR019092, 2017.

Margulis, S. A., Girotto, M., Cortés, G., and Durand, M.: A particle batch smoother approach to snow water equivalent estimation, Journal of Hydrometeorology, 16, 1752–1772, https://doi.org/10.1175/JHM-D-14-0177.1, 2015.

Mazzolini, M., Aalstad, K., Alonso-González, E., Westermann, S., and Treichler, D.: Spatio-temporal snow data assimilation with the ICESat-2 laser altimeter, The Cryosphere, 19, 3831–3848, https://doi.org/10.5194/tc-19-3831-2025, 2025.

Murphy, K.: Probabilistic Machine Learning: An Introduction, MIT Press, URL `http://probml.github.io/book1`, 2022.

Murphy, K.: Probabilistic Machine Learning: Advanced Topics, MIT Press, URL `http://probml.github.io/book2`, 2023.

Oberrauch, M., Cluzet, B., Magnusson, J., Schwaizer, G., and Jonas, T.: Combined assimilation of Sentinel-2 snow cover fraction and snow station data improves fully distributed snow simulation across multiple spatial scales, EGUsphere, 2026, 1–31, https://doi.org/10.5194/egusphere-2026-2725, 2026.

Sakov, P., Counillon, F., Bertino, L., Lisæter, K. A., Oke, P. R., and Korablev, A.: TOPAZ4: an ocean-sea ice data assimilation system for the North Atlantic and Arctic, Ocean Science, 8, 633–656, https://doi.org/10.5194/os-8-633-2012, 2012.

van Leeuwen, P., Künsch, H., Nerger, L., Potthast, R., and Reich, S.: Particle filters for high-dimensional geoscience applications: A review, Quarterly Journal of the Royal Meteorological Society, 145, 2335–2365, https://doi.org/10.1002/qj.3551, 2019.

van Leeuwen, P. J.: Particle Filtering in Geophysical Systems, Monthly Weather Review, 137, 4089 – 4114, https://doi.org/10.1175/2009MWR2835.1, 2009.