Authors' response to interactive comments by Referee #1

Black text: Reviewer's comment

Blue text: Authors' response

1. GENERAL

I discovered and read this paper with a great pleasure. Data assimilation of snow observation, in particular using the Particle Filter (PF) has been a growing topic in the snow/hydrology modelling community over the last five to ten years. However, this algorithm suffers from two strong limitations when it comes to large scale, data-scarce problems (which are the general case in the field). The first one is the curse of dimensionality, which can be solved by localizing the PF. But this solution generates a second strong limitation: individual members of localized PFs exhibit discontinuous spatial fields and noisy spatial correlation fields which are often detrimental for the PF performance itself.

Capitalizing on a localized PF variant using a spatial interpolation of the PF weights from Cantet et al., (2019), the authors efficiently introduce an approach coming from the hydrological modelling, the Schaake Shuffle to solve for both limitations: the localisation mitigates the curse of dimensionality, and the Schaake Shuffle is used to enforce the spatial structures of a deterministic run in the individual ensemble members in an elegant way.

Overall, I find that this paper is well within the scope of the journal, and has the potential to be a significant contribution to the snow/hydrology data assimilation community and beyond, because they address an important problem in a convincing and elegant way. I must admit, however, that I'm not satisfied with the theoretical justifications for the use of the interpolation of the weights within the PF from Cantet et al. (2019) (IDWPF), instead of the classical localised (LPF, e.g. Farchi and Bocquet, 2018). I would ask for more justifications, or a comparison between the LPF and the IDWPF.

The scientific quality of the writing is sometimes lacking rigor, especially in the Sections 1-3. I would ask for a significant effort on that. Nevertheless, I really appreciated the compactness of the paper and the efficiency of the results and discussion sections, which make a very clear and straightforward demonstration of the author's point.

To wrap up, I see a lot of potential in this paper, and despite my concerns, I am very confident that the authors will be able to address my comments in a revised version of the manuscript. Please find below some details on my main comments. I' m pleased to provide an annotated pdf version of the manuscript with comments and suggestions throughout.

We would like to thank Reviewer 1 for his detailed reading and his constructive commentaries and side notes. There is no doubt his review will help us improve the quality of our manuscript. We will try to address the different comments and to describe how this will affect the manuscript.

2. MAJOR comments

(1)

The authors an interpolation of the particle filter weights (IDWPF) from rather than a classical PF localisation (LPF) (see the review from Farchi et al., (2018)). I agree that the basic idea as formulated l 71-73 (a good-performing particle at a close location must also be good locally) resembles the theory. But in the classical LPF, based on Bayes theorem, and assuming equal prior weights, the posterior weights are computed by multiplying the likelihoods of the particles at the different (independent) observation sites (Eq. 27 from Farchi et al., (2018)) and then, normalising. Here, the IDWPF averages the normalised likelihoods of the particles at the different observation sites. There is a substantial conceptual difference here: averaging instead of multiplying. I suspect that this results in less sharp and potentially suboptimal (more conservative) PF analyses. For example, if a particle is given a zero likelihood at one location, the LPF will reject it, while there is still a chance for it to survive in the IDWPF.

Moreover, the arguments in Cantet et al., and the present manuscript used to justify the used of the IDWPF instead of the LPF failed to convince me: the LPF also proposes a ' tapering' method to smoothly reduce the influence of the observations with the distance (Eqs. 28-29 from Farchi et al., (2018)).

To wrap up, I'm not saying that the IDWPF method is wrong, and should be rejected. I can actually imagine that it could be more resilient to outliers in the observations, and its conservativeness could be an advantage. There may also be references in the literature to serve as base for the IDWPF. But in the present form, the justifications provided to substantially deviate from the main theory, the LPF, are too weak for me. I would suggest to make a considerable effort on justifications, or even to compare the IDWPF with the LPF. The latter would have the benefit of significantly increasing the potential impact of this paper thanks to the use of a more 'orthodox' method.

To help with the discussion, I'm pleased to provide a toy example comparing the IDWPF and the LPF in the form of a jupyter code attached or publicly available at:

https://github.com/bertrandcz/da_notebo

oks

We would like to underline that reviewer 1 summarised very accurately and very efficiently the main theoretical difference between the two kind of particle filter (PF). We would just like to stress out that IDWPF is also based and Bayes' theorem. At gauged sites, the posterior weights are also derived by multiplying the prior weights by the likelihood of the particle. Then the posterior weights are interpolated in space to be applied to ungauged sites (i.e. grid points). In a case where a null likelihood is given to a particle at all observation sites, then the posterior weights would also be 0 at all observation sites (through multiplication), so this particle would be given a 0 weight everywhere. Then, if a particle has a 0 likelihood at a site A, and a non zero likelihood at a site B, interpolated weight for ungauged site would vary from 0 around point A to higher values around point B. This is a desirable behavior, as in this situation the particle would carry some useful information in the area around point B. In the case of the LPF, assuming point A and B are in the same block, the particle would be given a weight of zero for the whole block. It is right to say LPF is more selective (sharp) as it would discard a particle as soon as one observation says it is unlikely (whatever the other observations around might say), while IDWPF would tend to conserve a particle as long as a single observation finds it likely.

We also believe that block design for a territory as vast as our study area would have been challenging, due to very different density of observations in different regions. Interpolation appears to be more flexible to deal with this situation.

To wrap up, we agree with reviewer 1 that more effort has to be given in explaining the specificities of IDWF and to compare it with the more standard LPF, especially in the introduction and methodology sections. We also agree that a detailed comparison of the two approaches with a common application would constitute a fascinating piece of work. Nevertheless, as the main purpose of our proposed manuscript is the introduction of the Schaake Shuffle to try to mediate the problem of spatial discontinuities in spatialized PF, we believe the comparison of IDWPF and LPF would be out of the scope of the paper. On this point, we would agree with reviewer 2 and consider the comparison for a future work and also suggest it as a work of interest in the conclusion.

(2)

The scientific quality and rigor of the writing is often not satisfactory in its present form, in particular in Secs. 1-3:

□ Even though there is no doubt that the authors have a deep understanding of the PF and its terminology, there is sometimes a lack of rigor in the terminology and approximations that make several sentences turn wrong, and arguments fall short (e.g. l.54, l.207-208, l. 218-219).

□ Even though the arguments are there, the logical formulation behind certain paragraphs is too loose to be convincing. I have no doubt that the authors can address that, but a considerable effort is required here.

- □ There are some approximations in the description of the literature which may induce the reader into having misconception on the references (e.g. l. 69, l. 227-228), and change the conclusions of some paragraphs

- □ the observation dataset and study area (Sec. 2) must be described with more details and rigor.

I'm pleased to provide several suggestions on these points in the attached pdf.

Again, the authors want to thank reviewer 1 for his detailed comments in the attached pdf. We will modify the text to integrate those comments, especially the ones regarding the theoretical justification of our approach. An effort will also be given in expanding the description of our datasets and study area.

3. MINOR COMMENTS:

(1) With the Schaake shuffle as used here, the authors enforce the spatial distribution of individual particles to match those of the model, (instead of historical observations): but by doing so, don't we miss the opportunity to adjust the ensemble to the observed spatial structures? I'd be curious about overlaying Figs 3. and 6 with observed (in-situ) values to assess that.

Our main challenge in this project was to identify the observed spatial structured. We do not have completely spatialised observations of SWE covering our study area (the available products are already the result of assimilation processes). It was intended to estimate a spatial structure from the local in-situ observations, but most of the data set is coming from manual snow surveys. SWE at different observation sites is not measured at the same date and estimating the spatial structure using data collected over a couple of weeks is tricky, as the snowpack keeps evolving in between measurements. Hopefully, the deployment of automatic sensors will help us solve this issue in the future.

For the same reason, the suggested figure would have a limited interest as only a limited number of observations are available for a given day.

(2) The abstract is lacking of a general scientific context to start with. I think that the start is too technical for the scope of TC, the notion of 'particles' should be introduced, and given the level of technicity of the paper, a brief sentence describing the particle filter might be required in the abstract, in particular to make the need for a reordering possible to understand. Details on the ensemble construction might be appreciated also. Would benefit from a more rigorous description of the observations and validation data sets.

We agree with reviewer 1 on this point. The abstract will be reworked to start from a more general perspective and provide more theoretical insights. Attention will also be given on methodology and data description. This will probably cause the abstract to be longer but will help less specialized readers to better understand the scope of the paper.

(3) When computing global metrics (Secs. 4.2 and 4.3, Fig 7,8 and 9), it could be fair and interesting to compare the assimilation products with their ensemble counterpart without assimilation, not only the deterministic run (called ‘open loop’ in the paper). Ensembles are often favored compared to deteministic runs in terms of RMSE and SWE, and it would enable the authors to put into perspective the impact of the assimilation in terms of spread-skill and CRPS.

This is an excellent point. We also believe an “ensemble open loop” simulation would be a better (and fairer) benchmark. It is expected it be characterized by a lower RMSE that is deterministic counterpart, but the spread of the members might become very large at the end of the accumulation season. Our intention is to integrate this suggestion. Nevertheless, we would like to stress out that this point is not a minor change, as it requires new simulations and in-depth modifications of most figures.

References:

Cantet, P., Boucher, M.-A., Lachance-Coutier, S., Turcotte, R., and Fortin, V.: Using a particle filter to estimate the spatial distribution of the snowpack water equivalent, Journal of Hydrometeorology, 20, 577–594, https://doi.org/10.1175/JHM-D-18-0140.1, 2019.

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/https://doi.org/10.5194/npg-25-765-2018, 2018.

https://github.com/bertrandcz/da_notebooks/

Please also note the supplement to this comment: https://tc.copernicus.org/preprints/tc-2021-322/tc-2021-322-RC1-supplement.pdf