

A Stochastic Covariance Shrinkage Approach to Particle Rejuvenation in the Ensemble Transform Particle Filter

Andrey A Popov¹, Amit N Subrahmanya¹, and Adrian Sandu¹

¹Computational Science Laboratory, Department of Computer Science, Virginia Tech, 2202 Kraft Drive, Blacksburg, VA, 24060, USA

Correspondence: Andrey A Popov (apopov@vt.edu)

Abstract. Rejuvenation in particle filters is necessary to prevent the collapse of the weights when the number of particles is insufficient to properly sample the high probability regions of the state space. Rejuvenation is often implemented in a heuristic manner by the addition of ~~stochastic samples that widen random noise that widens~~ the support of the ensemble. This work aims at improving canonical rejuvenation methodology by the introduction of additional prior information obtained from climatological samples; the dynamical particles used for importance sampling are augmented with samples obtained from stochastic covariance shrinkage. ~~The ensemble transport particle filter, and its second order variant, are extended with the proposed rejuvenation approach~~ A localized variant of the proposed method is developed. Numerical experiments with the Lorenz '63 model show that modified filters significantly improve the analyses for low dynamical ensemble sizes. Furthermore, localization experiments with the Lorenz '96 model show that the proposed methodology is extendable to larger systems.

1 Introduction

Ensemble-based data assimilation (Asch et al., 2016; Law et al., 2015; Reich and Cotter, 2015) aims to estimate ~~our uncertainty about~~ the state of some dynamical system in a Bayesian framework, and describes the uncertainty through an ensemble of possible states. ~~The assumed underlying distribution of these states is taken to be our uncertainty.~~

~~Oftentimes the ensemble cannot describe the distribution of our~~ Describing the distribution of state uncertainty to sufficient accuracy. ~~The requires very large ensembles, a phenomenon referred to as the curse of dimensionality (Tan et al., 2018) ensures that true descriptions of general high-dimensional distributions stay out of our reach (Tan et al., 2018; Snyder et al., 2008).~~ Several techniques such as the principal principle of maximum entropy (Jaynes, 2003) ~~exist in order to~~ attempt to alleviate ~~some of the this~~ burden by prescribing a distribution ~~to a set of data by constraining the choice with known or estimated quantities and qualities. For instance, the~~ constrained by known information. The ensemble Kalman filter (Burgers et al., 1998; Evensen, 1994, 2009), ~~could be thought of as an abuse of the principle of maximum entropy: by discarding information about the underlying dynamical system, and assuming that the ensemble only gives information about a mean (that lives in \mathbb{R}^n), and a covariance, the underlying distribution of our uncertainty is taken to be normal. In such a way, any~~ constraints the underlying distributions only by the ensemble mean and covariance; the application of Bayes' rule ~~has to transform our~~ transforms an assumed prior normal distribution into ~~our an~~ assumed posterior normal distribution. Recent attempts to apply particle filters to

25 high dimensional systems (Bocquet, 2018; Van Leeuwen, 2009; Van Leeuwen et al., 2019) have seen some success. However, particle filters are not yet competitive with other state-of-the-art methods such as the ensemble Kalman filter.

Previous work (Popov et al., 2020) has focused on augmenting the information represented by the ensemble with information derived from covariance shrinkage through a surrogate ensemble in the ensemble ~~transport~~ transform Kalman filter. In this paper, we extend this idea to the ensemble transport particle filter (ETPF) ~~(Reich, 2013). The ETPF attempts to use,~~
 30 Reich (2013)). ETPF transports a given ensemble that represents the posterior distribution using importance sampling (Liu, 2008), ~~not for a solution to the problem of Bayesian inference, but to simply transport a given ensemble into one that is equally sampled from some distribution to another, equally sampled, ensemble~~ whose moments, in the ~~ensemble limit~~ limit of infinitely many particles, approach the moments of the correct posterior distribution. ~~This means that the underlying distribution from which our posterior is sampled, could be potentially very far, in information distance, to the optimal posterior distribution, for a finite ensemble. Like all particle filters, ETPF is susceptible to weight collapse.~~

This work explores a new approach to particle rejuvenation, which is necessary to prevent weight collapse in particle filters. Rejuvenation in particle filters is a particular type of stochastic regularization (Musso et al., 2001), and is typically implemented in a heuristic manner. Instead of heuristics, ~~the~~ our approach makes use of prior information to enrich the ensemble subspace. ~~Our contributions~~ The new contributions of this work are as follows: (1) we introduce an alternative way of performing particle
 40 rejuvenation in ETPF by incorporating climatological covariance information. ~~We;~~ (2) we accomplish this by augmenting the dynamical (model) ensemble with synthetic anomalies with optimal scaling, accompanied by a statistically correct estimator. ~~We;~~ and (3) we show that this ~~method of performing particle rejuvenation~~ rejuvenation method significantly improves the analysis ~~RMSE~~ quality for low dynamical ensemble sizes.

This paper is organized as follows. Section 2 reviews the concept of Bayesian inference with the addition of prior infor-
 45 mation, and its ~~use in implementation in the context of~~ importance sampling. Section 3 ~~introduces~~ discusses the ensemble transform particle filter and its canonical rejuvenation heuristic. The concept of stochastic covariance shrinkage is proposed in Section 4, and ETPF is extended to make use of this shrinkage. Numerical experiment results are reported in Section 5. ~~Finally,~~ concluding Concluding remarks are drawn in Section 6.

2 Optimal coupling with prior information and the ensemble transform particle filter

50 Bayesian inference (Jaynes, 2003) aims at transforming prior information about the state of a system (represented by the distribution of a random variable X^f), additional qualitative and quantitative information (P), ~~which we will use to stand for,~~ and information obtained by observing the system (Y), into combined posterior information (X^a):

$$\pi_{X^a}(X | Y, P) = \frac{\pi_{Y|X}(Y | X, P) \pi_{X^f}(X | P)}{\pi_Y(Y | P)}, \quad (1)$$

where ~~$\pi_{X^f}(x|p)$~~ $\pi_{X^f}(X|P)$ represents the prior state probability density conditioned by all other relevant information, and
 55 ~~$\pi_{Y|X^f}(y|x,p)$~~ $\pi_{Y|X^f}(Y|X,P)$ is the observational likelihood conditioned by the ~~the~~ forecast and the prior information. Here

we consider the finite dimensional case where $X^f, X^a \in \mathbb{R}^n$, $Y \in \mathbb{R}^m$, ~~and~~ where the supports of the probability densities π_{X^f} and π_{X^a} are subsets of the respective spaces.

Classical particle filtering (Reich and Cotter, 2015) represents state distributions by collections of particles, i.e., ensembles of samples. Specifically, consider an ensemble of N^f particles ~~$\mathbf{X}^f = [\mathbf{X}_1^f, \dots, \mathbf{X}_{N^f}^f] \in \mathbb{R}^{n \times N^f}$~~ $\mathbf{X}^f = [\mathbf{x}_1^f, \dots, \mathbf{x}_{N^f}^f] \in \mathbb{R}^{n \times N^f}$. The
60 prior distribution density π_{X^f} is approximated weakly by the corresponding empirical measure

$$\pi_{\hat{X}^f}(X | P) = \sum_{\underline{i=j} \underline{j=1}}^{N^f} w_j^f \delta_{X - \mathbf{X}_j^f}, \quad (2)$$

where w_j^f for $1 \leq j \leq N^f$ are the prior importance weights associated with each particle. Similarly, the posterior density is approximated weakly by an empirical measure based on the same sample values (particle states) but with different posterior importance weights w_j^a for $1 \leq j \leq N^f$:

$$\pi_{\hat{X}^a}(\underline{x} \underline{X} | \underline{y} \underline{Y}, \underline{p} \underline{P}) = \sum_{\underline{i=j} \underline{j=1}}^{N^f} w_j^a \delta_{\underline{x} - \underline{\mathbf{X}}_j^f} \underline{X} - \underline{\mathbf{X}}_j^f, \quad (3)$$

The posterior importance sampling weights are obtained from eq. (1):

$$w_j^a \propto \pi_Y(\underline{y} \underline{Y} | X_j^f, P) \pi_P(P | X_j^f) = w_j^f \pi_Y(\underline{y} \underline{Y} | X_j^f, P). \quad (4)$$

The ensemble of weights is denoted by ~~$\mathbf{w} = [w_1, \dots, w_{N^f}]^T$~~ , ~~and \mathbf{w}^f and \mathbf{w}^a~~ $\mathbf{w} = [w_1, \dots, w_{N^f}]^T$, ~~and \mathbf{w}^f and \mathbf{w}^a~~ refer to the forecast and analysis weights respectively. Using (3) and (4) empirical estimates of the posterior mean and covariance,

$$\begin{aligned} \bar{\mathbf{x}}^a &= \sum_{j=1}^{N^f} w_j^a \mathbf{X}_j^f, \\ \Sigma_{X^a} &= \frac{1}{1 - \mathbf{w}^{a,T} \mathbf{w}^a} \mathbf{X}_j^f (\text{diag}(\mathbf{w}^a) - \mathbf{w}^a \mathbf{w}^{a,T}) \mathbf{X}_j^{f,T}, \end{aligned} \quad (5)$$

respectively, are obtained by the importance sampling approach (Liu, 2008).

~~The goal of particle filtering (with resampling)~~ Our goal is to find an analysis ensemble $\mathbf{X}^a \in \mathbb{R}^{n \times N^a}$ of N^a realizations of the random variable X^a that represents the posterior distribution π_{X^a} with equal weights. Specifically, the ~~the~~ posterior density is approximated weakly by the empirical measure

$$\pi_{\hat{X}^a}(\underline{x} \underline{X} | \underline{p} \underline{P}) = \sum_{j=1}^{N^a} \frac{1}{N^a} \delta_{\underline{x} - \underline{\mathbf{X}}_j^a} \underline{X} - \underline{\mathbf{X}}_j^a, \quad (6)$$




Figure 1. A visual representation of the continuous optimal transport procedure. The probability distribution π_{X^f} is transported into the distribution π_{X^a} through the optimal transport mapping \mathbf{T} . The discrete version of the solution is given by (9).

where the importance sampling weights are uniform and equal to $1/N^a$ (so as to be equally likely), and \hat{X}^a is the random variable corresponding to this measure. We impose that the empirical mean (5) is preserved by (6):

$$\bar{\mathbf{x}}^a = \sum_{j=1}^{N^a} \frac{1}{N^a} \mathbf{X}_j^a = \sum_{j=1}^N w_j^a \mathbf{X}_j^f = \mathbf{X}^f \mathbf{w}^a. \quad (7)$$

The optimal coupling (McCann et al., 1995; Reich and Cotter, 2015) between the prior empirical distribution eq. (2) and the posterior empirical distribution eq. (6), can be defined as an ensemble transformation,

$$\mathbf{X}^a = \mathbf{X}^f \mathbf{T}^*, \quad (8)$$

where $\mathbf{T}^* \in \mathbb{R}^{N^f \times N^a}$ is the solution to the optimal transport Monge-Kantorovich problem (Villani, 2003). It is important to see that each element of \mathbf{T}^* is positive. The discrete optimal transportation problem is

$$\mathbf{T}^* = \arg \min_{\mathbf{T}} \sum_{\substack{1 \leq j \leq N^f \\ 1 \leq k \leq N^a}} \mathbf{T}_{j,k} \|\mathbf{X}_j^f - \mathbf{X}_k^a\|_2^2 \quad (9)$$

subject to $\mathbf{T} \mathbf{1}_{N^a} = N^a \mathbf{w}^a$, $\mathbf{T}^T \mathbf{1}_{N^f} = \mathbf{1}_{N^a}$, $\mathbf{T}_{i,j} \geq 0$,

where the distance measure of squared Euclidean distance is taken for a provably unique solution to the Monge-Kantorovich problem to exist (McCann and Guillen, 2011). The vector of ones of size q is represented by $\mathbf{1}_q$. The problem eq. (9) is a linear programming problem. A visualization of the optimal transport between two continuous distributions is given in Figure 1.

The




Figure 2. A visual representation of \mathbf{R} -localization in the ETPF. The left panel represents the localization radius around the i th state variable represented by the black dot, with the observations represented by the open red dots. The right panel represents the decorrelation of \mathbf{R}^{-1} along the diagonal.

The standard ETPF (Reich, 2013) makes the assumption that the prior and posterior ensemble sizes are the same, $N := N^a = N^f$ in (9). The discrete optimal transport eq. (8) formulation begets a mapping $\mathbf{X}^a = \Psi_{N^f, N^a}(\mathbf{X}^f)$ that, in the limit of ensemble size ($N^f = N^a \rightarrow \infty$), converges weakly to a mapping Ψ , such that $X^a = \Psi(X^f)$ which has the exact desired distribution given by eq. (1) (Reich and Cotter, 2015, Theorem 5.19). We believe, but don't prove, that this is likely when N^f and N^a are not equal.

The standard ETPF (Reich, 2013) makes the assumption that the prior and posterior ensemble sizes are the same, $N := N^a = N^f$ in (Reich and Cotter, 2015, Theorem 5.19). A second order extension to the ETPF (which we will call ETPF2 here) (Acevedo et al., 2017) modifies the optimal transport equation (8) as follows:

$$\mathbf{X}^a = \mathbf{X}^f (\mathbf{T}^* + \mathbf{D}), \quad (10)$$

where the additional term \mathbf{D} is a matrix that ensures that the empirical covariance estimate Σ_{X^a} from (5) is preserved by (6).

2.1 Localization

In high dimensional geophysical problems, spatial error correlations decrease with increasing spatial distance between states. Due to the undersampled nature of the ensemble, these correlations may not be accurately approximated. Localization allows to strictly enforce the shrinking of correlations between distant states. For localization in the ETPF, we follow the \mathbf{R} -localization formulation given in (Reich and Cotter, 2015; Acevedo et al., 2017). Figure 2 provides an illustration of \mathbf{R} -localization, with the full procedure described below.

105 Typically, the observation error distribution is assumed to be unbiased and Gaussian, with the probability density used to compute the weights in (4) for particular realizations of the observation \mathbf{Y} , and of the state \mathbf{X} defined as

$$\pi_Y(\mathbf{Y}|\mathbf{X}, P) = \frac{1}{\sqrt{|2\pi\mathbf{R}|}} e^{-\frac{1}{2}(\mathbf{Y}-\mathcal{H}(\mathbf{X}))^T \mathbf{R}(\mathbf{Y}-\mathcal{H}(\mathbf{X}))}, \quad (11)$$

where \mathcal{H} is the observation operator. In this case it can be fully parametrized by the observation error covariance $\mathbf{R} \in \mathbb{R}^{m \times m}$, where m is the number of observations.

110 We assume that the observations are uncorrelated, making $\mathbf{R} = \text{diag}_j R_{j,j}$ a diagonal matrix. For the ℓ^{th} state variable $x^{(\ell)}$, we define the localized observation error covariance matrix \mathbf{R}_ℓ via:

$$\mathbf{R}_\ell^{-1} = \text{diag}_{j=1,\dots,m} \{ \rho(d(\ell, j)/r) \} \circ \mathbf{R}^{-1}, \quad (12)$$

where d is some distance function defined between the ℓ^{th} state space variable and the j th observation space variable, r is the localization radius, and ρ is a decorrelation function. In this work we use the Gaspari-Cohn decorrelation function (Gaspari and Cohn, 1999)

115 The localized inverse of the observation error covariance (12) is then used in the generation of the weight ensemble \mathbf{w}_ℓ^a similar to (4).

A different transform matrix is computed for each variable. Specifically, consider the ensembles the ℓ^{th} state space variables:

$$\mathbf{x}^{f,(\ell)} = \begin{bmatrix} x_1^{f,(\ell)} & \dots & x_{N^f}^{f,(\ell)} \end{bmatrix}, \quad \mathbf{x}^{a,(\ell)} = \begin{bmatrix} x_1^{a,(\ell)} & \dots & x_{N^a}^{a,(\ell)} \end{bmatrix}. \quad (13)$$

120 The Monge problem formulation in (9) is replaced by the localized formulation for the ℓ^{th} variable,

$$\mathbf{T}_\ell^* = \arg \min_{\mathbf{T}} \sum_{\substack{1 \leq j \leq N^f \\ 1 \leq k \leq N^a}} \mathbf{T}_{j,k} \left\| x_j^{f,(\ell)} - x_k^{a,(\ell)} \right\|_2^2 \quad (14)$$

$$\text{subject to } \mathbf{T} \mathbf{1}_{N^a} = N^a \mathbf{w}_\ell^a, \quad \mathbf{T}^T \mathbf{1}_{N^f} = \mathbf{1}_{N^f}, \quad \mathbf{T}_{i,j} \geq 0,$$

with the analysis ensemble of the ℓ^{th} state space variable given by:

$$\mathbf{x}^{a,(\ell)} = \mathbf{x}^{f,(\ell)} \mathbf{T}_\ell^*. \quad (15)$$

3 Particle Rejuvenation in ETPF

125 Particle and ensemble-based filters often underrepresent uncertainty (Asch et al., 2016) due to the relatively small number of samples when compared to the dimension of the state and data spaces. Over several data assimilation cycles multiple **partiele**

particles start carrying either unimportant or redundant information, which leads to weight collapse or to ensemble degeneracy (Strogatz, 2018). To alleviate these effects, methods such as inflation (Anderson, 2001; Popov and Sandu, 2020), rejuvenation (Reich, 2013), and resampling (Reich and Cotter, 2015) (Reich and Cotter, 2015; Attia and Sandu, 2015) have been developed.

In order to avoid ensemble collapse, ETPF employs a particle rejuvenation approach (Acevedo et al., 2017; Reich, 2013; Chustagulprom et al., 2016) that perturbs the analysis ensemble by a random sampling from the ensemble of prior anomalies,

$$\mathbf{X}^a \leftarrow \mathbf{X}^a + \sqrt{\frac{\tau}{N-1}} \mathbf{A}^f \boldsymbol{\eta} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T), \quad (16)$$

where $\boldsymbol{\eta} \sim (\mathcal{N}(0, 1))_{N \times N}$ is a matrix of i.i.d. samples from the standard normal distribution of size N , the rejuvenation factor τ is treated as (also called the bandwidth parameter) is a hyperparameter that controls the magnitude of the correction, and the ensemble anomalies

$$\mathbf{A}^f = \mathbf{X}^f (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T), \quad (17)$$

are defined as the ensemble of deviations from the sample mean. Of note is the fact that the extra term $(\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T)$ in (16) ensures that the introduction of the random matrix $\boldsymbol{\eta}$ does not modify the mean of \mathbf{X}^a . This is due to the fact that,

$$(\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \mathbf{1}_N = \mathbf{0}_N, \quad \mathbf{1}_N^T (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) = \mathbf{0}_N^T. \quad (18)$$

Notice that if we define the matrix, Note that defining the matrix

$$\mathbf{B} := \sqrt{\frac{\tau}{N-1}} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \boldsymbol{\eta} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T), \quad (19)$$

it is possible allows to write the ETPF with standard rejuvenation (16) as follows:

$$\begin{aligned} \mathbf{X}^a &= \mathbf{X}^f \mathbf{T}^* + \sqrt{\frac{\tau}{N-1}} \mathbf{A}^f \boldsymbol{\eta} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \\ &= \mathbf{X}^f \mathbf{T}^* + \sqrt{\frac{\tau}{N-1}} \mathbf{X}^f (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \boldsymbol{\eta} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \\ &= \mathbf{X}^f \left[\mathbf{T}^* + \sqrt{\frac{\tau}{N-1}} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \boldsymbol{\eta} (\mathbf{I}_N - N^{-1} \mathbf{1}_N \mathbf{1}_N^T) \right] \\ &= \mathbf{X}^f \tilde{\mathbf{T}} \quad \text{with} \quad \tilde{\mathbf{T}} := \mathbf{T}^* + \mathbf{B}. \end{aligned} \quad (20)$$

with the The matrix \mathbf{B} acting acts as a stochastic perturbation of the optimal transport operator \mathbf{T}^* ; such that. The choice of \mathbf{B} preserves the constraints linear constraints of the Monge-Kantorovich problem eq. (9) since $\mathbf{1}_N^T \mathbf{B} = \mathbf{0}_N^T$ and $\mathbf{B} \mathbf{1}_N = \mathbf{0}_N$ in eq. (9) because of the results in due to the property (18). This, of course, immediately calls into question the optimality of the




Figure 3. A representation of the mixing of the dynamical ensemble \mathbf{X}^f with the synthetic ensemble \mathcal{X}^f . The dynamical ensemble comes from the propagation of a dynamical system (here the Lorenz '63 system on the left) while the synthetic ensemble comes from a climatological distribution (here represented by the bell curve on the right).

transport for a finite ensemble, as adding this type of \mathbf{B} matrix is perturbing the transport mapping $\tilde{\mathbf{T}}$ away from the optimum \mathbf{T}^* .

150 4 Particle Rejuvenation Through Stochastic Shrinkage

In the context of ensemble methods, covariance shrinkage (Nino-Ruiz and Sandu, 2018, 2015; Ruiz et al., 2014) is used, similar to other canonical covariance tapering techniques such as inflation (Anderson, 2001; Popov and Sandu, 2020), localization (Anderson, 2012; Hunt et al., 2007; Nino-Ruiz and Sandu, 2017; Nino-Ruiz et al., 2015; Petrie, 2008; Zhang et al., 2010), to enrich the information represented by an undersampled covariance matrix.

155 From a Bayesian perspective, covariance shrinkage seeks to incorporate additional prior information on error correlations into the analysis, in order to enhance the inference. In many data assimilation models, climatological covariance information is often available, i.e., it is known prior information. Climatological covariances are typically precomputed or derived from climatological models and are often employed in variational data assimilation ~~Lorenc et al. (2015)~~ (Lorenc et al., 2015).

Following (Popov et al., 2020), we describe the stochastic covariance shrinkage technique. Instead of perturbing the trans-
 160 form matrix as in (20), we instead consider enhancing the dynamic ensemble $\mathbf{X}^f \in \mathbb{R}^{n \times N}$ with an M member synthetic ensemble $\mathcal{X}^f \in \mathbb{R}^{n \times M}$ of samples independent of the dynamical ensemble. ~~This leads to the total $N + M$ members ensemble~~
~~Each synthetic ensemble member is a biased sample distributed as,~~

$$\underline{\mathbf{X}^f = \begin{bmatrix} \mathbf{X}^f & \mathcal{X}^f \end{bmatrix} \in \mathbb{R}^{n \times (N+M)}}$$

with weights $\mathbf{w}^f = [w_1^f, \dots, w_{N+M}^f]$. We make the ansatz that each ensemble member is distributed as

$$165 \quad \mathbf{x}_{:,i}^f \sim \pi_{X^f}(X | P), \quad (21)$$

which is the full distribution of the forecast conditioned by the prior information that we have provided to the algorithm. Note that (21) is not the empirical measure distribution (2), that only has information from the ensemble members, but rather the ‘exaetenhanced’ distribution that is assumed to contain all the information from the forecast ~~· This is because we are now incorporating more prior and the prior climatological~~ information ~~P in the form of climatological information.~~ An illustration of this ensemble combination is shown in Figure 3.

Augmenting the dynamical with the synthetic ensembles leads to the total $N + M$ members ensemble:

$$170 \quad \mathbf{x}^f = \begin{bmatrix} \mathbf{X}^f & \mathcal{X}^f \end{bmatrix} \in \mathbb{R}^{n \times (N+M)} \quad (22)$$

with weights $\mathbf{w}^f = [w_1^f, \dots, w_{N+M}^f]$ with respect to $\pi_{|P|X^f}$.

Taking \mathbf{A}^f to be the anomalies of the dynamic ensemble (17), and

$$175 \quad \mathcal{A}^f = \mathcal{X}^f (\mathbf{I}_M - M^{-1} \mathbf{1}_M \mathbf{1}_M^T), \quad (23)$$

to be the anomalies of the synthetic ensemble, the ~~total empirical (unbiased) covariance~~ empirical covariance of the total ensemble can be written as,

$$\Sigma_{\mathbf{x}^f} = \Sigma_{\mathbf{X}^f} + \Sigma_{\mathcal{X}^f}, \quad (24)$$

where the constituent empirical (unbiased) covariances are defined in terms of the weights

$$180 \quad \Sigma_{\mathbf{X}^f} = \sum_{i=1}^N w_i^f \frac{N}{N-1} \mathbf{A}_{:,i}^f \mathbf{A}_{:,i}^{f,T}, \quad \Sigma_{\mathcal{X}^f} = \sum_{i=1}^M w_{N+i}^f \frac{M}{M-1} \mathcal{A}_{:,i}^f \mathcal{A}_{:,i}^{f,T}. \quad (25)$$

In the covariance shrinkage approach, to ensure that the sample mean of the augmented ensemble is the same as that of the dynamic ensemble, the synthetic ensemble is ~~assumed to be constructed with a mean equal to~~ the sample mean of the dynamic ensemble~~·~~.

$$\underline{M}^{-1} \mathcal{X}^f \mathbf{1}_M = \underline{N}^{-1} \mathbf{X}^f \mathbf{1}_N. \quad (26)$$

185 ~~by construction and, thus requiring that only~~ Thus, constructing the synthetic ensemble ~~anomalies need to be determined.~~ Taking a only requires sampling the synthetic anomalies. Consider a climatological covariance matrix \mathcal{P} ~~that represents a climatological estimate of the covariance, we sample the anomalies.~~ The synthetic anomalies are sampled from some unbiased

distribution with ~~a scaling, by a factor covariance $\mu\mathcal{P}$, where μ , of said covariance~~ is a scaling factor defined later. In the Gaussian case

$$190 \quad \mathcal{A}_{:,i}^f \sim \mathcal{N}(\mathbf{0}_n, \mu\mathcal{P}), \quad (27)$$

~~where μ is a scaling factor defined later~~. An alternate choice of distribution that we explore is the symmetric Laplace distribution (Kozubowski et al., 2013),

$$\mathcal{A}_{:,i}^f \sim \mathcal{L}(\mathbf{0}_n, \mu\mathcal{P}), \quad (28)$$

which is described by the pdf

$$195 \quad \pi(\mathbf{x}) = \frac{2}{[(2\pi)^n \mu\mathcal{P}]^{\frac{1}{2}}} \left(\frac{\mathbf{x}^T \mathcal{P}^{-1} \mathbf{x}}{2\mu} \right)^{\frac{2-n}{4}} K_{\frac{2-n}{2}} \left(\sqrt{\frac{2\mathbf{x}^T \mathcal{P}^{-1} \mathbf{x}}{\mu}} \right), \quad (29)$$

where K is the modified Bessel function of the second kind (Olver et al., 2010). The choice of Laplace distribution is motivated by robust statistics techniques (Rao et al., 2017). The resulting sampled covariance would therefore be an estimate of the scaled climatological covariance,

$$\Sigma_{\mathcal{X}^f} \approx \mu\mathcal{P}. \quad (30)$$

200 **Remark 1.** In order to stay consistent with the mean estimate, the sampled anomalies are replaced with their ~~sample~~-mean zero counterparts

$$\mathcal{A}^f \leftarrow \mathcal{A}^f (\mathbf{I}_M - M^{-1} \mathbf{1}_M \mathbf{1}_M^T). \quad (31)$$

The weights \mathbf{w}^f are divided into two classes: those that are associated with the dynamic ensemble, and those that are associated with the synthetic ensemble,

$$205 \quad w_i^f = \begin{cases} 1 - \gamma & 1 \leq i \leq N, \\ \gamma & N + 1 \leq i \leq N + M, \end{cases} \quad (32)$$

where the parameter γ is known as the covariance shrinkage factor.

One choice to calculate γ is the Rao-Blackwell Ledoit-Wolf (RBLW) estimator (Chen et al., 2009) (Nino-Ruiz and Sandu, 2017, equation (9)),

$$\gamma_{\text{RBLW}} = \min \left(\left[\frac{N-2}{N(N+2)} + \frac{(n+1)N-2}{\hat{U}(\mathcal{P}, \Sigma_{\mathcal{X}^f})N(N+2)(n-1)} \right], 1 \right), \quad (33)$$

210 where the sphericity factor,

$$\hat{U}(\mathcal{P}, \Sigma_{\mathbf{X}^f}) := \frac{1}{n-1} \left(\frac{n \operatorname{tr}(\mathcal{C}_i^2)}{\operatorname{tr}^2(\mathcal{C}_i)} \frac{n \operatorname{tr}(\mathcal{C}^2)}{\operatorname{tr}^2(\mathcal{C})} - 1 \right), \quad \text{with } \mathcal{C} := \mathcal{P}^{-\frac{1}{2}} \Sigma_{\mathbf{X}^f} \mathcal{P}^{-\frac{1}{2}}, \quad (34)$$

represents the mismatch between the climatological covariance (called the “target” in statistical literature) and sample covariance matrices. Note that if $\mathcal{P} = \Sigma_{\mathbf{X}^f}$, then $\hat{U}(\mathcal{P}, \Sigma_{\mathbf{X}^f}) = 0$. In such a framework, $\Sigma_{\mathbf{X}^f} \rightarrow \mathcal{P}$ then $\hat{U}(\mathcal{P}, \Sigma_{\mathbf{X}^f}) \rightarrow 0$ and, from (33), $\gamma_{\text{RBLW}} \rightarrow 1$. This represents a particular degenerate case whereby the dynamical ensemble is deemed to not be needed, and the climatological distribution is deemed to perfectly represent the forecast.

215

In this framework (Chen et al., 2009) the scaling parameter for the climatological covariance is defined to be

$$\mu = \frac{\operatorname{tr}(\mathcal{C})}{n}. \quad (35)$$

Remark 2. The RBLW estimator (33) makes the assumption that the underlying distribution of the dynamic ensemble is Gaussian. Typically this assumption is violated for dynamical systems of interest.

220 **Remark 3.** In statistical literature, the target covariance is often taken to be the identity, $\mathcal{P} = \mathbf{I}$, which implies that $\mathcal{C} = \Sigma_{\mathbf{X}^f}$ in (34). The assumption that the target is a climatological covariance is natural generalization in the specific context of sequential data assimilation.

Remark 4. The scaling of the target matrix \mathcal{P} is not of any consequence. Let $\tilde{\mathcal{P}} = \beta \mathcal{P}$ be a scalar scaling of the target matrix, then $\tilde{\mathcal{C}}_i = \frac{1}{\beta} \mathcal{C}_i$, implying that $\tilde{\mu}_i = \frac{1}{\beta} \mu_i$, rendering the matrix scaling inconsequential for computing μ . For γ_{RBLW} observe that the trace is linear operation, thus the scaling of \mathcal{C} plays no role in computing \hat{U} .

225

The Using the prior weight ensemble determined by (32), the importance sampling weights of the total ensemble \mathfrak{X}^f can be computed using (4), begetting the weight ensemble \mathfrak{w}^a . By leveraging this, the resulting analysis ensemble based on prior states and importance sampling weights of $N + M$ states is transported into an equally weighted posterior ensemble of N states through the transformation

$$230 \quad \mathbf{X}^a = \mathfrak{X}^f \mathbf{T}^*, \quad (36)$$

where the optimal transport matrix $\mathbf{T}^* \in \mathbf{R}^{(N+M) \times N}$ is computed by solving (9).

Recall that in the traditional method of rejuvenation (20), the optimal transport matrix is perturbed randomly into a nearby transport matrix; no new prior information is introduced. We take a fundamentally different approach by incorporating “unseen” prior information derived from a climatological covariance. To this end, before the Monge-Kantorovich problem (9) is solved,

235 we augment of the empirical measure distribution (2) with samples from the climatological distribution, to accommodate the

total ensemble (22),

$$\pi_{\hat{X}^f}(\underline{x} \underline{X} \mid \underline{p} \underline{P}) = \sum_{i=j=1}^{N^f} w_j^f \delta_{\underline{x} - \underline{x}_j^f} \underline{X} - \underline{x}_j^f, \quad (37)$$

240 ~~before the Monge-Kantorovich problem is solved~~ with \hat{X}^f being the random variable corresponding to this measure, estimating the distribution (21), and the weights (32) coming from the RBLW shrinkage factor (33). In effect ~~we are able attempt~~ to avoid ensemble collapse by enhancing the empirical measure distribution (37) with new prior information, as opposed to a reweighing of the old prior information. We ~~will~~ denote our method as FETPF, standing for ‘foresight’ ETPF, as we believe including prior information in the analysis procedure is a type of foresight. ~~When this procedure is combined with localization as described in section 2.1, we arrive at the Localized FETPF, or LFETPF, algorithm.~~

4.1 Convergence of the FETPF

245 ~~In this section we show that the FETPF reduces to (generalized) variants of the ETPF in two different ways: in the synthetic ensemble limit, and in the synthetic distribution limit.~~

~~Assume that the synthetic sample distribution is inexact in the mean and covariance, violating the assumption made in (21). As the dynamical ensemble size N increases, the shrinkage factor γ_{RBLW} in (33) approaches 0. This means that in the limit of an infinite dynamical ensemble, the FETPF reduces to the ETPF.~~

250 ~~Assume on the contrary that the synthetic sample distribution is exact, meaning that the climatology produces samples indistinguishable from the forecast, and that the assumption in (21) is fully satisfied. For a finite dynamical ensemble, the shrinkage factor γ_{RBLW} in (33) approaches 1, and the synthetic ensemble is taken as the forecast. This reduces to the ETPF in the case when the dynamical ensemble size is equal to the synthetic ensemble size $N = M$, and should result in an equivalent formulation when $M > N$.~~

255 ~~This leaves a gap however, as the shrinkage factor γ_{RBLW} only accounts for the covariance, thus if the synthetic ensemble distribution effectively predicts the covariance, but does not predict the higher order moments well, then the synthetic ensemble will still violate (21) even when it is treated as the forecast. Thus, an ideal shrinkage factor superior to (33) that takes into account more than just the covariance is required, though this is significantly outside the scope of this work.~~

Remark 5. ~~Most ensemble-based methods, including the ETPF, can produce physically unrealistic analysis ensemble members because of the linear nature of the optimal transport matrix. As the ETPF and FETPF perform inference that converges in distribution to the exact analysis distribution, as the dynamical ensemble size N grows, the algorithms produce physically realistic realizations with probability one.~~

4.2 Multiple Climatological Covariance Matrices

265 It is conceivable that multiple climatological models give rise to multiple climatological covariances, or alternatively multiple candidates for the most ‘common’ behavior of the model is to be chosen.

Given a collection of target covariances, $\{\mathcal{P}_j\}_{j \in \mathcal{J}}$, we must choose the appropriate covariance from which to sample. We consider the sphericity of the mismatch between the target and forecast covariances eq. (34). Based on authors' numerical experience, we select the target covariance that corresponds to the highest sphericity of the mismatch:

$$\mathcal{P}^* = \arg \max_{\mathcal{P}_j} \hat{U}(\mathcal{P}_j, \Sigma_{\mathbf{x}^f}), \quad (38)$$

270 We can justify this choice by realizing that the smaller the sphericity, the closer our samples are to that of canonical rejuvenation techniques. The aim of climatological shrinkage is to introduce unknown information into our inference procedure, thus the target covariance with the highest mismatch introduces the highest amount of outside information.

Remark 6. It is also possible to construct ‘multi-target’ shrinkage estimators (Lancewicki and Aladjem, 2014) that consider all target matrices simultaneously.

275 5 Numerical Experiments

~~Analysis RMSE versus dynamic ensemble size (N) of the Gaussian (G) and Laplace (L) covariance shrinkage approaches (FETPF) to particle rejuvenation for a synthetic ensemble size of $M=100$ and for various values of synthetic anomaly inflation, α , from , with respect to a canonical particle rejuvenation approach for first-order ETPF and second-order ETPF (denoted ETPF2) for an optimal rejuvenation factor $\tau=0.04$. The target covariance is taken to be .~~ We start with a short introduction to test problems configurations and the numerical experiments setups.

~~Analysis RMSE versus Dynamic ensemble size (N) of the Gaussian (G) and Laplace (L) covariance shrinkage approaches (FETPF) to particle rejuvenation with the multi-target covariances for a synthetic ensemble size of $M=100$ and for various values of synthetic anomaly inflation (α) with respect to a canonical particle rejuvenation approach for first-order ETPF and second-order ETPF (denoted ETPF2) for an optimal rejuvenation factor $\tau=0.04$. In order to stay in line with other particle~~ rejuvenation techniques, anomaly inflation is used as a heuristic to try and overcome deficiencies in the descriptive power of the synthetic ensemble. Formally,

$$\mathcal{A}^f \leftarrow \alpha \mathcal{A}^f, \quad (39)$$

which is equivalent to assuming an inflated scaling factor μ in (35). We therefore have two parameters that can be configured in the rejuvenation technique: M , the size of our synthetic ensemble, and α , the inflation applied to its realizations. It is

290 important that inflation only be applied to the synthetic ensemble, and not the dynamical ensemble, as to not violate the physical constraints of the dynamics.

~~Analysis RMSE of the covariance shrinkage approach to particle rejuvenation (FETPF) for different values of the synthetic ensemble size M and synthetic inflation factor α , and for a dynamical ensemble size of $N=5$. In our experiments we report the error of the analysis mean with respect to the truth (reference), measured by the spatio-temporal root mean square error~~

295 (RMSE):

$$\text{RMSE}(\mathbf{x}^t, \bar{\mathbf{x}}^a) = \sqrt{\frac{1}{nT} \sum_{i=1}^T \|\mathbf{x}_i^t - \bar{\mathbf{x}}_i^a\|_2^2}, \quad (40)$$

where T stands for the relevant measured timeframe of the experiments.

~~For all our experiments~~ We now describe the three variable Lorenz '63 model and the 40 variable Lorenz '96 model that are used in the experiments. We use the implementation of both these problems from the ODE test problem suite (Computational Science Labor

300

5.1 Lorenz '63 model

For first set of numerical experiments, we use the Lorenz '63 system (Lorenz, 1963):

$$\begin{aligned} x' &= \sigma(y - x), \\ y' &= x(\rho - z) - y, \\ z' &= xy - \beta z, \end{aligned} \quad (41)$$

with chaotic canonical parameter values $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. We observe the first component, with Gaussian unbiased observation error, with a very large variance of $\mathbf{R} = 8$. ~~The implementation is from our test problem suite (Computational Science Laborator~~

We perform 10,000 assimilation steps, but discard the first 1,000 that are used for spinup. The time interval between successive observations is $\Delta t = 0.12$. We perform 20 independent runs and take the mean of the results to obtain an accurate estimate of the expected error. All reported results are for differences that are two or more standard deviations apart.

310 This problem setup is challenging for the ensemble Kalman filter, which does not converge even for larger ensemble sizes. Therefore, this is a relevant test for non-Gaussian algorithms.

5.2 Lorenz '63 FETPF analysis results

As discussed previously, the canonical choice for the shrinkage covariance is the identity matrix. It has been the authors' experience that for most dynamical systems ~~the~~ this choice is poor. Moreover, the sequential data assimilation problem typically provides ways to calculate climatological approximations to the covariance. We take advantage of such techniques in this paper.

The first type of climatological covariance that we investigate is that of the distribution over that of the whole manifold of the dynamics. The (trace-state normalized) matrix that is obtained by taking the temporal covariance of ~~50000~~ 50,000 sample




Figure 4. Analysis RMSE versus dynamic ensemble size (N) of the Gaussian (G) and Laplace (L) covariance shrinkage approaches (FETPF) to particle rejuvenation for a synthetic ensemble size of $M = 100$ and for various values of synthetic anomaly inflation, α , from (39), with respect to a canonical particle rejuvenation approach for first order ETPF and second order ETPF (denoted ETPF2) for an optimal rejuvenation factor $\tau = 0.04$. The target covariance is taken to be (42). The baseline error is denoted by the ‘true’ SIR filter.

points on the attractor of the canonical Lorenz ’63 model is:

$$\mathcal{P} = \begin{bmatrix} 0.8616 & 0.8618 & -0.0148 \\ 0.8618 & 1.1149 & -0.0035 \\ -0.0148 & -0.0035 & 1.0234 \end{bmatrix}, \quad (42)$$

with condition number 15.88.

Our first round of experiments compares the canonical method of rejuvenation in ETPF and ETPF2 with a rejuvenation factor of $\tau = 0.04$ in (19) (see Acevedo et al. (2017)) to the stochastic covariance shrinkage technique for both Gaussian and Laplace samples. A dynamic ensemble size $M = 100$, the inflation factors $\alpha \in \{1.0, 1.2\}$, and the target covariance (42) are used. The baseline error is computed using a sequential importance resampling (SIR) filter with an ensemble of $N = 10^5$ particles.

The results for the first round of experiments are shown in Figure 4. From the results, all possible ETPF-based algorithms seem to converge to the SIR baseline for around $N = 100$ particles. The differences between the various algorithms only become apparent at smaller ensemble sizes. As reported in Acevedo et al. (2017), the second-order accurate ETPF2 performs better than the standard ETPF.

The FETPF without synthetic inflation performs worse than both the ETPF and ETPF2 for Gaussian synthetic samples, while it performs better when equipped with Laplacian synthetic samples. When the synthetic samples are inflated with inflation factor $\alpha = 1.2$, the FETPF performs significantly better than all other algorithms.




Figure 5. Analysis RMSE versus dynamic ensemble size (N) of the Gaussian (G) and Laplace (L) covariance shrinkage approaches (FETPF) to particle rejuvenation with the multi-target covariances (43) for a synthetic ensemble size of $M = 100$ and for various values of synthetic anomaly inflation (α) with respect to a canonical particle rejuvenation approach for first order ETPF and second order ETPF (denoted ETPF2) for the rejuvenation factor $\tau = 0.04$. The baseline error is denoted by the ‘true’ SIR filter.

The second round of experiments uses multiple values of the climatological covariance \mathcal{P} . The rest of the setup is identical to the previous experiment. For testing multiple covariances, we run an ETPF with $N = 100$ with 20000 evenly spaced samples over a time interval of 2400 time units and calculate the trace-state normalized forecast covariances. Under a square Frobenius norm distance, we cluster the empirical covariance matrices of the same ensemble at different times using the k -means algorithm (Tan et al., 2018) into two clusters. The collection of climatological covariances for the Lorenz ’63 thus consists of the centroids of each cluster,

$$\mathcal{P}_1 = \begin{bmatrix} 0.5017 & 0.5524 & -0.4587 \\ 0.5524 & 1.0731 & -0.6723 \\ -0.4587 & -0.6723 & 1.4252 \end{bmatrix}, \quad \mathcal{P}_2 = \begin{bmatrix} 0.5443 & 0.6830 & 0.4330 \\ 0.6830 & 1.2748 & 0.6318 \\ 0.4330 & 0.6318 & 1.1808 \end{bmatrix}, \quad (43)$$

with condition numbers 13.68 and 16.98 respectively. As can be seen, the clusters are mainly split by the correlation factors of z with respect to the other variables being positive or negative.

In order to stay in line with other particle rejuvenation techniques, a heuristic that we use is inflation. The second round of experiments are reported in Figure 5 using the climatological covariances (43). The analysis of the results is largely similar to the previous experiment, with the only difference being that the FETPF with multiple covariances does not seem to require synthetic inflation. As the covariance chosen depends on the dynamical ensemble, these results indicate that a more detailed climatological distribution that varies seasonally might induce an even greater decrease in error.

350 The results empirically show that supplementing the ensemble with additional synthetic information during assimilation is more effective than randomly perturbing the ensemble post-assimilation, for a small problem. The authors hypothesize that the results point strongly towards the need of intelligently, and adaptively choosing the target covariance matrices, and to the need for better operational calculation of the covariance shrinkage factor γ .

5.3 Lorenz '63 parameter search

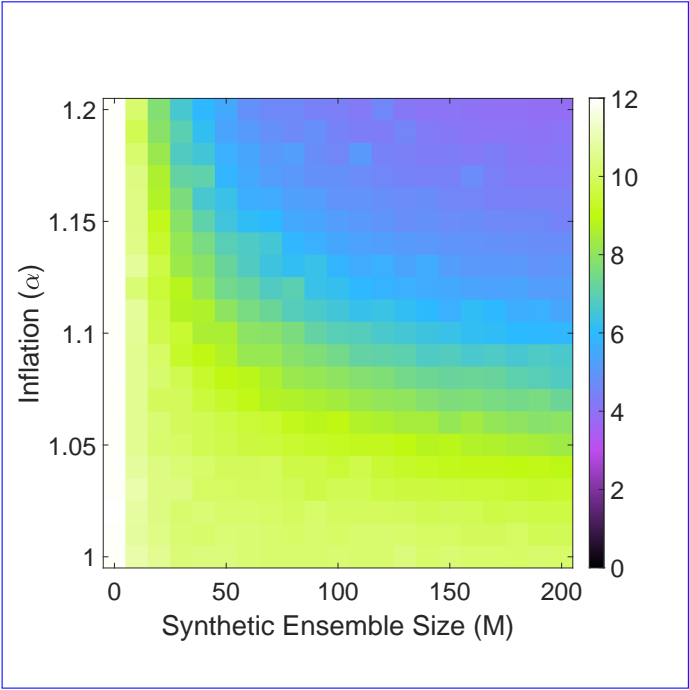


Figure 6. Analysis RMSE of the covariance shrinkage approach to particle rejuvenation (FETPF) for different values of the synthetic ensemble size M and synthetic inflation factor α , and for a dynamical ensemble size of $N = 5$.

Our third round of experiments with the Lorenz '63 system seeks to understand the effect of selecting the two free parameters, i.e., the synthetic ensemble size M and the synthetic ensemble, so that it is possible to overcome deficiencies in its descriptive power:-

355 $\mathcal{A}^f \leftarrow \alpha \mathcal{A}^f,$

which is equivalent to assuming an inflated scaling factor μ in-. We therefore have two parameters that we can configure in our rejuvenation technique:- synthetic ensemble inflation factor α . We keep the dynamic ensemble size to a small constant size of $N = 5$, and vary M , the size of our surrogate synthetic ensemble, in the range $[0, 200]$, with α varying in the range $[1, 1.2]$.

Figure 6 shows the spatio-temporal RMSE for various values of M and α , with Gaussian synthetic samples using the single target matrix (42). The results are not surprising. An increase in the synthetic ensemble size M corresponds to a decrease in error. Similarly, an increase in the inflation factor also corresponds to a decrease in error. Furthermore the factors are complementary, meaning that increasing both decreases the error even more significantly.

An interesting effect is that very large synthetic ensemble sizes are required to correspond to a noticeable decrease in error, relative to the dimension of the system. This might pose a challenge when this algorithm is utilized without further corrections such as localization.

5.4 Lorenz '63 γ_{RBLW} histograms




Figure 7. The distribution of the γ_{RBLW} parameter for the Lorenz '63 model. Both Gaussian (27) and Laplace (28) assumptions are used for the synthetic ensemble, with no synthetic inflation ($\alpha = 1$).

Our fourth experiment with the Lorenz '63 equations looks at the distribution of the inflation applied to its realizations. It is the authors' experience that inflation should only be applied to the synthetic ensemble, and not the dynamical ensemble; otherwise the shrinkage estimate values of the shrinkage parameter γ_{RBLW} from (33) does not lead to a stable algorithm that is obtained through the assimilation procedure. We test with the Gaussian (27) and Laplace (28) on the climatological distribution with no synthetic inflation (39) ($\alpha = 1$).

In our experiments we report the error of the analysis with respect to the truth, measured by the spatio-temporal RMSE, given by the formula,

$$\text{RMSE}(\mathbf{x}^t, \bar{\mathbf{x}}^a) = \sqrt{\frac{1}{nK} \sum_{i=1}^T \sum_{j=1}^n \left([\mathbf{x}_i^t]_j - [\bar{\mathbf{x}}_i^a]_j \right)^2},$$

where T stands for the relevant measured timeframe of the experiments. Figure 7 shows an approximation to the distributions of γ_{RBLW} for several choices of the dynamical ensemble size N with all other settings kept the same as in the previous experiments. As shown in section 4.1, the shrinkage factor γ_{RBLW} tends towards a distribution that starts resembling a degenerate distribution around zero as N increases. For Gaussian samples, this happens in a smooth fashion, but for Laplace samples something interesting occurs. For $N = 5$, the distribution of γ_{RBLW} is significantly much more significantly skewed towards smaller values, meaning that less confidence is placed in the synthetic ensemble. While the authors do not see a convincing explanation for this behavior, the effect does explain why the Laplace distributed synthetic samples were an improvement over Gaussian samples in 5.2.

Our first round of experiments reported in ?? compares the canonical method of rejuvenation in ETPF (and second order extensions) with the optimal rejuvenation factor of $\tau = 0.04$ in (computed by parameter search) to the stochastic covariance shrinkage technique for both Gaussian and Laplace samples.

5.5 Lorenz '96 model

For numerical experiments with localization, we use the Lorenz '96 system (Lorenz, 1996; van Kekem, 2018):

$$x'_i = -x_{i-1}(x_{i-2} - x_{i+1}) - x_i + F, \quad i = 1, \dots, 40, \quad F = 8. \quad (44)$$

with $x_0 = x_{40}$, $x_{-1} = x_{39}$, and $x_{41} = x_1$. The Lorenz '96 system provides a more challenging medium-dimension assimilation scenario. A dynamic ensemble size of $M = 100$, the inflation factors $\alpha = \{1.0, 1.2\}$. The target covariance is used. We perform 10000-5,000 assimilation steps, but discard the first 1000-500 that are used for spinup. The time interval between successive observations is $\Delta t = 0.12$ $\Delta t = 0.05$. We perform 20-12 independent runs and take the mean of the results to obtain an accurate estimate of the expected error. For localization, we take the Gaspari-Cohn decorrelation function with a radius of $r = 4$ variables.

Results in ?? show that We test with two observation operators. First, we consider a standard linear observation operator

$$\mathcal{H}(\mathbf{x}) = \mathbf{x}, \quad (45)$$

that observes all states with an observation covariance of $\mathbf{R} = \mathbf{I}_{40}$. Second, we use the nonlinear observation operator (Asch et al., 2016)

$$\mathcal{H}(\mathbf{x}) = \frac{\mathbf{x}}{2} \circ \left[1 + \left(\frac{|\mathbf{x}|}{10} \right)^{\circ(\omega-1)} \right], \quad (46)$$

where \circ stands for element-wise operations (multiplication and exponentiation, and $|\cdot|$ stands for element-wise absolute value, that observe all states through a non-linear fashion with the observation covariance $\mathbf{R} = \mathbf{I}_{40}$. We set the control parameter to $\omega = 5$ for a moderately non-linear system.

5.6 Lorenz '96 localization results




Figure 8. For the Lorenz '96 model with the linear observation operator (45): analysis RMSE versus dynamic ensemble size (N) of the Gaussian (G) localized covariance shrinkage approach (LFETPF) to particle rejuvenation for a synthetic ensemble size of $M = 100$ and for synthetic anomaly inflation, $\alpha = 1.05$, versus the LETKF (Hunt et al., 2007) method with inflation $\alpha = 1.05$.

For the Lorenz '96 experiments we aim to compare LFETPF to the LETPF, and to the Localized ensemble transform Kalman filter (LETKF). The matrix \mathcal{P} is computed in a similar way as (42) for the Lorenz '63 Model, and is not shown due to space limitations. We set the stochastic covariance shrinkage technique converges to the same 'correct' RMSE in the case of a large dynamic ensemble of $N = 100$. Moreover, in the small ensemble case of $N = 5$, methods that inflate the synthetic ensemble with factor $\alpha = 1.2$ perform significantly better than those that do not. Laplace distributed synthetic samples also

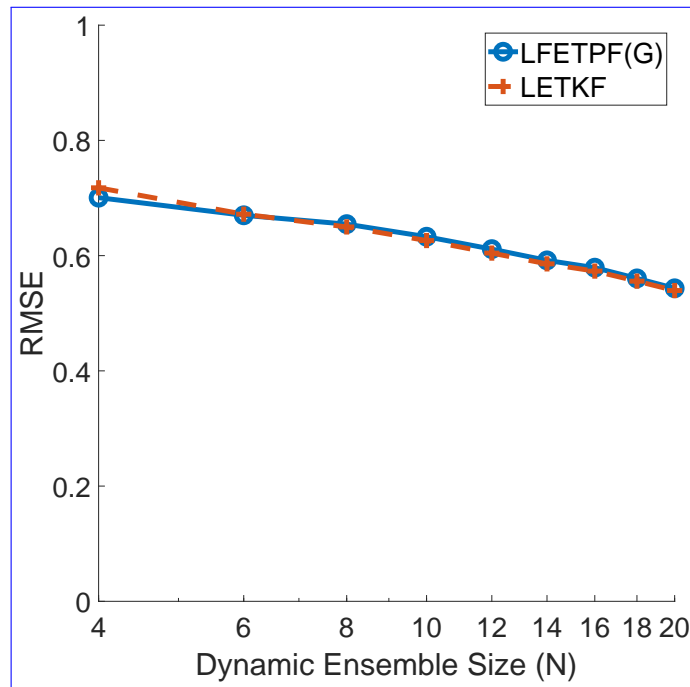


Figure 9. For the Lorenz '96 model with the nonlinear observation operator (46): analysis RMSE versus dynamic ensemble size (N) of the Gaussian (G) localized covariance shrinkage approach (LFETPF) to particle rejuvenation for a synthetic ensemble size of $M = 100$ and for synthetic anomaly inflation, $\alpha = 1.05$, versus the LETKF (Hunt et al., 2007) method with inflation $\alpha = 1.05$.

seem to slightly reduce the error. size to a constant $M = 100$ and the synthetic ensemble inflation factor to $\alpha = 1.05$. The inflation factor $\alpha = 1.05$ is also used in the LETKF. We vary the dynamical ensemble size in the range $N \in [2, 20]$, and plot the spatio-temporal analysis RMSE over the time interval after spinup.

Our second round of experiments reported in ?? makes use of For rejuvenation in LETPF multiple values of the climatological covariance \mathcal{P} , namely those in-. The rest of the setup is identical to the previous experiment-. Again rejuvenation factor τ were tried, but none lead to convergence.

For the linear observation operator (45), the results of this experiment can be seen in Figure 8. The proposed LFETPF did converge, for a large value of dynamic ensemble size $N = 100$, the stochastic covariance shrinkage approach attains the correct error statistics. For a low ensemble size, however, there is virtually no difference between the Gaussian, Laplace, inflation, and no-inflation stochastic covariance shrinkage methods as compared to the ETPF dynamical ensemble size of as little as $N = 4$, however it provides an analysis less accurate than the state-of-the-art LETKF. The authors believe that this is likely due to the highly Gaussian nature of the problem setup, coupled with the inexactness that localization brings.

Our final round of experiments seeks to understand the effect of selecting the two free parameters, i. e., the synthetic ensemble size M and the synthetic ensemble inflation factor α . ?? shows the spatio-temporal RMSE of a small dynamic ensemble ($N = 5$) for various values of M and α , with Gaussian synthetic samples using the single target matrix-. The results clearly show that in many operationally useful cases, it is necessary to have a sufficiently expressive synthetic ensemble,

425 ~~whose anomalies are sufficiently inflated. For the non-linear observation operator (46) the results are plotted in Figure 9. This~~
~~time, the LFETPF and LETKF are nearly identical in error, suggesting that the problem setup is highly Gaussian, and can be~~
~~well-accommodated by synthetic Gaussian samples.~~

~~The results of the first two experiments show that adding additional synthetic information during assimilation is more~~
~~effective than randomly perturbing the ensemble post-assimilation. The authors hypothesize that the results point strongly~~
430 ~~towards the need of intelligently, and adaptively choosing the target covariance matrices, and to the need for better operational~~
~~calculation of the covariance shrinkage factor.~~ Methods such as hybrid filters Acevedo et al. (2017) might provide for a happy
medium between the LFETPF and the LETKF, though the exploration of this idea is outside the scope of this work. The authors
believe that these results show great promise in the applicability of the LFETPF for medium-scale data assimilation workloads.

6 Conclusions

435 This paper introduces a stochastic covariance shrinkage-based particle rejuvenation technique for the ensemble transport par-
ticle filter. Instead of reweighing existing prior information, the approach incorporates additional prior information into the
ensemble through the use of synthetic anomalies. These ~~anomalies come from~~ synthetic anomalies are sampled from from any
chosen distribution family, such that they are consistent with the climatological covariance information. Numerical experiments
show that the use of climatological prior information to perform rejuvenation leads to reduced analyses errors for significantly
440 smaller dynamical ensemble sizes than the original rejuvenation approach. Further research includes testing FETPF on higher
dimensional operational problems, which could require large synthetic ensemble sizes. This would need computationally viable
sampling methods.

~~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.~~

445 Acknowledgements

This work was supported by ~~awards NSF CDS&E-MSS-1953113, DOE ASCR DE-SC0021313, DOE through award ASCR~~
DE-SC0021313, by NSF through award CDS&E-MSS 1953113, and by the Computational Science Laboratory at Virginia
Tech.

References

- 450 Acevedo, W., de Wiljes, J., and Reich, S.: Second-order accurate ensemble transform particle filters, *SIAM Journal on Scientific Computing*, 39, A1834–A1850, 2017.
- Anderson, J. L.: An ensemble adjustment Kalman filter for data assimilation, *Monthly weather review*, 129, 2884–2903, 2001.
- Anderson, J. L.: Localization and sampling error correction in ensemble Kalman filter data assimilation, *Monthly Weather Review*, 140, 2359–2371, 2012.
- 455 Asch, M., Bocquet, M., and Nodet, M.: *Data assimilation: methods, algorithms, and applications*, SIAM, 2016.
- Attia, A. and Sandu, A.: A hybrid Monte–Carlo sampling filter for non-Gaussian data assimilation, *AIMS Geosciences*, 1, 41–78, <https://doi.org/10.3934/geosci.2015.1.41>, 2015.
- Bocquet, M.: Comparison of local particle filters and new implementations, *Nonlinear Processes in Geophysics*, 2018.
- Burgers, G., van Leeuwen, P. J., and Evensen, G.: Analysis scheme in the ensemble Kalman filter, *Monthly weather review*, 126, 1719–1724, 1998.
- 460 Chen, Y., Wiesel, A., and Hero, A. O.: Shrinkage estimation of high dimensional covariance matrices, in: *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 2937–2940, IEEE, 2009.
- Chustagulprom, N., Reich, S., and Reinhardt, M.: A hybrid ensemble transform particle filter for nonlinear and spatially extended dynamical systems, *SIAM/ASA Journal on Uncertainty Quantification*, 4, 592–608, 2016.
- 465 Computational Science Laboratory: ODE Test Problems, <https://github.com/ComputationalScienceLaboratory/ODE-Test-Problems>, 2020.
- 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, 1994.
- Evensen, G.: *Data assimilation: the ensemble Kalman filter*, Springer Science & Business Media, 2009.
- Gaspari, G. and Cohn, S. E.: Construction of correlation functions in two and three dimensions, *Quarterly Journal of the Royal Meteorological Society*, 125, 723–757, 1999.
- 470 Hunt, B. R., Kostelich, E. J., and Szunyogh, I.: Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter, *Physica D: Nonlinear Phenomena*, 230, 112–126, 2007.
- Jaynes, E. T.: *Probability theory: The logic of science*, Cambridge university press, 2003.
- Kozubowski, T. J., Podgórski, K., and Rychlik, I.: Multivariate generalized Laplace distribution and related random fields, *Journal of Multivariate Analysis*, 113, 59–72, 2013.
- 475 Lancewicki, T. and Aladjem, M.: Multi-target shrinkage estimation for covariance matrices, *IEEE Transactions on Signal Processing*, 62, 6380–6390, 2014.
- Law, K., Stuart, A., and Zygalakis, K.: *Data assimilation: a mathematical introduction*, vol. 62, Springer, 2015.
- Liu, J. S.: *Monte Carlo strategies in scientific computing*, Springer Science & Business Media, 2008.
- 480 Lorenc, A. C., Bowler, N. E., Clayton, A. M., Pring, S. R., and Fairbairn, D.: Comparison of hybrid-4DEnVar and hybrid-4DVar data assimilation methods for global NWP, *Monthly Weather Review*, 143, 212–229, 2015.
- Lorenz, E. N.: Deterministic nonperiodic flow, *Journal of the atmospheric sciences*, 20, 130–141, 1963.
- Lorenz, E. N.: Predictability: A problem partly solved, in: *Proc. Seminar on predictability*, vol. 1, 1996.
- McCann, R. J. and Guillen, N.: Five lectures on optimal transportation: geometry, regularity and applications, *Analysis and geometry of metric measure spaces: lecture notes of the séminaire de Mathématiques Supérieure (SMS) Montréal*, pp. 145–180, 2011.
- 485

- McCann, R. J. et al.: Existence and uniqueness of monotone measure-preserving maps, *Duke Mathematical Journal*, 80, 309–324, 1995.
- Musso, C., Oudjane, N., and Le Gland, F.: Improving regularised particle filters, in: *Sequential Monte Carlo methods in practice*, pp. 247–271, Springer, 2001.
- Nino-Ruiz, E. and Sandu, A.: Ensemble Kalman filter implementations based on shrinkage covariance matrix estimation, *Ocean Dynamics*, 65, 1423–1439, <https://doi.org/10.1007/s10236-015-0888-9>, 2015.
- Nino-Ruiz, E. D. and Sandu, A.: An ensemble Kalman filter implementation based on modified Cholesky decomposition for inverse covariance matrix estimation, *SIAM Journal on Scientific Computing*, submitted, <https://arxiv.org/abs/1605.08875>, 2017.
- Nino-Ruiz, E. D. and Sandu, A.: Efficient parallel implementation of DDDAS inference using an ensemble Kalman filter with shrinkage covariance matrix estimation, *Cluster Computing*, pp. 1–11, <https://link.springer.com/article/10.1007/s10586-017-1407-1>, 2018.
- Nino-Ruiz, E. D., Sandu, A., and Deng, X.: A parallel ensemble Kalman filter implementation based on modified Cholesky decomposition, in: *Proceedings of the 6th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*, vol. Supercomputing 2015 of *Scala '15*, Austin, Texas, <https://doi.org/10.1145/2832080.2832084>, 2015.
- Olver, F. W., Lozier, D. W., Boisvert, R. F., and Clark, C. W.: *NIST handbook of mathematical functions* hardback and CD-ROM, Cambridge university press, 2010.
- Petrie, R.: *Localization in the ensemble Kalman filter*, MSc Atmosphere, Ocean and Climate University of Reading, 2008.
- Popov, A. A. and Sandu, A.: An Explicit Probabilistic Derivation of Inflation in a Scalar Ensemble Kalman Filter for Finite Step, Finite Ensemble Convergence, *arXiv preprint arXiv:2003.13162*, 2020.
- Popov, A. A., Sandu, A., Nino-Ruiz, E. D., and Evensen, G.: A Stochastic Covariance Shrinkage Approach in Ensemble Transform Kalman Filtering, 2020.
- Rao, V., Sandu, A., Ng, M., and Nino-Ruiz, E. D.: Robust Data Assimilation using L_1 and Huber norms, *SIAM Journal on Scientific Computing*, 39, B548–B570, 2017.
- Reich, S.: A nonparametric ensemble transform method for Bayesian inference, *SIAM Journal on Scientific Computing*, 35, A2013–A2024, 2013.
- Reich, S. and Cotter, C.: *Probabilistic forecasting and Bayesian data assimilation*, Cambridge University Press, 2015.
- Roberts, S., Popov, A. A., and Sandu, A.: ODE Test Problems: a MATLAB suite of initial value problems, *CoRR*, [abs/1901.04098](https://arxiv.org/abs/1901.04098), <http://arxiv.org/abs/1901.04098>, 2019.
- Ruiz, E. N., Sandu, A., and Anderson, J.: An efficient implementation of the ensemble Kalman filter based on an iterative Sherman–Morrison formula, *Statistics and Computing*, pp. 1–17, <https://doi.org/10.1007/s11222-014-9454-4>, 2014.
- Snyder, C., Bengtsson, T., Bickel, P., and Anderson, J.: Obstacles to high-dimensional particle filtering, *Monthly Weather Review*, 136, 4629–4640, 2008.
- Strogatz, S. H.: *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*, CRC press, 2018.
- Tan, P.-N., Steinbach, M., Karpatne, A., and Kumar, V.: *Introduction to Data Mining* (2nd Edition), Pearson, 2nd edn., 2018.
- van Kekem, D. L.: *Dynamics of the Lorenz-96 model: Bifurcations, symmetries and waves*, Ph.D. thesis, University of Groningen, 2018.
- Van Leeuwen, P. J.: Particle filtering in geophysical systems, *Monthly Weather Review*, 137, 4089–4114, 2009.
- Van Leeuwen, P. J., Künsch, H. R., Nerger, L., Potthast, R., and Reich, S.: Particle filters for high-dimensional geoscience applications: A review, *Quarterly Journal of the Royal Meteorological Society*, 145, 2335–2365, 2019.
- Villani, C.: *Topics in optimal transportation*, 58, American Mathematical Soc., 2003.

Zhang, Y., Liu, N., and Oliver, D. S.: Ensemble filter methods with perturbed observations applied to nonlinear problems, Computational Geosciences, 14, 249–261, 2010.