Large-scale snow data assimilation using a spatialized particle filter: recovering the spatial structure of the particles

Jean Odry[1], Marie-Amélie Boucher[1], Simon Lachance-Cloutier[2], Richard Turcotte[2], and Pierre-Yves St-Louis[2]

[1]Department of Civil and Building Engineering, Université de Sherbrooke, Sherbrooke, Canada
[2]Quebec Ministère de l'Environnement et de la Lutte contre les Changements Climatiques, Quebec, Canada

Correspondence: Jean Odry (jean.odry@usherbrooke.ca)

Abstract. Data assimilation is an essential component of any hydrological forecasting systems. Its purpose is to incorporate some observations from the field when they become available in order to correct the state variables of the model prior to the forecasting phase. The goal is to ensure that the forecasts are initialized from state variables that are as representative of reality as possible, and also to estimate the uncertainty of the state variables. There exist several data assimilation methods, and particle
5 filters are increasingly popular because of their minimal assumptions. The baseline idea is to produce an ensemble of scenarios (i.e. the particles) using perturbations of the forcing variables and/or state variables of the model. The different particles are weighted using the observations when they become available. However, implementing a particle filter over a domain with large spatial dimensions remains challenging, as the number of required particles rises exponentially as domain size increases. Such situation is referred to as the "curse of dimensionality", or a "dimensionality limit". A common solution to overcome this curse
10 is to localize the particle filter. This consists in dividing the large spatial domain into smaller portions, or « blocks », and applying the particle filter separately for each block. This can solve the above-mentioned dimensionality problem because it reduces the spatial scale on which each particle filter must be applied. However, it can also cause spatial discontinuities when the blocks are reassembled to form the whole domain. This issue can become even more problematic when additional data is assimilated. The purpose of this study is to test the possibility of remedying the spatial discontinuities of the particles by
15 locally reordering them.

We implement a spatialized particle filter to estimate the snow water equivalent (SWE) over a large territory in eastern Canada by assimilating local SWE observations from manual snow surveys. We apply two reordering strategies based on 1) a simple ascending order sorting and 2) the Schaake Shuffle and evaluate their ability to maintain the spatial structure of the particles. To increase the amount of assimilated data, we investigate the inclusion of a second data set (SR50), in which SWE is
20 indirectly estimated from automatic measurements of snow depth using sonic sensors. The two reordering solutions maintain the spatial structure of the individual particles throughout the winter season, which significantly reduces the spatial random noise in the distribution of the particles and decreases the uncertainty associated with the estimation. The Schaake Shuffle proves to be a better tool for maintaining a realistic spatial structure for all particles, although we also found that sorting provides a simpler and satisfactory solution. The assimilation of the secondary data set improved SWE estimates in ungauged

sites when compared with the deterministic model, but we noted no significant improvement when both snow courses and the SR50 data were assimilated.

1 Introduction

The accumulation and melting of snow dominate the hydrology of Nordic and mountainous regions (Doesken and Judson, 1997; Barnett et al., 2005; Hock et al., 2006). In these regions, accurate information about snow water equivalent (SWE) is crucial for streamflow forecasting (Li and Simonovic, 2002) and reservoir management (Schaefli et al., 2007). Over large territories or catchments, the spatial distribution of SWE can be assessed using remote sensing (Goïta et al., 2003) or snow modeling (Marks et al., 1999; Ohmura, 2001; Essery et al., 2013). The advantages of snow modeling include the possibility to generalize over large territories and ensure a complete temporal coverage with a resolution determined by the user. However, as with any kind of model, snow modeling is subject to uncertainty and inaccuracy stemming from errors in the inputs or the model structure itself (structural uncertainty). The origin of structural uncertainty includes both epistemic uncertainty and one's perceptual model (Beven, 2012). Those factors, in addition to tradition (or "legacy", see for instance Addor and Melsen, 2019) influence one's choice of model algorithmic structure, level of complexity and process description (e.g., Clark et al., 2011)

Given that SWE is a cumulative variable, these errors obviously increase in importance throughout the winter, making SWE estimates highly uncertain at the beginning of the melting season. Nonetheless, a correct assessment of the spatial distribution of SWE over large domains is often required for diverse applications, including flood forecasting in large catchments. Therefore, any data assimilation technique for snow that can be applied to such a large spatial domain with satisfactory results is of very high practical importance for operational hydrology, as it could potentially improve flood forecasting and reservoir management.

Data assimilation can be used to limit the accumulation of uncertainty and error from input data and models. The general idea is to correct the state variables of the model using the observations when they become available. Several elaborate data assimilation schemes have been proposed in the literature to deal with these limitations, including optimal interpolation (Burgess and Webster, 1980), ensemble Kalman filtering (Evensen, 2003), and particle filtering (Gordon et al., 1993; Moradkhani et al., 2005). The latter relies on fewer hypotheses than the other schemes in relation to the nonlinearity and non-normality of the uncertainty distributions of the model's input variables, state variables, and outputs. The particle filter (PF) is based on the simulation and weighting of different scenarios (particles) rather than modifying the state variables of the model; thus, the physical consistency of state variable is ensured and the disruption of cross-correlations between state variables is avoided. These scenarios are generated within the range of the uncertainty of the input data and/or the model structure. Assessing or defining the range of uncertainty for each variable is not a trivial task, but it is also not unique to the particle filter (see, for

instance, Thiboult and Anctil, 2015). The particles are weighted using the observations when they become available. These weights can be used to build the distribution of the variable of interest (model output).

A common limitation of the PF is the risk of particle degeneracy. After a certain number of time steps, the evolution of most particles brings them so far from the observation that they are associated with a weight close to zero. This can be seen as an impoverishment of the filter with fewer and fewer useful particles. At some point, a single particle concentrates all the weight, and a distribution can no longer be built. This degeneracy can be avoided using some kind of particle resampling; the particles with the lowest weights are discarded, whereas the better ones are duplicated. Several resampling algorithms have been proposed in the literature, such as those of Gordon et al. (1993); Douc and Cappe (2005); Leisenring and Moradkhani (2011).

The PF has been successfully used for the local (onsite) assimilation of snow data (Leisenring and Moradkhani, 2011; Magnusson et al., 2017). The application of PF over large territories, however, remains challenging. The number of particles must grow exponentially with the dimension of the observed space to avoid filter degeneracy (Snyder et al., 2015). This "curse of dimensionality" (Bengtsson et al., 2008) is a severe drawback for the application of the PF over large spatial domains with a consequent number of observations.

Several procedures have been proposed to overcome the curse of dimensionality. A first approach, known as error inflation (Larue et al., 2018), consists of increasing the error associated with the observation to make the filter more tolerant. Cluzet et al. (2020) proposed a generalization of this approach to apply it over large domains. The most common approach to overcome the curse of dimensionality is localization (Farchi and Bocquet, 2018). This family of methods assumes that points separated by a sufficiently large distance are independent. Then, given this independence, the assimilation is performed locally, that is to say only observations close enough to a given site of interest are considered. This can be achieved with techniques such as spatially limited block (where the space is divided in fixed spatial tiles) or localisation radius (where the selection of observation is proper to each site of interest). In this way, the impact of each observation is limited in space, and the assimilation is performed on several subregions of lower spatial dimensions.

Recent work involving spatialization or the PF for assimilating SWE over large spatial domains includes that of Cluzet et al. (2020) who developed the k-local framework to select the appropriate data set to be assimilated for each site of interest; this selection is based on the cross-correlation between ensemble values at the site of interest and observations sites.Poterjoy (2016) introduced a Bayesian updating of particles near observation sites while preserving particles located far away from observations. This method for locally updating the particle is similar to covariance localization used in Ensemble Kalman filtering. Cantet et al. (2019) proposed a spatialized particle filter using the localization approach. They used spatially structured perturbations to generate the particles, computed the weights of the particles locally for each observation site, then interpolated the weights in space. The underlying assumption of their method is that for any spatially coherent particle (i.e., model simulation), sites that are close to each other in space should be characterized by similar weights. The method was successfully applied to the assimilation of SWE data from manual snow surveys in the snow module of the HYDROTEL model (Turcotte et al., 2007) over the province of Quebec in eastern Canada. The spatialized particle filter from Cantet et al. (2019) is now fully operational

for SWE assimilation in the context of flood forecasting in Quebec. This approach is then considered as the baseline from an operational perspective in Quebec.

Although localization is an effective means of circumventing the curse of dimensionality, this approach can create physically unrealistic spatial discontinuities because of the local resampling implemented independently at different sites and because of the resulting noise (Farchi and Bocquet, 2018). This can occur when a global particle is formed (post-resampling) with a given particle on block a and with a different particle on the adjacent block b. The apparition of these discontinuities can be mitigated by improving the resampling algorithm. For instance, Farchi and Bocquet (2018) proposed the use of the optimal transport theory to minimize the movements of local particles during resampling. Such a technique aims to prevent the apparition of discontinuities. Nevertheless, when resampling must be performed frequently, it cannot prevent every discontinuities. Rather than (or in addition to) preventing the discontinuities, it is possible to mitigate them. Such a process can be called particle gluing. Some authors suggest smoothing out the discontinuities by space averaging the resampled fields (Penny and Miyoshi, 2016). This appears to be an effective but rather subjective way to remove the discontinuities, as this smoothing is not related to the underlying spatial structure of the physical variable. A solution for gluing the particles back together and solving the spatial discontinuities becomes necessary when the number of observation sites or the temporal frequency of the observations increases. Whereas being able to assimilate more data is always appealing to better capture the spatial and temporal variability of the snowpack properties, a greater amount of resampling is expected when more observations are used to evaluate the particles of the PF. Unfortunately, a greater number of resampling iterations is likely to create new spatial discontinuities.

In Cantet et al. (2019), only the data from manual snow surveys were assimilated in the snow model. The fact that this type of measurement is not available continuously in time reduces the risk of spatial discontinuities appearing because of resampling. However, a large number of SR50 sonic sensors, which measure snow depth at an hourly frequency, have been deployed across Quebec since 2000. Furthermore, Odry et al. (2020) recently proposed a machine-learning model to evaluate SWE from snow depth, temperature, and precipitation indicators in Quebec. On the one hand, there is potential added value in including multiple types of snow data in the particle filter. Those advantages include the addition of new observation sites, but also a finer consideration of the temporal variation of the snowpack (snowfall, detection of the beginning of the snowmelt,...). On the other hand, including data that is available continuously and performing data assimilation more frequently increases the risk of spatial discontinuity appearing.

In a context where snow observations are increasingly automated, it becomes crucial to remedy the curse of dimensionality. Here we present an innovative means of gluing the particles of a localized particle filter back together using the Schaake Shuffle (Clark et al., 2004). The Schaake Shuffle was introduced to rebuild the spatial structure in forecasted ensemble meteorological fields. It uses a set of spatial observations of the variable to reorganize the members of an ensemble cell by cell over a spatial grid to bring their spatial structures closer to the observations. It is an actual shuffling of the members, and there is no alteration in the distribution of the local ensemble. Considering that a localized setting of the particle filter is necessary to overcome the curse of dimensionality, and considering that spatial discontinuities can arise from this localized setting, we hypothesize that the Schaake Shuffle can be used within a localized particle filter to glue back the resampled particles together and restore the

spatial correlation. To rebuild the short range correlation among the particles using an alternative and simpler solution, we also implement a simple sorting, in ascending order, of the particles' SWE.

125   To verify this hypothesis, we propose using the spatial particle filter developed by Cantet et al. (2019) and implement the Schaake Shuffle after each resampling of the particles. For comparison purposes, we use a study area that includes that of Cantet et al. (2019) and the same model. The reference set of the Schaake Shuffle is simulated over a historical period using a deterministic run of the model without any assimilation, as no reliable spatial observation of snow is available for this region.

The availability of SR50 data provides a great opportunity to test both particle reordering when the amount of data is in-
130   creased and the ability of our approach to assimilate other types of data having different uncertainties in the spatial particle filter (SPF). In this context, and as a secondary objective of this study, we also hypothesize that the assimilation of this secondary SWE data set improves the estimates of SWE compared with the deterministic simulation and that the addition of data in the assimilation scheme improves the overall quality of the SWE estimates at ungauged sites.

The next section describes the study area and the data. The methodology is outlined in a third section, and the results and
135   related discussion are presented in the fourth part.

2   Experimental setup

2.1   Study area and snow data

Quebec is a province in eastern Canada. For most catchments, the hydrology of rivers within the province is largely influenced by the marked accumulation of snow during the winter, with an annual maximum peak flow occurring during the spring freshet.
140   For this reason, the Quebec government has put much effort into monitoring and simulating SWE across the province. Our research form a part of this strategy.

Historically, SWE monitoring has occurred through manual snow surveys using standard Federal Snow Samplers. The measurements are undertaken manually at various sites along predetermined snow lines at differing temporal frequencies (monthly to bi-weekly). This sampling strategy provides information regarding SWE, snow depth, and snow bulk density.
145   Since the beginning of the 2000s, automated snow sensors have been installed across Quebec. These automatic sonic sensors (SR50) measure snow depth continuously at an hourly time step. Although the temporal continuity of these snow depth series is a clear advantage, the absence of information regarding SWE is a drawback. This limitation is overcame by deriving daily time series of SWE from the SR50 data using an artificial neural network model developed for this purpose and trained and validated against the manual snow survey data set (Odry et al., 2020). Because this is an indirect estimation of SWE, it is
150   assumed that this second SWE data set is characterized by greater uncertainty than the data from the manual snow surveys. To be able to make use of both data sets, we focus on the 2011–2015 period for which both data sets have a relatively higher and more stable number of SR50 sensors.

All snow data used in this research was provided by the Réseau de surveillance du climat du Québec (MELCC, 2019), operated by the Quebec Ministère de l'Environnement et de la Lutte contre les Changements Climatiques (MELCC), from the
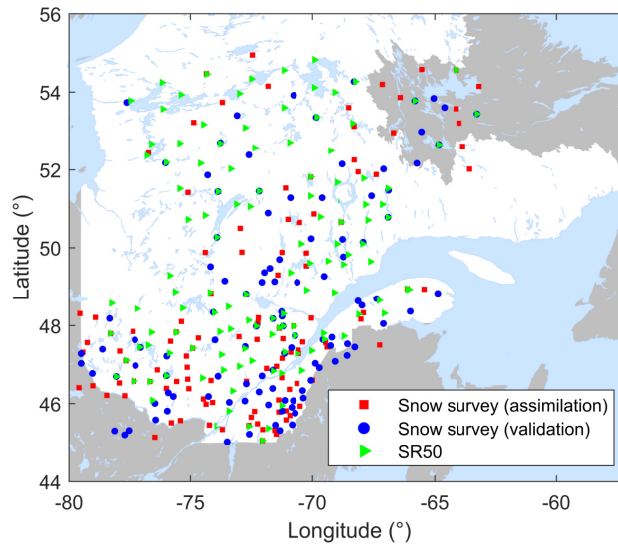
Figure 1. Snow observation sites in and around Quebec (Canada)

other private and public sectors members of the Réseau météorologique coopératif du Québec (RMCQ) and from their partners outside Quebec. The spatial distribution of observation sites is presented in Figure 1.

Figure 2 provides the observed SWE values from the entire manual snow survey dataset averaged by month or by year, for the individual sites and the whole area. This information has to be taken with caution, as the number of observations per site and year as well as the date of observations can change throughout the period. Nevertheless, it is possible to appreciate that the maximum SWE value is reached in March for most stations (i.e. the south of the province), but only in April for stations with the largest SWE values (i.e. the northern part of the province). No general trend is visible in the evolution of the monthly average. Nevertheless, it is possible to see the the mean April SWE used to be higher in the 70s, which is a known reality in Quebec hydrology.

This research focuses on a portion of the province located south of 53°N, where most of the population is concentrated. Northern Quebec differs markedly from the more southern region in terms of the amount of available snow data and snow characteristics; therefore, its inclusion would require a separate study.

2.2 Snow model and meteorological forcing

We use the HYDROTEL snowpack model (HSM) in this study (Turcotte et al., 2007), which is used operationally for real-time SWE estimation at the MELCC. HSM is a temperature index–based model and therefore provides a simplified representation of the water and energy budgets of the snowpack. It is a rather simple snow model that only requires daily precipitation and minimum and maximum temperatures as inputs. The implementation and parameterization of the model are those proposed by Cantet et al. (2019). More detailed specifications about the model structure and physics can be found in Turcotte et al. (2007).
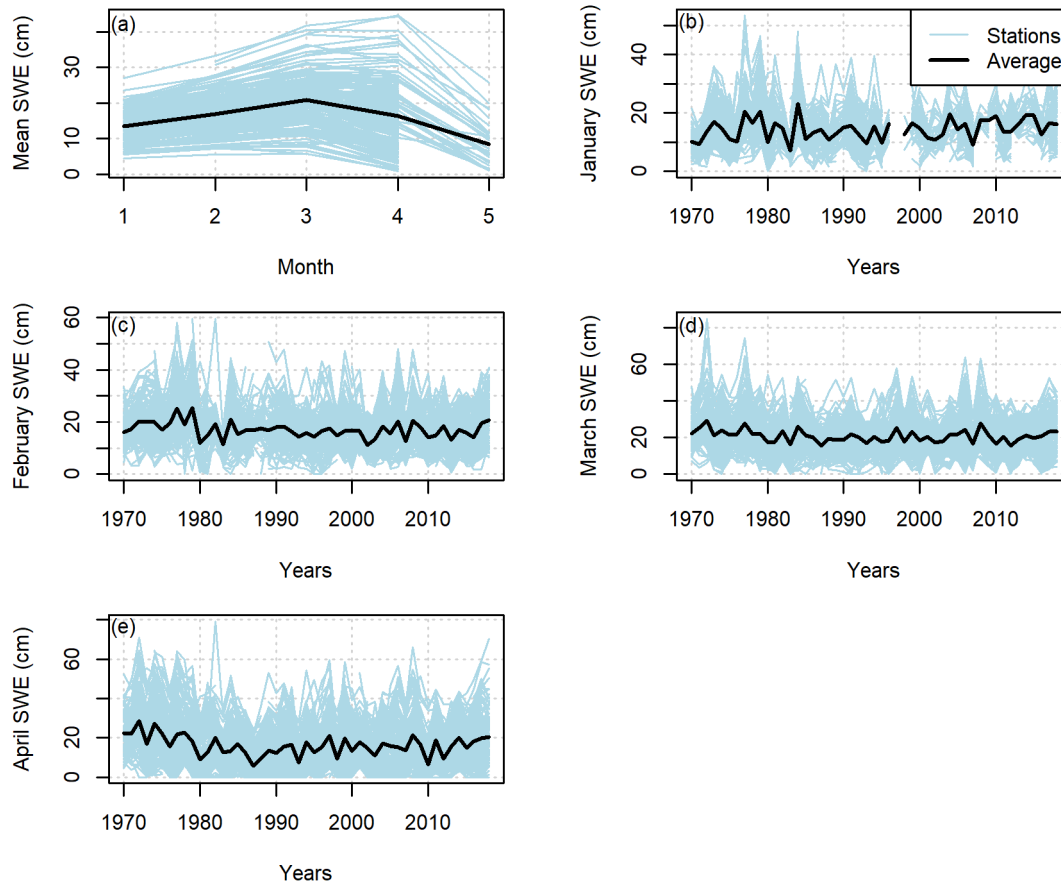
Figure 2. Averaged SWE observations in the study area between 1961–2018, (a) over the entire year; (b) in January; (c) in February; (d) in March; (e) in April

To force the HSM, daily grids of precipitation and minimum and maximum temperatures are used. They are produced by the MELCC by kriging in situ measurements from local meteorological ground stations (Bergeron, 2017) on a regular 0.1° resolution grid. Along with the interpolated variables, the corresponding kriging variances are also produced. They cover the entire 1961–2018 period.

3 Methodology

3.1 Spatialized particle filter

The particle filter implemented in this study is the one developed and tested by Cantet et al. (2019). This method was developed to assimilate at-site snow data into a large-scale modeling framework. The main feature is to generate particles characterized

by a realistic temporal and spatial correlation so that each particle can be seen as a reasonable scenario. The observations are then used to weight the different particles when and where they are available. In this context, it is hypothesized that the weights should be spatially correlated, given that the particles are also spatially correlated. In a second step, the weights estimated at observation sites are interpolated in space. To avoid filter degeneracy, Cantet et al. (2019) implemented the sampling importance resampling algorithm from Gordon et al. (1993).

The spatialized particle filter proposed by Cantet et al. (2019), differs from more traditional approaches such as the ones described by Farchi and Bocquet (2018). In Cantet et al. (2019), the likelihood of the particles is evaluated locally at each observation site. The likelihood is then multiplied by the prior weights to derives the posterior weights. The posterior weights at observation sites are then interpolated in space. In traditional localized approach, the likelihood is evaluated over a certain area (block or radius around an observation), and the posterior weights are directly derived for this area by multiplying the prior by the likelihood. However, the interpolation step in the approach of Cantet et al. (2019) could tend to make it less selective and too keep particles as long as one observation provides it with a sufficiently high weight.

Following Cantet et al. (2019), we generated 500 particles using perturbations of the meteorological input data and the predicted SWE. Perturbations of the precipitation and temperature inputs should represent the uncertainty associated with the forcing. Given that these inputs are created through a spatial interpolation, they are perturbed using an additive noise following a Gaussian distribution for which the variance is set to the interpolation variance. The simulated SWE is perturbed with a multiplicative uniform noise proposed by Clark et al. (2008) and used in Leisenring and Moradkhani (2011); this perturbation process represents the structural uncertainty associated with the snow model. The temporal correlation of the noise is maintained using the formulation (Evensen, 1994):

$$s_t = \alpha s_{t-1} + \sqrt{1-\alpha^2}\eta_t, \qquad \text{with} \qquad s_0 \sim \mathcal{N}(0,1), \tag{1}$$

where s_t is the random noise at time t, α represents the temporal correlation and is set at 0.95, and η_t is a random white noise following a standard normal distribution.

The spatial correlation of the noise is built up using the exponential model (Uboldi et al., 2008):

$$r = \exp\frac{-d^2}{L^2}, \tag{2}$$

where r is the spatial correlation between two points separated by a distance d, and L is the range of the correlation, set at 200 km (Cantet et al., 2019).

In our study, local particle filters are implemented at the observation sites. The weight of a given particle at an observation site at a specific time when an observation is available depends on the likelihood of the measured SWE value given the SWE associated with this particle. This likelihood is represented by a Gaussian process:

$$y_t|x_t^i \sim \mathcal{N}(y_t - x_t^i, \sigma^2), \tag{3}$$

where y_t represents the observed SWE at time t, x_t^i is the SWE simulated by the i[th] particle, and σ is the standard deviation of the Gaussian process and is related to the uncertainty associated with each type of SWE observation and is expressed in

SWE units. For manual snow surveys, σ should represent both the uncertainty of the measurement itself and the error of spatial representativeness, i.e., how well a local measurement can represent the SWE averaged over a grid cell. σ is tuned to $a * (0.1 * y_t + 1)$, where a is a unitless parameter set at 3 for the snow survey observations (by trial and error by Cantet et al. (2019)) and to be determined for the SR50. In the case of SWE estimated from the SR50 snow depth measurements, σ should also represent the structural uncertainty of the artificial neural network model (Odry et al., 2020) used to estimate SWE. Moreover, although manual snow surveys average 10 measurements along a 100–300 m snow line, SR50 sensors provide a purely at-point measurement. As a consequence, the error for spatial representativity should be larger for SR50 than for manual snow surveys.

Under these assumptions, when a given observation becomes available at time t, the weight x_t^i associated with the i[th] particle at the observation site can be updated as:

$$w_t^i = \frac{w_{t-1}^i p(y_t | x_t^i)}{\sum_{i=1}^{N_P} w_{t-1}^i p(y_t | x_t^i)}, \tag{4}$$

where N_P is the number of particles. One can notice that the weight of a particle at a given time step depends on the weight at the previous time step. Therefore, the assimilation of continuous data with a time correlation can lead to assimilating the same information several times and to give too much weight to this information. Here, it is chosen to assimilate the SR50 at a weekly frequency. In operational real-time simulations, the model is restarted each day from seven days in the past to include all observations obtained over that one-week period. It would thus be possible to assimilate only the most recent SR50 observation and have an assimilation frequency close to weekly.

At a given time step, the weights are updated for any newly available observation using Equation 4. The weights at the observation sites are spatially interpolated using a simple inverse distance-weighted method to estimate weights for each grid cell. This process makes it possible to include different observations obtained inside the same grid cell.

Once the weights have been updated and interpolated, it is possible to investigate the risk of filter degeneracy. Degeneracy occurs when a few particles concentrate all the weight. For each observation site and grid cell, N_t^{eff}, the effective sample size at time t is computed as:

$$N_t^{eff} = \frac{1}{\sum_{i=1}^{N_P} (w_t^i)^2}. \tag{5}$$

If the effective sample size falls below the $0.8 N_P$ threshold (Magnusson et al., 2017), the particles are then locally resampled. It is important to highlight that whereas the particles are generated with a spatial structure, they are evaluated and resampled locally. The resampling is performed using the SIR algorithm (Gordon et al., 1993) and consists of deleting the particles with the lowest weights and duplicating those with the largest weights. The objective of this resampling is to obtain a set of particles associated with equal weights of $\frac{1}{N_P}$ that describe the same distribution as the weighted particles before resampling. Through this process, the filter retains a large number of meaningful particles, while the posterior distribution of SWE is marginally

affected (the resampling algorithm should be selected to limit the impact on the posterior distribution). Moreover, the SWE values and other state variables of the different particles are not modified.

245   Although the resampling aims to preserve the filter from degeneracy, it is also responsible for the risk of spatial discontinuity, as it is performed independently for each site and grid cell. The risk of discontinuity is managed using a reordering of the particles.

3.2   Recovering the spatial structure by reordering the particles

We implemented two alternative procedures to maintain the spatial structure of each particle despite resampling: the Schaake
250   Shuffle, proposed by Clark et al. (2004), and a simple sorting of the particles at each site. The use of any of those strategies in the context of particles filtering constitutes the novelty of this piece of research.

The Schaake Shuffle is a reordering method that aims to recover the space–time structure in an ensemble of precipitation or temperature forecast fields. In the original application of the Schaake Shuffle, the spatial structure of the ensemble members is lost during a downscaling procedure. The idea behind the Schaake Shuffle is to rank the members of the ensemble and
255   match these ranks with those of past observations selected randomly from similar dates in the historical records. Although the Schaake Shuffle was not initially developed to be applied to the field of snow or data assimilation, there is no technical difficulty in transferring its application from an ensemble of meteorological forecasts to a set of particles describing a snowpack.

For a given time and location, let $\mathbf{X}$ be the vector of length N_P of SWE values for all the particles. A vector of equal size $\mathbf{Y}$ is built from historical record at the same location and similar date. Considerations about the construction of $\mathbf{Y}$ are provided
260   in the following sections. The vectors $\boldsymbol{\chi}$ and $\boldsymbol{\gamma}$ are the sorted versions of $\mathbf{X}$ and $\mathbf{Y}$, respectively:

$$\mathbf{X} = (x_1, x_2, ..., x_{N_P}), \tag{6}$$

$$\boldsymbol{\chi} = (x_{(1)}, x_{(2)}, ..., x_{(N_P)}), \qquad \text{where} \qquad x_{(1)} \leq x_{(2)} \leq ... \leq x_{(N_P)}, \tag{7}$$

$$\mathbf{Y} = (y_1, y_2, ..., y_{N_P}), \tag{8}$$

$$\boldsymbol{\gamma} = (y_{(1)}, y_{(2)}, ..., y_{(N_P)}), \qquad \text{where} \qquad y_{(1)} \leq y_{(2)} \leq ... \leq y_{(N_P)}. \tag{9}$$

265   Moreover, $\mathbf{B}$ is taken as the vector of indices describing the positions of the elements of γ in the original vector $\mathbf{Y}$. That is to say the i[th] element of $\mathbf{B}$ describes where to find the i[th] element of γ in $\mathbf{Y}$. Finally, if $\mathbf{Z}$ is the reordered version of $\mathbf{X}$, then:

$$\mathbf{Z} = (z_1, z_2, ..., z_{N_P}), \qquad \text{with} \tag{10}$$

$$z_q = x_{(r)}, \tag{11}$$

$$q = \mathbf{B}[r], \qquad \text{and} \tag{12}$$

270    $$r = 1, ..., N_P. \tag{13}$$

This reordering must be applied at each site or grid cell. A more detailed procedure and a visual example are provided in (Clark et al., 2004).

As it appears in equations 7–13, the main limitation to the application of the Schaake Shuffle is the availability of historical records for the construction of vector $\mathbf{Y}$. It is necessary to have records covering each site (observation sites and grid cells) and be of sufficient length to be able to draw N_P values from similar periods for each site. Given that an extensive data set does not exist for SWE in Quebec, we propose to use a record of simulations using the HSM in its deterministic configuration without assimilation. This selection relies on the assumption that the model without data assimilation can correctly estimate the spatial structure of SWE. This is a fairly strong assumption, nevertheless in the absence of a spatially distributed SWE product covering the whole study area for a long period of time, and considering the spatial structure in the HSM deterministic simulation is controlled by the meteorological inputs and not a pre-imposed spatial structure, this choice is considered reasonable. One can notice that the actual SWE values in the historical records are not used in the Schaake Shuffle, only their respective rank is used to reorder the particles, so the estimation errors and biases from the model are not of prime importance. In our application, the vector $\mathbf{Y}$ is constructed from a reference set of simulations produced with a deterministic run of the HSM for the years 1961 to 2004. It is assumed that the overall spatial structure of SWE remains the same during this period. This assumption is difficult to verify as there is no spatial observation of SWE covering the whole area and period, but Figure 2 does not show any strong tendency over the period. Following Clark et al. (2004), only values associated with a date within a window of ± 7 days around the assimilation date (but with different years) are used. The vector $\mathbf{Y}$ is built by randomly drawing 500 values from the 660 available records.

As a simpler alternative to the Schaake Shuffle, we also implemented a sorting of the particles. In this case, at each site, all 500 particles are sorted in ascending order of SWE values. The idea is to rebuild the correlation between the sets of particles in neighboring sites—in the sense of the rank correlation. Although the sorting procedure is much simpler to implement and does not require any reference record, it also removes the possibility to consider each particle as a reasonable snow scenario at a large scale. Compared to the Schaake-Shuffle, the sorting can be seen as a way of trading the large-scale structure for a short-scale one.

3.3 Validation procedure and metrics

To evaluate the impact of the different reordering strategies on the spatial structure of the particles and the final SWE estimation, we use experimental variograms. These plots provide the evolution of the variance between the sites as a function of the separating distance. Because the variograms are only computed for a single date, the Pearson correlation between the 500 SWE values (from each particle) at two nearby sites is also used and computed at each time step so that the temporal evolution of the spatial structure can also be assessed. When investigating the spatial structure, all snow survey sites are used for assimilation.

To assess the ability of the overall method to estimate SWE, we compute two deterministic scores, the root mean squared error (RMSE) and the mean bias error (MBE), as well as two ensemble-based scores, the continuous ranked probability score (CRPS, Matheson and Winkler, 1976) and the skill-to-spread ratio (Fortin et al., 2014). The CRPS considers two, sometimes competing, aspects of ensemble quality. The first aspect is resolution. An ensemble has a good resolution (or sharpness)

Table 1. Performance metrics

Metric	Focus	Perfect value	Formula
RMSE	Predictive accuracy	0	$RMSE = \sqrt{\frac{1}{n}\sum_{k=1}^{n}(\hat{y_k} - y_k)^2}$
MBE	Systematic bias	0	$MBE = \frac{1}{n}\sum_{k=1}^{n}(\hat{y_k} - y_k)$
CRPS	Global ensemble evaluation	0	$CRPS = \frac{1}{n}\sum_{k=1}^{n}\int_{-\infty}^{+\infty}(F_k^e(u) - F_k^o(u))^2 du$
Skill-to-spread ratio	Adequacy of spread and skill	1	$r = \dfrac{RMSE}{\sqrt{\frac{1}{n}\sum_{k=1}^{n} var_k}}$

n is the number of points, x_k and $\hat{x_k}$ are respectively the estimated and observed SWE values, var_k is the variance of the particles of the k^{th} point, F_k^e is the cumulative distribution function described by the weighted particles of the k^{th} point, and F_k^o is the Heaviside function associated with the observation.

when it discriminates events. For instance, a climatological ensemble that covers the whole range of possibilities, has no resolution. A deterministic simulation series that follows the observed series very closely has an excellent resolution. The second aspect is reliability. Reliable ensembles display exactly the good amount of variability. For instance, a 95% confidence interval computed from the daily ensemble is reliable if it contains the observation 95% of the time on average. The RMSE is a measure of accuracy, which is related to resolution. If two different series of ensemble simulations have the same level of resolution (similar RMSE) but very different reliability values, the more reliable one will obtain a lower (better) CRPS.

The metrics are summarized in Table 1. As the purpose of the spatialized particle filter is mainly to produce estimates at ungauged sites, only 50% of the manual snow survey sites are used for the assimilation; the other half is kept as a validation subset. This sub-sampling strategy is presented in Figure 1. All provided scores are computed only on the validation subset.

4 Results and discussion

4.1 Spatial structure of the particles

The variograms of the estimated SWE for February 28[th], 2012, for the different configurations (deterministic run, ensemble open loop, assimilation without reordering, assimilation with sorting of the particles, and assimilation with the Schaake Shuffle) are provided in Figure 3. We selected February 28[th] because February is considered to be the last month of the accumulation season in Quebec. The deterministic run corresponds to the deterministic simulation of the HSM without any assimilation, whereas the three other curves correspond to the weighted average of the 500 particles. The five variograms are very similar, which demonstrates that the different assimilation strategies lead to final SWE estimates that are comparable to the deterministic run of the HSM, with similar spatial structures. This similarity is a desired behavior, as the spatial structure should be guided
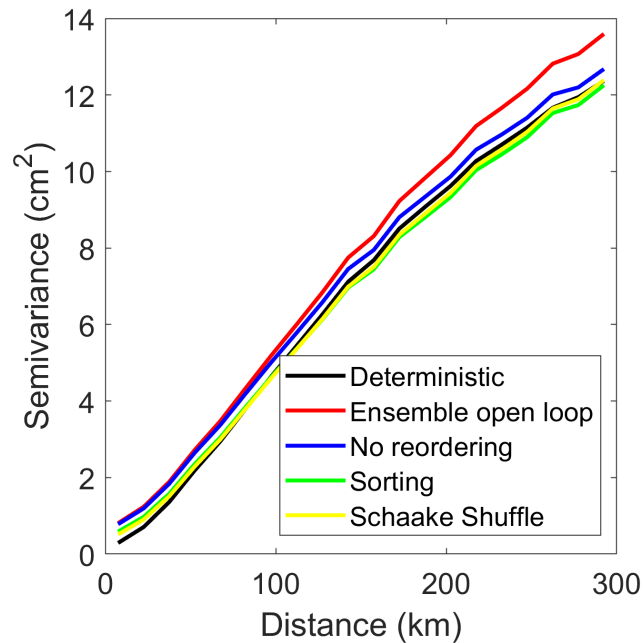
Figure 3. Variograms of the estimated snow water equivalent (SWE) on February, 28, 2012

by the snow model and its meteorological forcings. Especially, Figure 3 demonstrates that particle reordering does not affect the overall spatial structure of the final SWE estimate.

325   The variograms in Figure 3 have a semi-variance below 1 cm^2 for separation distances close to 0. The semi-variance that corresponds to a null distance (also called the nugget effect) is important, as it describes the spatial roughness or local noise. The presence of spatial discontinuities would create a noticeable nugget effect. As expected, the deterministic model is characterized by the lowest nugget effect (around 0 cm^2), which means an absence of local random noise. In addition, it is interesting to note that the assimilation without reordering has the largest nugget effect, with a value around 1 cm^2. This large nugget effect can

330 be a symptom of some noise affecting the correlation between neighboring sites. Finally, the ensemble open loop displays the largest long distance semi-variance, which is a result of a larger dispersion of the particles. This is expected in the absence of data assimilation. Despite those small differences, we conclude that the different configurations result in comparable spatial structures.

    The maps corresponding to the variograms in Figure 3 are provided in Figure 4. The five maps exhibit a similar spatial

335 pattern, although the map from the deterministic simulation (panel a) is the most distinct from the other ones. This was expected because there is no assimilation involved in this simulation. The deterministic simulation also provides a smoother map compared with the other three, which is compatible with what is observed in the variograms. The overall similarities of the different maps was expected, considering that they are characterized by similar variograms. Figure 4 evidences that the assimilation process as well as the different reordering methods do not impact the spatial structure of the final SWE estimate.
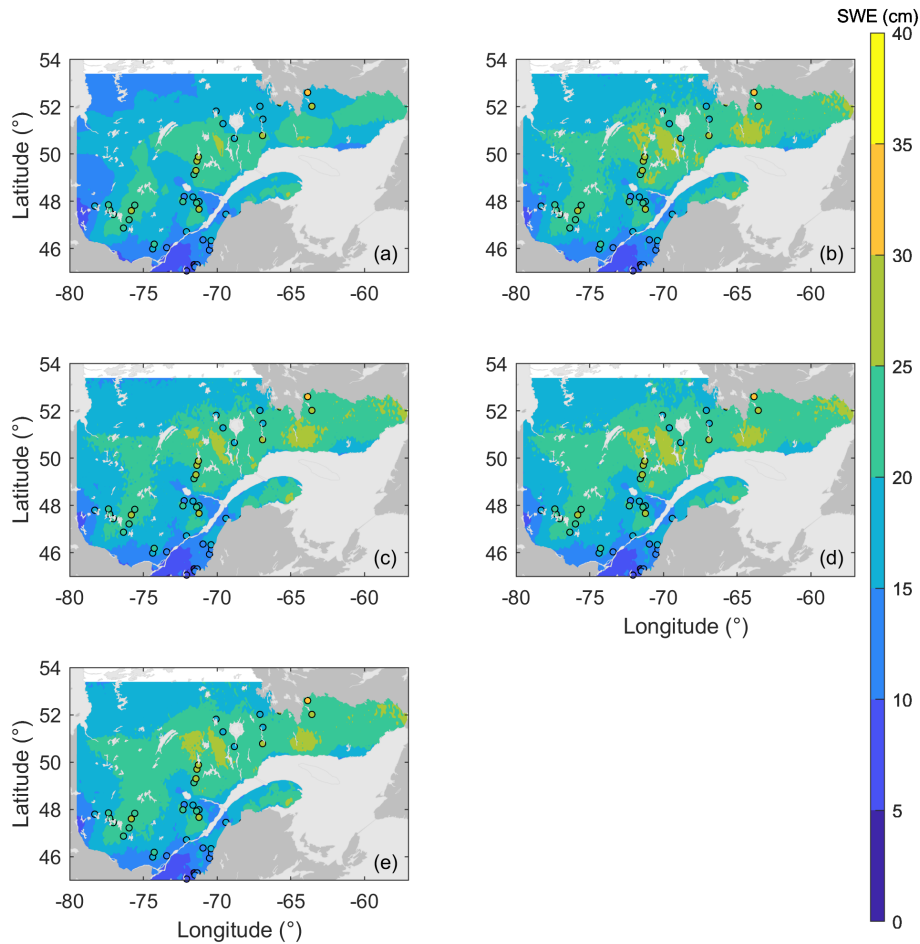
Figure 4. Estimated snow water equivalent (SWE) maps for February, 28, 2012; (a) deterministic run, (b) ensemble open loop (c) no reordering, (d) sorting, (e) Schaake Shuffle; points represent the locations of the available snow surveys on this day

340   The behaviour was expected since the overall structure is driven by the forcing data and the spatially correlated perturbations. Figure 4 shows that the reordering process does not affect the overall spatial structure of the estimated SWE.

     Nevertheless, although it is important to assess the spatial structure of the final SWE estimates, the spatialized particle filter relies on the assumption that the weights of the particles can be interpolated in space. Therefore, the spatial structure of each particle must be preserved. To verify this, we plotted the variograms of the 500 particles for the three reordering options:

345   no reordering, sorting the particles in ascending order and the Schaake Shuffle (Figure 5). In Figure 5, the thin gray curves represent the variograms of individual particles, whereas the bold colored lines are identical to those in Figure 3. Note that the vertical axes do not share the same scale. Moreover, the maps corresponding to two particles (number 250 and number 500) are provided in Figure 6 to illustrate the link between the variograms and the overall spatial structure. First, for both the ensemble open loop and the assimilation without reordering (Figure 5 a and b), it is clear that the individual spatial structure of
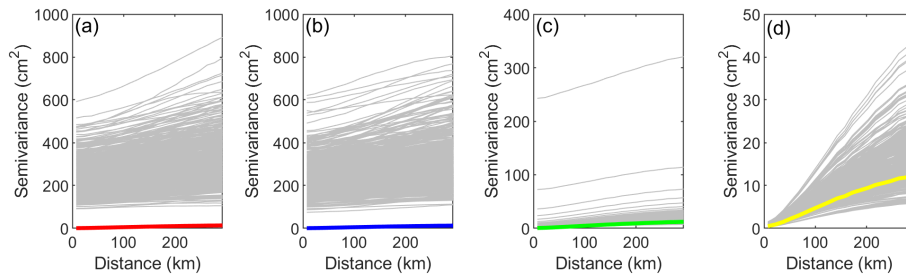
Figure 5. Variograms of the particles for the different reordering strategies for February, 28, 2012; (a) ensemble open loop; (b) no reordering; (c) sorting; (d) Schaake Shuffle. Solid lines represent the variograms of the estimated SWE as displayed on Figure 3 and gray lines represent the variograms of individual particles.

350 each particle is different from the weighted average. More specifically, all particles are characterized by a much larger nugget effect, ranging from approximately 100 cm^2 to around 600 cm^2. Therefore, there is significant short-range random variability. This nugget effect can also be observed on Figure 6 a and b as well as c and d, where the maps display a very granular aspect. The fact that individual particles exhibit an important nugget effect whereas their weighted average does not exhibit this effect indicates a large amount of random noise associated to each particle, which is in disagreement with the underlying assumptions

355 of the spatialized particle filter.

Figure 5c shows the variograms of the individual particles when the particles are sorted in ascending order following each resampling. This sorting attenuates the random noise associated to individual particles very effectively. Following this sorting, most particles fall around the final estimate. Nevertheless, some still exhibit a very large nugget effect. This granularity is associated with the accumulation of small random differences between neighbouring cells throughout the winter. Even if the

360 perturbations (on precipitation, temperature and SWE) are spatially correlated, they are not perfectly identical. This nugget effect on individual particles dissipates when the particles are averaged, as seen on Figures 3, 4 and 5. With reordering the particles in ascending order, we expected that some of those particles would display a disturbed spatial structure. Panels e and f of Figure 6 provide more insight regarding this expectation. The maps represent the state of two individual particles (#250 and #500) for a given moment in time. The two particles present very different patterns. This difference was expected, as particle

365 #250 is located in the middle of the ensemble, whereas particle #500 gathers all the highest SWE values. Such differences illustrate that when particles are sorted in ascending order, the individual particles cannot be considered as potential SWE map scenarios anymore. Sorting can rebuild the short-range correlation, which is overall necessary for the spatialized particle filter, but that same sorting can jeopardize the long-range correlation and the general spatial pattern.

Reordering the particles with the Schaake Shuffle appears to be a better option to fix the spatial structure of the particles.

370 Figure 5d illustrates that all particles have a near-zero nugget effect, and their variograms lie around that of the final estimate. This is compatible with the idea that the particles are alternative SWE scenarios, each providing a reasonable spatial structure. The same conclusions can be drawn from Figure 6h and g.
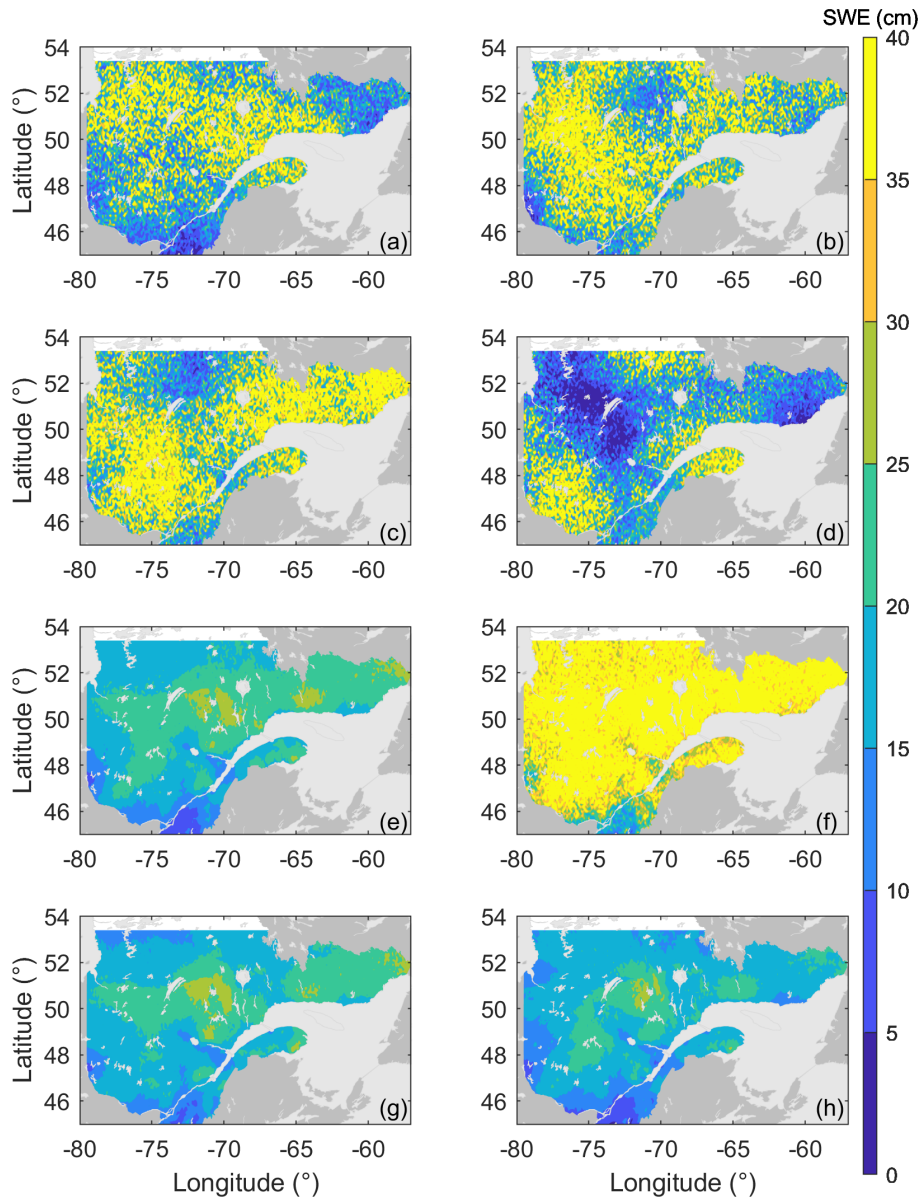
Figure 6. Snow water equivalent (SWE) for two particles on February, 28, 2012: particle #250 (a,c,e,g) and particle #500 (b,d,f,h); ensemble open loop (a and b), no reordering (c and d), sorting (e and f), and Schaake Shuffle (g and h)

As a summary, Figures 5 and 6 demonstrate that the individual particles are characterized by a non realistic spatial structure in the ensemble open loop, and when the assimilation is performed without any reordering. In both cases, the particles are characterized a strong nugget effect, which is also associated to granular SWE maps. This behavior is incompatible with the underlying assumptions of the spatialized particle filter. On the contrary, the two reordering methods result in more re-
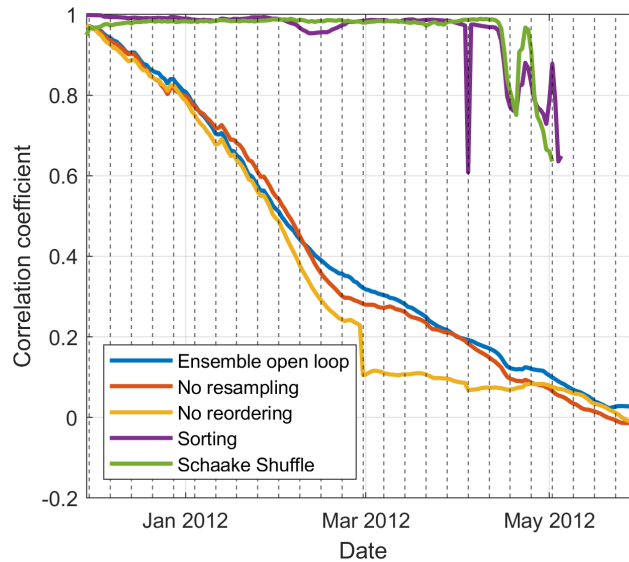
Figure 7. Evolution of the correlation between between two adjacent grid cells over the winter 2011–2012

alistic spatial structures for the individual particles. However, simply sorting the particles in ascending order produces some problematic particles. Only the Schaake Shuffle is able to maintain an acceptable spatial structure for all individual particles.

To evidence the temporal evolution of the spatial structure, we refer to Figure 7, which shows the evolution of the correlation coefficient between the sets of particles at two given neighboring sites for winter 2011–2012. On this figure, the vertical dotted lines represent assimilation dates. Due to the proximity of the two sites, a strong correlation between both sets of particles is to be expected, as SWE is correlated in space at short distance in reality.

Figure 7 also shows the results for data assimilation without any resampling, in order to assess the impact of resampling on the spatial structure. Correlation decreases continuously with time for the assimilation without resampling (nor reordering) and without reordering. The main difference between those two options is a sudden drop of correlation values in late February for the case of resampling without reordering. This drop is associated with a resampling. Whereas resampling can be associated with a large and immediate drop in correlation because of a disassociation of particles at different sites, resampling cannot explain the continuous decrease in correlation. This decrease can be explained, however, by the cumulative nature of the snowpack and the perturbations. In this study and following Cantet et al. (2019), precipitation, temperature, and SWE itself are perturbed independently. The perturbations are generated with a spatial structure, meaning that perturbations applied to neighboring sites are strongly correlated but are not identical. The simulated snowpack accumulates these differences throughout the winter, thus explaining the continuous decrease in correlation. Therefore, it appears that the assumption used to justify the interpolation of particle weights is not supported after a certain number of time steps without reordering, and resampling only aggravates the problem.

395   In Figure 7, the two reordering strategies maintain a very high correlation between neighboring sites. Some momentary correlation losses are observed in April, which can be associated with the disappearing snow (some grid cells and particles having no snow).

As a consequence, it appears that both reordering strategies can maintain the spatial structure of the particles. In the present context, such reordering is necessary to verify the underlying assumption of the spatial particle filter and to ensure that the

400   weights of the particles can be interpolated meaningfully. From a more general point of view, reordering the particles appears to be an effective way to deal with the spatial discontinuities created by the resampling step of the particle filter. Reordering by sorting the particles in ascending order is easy to implement, as it does not require any additional data set, but in this case, the individual particles cannot be seen as scenarios each having a reasonable long-range correlation. In contrast, the Schaake Shuffle preserves the idea of each particle representing a potential realistic scenario. Consequently, because the spatialized

405   particle filter relies largely on spatially structured perturbations and particles, the Schaake Shuffle appears to be a better solution. Nonetheless, the Schaake Shuffle requires a long spatialized record to build the reference set. Here, we demonstrate that a simulated deterministic run can be used as a pseudo-historic climatology of SWE, if we assume that the model can simulate a realistic spatial structure. Still, in situations where a reference set cannot be acquired or created, the reordering by particles sorting them in ascending order can be an acceptable alternative to maintain the short-range correlation of the particles.

410   ## 4.2   Snow water equivalent estimation

In this and the following sections, all scores are computed solely on the validation subset (see Figure 1); the SWE data from the validation sites are never assimilated. The calculated scores intend to show the ability of the spatialized PF to estimate SWE in ungauged sites.

Figure 8 presents the cumulative distribution of the different performance metrics over the validation sites when only manual

415   snow surveys are assimilated. Table 2 provides the scores averaged over all the sites in the validation sample. The shift from the deterministic to ensemble open loop configuration as well as the data assimilation sensibly reduce the RMSE (panel a) and the MBE (panel b) relative to the deterministic run simulation. The ensemble open loop simulation is characterized by as much bias as the the simulations with data assimilation, but the open loop ensemble simulation obtains a higher (worse) RMSE. Nonetheless, RMSE and MBE values are very comparable among the three assimilation and reordering strategies.

420   There seems to be only a slight decrease in RMSE when using the Schaake Shuffle. Differences are more noticeable for ensemble-based performance metrics. First and foremost, Figure 8c exhibits a large decrease in CRPS when using one of the reordering methods. As the RMSE values are comparable, this improvement of the CRPS is necessarily associated with a narrowing of the posterior distribution described by the particles. This result can be linked with the reduced amount of random noise associated with the reordering of the particles (see Section 4.1). The reduction of the ensemble spread is also apparent in

425   Figure 8d. Because the skill (evaluated using the RMSE) is not sensitive to the reordering, the increase in the skill-to-spread ratio can only be explained by a reduction of the spread. The ratio is also systematically higher with the Schaake Shuffle than with the sorting, meaning that the Schaake Shuffle makes the spread of the ensemble of particles even narrower. In the case of the simple sorting, 80% of the sites have a ratio below 1; thus, the particles are over-dispersive. The situation improves slightly
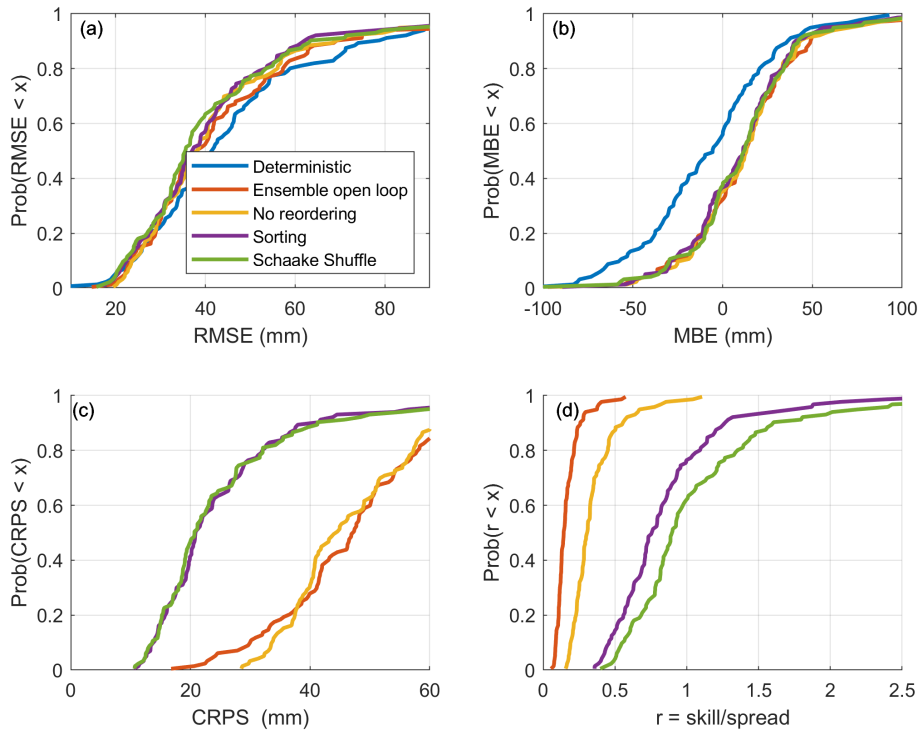
Figure 8. Cumulative distribution of the scores calculated for the 50% validation snow courses

with use of the Schaake Shuffle. As expected, the ensemble open loop simulation is characterised by the worst ensemble scores.
430 In the absence of assimilation, the particles are free to evolve, and the dispersion of the ensemble becomes very large.

 Figure 8 shows that the gain in deterministic error reduction from particle reordering is very limited. Nevertheless, reordering the particles largely improves the ensemble scores, which means that the posterior distribution of the particles is a better estimate of the actual uncertainty. Considering that data assimilation techniques in hydrology are mainly used for forecast applications, a better quantification of the uncertainty is a very important achievement. According to those results, the simple
435 sorting of the particles appears to be as effective as the more elaborate Schaake Shuffle to maintain the spatial structure of the particles.

4.3 Inclusion of the data from the SR50 sensors within the assimilation scheme

The inclusion of the data from SR50 sensors in the experiment aims at improving the spatio-temporal coverage of the territory in terms of observations. This new dataset includes new observation sites and continuous time-series. The first step before
440 including the SR50 data is to estimate the uncertainty associated with the indirect observations. The total uncertainty is composed of the measurement uncertainty, the depth-to-SWE conversion error, and the spatial representativeness error. Here, the spatial representativeness error represents the difference between the local in situ observation and the simulated value, which is averaged over a $0.1°$ resolution grid cell.

Table 2. Averaged scores over the 50% validation snow courses

Score	Deterministic	Ensemble open loop	No reordering	Sorting	Schaake Shuffle
RMSE (mm)	45.8	44.6	43.3	41.8	41.9
MBE (mm)	-8.5	12.9	12.6	10.1	11.2
CRPS (mm)	NA	47.8	47.2	25.8	26.0
skill/spread	NA	0.17	0.35	0.85	1.06

We assume that the uncertainty associated with the SR50 data follows specifications provided in section 3.1. Figure 9
displays the cumulative distribution of the performance metrics on validation sites when only the SR50 data is assimilated for
different a values. The first observation is that the lowest (worst) scores are achieved for an a-value of 4. Higher (better) scores
are attained at higher a-values, which confirms that the uncertainty associated with the SR50 must be greater than that of the
snow surveys. It also appears that if a is greater than 4, there are no large differences in the distribution of RMSE, MBE, or
CRPS. Thus, the performances are not very sensitive to the a-value as long as it is sufficiently high (greater than 4). In Figure
9d, the skill-to-spread ratio increases with the a-values of the uncertainty associated with the SR50 observations. This result
indicates that greater uncertainty for these observations creates a larger spread of the particles, which is expected as a higher
observation uncertainty provides greater flexibility to the particles around the observation and gives more weight to the prior
distribution. Here, in the end and considering the observations above, we selected an a-value of 10 for the SR50 data as a
reasonable compromise.

Figure 10 displays the cumulative distributions of the performance metrics over the validation sites, according to the type of
data that is assimilated: snow surveys only, SR50 only, or both. First, all assimilation schemes significantly reduce the RMSE
(panel a) for most sites compared to the deterministic run. These results, similar to those of Cantet et al. (2019), demonstrate
that the spatialized particle filter can assimilate local observations and use them to improve the global estimation of SWE. On
panel b, the assimilation schemes increase the MBE values in all sites relative to those of the deterministic simulation, which
confirms that assimilation tends to correct the underestimation of the deterministic simulation. When comparing the RMSE
(panel a) and CRPS (panel c), depending on the type of data that is assimilated, improved results are obtained when only the
snow surveys are assimilated. The assimilation of SR50 observations only or of both types of observations simultaneously
displays comparable performances, although still higher (better) than that of the deterministic run. Thus, in absence of snow
surveys, it is possible and even desirable to assimilate estimated SWE data from SR50 observations to improve the simulation
of SWE. Nonetheless, at this point, the validation process reaches its limitations. In our study, 50% of the snow survey sites
are used in validation; therefore, it is assumed that these observations represent the actual amount of SWE. Nevertheless,
the model simulates SWE over a grid having a 0.1° resolution, thus local observations may not be representative. Moreover,
SR50 sensors and snow survey sites tend to not be installed in the same environment; SR50 sensors are generally installed
as part of ground meteorological stations, which are located in more open environments as per the World Meteorological
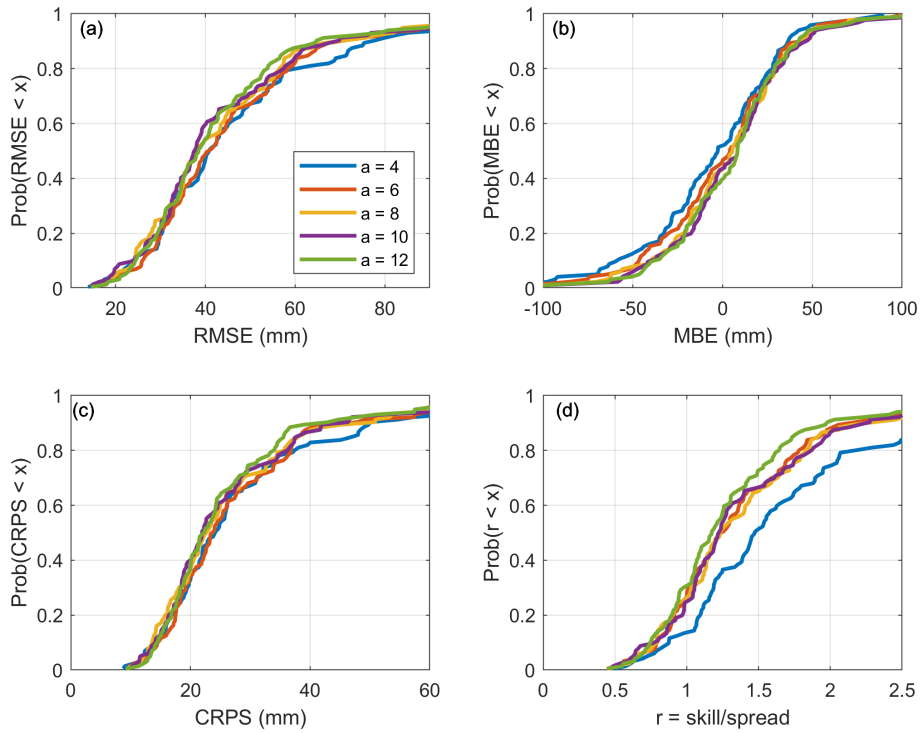
Figure 9. Cumulative distribution of the scores for the validation sites when assimilating SR50 data, according to the applied uncertainty model

Organization's guidelines (World Meteorological Organization, 2018), whereas snow surveys are almost always conducted in a forested environment. According to the official guidelines, a radius of 3m around each sampling point must be cleared from trees and other vegetation. Therefore, while manual snow survey sampling points are not directly under the canopy, they are protected from the wind. Consequently, the validation of snow survey sites only may produce a bias. Finally, in Figure 10d, we note that the type of assimilated data has a large effect on the skill-to-spread ratio. The assimilation of SR50 data greatly increases the ratio compared with the assimilation of snow surveys alone, whereas the assimilation of both types of data produces the highest values of skill-to-spread ratio. Given that the level of skill (RMSE in panel a) cannot explain this difference, the assimilation of SR50 data only or of both types of data therefore decreases the ensemble spread. These results may be explained by the amount of assimilated data, as the more observations are assimilated, the more the particles are constrained. The number of assimilated snow survey sites and SR50 sites are similar, but even with a weekly frequency, the number of observations is much higher for the SR50, which explains the decrease in spread. Obviously, the number of observations is even higher when assimilating both types of observations.
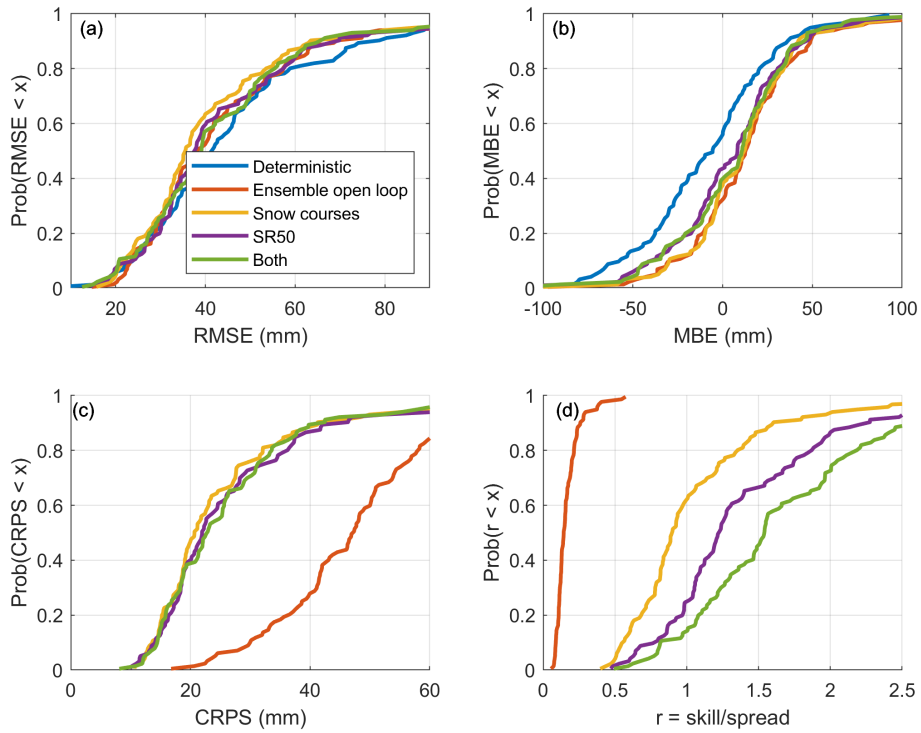
Figure 10. Cumulative distribution of the scores for the validation sites according to the nature of the assimilated data

5 Conclusions

In this study, we propose an improvement of the spatialized particle filter introduced by Cantet et al. (2019). By reordering the particles of the filter it was possible to slighlty reduce the RMSE associated with SWE estimation at validation sites as well as largely reducing the CRPS which means that the uncertainty associated to the SWE estimation in ungauged site is significantly reduced. In a more general context, particle reordering is a good remedial solution to solve spatial discontinuity problems arising from resampling and localization when applying a particle filter to large spatial areas. Particle reordering may then be a new solution to the curse of dimensionality for particle filtering. To confirm this result, integration of the reordering procedures in different kind of localized particles filter should be considered. While it has been shown that the spatialized particle filter described in this paper improved SWE estimation and provided better uncertainty estimates than previous versions or the open loop simulation, it would be beneficial to undertake a more general comparison of different localized particles filters (such as the ones described by Farchi and Bocquet (2018) or Cluzet et al. (2020)).

The reordering of particles using the Schaake Shuffle technique or a sorting procedure can help maintain the spatial structure of the particles within the ensemble. The Schaake Shuffle appears to be more effective in recreating the overall spatial structure of all particles. This solution makes it more reasonable to assume that it is possible to interpolate the weights of the particles

in space; this is the basic assumption behind the spatialized particle filter. In cases where a sufficiently long reference record cannot be acquired, a reordering by sorting the particles can provide a good alternative.

This research was also an opportunity to test the possibility of assimilating SWE estimates derived from automatic snow depth observations (SR50) and an ensemble of artificial neural networks (Odry et al., 2020). We demonstrated that the assimilation of this additional data set alone outperforms the deterministic simulation. This observation confirms the relevance of this new SWE data set. However, the assimilation of both SR50 data and snow surveys did not improve the simulation when compared with the assimilation of only the manual snow surveys. We attribute this result to the lower quality of the SWE estimates from the SR50 and also to the validation process itself, which involves preserving a portion of snow survey sites for validation. Nevertheless, the exercise demonstrated the possibility of assimilating different types of data together in the spatialized particle filter, using the uncertainty associated to each type of snow observation to weight their relative influence. In this context, the current ongoing deployment in Quebec of automatic sensors capable of measuring SWE (rather than snow depth) by using natural gamma radiation constitutes a great opportunity. Not only could this third source of information be assimilated into the particle filter, but alternatively, it could also be used as an improved validation set, as these new gamma ray–based sensors provide sub-daily observations, and the added value of the SR50 data in terms of temporal representativeness may be better captured. The ability of the spatialized particle filter to assimilate different types of data could also be compared to other operational assimilation framework such as the one decribed by Zhang and Yang (2016).

Code and data availability. The code for the SWE simulation using the spatialized PF with particle reordering is provided on Github https://github.com/TheDroplets/Snow_spatial_particle_filter (Odry, 2021). The data set used in this study, including the meteorological data and a part of the snow data that we are authorized to share can be accessed though the Harvard Dataverse, at https://dataverse.harvard.edu/dataverse/Odry_etal2021_largeScaleSpatialStructure. (Boucher, 2021a, b, c, d). There is a mandatory form to fill to access the data, and this form can also be found in the same Dataverse.

Author contributions. Jean Odry performed all the computations, testing, and primary analysis and also wrote the first draft of the manuscript. Marie-Amélie Boucher initially proposed this study and supervised the work throughout the research and writing phases. Simon Lachance-Cloutier, Richard Turcotte, and Pierre-Yves St-Louis offered guidance, technical support, and expertise regarding data, snow models, and operational reality. They also participated in the editing of the manuscript.

Competing interests. The authors declare they have no competing interests.

Acknowledgements. The authors acknowledge the MELCC of Quebec for its financial support in the undertaking of this work. The authors thank Philippe Cantet for sharing the original code for the spatialized particle filter. Finally, they also want to thank the members of the

Réseau météorologique coopératif du Québec for sharing the snow data: Quebec's Ministère de l'Environnement et de la Lutte contre les Changements Climatiques, Quebec's Ministère des Forêts, de la Faune et des Parcs, Quebec's Société de protection des forêts contre le feu, Environment and Climate Change Canada, Hydro-Québec, Rio Tinto, as well as their partners from outside Quebec: Ontario Power Generation, the Ministry of Northern Development, Mines, Natural Resources and Forestry of Ontario, Churchill Falls Corporation, Environment New Brunswick, Maine Cooperative Snow Survey.

References

530   Addor, N. and Melsen, L.: Legacy, Rather Than Adequacy, Drives the Selection of Hydrological Models, Water Resources Research, 55, 378–390, https://doi.org/doi.org/10.1029/2018WR022958, 2019.

Barnett, T. P., Adam, J. C., and Lettenmaier, D. P.: Potential impacts of a warming climate on water availability in snow-dominated regions, Nature, 438, 303, https://doi.org/10.1038/nature04141, 2005.

Bengtsson, T., Bickel, P., and Li, B.: Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems, Institute of
535   Mathematical Statistics, https://doi.org/10.1214/193940307000000518, 2008.

Bergeron, O.: Grilles climatiques quotidiennes du Programme de surveillance du climat du Québec, version 2 – Guide d'utilisation,, Tech. rep., Ministère du Développement durable, de l'Environnement et de la Lutte contre les changements climatiques, Direction du suivi de l'état de l'environnement,, Québec, 2017.

Beven, K.: Rainfall-Runoff Modelling: a primer, 2nd edition, 2012.

540   Boucher, M.-A.: Meteorological inputs, https://doi.org/10.7910/DVN/BXXRHL, 2021a.

Boucher, M.-A.: HYDROTEL Snow Model Parameters, https://doi.org/10.7910/DVN/RJSZIP, 2021b.

Boucher, M.-A.: Historical Snow Simulation (Open Loop), https://doi.org/10.7910/DVN/CJYMCV, 2021c.

Boucher, M.-A.: Snow Observation Data, https://doi.org/10.7910/DVN/NPB1JY, 2021d.

Burgess, T. M. and Webster, R.: Optimal Interpolation and Isarithmic Mapping of Soil Properties, Journal of Soil Science, 31, 315–331,
545   https://doi.org/https://doi.org/10.1111/j.1365-2389.1980.tb02084.x, 1980.

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.

Clark, M., Gangopadhyay, S., Hay, L., Rajagopalan, B., and Wilby, R.: The Schaake Shuffle: A Method for Reconstructing Space–Time Variability in Forecasted Precipitation and Temperature Fields, Journal of Hydrometeorology, 5, 243–262, https://doi.org/10.1175/1525-
550   7541(2004)005<0243:TSSAMF>2.0.CO;2, 2004.

Clark, M., Kavetski, D., and F., F.: Pursuing the method of multiple working hypotheses for hydrological modeling, Water Resources Research, 47, https://doi.org/doi:10.1029/2010WR009827, 2011.

Clark, M. P., Rupp, D. E., Woods, R. A., Zheng, X., Ibbitt, R. P., Slater, A. G., Schmidt, J., and Uddstrom, M. J.: Hydrological data assimilation with the ensemble Kalman filter: Use of streamflow observations to update states in a distributed hydrological model, Advances in
555   Water Resources, 31, 1309–1324, https://doi.org/10.1016/j.advwatres.2008.06.005, 2008.

Cluzet, B., Lafaysse, M., Cosme, E., Albergel, C., Meunier, L.-F., and Dumont, M.: CrocO_v1.0 : a Particle Filter to assimilate snowpack observations in a spatialised framework, Geoscientific Model Development Discussions, pp. 1–36, https://doi.org/https://doi.org/10.5194/gmd-2020-130, 2020.

Doesken, N. J. and Judson, A.: The snow booklet: A guide to the science, climatology, and measurement of snow in the United States,
560   Colorado State University Publications & Printing, 1997.

Douc, R. and Cappe, O.: Comparison of resampling schemes for particle filtering, in: ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005., pp. 64–69, https://doi.org/10.1109/ISPA.2005.195385, iSSN: 1845-5921, 2005.

Essery, R., Morin, S., Lejeune, Y., and B Ménard, C.: A comparison of 1701 snow models using observations from an alpine site, Advances
565   in Water Resources, 55, https://doi.org/10.1016/j.advwatres.2012.07.013, 2013.

Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, Journal of Geophysical Research: Oceans, 99, 10 143–10 162, https://doi.org/https://doi.org/10.1029/94JC00572, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/94JC00572, 1994.

Evensen, G.: The Ensemble Kalman Filter: theoretical formulation and practical implementation, Ocean Dynamics, 53, 343–367, https://doi.org/10.1007/s10236-003-0036-9, 2003.

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.

Fortin, V., Abaza, M., Anctil, F., and Turcotte, R.: Why Should Ensemble Spread Match the RMSE of the Ensemble Mean?, Journal of Hydrometeorology, 15, 1708–1713, https://doi.org/10.1175/JHM-D-14-0008.1, 2014.

Gordon, N. J., Salmond, D. J., and Smith, A. F. M.: Novel approach to nonlinear/non-Gaussian Bayesian state estimation, IEE Proceedings F (Radar and Signal Processing), 140, 107–113, https://doi.org/10.1049/ip-f-2.1993.0015, 1993.

Goïta, K., Walker, A. E., and Goodison, B. E.: Algorithm development for the estimation of snow water equivalent in the boreal forest using passive microwave data, International Journal of Remote Sensing, 24, 1097–1102, https://doi.org/10.1080/0143116021000044805, 2003.

Hock, R., Rees, G., Williams, M. W., and Ramirez, E.: Contribution from glaciers and snow cover to runoff from mountains in different climates, Hydrological Processes, 20, 2089–2090, https://doi.org/10.1002/hyp.6206, 2006.

Larue, F., Royer, A., Sève, D. D., Roy, A., and Cosme, E.: Assimilation of passive microwave AMSR-2 satellite observations in a snowpack evolution model over northeastern Canada, Hydrology and Earth System Sciences, 22, 5711–5734, https://doi.org/https://doi.org/10.5194/hess-22-5711-2018, 2018.

Leisenring, M. and Moradkhani, H.: Snow water equivalent prediction using Bayesian data assimilation methods, Stochastic Environmental Research and Risk Assessment, 25, 253–270, https://doi.org/10.1007/s00477-010-0445-5, 2011.

Li, L. and Simonovic, S. P.: System dynamics model for predicting floods from snowmelt in North American prairie watersheds, Hydrological Processes, 16, https://doi.org/10.1002/hyp.1064, 2002.

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.

Marks, D., Domingo, J., Susong, D., Link, T., and Garen, D.: A spatially distributed energy balance snowmelt model for application in mountain basins, Hydrological Processes, 13, 1935–1959, https://doi.org/10.1002/(SICI)1099-1085(199909)13:12/13<1935::AID-HYP868>3.0.CO;2-C, 1999.

Matheson, J. E. and Winkler, R. L.: Scoring rules for continuous probability distributions, Management science, 22, 1087–1096, 1976.

MELCC: Données du réseau de surveillance du climat du Québec, Tech. rep., Direction générale du suivi de l'état de l'environnement, Québec, 2019.

Moradkhani, H., Hsu, K.-L., Gupta, H., and Sorooshian, S.: Uncertainty assessment of hydrologic model states and parameters: Sequential data assimilation using the particle filter, Water Resources Research, 41, https://doi.org/https://doi.org/10.1029/2004WR003604, 2005.

Odry, J.: TheDroplets/Snow_spatial_particle_filter: First release of the codes for the spatial particle filter, https://doi.org/10.5281/zenodo.5531771, 2021.

Odry, J., Boucher, M. A., Cantet, P., Lachance-Cloutier, S., Turcotte, R., and St-Louis, P. Y.: Using artificial neural networks to estimate snow water equivalent from snow depth, Canadian Water Resources Journal / Revue canadienne des ressources hydriques, 0, 1–17, https://doi.org/10.1080/07011784.2020.1796817, 2020.

Ohmura, A.: Physical Basis for the Temperature-Based Melt-Index Method, Journal of Applied Meteorology and Climatology, 40, 753–761, https://doi.org/10.1175/1520-0450(2001)040<0753:PBFTTB>2.0.CO;2, 2001.

605   Penny, S. G. and Miyoshi, T.: A local particle filter for high-dimensional geophysical systems, Nonlinear Processes in Geophysics, 23, 391–405, https://doi.org/10.5194/npg-23-391-2016, 2016.

Poterjoy, J.: A Localized Particle Filter for High-Dimensional Nonlinear Systems, Monthly Weather Review, 144, 59 – 76, https://doi.org/10.1175/MWR-D-15-0163.1, 2016.

Schaefli, B., Hingray, B., and Musy, A.: Climate change and hydropower production in the Swiss Alps: quantification of potential impacts

610   and related modelling uncertainties, Hydrology and Earth System Sciences, 11, 1191–1205, https://doi.org/https://doi.org/10.5194/hess-11-1191-2007, 2007.

Snyder, C., Bengtsson, T., and Morzfeld, M.: Performance Bounds for Particle Filters Using the Optimal Proposal, Monthly Weather Review, 143, 4750–4761, https://doi.org/10.1175/MWR-D-15-0144.1, 2015.

Thiboult, A. and Anctil, F.: On the difficulty to optimally implement the ensemble Kalman filter: An experiment based on many hydrological

615   models and catchments, Journal of Hydrology, 529, 1147–1160, 2015.

Turcotte, R., Fortin, L.-G., Fortin, V., Fortin, J.-P., and Villeneuve, J.-P.: Operational analysis of the spatial distribution and the temporal evolution of the snowpack water equivalent in southern Québec, Canada, Hydrology Research, 38, https://doi.org/10.2166/nh.2007.009, 2007.

Uboldi, F., Lussana, C., and Salvati, M.: Three-dimensional spatial interpolation of surface meteorological observations from

620   high-resolution local networks, Meteorological Applications, 15, 331–345, https://doi.org/https://doi.org/10.1002/met.76, _eprint: https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/met.76, 2008.

World Meteorological Organization: Guide to Instruments and Methods of Observation. Volume I –Measurement of Meteorological Variables. 2018 edition, 2018.

Zhang, Y.-F. and Yang, Z.-L.: Estimating uncertainties in the newly developed multi-source land snow data assimilation system, Journal of

625   Geophysical Research: Atmospheres, 121, 8254–8268, https://doi.org/https://doi.org/10.1002/2015JD024248, 2016.