Response to Referees

Andrey A Popov et al.

January 31, 2022

1 Misc. Changes

Three figures have been added. One, illustrating the optimal transport methodology in the continuous one dimensional case. We hope that this figure would help the reader visually understand the optimal transport methodology. One illustrating the Stochastic shrinkage approach. One illustrating R-Localization.

2 Response to Referee 1

2.1 General Comments

> Throughout the manuscript, the notation is inconsistent. For example, in Eq. (2) the argument if X|P while in Eq. (3) it is x|y, p. Using a consistent notation would really make the manuscript clearer and hence help the reader. Furthermore, I strongly recommend to follow the usual conventions of the data assimilation community (which, if I am not mistaken, also coincide with the journal conventions): • bold face uppercase for matrices (ex M); • bold face lowercase italic for vectors (ex v); • lowercase italic or greek letters for scalar quantities (ex n or ...); 1 • uppercase italic for sizes (ex N).

We have changed as much notation as possible without interrupting the flow of what we feel the text represents. We have left some of the more novel notation alone (such as the use of caligraphy and fraktur) as it is in line with previous work, which we strongly believe is important for consistency.

> 1.2 Regularisation or particle rejuvenation The entire method derived by the authors is designed as a sort of extension of the ETPF of Reich (2013), therefore I am not surprised that the authors adopt the same terminology. Nevertheless, one should keep in mind that ”particle rejuvenation” is not a new method invented by Reich and colleagues, it is just a new fancy name for one of the regularisation methods that have been introduced in the 2000s by Musso and colleagues. See, in particular, the chapter ”Improving Regularized Particle Filters” by Musso et al. in the book ”Sequential Monte Carlo Methods in Practice” by Doucet, Freitas, and Gordon (isbn: 978-0-387-95146-1). This historical perspective does not appear in the manuscript and I think that this is missing. In addition, I would like to mention that in my opinion, ”regularisation” is a better name than ”particle rejuvenation”, because I

> think that it better describes what is actually happening (in practice the posterior pdf is regularised). Of course, I acknowledge that the authors should have the right to choose which name they use!

We greatly appreciate the historical context given by the reviewer. While we did not make the connection in the manuscript before, we are aware of the strong connections between regularization and rejuvenation. We however believe that they are not one and the same. From our point of view regularization is an action on the probability distribution, while rejuvenation is an action on the samples. We added two lines about this in the introduction of the paper.

> 1.3 Particle filter and localisation The curse of dimensionality, mentioned in the introduction, is one of the main obstacles to the application of PF algorithms to high-dimensional problems (see Snyder et al., 2008, doi: 10.1175/2008MWR2529.1). More recently, a lot of studies have tried to apply localisation techniques in the PF to circumvent the curse of dimensionality (see in particular the review by Farchi and Bocquet, 2018, doi: 10.5194/npg-25-765-2018), which has lead to successful applications of PF algorithms to high-dimensional problems. In the manuscript, there is no discussion about the scalability of the new method, at a point that a naive reader could think that this new method could be applied as is to high-dimensional problems. This aspect must be clarified in the manuscript. In particular, I think that the author should explain whether the new method should (i) replace localisation or (ii) be used in conjunction with localisation.

We have cited both the papers provided, and have included a short discussion of high-dim particle filters in the introduction.

> 1.4 Numerical experiments and conclusions After reading the article, I am left with the impression that the numerical experiments are very brief. Nowadays, a couple of experiments with the Lorenz 1963 model is not enough to publish an article in a data assimilation journal like NPG. Therefore, I think that Section 5 (with the numerical experiments) has to be extended. At the very least, I think that a test series with the Lorenz 1996 model should be included, which would help illustrate a potential discussion about the scalability of the new method. Of course, implementing a PF algorithm with 40 variables (the classical size of the Lorenz 1996 model) is a challenge without localisation, which is why it is probably more reasonable to start with only 10 or 8 variables in the Lorenz 1996 model. In addition, the results of the numerical experiments are barely discussed in the manuscript, and the conclusions are very short (only two paragraphs!). I think that this is clearly not enough and that the authors should provide appropriate discussion of the results and conclusions.

In line with this notion and the next, we have added a localized FETPF, which we name the LFETPF. We have added experiments with the Lorenz ’96 equations, and have shown that the new methodology can produce good results with as few as 4 dynamical ensemble members.

We are very grateful to the reviewer for pushing us in this direction, as we believe this experiment makes the results provided in the manuscript stronger.

3 Technical Corrections

> L. 12-13 ”The curse of dimensionality”. At this point, the authors could cite Snyder et al. (2008).

We agree and have cited the paper.

> L. 17 ”by discarding information about the underlying dynamical system”. In the EnKF, the information about the dynamical system is taken into account through the use of the dynamical model (for forecast) and the observation operator. Could the authors explain what they mean here?

This is a largely philosophical point, which we believe that is worth making. Most dynamical systems of interest have compact support on complex and hard-to-describe manifolds. This is the "information" that the agent performing the inference has in possession, and must use for the inference to be Bayesian. This "information" is discarded when the ensemble of realizations is assumed Gaussian. Thus, the information about the dynamical system is only partially taken into account. We have modified this line to say "by partially discarding".

> L. 18 ”that lives in Rn”. At this point, n is not defined. I would delay this aspect until section 2 where the different spaces are introduced.

We have removed the reference to n and describe it as the "state space".

> LL. 19-20 ”into our assumed posterior normal distribution”. If both the prior and the likelihood are assumed Gaussian, then the posterior is Gaussian (this is not an assumption).

This is again a matter of philosophy. As we are assuming a Gaussian prior, and we know that in actuality the distribution is not and cannot be Gaussian, we know that we are not performing Bayesian inference. Thus, Bayes' theorem no longer applies to what we are doing. We therefore have to make assumptions about the posterior.

We recognize the point made by the reviewer, and we understand that his point of view would be more common in the community. However, we would like to be technically correct rather than conventionally correct.

> L. 21 In the reference list, two elements match the key ”Popov et al., 2020”. This should be corrected.

In our understanding this is an issue with the Copernicus bibliography style. As we have no way of fixing this without modifying the provided style, we do not have a way of modifying this.

> L. 23 I would suggest the following stylistic transformation ”(ETPF) (Reich, 2013)” $\rightarrow$ ”(ETPF, Reich, 2013)”.

We agree with the reviewer and have fixed it as best as is possible within the Copernicus style convention.

> L. 25 ”the ensemble limit”. To my knowledge, this is not clearly defined in the data assimilation community (even though I agree that this is understandable). I would recommend to explicitly define this term with, for example: ”in the limit of an infinite ensemble size”.

We agree with the reviewer and have made this change.

> L. 26 ”This means that”. The logical connection is incorrect here.

We changed the wording to add the modal "possible".

> LL. 45-46 ”and the supports of the probability densities $\pi(Xf)$ and $\pi(Xa)$ are subsets of the respective spaces”. This could be reformulated because as is, one could understand that it is possible that the support of these pdfs are not subsets of the respective spaces.

We have reworded this in a clearer way.

> LL. 50 and 54 The sum’s limits are incorrect in Eq. (2) and (3).

We have corrected this.

> L. 57 ”ensemble of weights”. I would suggest to name it ”weight vector”.

While we agree that notationally this should be a vector, we prefer using ensemble of weights rather than weight vector to make the link more obvious between the ensemble and the corresponding weights of each ensemble member.

> L. 58 ”Using (3) and (4) empirical estimates of the posterior mean and covariance”. This seems weird. I would suggest a reformulation.

We have reformulated this.

> L. 59 In equation (5), the authors use the prefactor N/(N - 1) to debias the sample covariances. However, for weighted sample the prefactor to debias the sample covariances is 1/(1 - ww), which is equal to N/(N - 1) only in the case where the weights are uniform w = 1N/N. Could the authors justify this choice?

This formula has been corrected. As we do not use this formula in the implementation, this had no consequences on the results.

> L. 61 ”The goal of particle filtering (with resampling)”. I would rather say that this is the goal of resampling.

Since what we are referring to is the goal of particle filtering and also resampling (which is one solution to particle filtering), we have reworded this as "Our goal".

> L. 62 ”the the posterior...” → ”the posterior...”

We have made this correction.

> LL. 65-66 ”We impose that the empirical mean (5) is preserved by (6)”. This is in general not possible with classical resampling algorithms. Of course, this is possible when using a linear ensemble transformation like Eq. (8) but at this point in the manuscript, Eq. (8) is not yet introduced!

This has been changed to "our goal" in the begining, thus circumventing the need for a discussion about resampling.

> L. 71 ”...” At this point, it could be interesting to remind the reader that T has positive coefficients.

Agreed. We have done this change.

> L. 73 L. 78

Both these formulae have been changed.

> L. 94 ”the factor τ”. The authors could mention that τ is usually called the bandwidth.

We have added this.

> L. 99 The second line of Eq. (13) is just the transpose of the first line, or did I miss anything?

It is indeed the case. We include it for the sake of completeness. We believe that it does not leave any guesswork for the reader.

> L. 97-99 It is true that the extra term ensures that the regularisation noise has zero mean. However, the author should mention that this extra term does modify the sample covariance of the noise. The same holds for Eq. (26).

Taking out the sample mean does not change the sample covariance. We are not sure what the reviewer means by this.

> L. 125 ”THis” → ”This”

This has been corrected.

> LL. 125-126 ”This is because we are now incorporating more prior information P in the form of climatological information”. From what I understand, such additional information was missing until now. Hence Eq. (2) – and the few following equations – should be written ”π Xf(X)” instead of ”π Xf(X|P)”, right?

This “additional information” is always present. Bayesian inference needs to take into account all information. Most of the time in practice, the vast majority of information is ignored, (meaning that almost all practical inference is in fact non-Bayesian) but that does not mean it is not there in the ideal case. Instead of ignoring this information, we recognize that it is always there. We refer the reviewer to the book “Probability theory: The logic of science” by Jaynes for more information.

> LL. 134-135 ”is assumed to be the sample mean of the dynamic ensemble” I would mention here that this choice is necessary to preserve the mean of the augmented ensemble.

We agree and have fixed this.

> L. 136 ”by construction and (13), thus requiring that only the synthetic ensemble anomalies need to be determined”. This is hardly understandable. I would suggest a reformulation.

We have reformulated this.

> L. 158 Please define the $\wedge$ symbol in Eq. (28).

We have used $\min$ instead of the AND operation.

> L. 162 ”Note that if P = ...”. How is the division by zero in Eq. (28) handled in this case?

We have clarified this.

> L. 162 ”In such a framework the scaling parameter...”. I think that a separation is needed here to indicate that this does not apply to P = ...

A paragraph separation was added.

> LL. 170-171 I understand that the scaling of P has no impact on the definition of ... defined by Eq. (25), but does it have an impact on γ defined by Eq. (28)?

We have expanded the remark to address that γ is not scaled by scaling P.

> L. 175 ”where the optimal transport matrix ... is computed by solving (9)”. This is a very concise description and I think that additional description is needed, because this is one of the core elements of the new method. At the very least, it should be mentioned that a posterior weight is needed for all members of the augmented ensemble, and that the formulation of (9) needs to be adjusted to take into account non-uniform prior weights. In addition, note that ”R” should be replaced by ”R”.

This paragraph has been expanded to mention the prior information weights and the importance sampling procedure.

> LL. 181-182 ”In effect we are able to avoid ensemble collapse by enhancing the empirical measure distribution (32) with new prior information”. From a theoretical perspective, this has not been proven. This is only illustrated in Section 5 using numerical illustration with a 3-variable model. At this point, the lack of discussion about scalability really hurts (see general comment in section 1.3).

We have changed this statement to be much weaker. “We attempt to avoid”.

> LL. 227 Why not use the time-averaged RMSE, defined as ...

We looked at the four major books on data assimilation:

- Evensen (2009) uses spatio-temporal RMSE
- Reich and Cotter equation (1.13) is also spatio-temporal RMSE (it is mis-named in the text, but it is spatio-temporal). See also example 8.9 again given the spatio-temporal RMSE equation.
- Law Stuart and Zygalakis uses time-averaged RMSE
- Asch, Bocquet and Nodet use time-averaged RMSE

We hope the reviewer agrees that it seems like there is no clear consensus in the literature. We prefer the spatio-temporal RMSE.

> LL. 230 ”with the optimal rejuvenation factor of $\tau = 0.04$”. I highly doubt that $\tau = 0.04$ is optimal for all values of the ensemble size N. In order to make a fair comparison, the value of τ should be optimally tuned for each ensemble size N. If not, we give an unfair advantage to any of the method.

We have removed the "optimal" word and instead cited the ETPF paper where the author used the same value for all experiments. Note that since our rejuvenation is a modification of the cited paper, the parameter is a square of that used in the cited paper.

> LL. 235 ”Results in Figure 1 show...”. I think that this test series (and the second one as well) is lacking a baseline. For this small 3-variable model, the baseline could be the score obtained with a classical PF (for example the SIR or Bootstrap filter) with a very large ensemble (typically more than 103 particles) and with optimally tuned regularisation.

We have added SIR baselines to the first two experiments.

> LL. 241-243 ”For a low ensemble size... as compared to the ETPF”. From what can be seen in Fig. 2, there is a difference between, for example, $\alpha = 1$ and $\alpha = 1.2$ (the latter being more accurate). It is possible that this difference falls into the variability between the 20 independent runs, but this is not clearly explained in the text or in the figure.

Both points are outside of 3 standard deviations of each other. We have added a comment that all reported differences are for values more than 3 standard deviations apart.

> LL. 259-260 ”We believe that the stochastic covariance shrinkage approach to importance sampling can be used not just for particle rejuvenation in the ETPF, but in other particle filters as well”. Let us take the example of the most basic PF, the SIR filter. During the resampling step, some particles (typically those with low weight) are discarded and replaced by other particles (typically those with high weight). If applying the new proposed method, this would unavoidably lead to replacing original particles by synthetic particles, which is probably not something that is desirable. With this small example, I hope that I have convinced the authors that more discussion is needed here regarding the application of the new method to other PFs than the ETPF.

We agree and this piece of text has been removed.

> Figs. 1 and 2 Some of the lines cannot be distinguished when the manuscript is printed in black and white. This should be corrected. In addition, I would recommend a log scale for the x-axis and I would recommend to show the grid for clarity.

We have remade all the figures with these suggestions in mind.

4 Response to Referee 2

> 1. The discussion of convergence of the method is rather short (Ln 235). In principle, for a large number the dynamical particles the estimator will give weight 0 to the synthetic particles.Is that correct? Is this the reason of the convergence? The authors should discuss how the estimator varies in the experiments as a function of the number of particles. In a regime with small/medium number of particles, there should be two effects that the methodology should lead to a sub-optimal filter. The synthetic particles are sampled from a Gaussian distribution which may deteriorate the performance of the filter for non-Gaussian forecast predictions. In several dynamical systems the prediction covariance varies with time, then the use of prior information with a climatological covariance should give a sub-optimal filter. The authors should discuss these limitations and evaluate them in the experiments.

We have added a section on convergence, and showed that the FETPF converges in the limit of dynamical ensemble size (if some conditions are met).

> 2. A plot and/or a discussion in the experiments about the values given by the shrinkage estimator (which is used as weight in the dynamical and synthetic particles) are required.

We have added histograms of the distributions in a new plot.

> 3. To my understanding of the methodology, the authors are not using the sampled particles as a rejuvenation of the sequential filter but just to improve the inference step in the ETKF. The synthetic particles are not used in the prediction step,is this correct?. I had in mind that in the ETPF, the rejuvenated particles were used in the prediction step? Could the authors discuss this point in the manuscript?

Rejuvenation for the ETPF refers to a stochastic regularization by adding random perturbations to the existing ensemble after the inference step. Purely random data is not predicted upon. In the FETPF the only difference is the synthetic data is combined during the inference step instead of after, in a more informed manner. We have added a discussion of regularization that hopefully clarifies this a bit more in the text.

> 4. Sampling perturbations from climatological covariances in geophysical applications may give physically unrealistic states. The authors should comment how the method could be extended to be applied in realistic applications.

We have added a remark about physically unrealistic realization in the convergence section. In both the ETPF and FETPF, in the limit of ensemble size, physically realistic realization are generated with probability one. This is now explicitly stated.

> 5. Experiments are rather insufficient. An experiment with a nonlinear observational operator or any other configuration that result in a non-Gaussian posterior distribution would also be illustrative. Experiments with different observational errors (particularly smaller ones) are also required.

We have significantly expanded on the experiments. We have added two experiments with Lorenz '96 with linear and non-linear observations.

As for small observation error: if we have a small Gaussian observation error, then the analysis distribution would tend towards being Gaussian. As we wish to examine the case of a non-Gaussian analysis distribution, a large observation error is chosen.

> 6. The authors may compare the performance of the experiment with a standard inflation technique (with optimal inflation factor) as a baseline. They mentioned in passing that the ETKF deteriorates with inflation, a deeper examination of this point could be useful to further motivate the proposed method.

We are not sure as to what the reviewer is referring. No inflation is performed on the dynamical ensemble anywhere in the paper. Additionally, the Lorenz '63 cases discussed are sufficiently non-Gaussian such that no ensemble Kalman-based filter ever converges for the given problem.

New experiments with the Lorenz '96 equations do look at the ETKF, and thus approximately optimal inflation factors are used.

> 7. The experiments comparing the Laplacian with Gaussian sampling, and the ones comparing two climatological covariances versus one do not appear to have a conclusion.

All conclusions have been rewritten and the experiments part significantly expanded. This is hopefully not an issue anymore.

> 8. I had in mind covariance shrinkage as a covariance regularization method for small samples. However, the low-dimensional example shown in the manuscript does not appear to evaluate the regularization of the long-distance correlations. The authors should evaluate at least in a 40-dimension Lorenz-96 the performance of the methodology that they are proposing.In principle one expects a larger impact of the covariance shrinkage estimator in higher dimensional systems.

40-variable localized Lorenz '96 experiments have been added, with both linear and non-linear observation operators. We hope that these new experiments satisfy all concerns.