Hydrol. Earth Syst. Sci. Discuss.,
https://doi.org/10.5194/hess-2018-121-RC2, 2018

Hydrology and Earth System Sciences
Open Access
Discussions

Interactive comment on "Covariance resampling for particle filter – state and parameter estimation for soil hydrology" *by* Daniel Berg et al.

J. Vrugt (Referee)

jasper@uci.edu

Received and published: 26 June 2018

Particle filters (PFs) have found widespread application and use for state and/or parameter estimation of dynamic system models. The premise of such filters is that they provide an exact approximation of the state forecast distribution. Yet, particle filters are not necessarily efficient as they may require a very large number of ensemble members (so-called particles) to approximate closely the evolving state distribution. This is particularly true in high-dimensional state spaces and complicates significantly the practical and/or real-time application of particle filters. What is more, particle filters are prone to sample impoverishment, that is, after a number of so-called assimilation steps, a very large number of the particles receives a negligible weight. These particles thus contribute little to the state forecast distribution and should be discarded/eliminated to

a) refocus the thrust of the filter on the high-density region of the state space, and b) maintain an adequate filter efficiency and use of CPU resources. In the past decades, different resampling methods have been proposed and/or used to periodically rejuvenate the particle ensemble and ensure an adequate tracking of the evolving state distribution. Of these, Sequential Importance Resampling (SIR) has found most application and use. This method re-samples the particle ensemble using the computed weights of the N particles. These weights are simply equivalent to the product of the prior and the likelihood of the particle's simulated trajectory. Whereas this SIR method is computationally efficient, it typically leads to a resampled ensemble with many copies of only a few of the "best" members (those with highest weights). In theory, this should not necessarily be a problem as the model operator (transition density) would disperse identical copies of the initial states as a result of the stochastic model error. This would work well in practice if the transition density of the state vector closely approximates the underlying system behavior. Unfortunately, even a modest deviation of the model operator from the actual system dynamics would deteriorate the particle ensemble to a point that most particles receive only a negligible weight. Thus, resampling is of crucial importance to periodically rejuvenate the particle ensemble and make sure that the simulated state PDF mimics closely the observed system behavior. Note that Ensemble Kalman filters do not suffer this same problem with sample impoverishment as they use a state analysis step to update the state forecasts of the N ensemble members each time an observation is becoming available.

In this paper, the authors present a new resampling method to improve the efficiency and practical application of particle filters. This resampling method stores only a single copy of the M "best" particles determined with a standard resampling method, say SIR, and simulates the $N - M$ "open spots" by drawing from a m$-variatenormaldistributionwithmeanandcovariancematrixderivedfromthe$m$-dimension$

This paper considers an important practical (and theoretical) problem in hydrologic data assimilation. This topic is relevant to HESS and should be of interest as well to an audience outside hydrology as it involves improvements to an existing method. The paper is generally well written but would benefit from careful editing. I now list my main comments.

1. The authors should consider a more realistic or appealing case study. Indeed, the present one-dimensional Richards' type flow problem with two horizontal layers is too simple to really demonstrate the advantages of the proposed covariance resampling methodology. The authors should consider a range of different state dimensionalities, m, to demonstrate that their method does not suffer from particle impoverishment. The authors should consider the Lorenz96 model with $m = 40$ state variables - this would demonstrate (or not) that the proposed resampling method works well in higher dimensional state spaces and would track closely the observed system dynamics. Such case study would make the paper much stronger and more appealing to those interested in methodological developments.

2. The authors refer to Figure 2 for a demonstration of the proposed covariance resampling method. I do not necessarily find this illustration to be particularly informative - that is - I think the authors can do a better job in detailing the proposed resampling method. The present animation assumes as if the target distribution is already well described with the present forecast distribution. In practice, this is often not true, certainly in higher dimensional state spaces. Also, I think the authors should differentiate between the state forecast density and the "true" or "unobserved" state forecast PDF. Then detail how the resampling works in practice. Personally, I always enjoy reading well-crafted algorithmic recipes (and associated coding) as those detail a step-by-step plan of how to implement the steps detailed in the main text.

3. On Page 5 (top part) the authors list some previous approaches that have been used to resample the particle ensemble. I think the authors should mention whether each of these listed approaches leave the target state PDF invariant - that is - they lead to an exact approximation of the evolving target PDF. This may not necessarily be of concern to most hydrologists but is a requirement for methods to find widespread

application and use. The same comment applies to the Introduction Section of the paper. In other words, I think it is good to emphasize that ad-hoc methods may provide results - but that such methods may not enjoy statistical underpinning.

4. I think the paper will be better if the authors replace Equation (1) with a recursive implementation of Bayes Law. This will make clear the relationship between the prior and posterior state PDF, how Equation (8) ties into this, and defines the importance weight, incremental importance weight and normalized importance weight. As it stands the current theory section omits completely the dimension of time - and this is key to state estimation.

5. I think it may be worthwhile to tie Equation (2) to the marginal likelihood. This is what you want to maximize with parameter estimation - but is of no real concern/importance for state estimation.

6. Do not understand the need for Equation (6) - and also do not necessarily directly understand how the normalized weights lead to the normalization constant. This denominator, or evidence, does not require the importance weights to add up to unity, right?

7. The authors assign a weight of 1/N to the samples drawn from the m-variate normal distribution. I am not sure whether this leaves the state PDF invariant. The authors treat as if the samples from the multivariate normal have an equal weight - this is fine if ALL N samples were drawn from the multivariate normal PDF - but the state vectors drawn from this normal PDF are combined with the existing M "best" particles of the state forecast distribution - and those latter ones do not have a weight of 1/N. This cannot be justified theoretically. So, it is of crucial importance to demonstrate that the proposed resampling method leads to the exact target PDF.

8. The present resampling method relies heavily on the simulated state forecast distribution. If this distribution does not properly approximate the actual target PDF then resampling will provide N unique samples but those state vectors are not expected

to produce a proper forecast PDF at the next time when a subsequent measurement becomes available. In other words, the present resampling method assumes that the transition density (model operator) approximates closely the true system dynamics. Once the state forecast PDF is systematically biased (likely to happen in real-world application) then the present resampling method may not necessarily enhance particle filtering results.

9. The authors use a perturbation factor, gamma, to inflate or deflate the covariance matrix of the normal resampling PDF. There is no justification for this - that is - its value is entirely subjective - indeed, one can tune gamma to provide appealing results, yet the value of gamma should guarantee an exact approximation of the target PDF (see Vrugt et al., 2013).

10. The authors use state augmentation to estimate jointly the model's state variables and parameter values. Do I conclude correctly that the authors use the normal resampling PDF to generate new parameter vectors? So, the resampling method assumes that the state variables and parameter values are multivariate normal. Would it not make sense to implement a mixture distribution instead - and estimate this distribution from the forecasted state/parameter distribution of the N particles?

11. The synthetic case study presented in the paper satisfies the assumption of a perfect model and thus transition density; in other words, the presented resampling method should work well as the forecast PDF (state/parameter) is not expected to deviate systematically from the observed data and system behavior. A real-world case study with actual measured data would provide a much stronger test of the proposed method. For example, one can use the Lorenz96 model to create an artificial data set - and then use an alternative model formulation to test, evaluate and benchmark the proposed covariance resampling method.

For now, I'll leave it with these comments. In summary, I think the authors have to investigate in more depth the statistical rigor of their resampling methodology. I have

serious concerns about the statistical validity of the proposed method - as far as I understand the details I do not think that the proposed resampling method leads to an exact approximation of the evolving target PDF. What is more, the authors should consider a more complex case study involving a higher dimensional state space and a transition density (model operator) that cannot track exactly observed system dynamics. If the authors can show that their method works well with a relatively inferior model then this would demonstrate in strongest possible terms the advantages of the proposed resampling method. Indeed, it is then when particle filters break down. Also, the authors should provide theoretical justification for the use of a multiplicative factor to scale the covariance matrix of the multivariate normal resampling PDF. This factor is entirely subjective and can be tuned so that one achieves desired results - but what is the value of this scalar for other studies? How do we know what value for gamma to take in practice?

I do not want to discourage the authors, but I believe that proper resampling necessitates the use of MCMC simulation and re-simulation of part of the historic trajectory to determine whether to accept a proposal or not. Such re-simulation is not particularly appealing, yet required to track properly the evolving state PDF. I elude to the work we published in Vrugt et al. (2013) (particle filtering with DREAM resampling/resimulation) - which ultimately led me to conclude that particle filters are not particularly useful for real-world application to complex systems - unless you have at your disposal sufficient CPU resources and can afford the use of an excessively large particle ensemble. This may then guarantee a sufficient coverage of the state space so that particle resampling can rapidly rectify systematic deviations between the forecast PDF and the "measured" state PDF as expressed in the measured data.

I hope my comments are useful to further improve this paper. As always, my comments/interpretations may be wrong! As usual, I welcome dialogue with the authors on this and/or related topics. Jasper A. Vrugt Irvine, June 26, 2018

Interactive comment on Hydrol. Earth Syst. Sci. Discuss., https://doi.org/10.5194/hess-2018-121, 2018.