

A Stochastic Covariance Shrinkage Approach in Ensemble Transform Kalman Filtering


Andrey A Popov, Adrian Sandu, Elias D. Nino-Ruiz,
Geir Evensen

Computational Science Laboratory Report
CSL-TR-20-3
June 6, 2022

Computational Science Laboratory
“Compute the Future!”

Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, VA 24060
Phone: (540)-231-2193
Fax: (540)-231-6075

Email: apopov@vt.edu, sandu@cs.vt.edu, enino@uninorte.edu.co,
geev@norceresearch.no
Web: <http://csl.cs.vt.edu>



A Stochastic Covariance Shrinkage Approach in Ensemble Transform Kalman Filtering

Andrey A Popov^{*} and Adrian Sandu[†]

Computational Science Laboratory, Department of Computer Science, Virginia Tech

Elias D. Nino-Ruiz[‡]

Applied Math and Computer Science Lab, Universidad del Norte, Colombia

Geir Evensen[§]

Norwegian Research Center (NORCE) and Nansen

Environmental and Remote Sensing Center (NERSC), Norway

(Dated: June 6, 2022)

The Ensemble Kalman Filters (EnKF) employ a Monte-Carlo approach to represent covariance information, and are affected by sampling errors in operational settings where the number of model realizations is much smaller than the model state dimension. To alleviate the effects of these errors EnKF relies on model-specific heuristics such as covariance localization, which takes advantage of the spatial locality of correlations among the model variables. This work proposes an approach to alleviate sampling errors that utilizes a locally averaged-in-time dynamics of the model, described in terms of a climatological covariance of the dynamical system. We use this covariance as the target matrix in covariance shrinkage methods, and develop a stochastic covariance shrinkage approach where synthetic ensemble members are drawn to enrich both the ensemble subspace and the ensemble transformation. We additionally provide for a way in which this methodology can be localized similar to the state-of-the-art LETKF method, and that for a certain model setup, our methodology significantly outperforms it.

1. INTRODUCTION

The ensemble Kalman filter [7, 12, 13], one of the most widely applied data assimilation algorithms [4, 24, 38], uses a Monte Carlo approach to provide a non-linear approximation to the Kalman filter [22]. In the typical case of an undersampled ensemble the algorithm requires correction procedures such as inflation [1], localization [2, 19, 31, 33, 35, 45], and ensemble subspace

^{*} apopov@vt.edu

[†] sandu@cs.vt.edu

[‡] enino@uninorte.edu.co

[§] geev@norceresearch.no

enrichment [30, 32, 40].

Hybrid data assimilation [18] is typically an umbrella term for assimilation techniques that combine both offline-estimated climatological covariances with their online-estimated statistical counterparts. These methods are often thought of as heuristic corrections, but in fact stem from statistically rigorous covariance shrinkage techniques.

This work is based on enriching the ensemble subspace through the use of climatological covariances. Previous work [30, 32] proposed augmenting the covariance estimates derived from the ensemble by a full rank shrinkage covariance matrix approximation. In this work we consider augmenting the physical ensemble with synthetic members drawn from a normal distribution with a possibly low rank covariance matrix derived from *a priori* information such a climatological information or method of snapshots. We show that this is equivalent to a stochastic implementation of the shrinkage covariance matrix estimate proposed in [30, 32], and therefore augmenting the physical ensemble with synthetic members enriches the rank of the covariance matrix, and nudges the resulting covariance estimate toward the true covariance.

2. BACKGROUND

Our aim is to understand the behavior of an evolving natural phenomenon. The evolution of the natural phenomenon is approximated by an imperfect dynamical model:

$$X_i = \mathcal{M}_{i-1,i}(X_{i-1}) + \boldsymbol{\xi}_i, \quad (2.1)$$

where X_{i-1} is a random variable (RV) whose distribution represents our uncertainty in the state of the system at time $i - 1$, $\mathcal{M}_{i-1,i}$ is the (imperfect) dynamical model, $\boldsymbol{\xi}_i$ is a RV whose distribution represents our uncertainty in the additive modeling error, and X_i is the RV whose distribution represents our uncertainty in the (forecasted) state at time i .

One collects noisy observations of the truth:

$$\mathbf{y}_i^o = \mathcal{H}_i(\mathbf{x}_i^t) + \boldsymbol{\eta}_i, \quad (2.2)$$

where $\mathbf{x}^t$ represents the true state of nature represented in model space, $\mathcal{H}_i$ is the (potentially non-linear) observation operator, $\boldsymbol{\eta}_i$ is a RV whose distribution represents our uncertainty in the observations, and $\mathbf{y}_i^o$ are the observation values, assumed to be realizations of an observation RV Y_i . Take n to be the dimension of the state-space, and m to be the dimension of the observation

space.

The goal of data assimilation is to find the *a posteriori* estimate of the state given the observations, which is typically achieved through Bayes' theorem. At time i we have:

$$\pi(X_i|Y_i) \propto \pi(Y_i|X_i) \pi(X_i). \quad (2.3)$$

In typical Kalman filtering the assumption of Gaussianity is made, whereby the states at all times, as well as the additive model and observation errors, are assumed to be Gaussian and independently distributed. Specifically one assumes $\boldsymbol{\xi}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_i)$ and $\boldsymbol{\eta}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_i)$.

In what follows we use the following notation. The *a priori* estimates at all times are represented with the superscript $\square^f$, for forecast (as from the imperfect model), and the *a posteriori* estimates are represented with the superscript $\square^a$, for analysis (through a DA algorithm).

2.1. Ensemble Transform Kalman Filter

Forecasting with an ensemble of coarse models has proven to be a more robust methodology than forecasting with a single fine model [23]. Ensemble Kalman filtering aims to utilize the ensemble of forecasted states to construct empirical moments and use them to implement the Kalman filter formula. The Ensemble Transform Kalman Filter (ETKF) [5] computes an optimal transformation of the prior ensemble member states to the posterior member states; for Gaussian distributions the optimal transform is described by a symmetric transform matrix.

We now describe the standard ETKF. Let $\mathbf{X}_{i-1}^a = [\mathbf{x}_{i-1}^{(1),a}, \dots, \mathbf{x}_{i-1}^{(N),a}]$ represent the N -members analysis ensemble at time $i-1$. The forecast step is:

$$\mathbf{x}_i^{(k),f} = \mathcal{M}_{i-1,i}(\mathbf{x}_{i-1}^{(k),a}) + \boldsymbol{\xi}_i^{(k)}, \quad k = 1, \dots, N, \quad (2.4)$$

where $\boldsymbol{\xi}_i^{(k)}$ is a random draw from $\mathcal{N}(\mathbf{0}, \mathbf{Q}_i)$.

The ETKF analysis step reads:

$$\mathbf{A}_i^a = \mathbf{A}_i^f \mathbf{T}_i, \quad (2.5a)$$

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{A}_i^a \mathbf{Z}^{a,\top} \mathbf{R}^{-1} \mathbf{d}_i, \quad (2.5b)$$

where

$$\mathbf{T}_i = \left(\mathbf{I} - \mathbf{Z}_i^{\text{f}\top} \mathbf{S}_i^{-1} \mathbf{Z}_i^{\text{f}} \right)^{\frac{1}{2}}, \quad (2.6a)$$

$$\mathbf{S}_i = \mathbf{Z}_i^{\text{f}} \mathbf{Z}_i^{\text{f}\top} + \mathbf{R}_i, \quad (2.6b)$$

$$\mathbf{A}_i^{\text{f}} = \frac{1}{\sqrt{N-1}} \left(\mathbf{X}_i^{\text{f}} - \bar{\mathbf{x}}_i^{\text{f}} \mathbf{1}^\top \right), \quad (2.6c)$$

$$\mathbf{Z}_i^{\text{f}} = \frac{1}{\sqrt{N-1}} \left(\mathcal{H}(\mathbf{X}^{\text{f}}) - \overline{\mathcal{H}(\mathbf{X}^{\text{f}})} \mathbf{1}^\top \right), \quad (2.6d)$$

$$\mathbf{d}_i = \mathbf{y}_i^{\text{o}} - \overline{\mathcal{H}(\mathbf{X}_i^{\text{f}})}, \quad (2.6e)$$

$$\bar{\mathbf{x}}_i^{\text{f}} = \frac{1}{N} \sum_{k=1}^N \mathbf{X}_i^{\text{f},(k)}, \quad (2.6f)$$

$$\overline{\mathcal{H}(\mathbf{X}_i^{\text{f}})} = \frac{1}{N} \sum_{k=1}^N \mathcal{H}(\mathbf{X}_i^{\text{f},(k)}). \quad (2.6g)$$

Here the unique symmetric square root of the matrix is used, as there is evidence of that option being the most numerically stable [42].

The empirical forecast covariance estimate

$$\widetilde{\Sigma}_{X_i^{\text{f}} X_i^{\text{f}}} = \mathbf{A}_i^{\text{f}} \mathbf{A}_i^{\text{f}\top} \quad (2.7)$$

is inexact due to a multitude of deficiencies. One method to improve the empirical covariance estimate is inflation [1], which is applied to the ensemble anomalies,

$$\mathbf{A}_i^{\text{f}} \leftarrow \alpha \mathbf{A}_i^{\text{f}}, \quad (2.8)$$

before any other computation is performed (meaning that it is also applied to the observation anomalies, $\mathbf{Z}_i^{\text{f}}$ as well). The inflation parameter $\alpha > 1$ is known to be a requirement for the EnKF to converge for linear models [37].

2.2. Covariance localization

Traditional state-space localization of the empirical covariance (2.7) is done by tapering, i.e., by using a Schur product of the empirical covariance with a localization matrix $\boldsymbol{\rho}_i$:

$$\mathbf{B}_i^{\text{f}} = \boldsymbol{\rho}_i \circ \widetilde{\Sigma}_{X_i^{\text{f}} X_i^{\text{f}}}, \quad (2.9)$$

where $\boldsymbol{\rho}_i$ contains entries that are progressively smaller as the (physically-relevant) distance between the corresponding variables increases.

The localized ETKF (LETKF) [19] is an efficient implementation of localization for ETKF. The LETKF and its variants are considered to be one of the state-of-the-art EnKF methods. In the state-space approach to the LETKF, the j -th state space variable $\mathbf{x}_{i,[j]}$ is assimilated independently of all others, with the observation space error covariance inverse replaced by

$$\mathbf{R}_i^{-1} \leftarrow \boldsymbol{\rho}_{i,[j]} \circ \mathbf{R}_i^{-1}, \quad (2.10)$$

where $\boldsymbol{\rho}_{i,[j]}$ is a diagonal matrix, with diagonal entries representing the decorrelation factors between all observation space variables and the j -th state space variable. Each diagonal element represents a tapering factor, and is often chosen to be a function of the distance from the state-space variable $\mathbf{x}_{i,[j]}$ being assimilated and the corresponding observation-space variable. The implicit assumption is that all observations are independent of each other, in both the observation error ($\mathbf{R}_i$ is diagonal), and forecast error ($\mathbf{Z}^f \mathbf{Z}^{f,T}$ is assumed to be diagonal).

2.3. Covariance shrinkage

In the statistical literature [8–10, 25] covariance shrinkage refers to the methodology under which an empirical covariance is made to approach the “true” covariance from which the set of samples is derived. For the vast majority of statistical applications, there is no additional apriori knowledge about the distribution of the samples, thus assumptions such as Gaussianity and sphericity are made. In data assimilation applications, however, climatological estimates of covariance are commonplace.

Assume that one has access to a target covariance matrix $\mathbf{P}$ that represents the *a priori* knowledge about the error covariances. This matrix can be a climatological estimate of the covariance, or can be chosen through some ergodic assumption (with localization) using previous data. In existing data assimilation algorithms, such estimates most often exist for 4D-Var methods, and take the form of a static known background covariance that is an independent estimate from the current state.

We seek to combine this offline estimate of the covariance containing prior knowledge with the online estimate of the covariance obtained from the EnKF ensemble. In this paper we focus on an additive shrinkage covariance structure which is a linear combination of the the target covariance

matrix and the empirical covariance (2.7):

$$\mathbf{B}_i^f = \gamma_i \mu_i \mathbf{P} + (1 - \gamma_i) \tilde{\Sigma}_{X_i^f X_i^f}, \quad (2.11)$$

with γ_i represents the shrinkage factor (the linear combination coefficient) and μ_i represents a scaling factor. The choice of γ_i is extremely important. In contemporary data assimilation literature (e.g., see [4]) this factor is taken as a hyper-parameter whose optimal tuning could lead to significant reduction in error; however, the empirical tuning by trial and error is costly, and poor choices can offset possible improvements.

By employing a general invertible target matrix $\mathbf{P}$, and optimizing for a 2-norm distance over the “true” covariance, a closed-form expression to compute the shrinkage factor γ_i is proposed in [44, 46]. In this derivation, weights are computed as follows:

$$\gamma_i = \min \left(\frac{\frac{1}{N^2} \sum_{k=1}^N \|\mathbf{x}_i^{(k),f} - \bar{\mathbf{x}}_i^f\|^4 - \frac{1}{N} \|\tilde{\Sigma}_{X_i^f X_i^f}\|^2}{\|\tilde{\Sigma}_{X_i^f X_i^f} - \mathbf{P}\|^2}, 1 \right). \quad (2.12)$$

Since the estimate (2.12) is expensive to compute in an operational setting, here we will settle for a more computationally inexpensive method. No assumptions about the structure of $\mathbf{P}$ are made to compute γ_i . The general form (2.11) can be reduced to a standard form where the target matrix is the (scaled) identity by defining:

$$\mathbf{C}_i := \mathbf{P}^{-\frac{1}{2}} \tilde{\Sigma}_{X_i^f X_i^f} \mathbf{P}^{-\frac{1}{2}}, \quad \mu_i = \frac{\text{tr}(\mathbf{C}_i)}{n}, \quad (2.13)$$

where the new target matrix $\mu_i \mathbf{I}_{n \times n}$ represents a spherical climatological assumption on $\mathbf{C}_i$. Equivalently we can write

$$\mathbf{P}^{-\frac{1}{2}} \mathbf{B}_i^f \mathbf{P}^{-\frac{1}{2}} = \gamma_i \mu_i \mathbf{I}_{n \times n} + (1 - \gamma_i) \mathbf{C}_i. \quad (2.14)$$

As is traditional with Kalman type methods, we make the assumption that all our samples are drawn from an underlying Gaussian distribution. This assumption allows for a simpler computation of γ_i . The Rao-Blackwellized Ledoit-Wolf (RBLW) estimator [9] [31, equation (9)]:

$$\gamma_{i,\text{RBLW}} = \min \left[\frac{N-2}{N(N+2)} + \frac{(n+1)N-2}{\hat{U}_i N(N+2)(n-1)}, 1 \right], \quad (2.15)$$

is the optimal estimate of the covariance shrinkage factor under Gaussian assumptions. The com-

putationally dominant (and interesting) term in (2.15) is the sphericity factor

$$\hat{U}_i = \frac{1}{n-1} \left(\frac{n \operatorname{tr}(\mathbf{C}_i^2)}{[\operatorname{tr}(\mathbf{C}_i)]^2} - 1 \right), \quad (2.16)$$

which measures how similar the correlation structures of the sample and the target covariance are. For example if the both the target matrix $\mathbf{P}$ and the empirical covariance matrix $\tilde{\Sigma}_{X_i^f X_i^f}$ are diagonal ($\mathbf{C}_i$ is diagonal in (2.14)), then $\hat{U}_i = 0$ in (2.16), meaning that the RBLW estimate (2.15) would be $\gamma_{i,\text{RBLW}} = 1$. If, on the other hand, there is a large difference between the structures of $\mathbf{P}$ and $\tilde{\Sigma}_{X_i^f X_i^f}$ ($\mathbf{C}_i$ has large off-diagonal elements in (2.14)), then the sphericity factor is close to 1, forcing the RBLW estimate to be small (e.g., for $n = 10^{10}$ and $N = 50$, with $\hat{U}_i = 1$, one has $\gamma_{i,\text{RBLW}} = 0.038$).

Note that if our samples are also used to calculate the sample mean, the effective sample size of the sample covariance is smaller by one, therefore for most practical applications one replaces N by $N - 1$ in (2.15).

A drawback of the RBLW estimate is its reliance on the Gaussian assumption. A second drawback is that it is only valid for an over sampled ensemble with $N > n$, meaning that in the typical undersampled regime of EnKF with $N \ll n$, the factor is technically not well-defined. It is nonetheless still useful, in a similar fashion to ill-defined covariance estimates in the EnKF.

Aside from the inherent issues with the RBLW estimator, there are two major issues with its application in the EnKF, both related to its reliance on the sphericity of $\mathbf{C}_i$. First, when operating in the undersampled regime $N \ll n$, the estimate $\mathbf{C}_i$ (2.13) is also undersampled, and the problem of “spurious correlations” will affect the measure of sphericity (2.16). The second related issue regards the climatological estimate $\mathbf{P}$. Unless the climatological estimate accurately measures the correlation structure of the sample covariance, the shrinkage estimate (2.11) could potentially not be representative of our current uncertainty. The long-term accuracy of the climatological estimate to the covariance is thus of great importance.

Note that there are alternatives for non-invertible $\mathbf{P}$. Commonly, a reduced spectral version of $\mathbf{P}$ is known, $\mathbf{P} = \mathcal{U} \mathcal{L} \mathcal{U}^*$, with the $\mathcal{L}$ being a diagonal matrix of $r \ll n$ spectral coefficients, and $\mathcal{U}$ being an $n \times r$ matrix of orthonormal coefficients. The canonical symmetric pseudo-inverse square-root of $\mathbf{P}$ would therefore be $\mathbf{P}^{-\frac{1}{2}} = \mathcal{U} \mathcal{L}^{-\frac{1}{2}} \mathcal{U}^*$. If σ_k is the k -th singular value of $\mathbf{P}^{-1/2} \mathbf{A}^f$, then the traces appearing in (2.16) can be computed as follows:

$$\operatorname{tr}(\mathbf{C}_i) = \sum_{k=1}^{N-1} \sigma_k^2, \quad \operatorname{tr}(\mathbf{C}_i^2) = \sum_{k=1}^{N-1} \sigma_k^4.$$

Note that only the first $N - 1$ singular values are required for computation, even if $n \gg r \gg N - 1$.

The choice of a suitable target matrix $\mathbf{P}$ is very much an open question, and depends entirely on the problem at hand and on the available data. Some of the possible options include:

- matrices that are used in variational data assimilation methods,
- localized (through (2.9)) estimates from historical data, such as from a previous cycle in quasi-periodic models, and
- estimates derived from more long-term models (such as climate models).

The above is a non-exhaustive list, and it would be up to the practitioner to decide the validity of one estimate over another. The mismatch of the target covariance with the covariance estimate derived from the dynamical ensemble through the sphericity factor (2.16) could also be used in an online manner to determine the utility of the target $\mathbf{P}$, as a poor choice of the target matrix could significantly decrease the overall accuracy of the method, in the author's experience.

3. ETKF IMPLEMENTATION WITH STOCHASTIC SHRINKAGE COVARIANCE ESTIMATES

In ensemble-based methods our uncertainty is represented by an ensemble of samples of the underlying probability distribution. We wish to augment this representation of our uncertainty by augmenting the ensemble of samples with historical (climatological) samples of said information, as the application of Bayes' rule requires that all available information is used [20].

A naive approach to augmenting the ensemble would simply involve sampling from some known climatological distribution, for example sampling synthetic anomalies from a mean-zero Gaussian with known covariance, and appending this ensemble to our existing dynamical ensemble. This would, however, not be statistically sound, as the coupling between the two distributions would not be explicitly utilized. We therefore attempt to make use of the covariance shrinkage estimate (2.11) to the covariance in order to couple the dynamical and synthetic ensembles correctly.

We build on previous work by Nino-Ruiz and Sandu [30, 32] who proposed to replace the empirical covariance in EnKF with a shrinkage covariance estimator (2.11). They showed that this considerably improves the analysis at a modest additional computational cost. Additional, it was shown that synthetic ensemble members drawn from a normal distribution with covariance $\mathbf{B}^f$ are used to decrease the sampling errors.

In this work we develop an implementation of ETKF with a stochastic shrinkage covariance estimator (2.11). Rather than computing the covariance estimate (2.11), we build a synthetic ensemble by sampling directly from a distribution with covariance $\mu_i \mathbf{P}$. The anomalies of this synthetic ensemble are independent of the anomalies of the forecast EnKF ensemble.

Our approach works in a similar manner, but instead of simply augmenting the ensemble in a naive manner, we attempt to augment the ensemble in a statistically consistent manner by utilizing the theory behind optimal shrinkage estimators. If the dynamical system is locally (in time) stationary, climatologies about the local time roughly describe a measure of averaged-in-space uncertainty.

To be specific, let $\mathcal{X}^f \in \mathbb{R}^{n \times M}$ be a synthetic ensemble with M members (as opposed to the dynamic ensemble $\mathbf{X}_i^f$ with N members) drawn from a climatological probability density. We denote the variables related to the synthetic ensemble by calligraphic letters.

An important issue is the choice of the climatological distribution. As sampling from the dynamical manifold is impractical, heuristic assumptions are made about the distributions involved. A useful known heuristic is the principle of maximum entropy (PME) [20]. Assume that the mean and covariance of the distribution are known (through sampling), and that the distribution is supported over all of $\mathbb{R}^n$. The synthetic ensemble distribution of maximum entropy consistent with these assumptions is Gaussian:

$$\mathcal{X}_i^f \sim \mathcal{N}(\bar{\mathbf{x}}_i^f, \mu_i \mathbf{P}). \quad (3.1)$$

The synthetic ensemble anomalies in the state and observation spaces are:

$$\begin{aligned} \mathcal{X}_i^f &= \frac{1}{\sqrt{M-1}} \left(\mathcal{X}_i^f - \bar{\mathcal{X}}_i^f \mathbf{1}_M^\top \right) \in \mathbb{R}^{n \times M}, \\ \mathcal{Z}_i^f &= \frac{1}{\sqrt{M-1}} \left(\mathcal{H}(\mathcal{X}_i^f) - \overline{\mathcal{H}(\mathcal{X}_i^f)} \mathbf{1}_M^\top \right) \in \mathbb{R}^{m \times M}. \end{aligned} \quad (3.2)$$

The shrinkage estimator (2.11) of the forecast error covariance for $\mathbf{B}_i^f$ is represented in terms of synthetic and forecast anomalies as follows:

$$\tilde{\mathbf{B}}_i^f = \gamma_i \mathcal{A}_i^f \mathcal{A}_i^{f\top} + (1 - \gamma_i) \mathbf{A}_i^f \mathbf{A}_i^{f\top}. \quad (3.3)$$

The Kalman filter formulation [22] yields the following analysis covariance matrix:

$$\tilde{\mathbf{B}}_i^a = \tilde{\mathbf{B}}_i^f - \tilde{\mathbf{B}}_i^f \mathbf{H}_i^\top \mathbf{S}_i^{-1} \mathbf{H}_i \tilde{\mathbf{B}}_i^f, \quad (3.4)$$

where $\mathbf{S}_i$ will be discussed later.

Using the forecast error covariance estimate (3.3) in (3.4) leads to the following analysis covariance:

$$\begin{aligned} \tilde{\mathbf{B}}_i^a &= \gamma_i \mathcal{A}_i^f \mathcal{A}_i^{f,\top} + (1 - \gamma_i) \mathbf{A}_i^f \mathbf{A}_i^{f,\top} \\ &\quad - \left(\gamma_i \mathcal{A}_i^f \mathcal{Z}_i^{f,\top} + (1 - \gamma_i) \mathbf{A}_i^f \mathbf{Z}_i^{f,\top} \right) \mathbf{S}_i^{-1} \left(\gamma_i \mathcal{Z}_i^f \mathcal{A}_i^{f,\top} + (1 - \gamma_i) \mathbf{Z}_i^f \mathbf{A}_i^{f,\top} \right), \end{aligned} \quad (\text{goal-cov})$$

which we refer to as the “goal” analysis covariance formula. Where in this paper the factor the factor γ_i is chosen to be the RBLW estimate in (2.15), unless otherwise specified.

The ensemble goal of our modified ensemble Kalman filter is to construct an N -member analysis ensemble such that the anomalies $\mathbf{A}_i^a$ (2.5a) represent the (goal-cov) analysis covariance as well as possible:

$$\text{Find } \mathbf{A}_i^a \in \mathbb{R}^{n \times N} \text{ such that : } \mathbf{A}_i^a \mathbf{A}_i^{a,\top} \approx \tilde{\mathbf{B}}_i^a. \quad (\text{goal-an})$$

In the proposed method, we enrich our forecast ensemble in a way that closely approximates the shrinkage covariance (2.11).

3.1. The stochastic shrinkage implementation

We enrich the ensembles of forecast anomalies with synthetic anomalies (3.2):

$$\begin{aligned} \mathcal{A}_i^f &= \begin{bmatrix} \sqrt{1 - \gamma_i} \mathbf{A}_i^f & \sqrt{\gamma_i} \mathcal{A}_i^f \end{bmatrix} \in \mathbb{R}^{n \times (N+M)}, \\ \mathcal{Z}_i^f &= \begin{bmatrix} \sqrt{1 - \gamma_i} \mathbf{Z}_i^f & \sqrt{\gamma_i} \mathcal{Z}_i^f \end{bmatrix} \in \mathbb{R}^{m \times (N+M)}. \end{aligned} \quad (3.5)$$

Next, we define a transform matrix (2.5a) that is applied to the enriched ensemble (3.5), and leads to an analysis ensemble that represents the target analysis covariance (goal-an). Specifically, we search for a transform matrix $\mathcal{T}_i$ such that:

$$\tilde{\mathbf{B}}_i^a = \mathcal{A}_i^f \mathcal{T}_i \mathcal{T}_i^\top \mathcal{A}_i^{f,\top}. \quad (3.6)$$

Using the extended ensembles (3.6) the (goal-cov) becomes

$$\tilde{\mathbf{B}}_i^a = \mathcal{A}_i^f \left(\mathbf{I}_{(N+M) \times (N+M)} - \mathcal{Z}_i^{f,\top} \mathbf{S}_i^{-1} \mathcal{Z}_i^f \right) \mathcal{A}_i^{f,\top}, \quad (3.7)$$

where, from (2.6b),

$$\mathbf{S}_i = \mathbf{Z}_i^f \mathbf{Z}_i^{f,\top} + \mathbf{R}_i. \quad (3.8)$$

The transform matrix (2.6a) is a square root of (3.6):

$$\mathcal{T}_i = \left(\mathbf{I}_{(N+M) \times (N+M)} - \mathbf{Z}_i^{f,\top} \mathbf{S}_i^{-1} \mathbf{Z}_i^f \right)^{\frac{1}{2}}. \quad (3.9)$$

We compute the analysis mean using the shrinkage covariance estimate. From (2.5b):

$$\bar{\mathbf{x}}_i^a = \bar{\mathbf{x}}_i^f + \mathcal{A}_i^f \mathcal{T}_i \mathcal{T}_i^\top \mathbf{Z}_i^{f,\top} \mathbf{R}_i^{-1} \mathbf{d}_i, \quad (3.10)$$

where the full analysis covariance estimate (3.7) is used. In addition, we achieve the (goal-an) by keeping the first N members of the transformed extended ensemble, or equivalently, the first N columns of the symmetric square root (3.9). From (2.5a) we have:

$$\mathbf{A}_i^a = \frac{1}{\sqrt{1-\gamma_i}} [\mathcal{A}_i^f \mathcal{T}_i]_{:,1:N} = \frac{1}{\sqrt{1-\gamma_i}} \mathcal{A}_i^f \check{\mathcal{T}}_i, \quad \check{\mathcal{T}}_i = [\mathcal{T}_i]_{:,1:N}. \quad (3.11)$$

An alternative approach to achieve the (goal-an) is to look for a low-rank, approximate square root instead of the symmetric square root (3.9). Specifically, we seek a transformation matrix $\widehat{\mathcal{T}}_i$ such that:

$$\widehat{\mathcal{T}}_i \in \mathbb{R}^{(N+M) \times N}, \quad \widehat{\mathcal{T}}_i \widehat{\mathcal{T}}_i^\top \approx \mathbf{I}_{(N+M) \times (N+M)} - \mathbf{Z}_i^{f,\top} \mathbf{S}_i^{-1} \mathbf{Z}_i^f. \quad (3.12)$$

The calculation of the symmetric square root (3.9) requires an SVD of the right hand side matrix.

With the same computational effort one can compute the low rank transformation:

$$\begin{aligned} \mathbf{U} \mathbf{\Sigma} \mathbf{U}^\top &= \left(\mathbf{I} - \mathbf{Z}_i^{f,\top} \mathbf{S}_i^{-1} \mathbf{Z}_i^f \right), \quad \mathbf{U}, \mathbf{\Sigma} \in \mathbb{R}^{(N+M) \times (N+M)}, \\ \check{\mathcal{T}}_i &= \mathbf{U} \mathbf{\Sigma}^{1/2} [\mathbf{U}_{1:N,:}]^\top \in \mathbb{R}^{(N+M) \times N} \quad (\text{symmetric square root (3.9)}); \\ \widehat{\mathcal{T}}_i &= \mathbf{U}_{:,1:N} \mathbf{\Sigma}_{1:N,1:N}^{1/2} \in \mathbb{R}^{(N+M) \times N} \quad (\text{low rank square root (3.12)}). \end{aligned} \quad (3.13)$$

The mean calculation (3.10) is the same. The ensemble transform produces N transformed ensemble members that contain “mixed” information from both the physical and the synthetic ensembles:

$$\mathbf{A}_i^a = \frac{1}{\sqrt{1-\gamma_i}} \mathcal{A}_i^f \widehat{\mathcal{T}}_i.$$

3.2. Localization

It is possible to combine the proposed stochastic shrinkage approach with traditional localization. The LETKF implementation [19] computes transform matrices for subsets of variables, corresponding to localized spatial domains. In a similar vein one can combine our shrinkage algorithm with classical localization, as follows. Subsets of variables of the enriched ensembles (3.5) are used to compute local transform matrices (3.9) or (3.12), which are then applied to transform the corresponding local subsets, i.e. to compute the corresponding rows in equations (3.11) or (3.11), respectively. Utilizing the Sherman-Morrison-Woodbury identity [34], it is possible to decompose the inverse of (3.8) into,

$$\mathbf{S}_i^{-1} = \mathbf{R}_i^{-1} - \mathbf{R}_i^{-1} \mathbf{Z}_i (\mathbf{Z}_i^T \mathbf{R}_i^{-1} \mathbf{Z}_i + \mathbf{I}_{N+M})^{-1} \mathbf{Z}_i^T \mathbf{R}_i^{-1}, \quad (3.14)$$

where the inverse observation covariance is replaced by the localized variant in (2.10).

4. NUMERICAL EXPERIMENTS

In the numerical experiments we aim to assess the performance of the methodology in three different regimes: (i) a small scale model (Lorenz '96) to empirically test the performance of the optimally estimate the covariance shrinkage factors γ (2.15) constants against hand-picked values, (ii) a medium scale model (Quasi-geostrophic equations) with small observation errors to test the unlocalized shrinkage covariance methodology against the state-of-the-art LETKF, (iii) a geophysical model (shallow water on the sphere) with large observation errors to test our localized methodology against the LETKF.

All test problem implementations are available in the ‘ODE Test Problems’ suite [11, 39].

4.1. The Lorenz'96 model (L96)

We first consider the 40-variable Lorenz '96 problem [26],

$$[y]_i' = -[y]_{i-1} ([y]_{i-2} - [y]_{i+1}) - [y]_i + F, \quad i = 1, \dots, 40, \quad F = 8. \quad (4.1)$$

Assuming (4.1) is ergodic (thus having a constant spatio-temporal measure of uncertainty on the manifold of the attractor), we compute the target covariance matrix $\mathbf{P}$ as the empirical covariance




FIG. 1: Results for the L96 problem with dynamic ensemble sizes of $N = 5$ and $N = 14$, inflation factor $\alpha = 1.1$, and different synthetic ensemble sizes M . We compute the KL divergence of the rank histogram (4.2) and the RMSE (4.3) for the methods. Error bars show two standard deviations.

from 10,000 independent ensemble members run over 225 days in the system (where 0.05 time units corresponds to 6 hours), with an interval of 6 hours between snapshots. This system is known to have 13 positive Lyapunov exponents, with a Kaplan-Yorke dimension of about 27.1 [36].

The time between consecutive assimilation steps is $\Delta t = 0.05$ units, corresponding to six hours in the system. All variables are observed directly with an observation error covariance matrix of $\mathbf{R}_i = \mathbf{I}_{40}$. The time integration of the model is performed with RK4 the fourth order Runge-Kutta scheme RK4 [16] with a step size $h = \Delta t$. The problem is run over 2200 assimilation steps. The first 200 are discarded to account for model spinup. Twenty independent model realizations are performed in order to glean statistical information thereof.




FIG. 2: Results for the L96 problem. The left panel presents the analysis RMSE for various values of the dynamic and synthetic ensemble sizes. The right panel presents the shrinkage factor γ (2.15) for a synthetic ensemble size of $M = 100$ over a number of assimilation steps, with error bars showing two standard deviations.

4.2. L96 assimilation results

We assess the quality of the analysis ensembles using a rank histogram [17], cumulative over 20 independent runs. For a quantitative metric we compute the KL divergence from Q to P ,

$$D_{KL}(P||Q) = - \sum_k P_k \log \left(\frac{P_k}{Q_k} \right), \quad (4.2)$$

where P is the uniform distribution and Q is our ensemble rank histogram, and P_k & Q_k are the discrete probabilities associated with each bin. A low KL divergence would indicate that our rank histogram is close to uniform, and thus the ensemble is representative of the truth.

Additionally, for testing the accuracy of all our methods we compute the spatio-temporal analysis RMSE,

$$\sqrt{\frac{1}{Kn} \sum_{i=1}^K \sum_{j=1}^n [x_i^a - x_i^t]_j^2}, \quad (4.3)$$

with K representing the amount of snapshots at which the analysis is computed, and $[x_i]_j$ is the j th component of the state variable at time i .

For the given settings of a severely undersampled ensemble ($N = 5$) and mild inflation ($\alpha = 1.1$), we compare the Gaussian sampling methodology coupled to the RBLW formulation for the shrinkage factor γ (2.15), with the optimal static $\gamma = 0.85$ shrinkage factor. For a dynamic ensemble that captures the positive error growth modes ($N = 14$) we will compare the RBLW estimator with the optimal static $\gamma = 0.1$. We will compare the mean and variance of the KL divergence of

the rank histogram of the variable $[y]_{17}$ from the uniform, and the statistics of the spatio-temporal RMSE.

The results are reported in Figure 1. For both an undersampled and sufficient ensemble, the optimal shrinkage factor has a smaller mean error, and a smaller KL divergence with less variance (top left, top middle, bottom left, and bottom middle panels). In the undersampled case, the RBLW estimator induces more variance into the RMSE (top middle panel). For the sufficiently sampled ensemble, however, the optimal static shrinkage value induced significantly more variance into the error, with the RBLW estimator reducing the error significantly (bottom middle panel).

It is possible that a better estimator than RBLW may get the ‘best of both worlds’ and induce low error with low variance, though this is as-of-yet out of reach. This is to be expected as the RBLW estimate is only accurate in the limit of ensemble size, and there is no theory about its accuracy in the undersampled case. In the authors’ experience other estimators such as OAS, while having the theoretically desired properties, perform empirically worse in conjunction with ensemble methods. Currently, for a modest reduction in accuracy, one of the hyperparameters can be estimated online by the methodology.

For the second round of experiments with Lorenz ’96, reported in Figure 2, we compare analysis errors when the synthetic and dynamic ensemble sizes are modified (left panel). It is evident that increases in both dynamic and synthetic ensemble size lead to lower error. In the right panel we also compare dynamic ensemble size to the values of γ that are produced. It is clear that an increase in dynamic ensemble size decreases the need for shrinkage.

4.3. The Quasi-Geostrophic model (QG)

We follow the QG formulations given in [28, 43]. We discretize the equation

$$\begin{aligned} \omega_t + J(\psi, \omega) - Ro^{-1} \psi_x &= Re^{-1} \Delta \omega + Ro^{-1} F, \\ J(\psi, \omega) &\equiv \psi_y \omega_x - \psi_x \omega_y, \end{aligned} \tag{4.4}$$

where ω stands for the vorticity, ψ stands for the stream function, Re is the Reynolds number, Ro is the Rossby number, J is the Jacobian term, and F is a constant (in time) forcing term.

The relationship between stream and vorticity, $\omega = -\Delta \psi$ is explicitly enforced in the evaluation of the ODE. The forcing term is a symmetric double gyre,

$$F = \sin(\pi(y - 1)). \tag{4.5}$$




FIG. 3: Analysis RMSE results for the QG model. The experiments use a synthetic ensemble size $M = 100$ and Gaussian samples. Results are compared against LETKF with the Gaspari-Cohn decorrelation function (GC)

Homogeneous Dirichlet boundary conditions are enforced on the spatial domain $[0, 1] \times [0, 2]$. The spatial discretization is a second order central finite difference for the first derivatives, and the Laplacian, with the Arakawa approximation [3] (a pseudo finite element scheme [21]) used for computing the Jacobian term. All spatial discretizations exclude the trivial boundary points from explicit computation.

The matrix $\mathbf{P}$ is approximated from 700 snapshots of the solution about 283 hours apart each, with Gaspari-Cohn localization applied, so as to keep the matrix sparse. The true model is run outside of time of the snapshots so as to not pollute the results. Nature utilizes a 255×511 spatial discretization, and the model a 63×127 spatial discretization. Observations are first relaxed into the model space via multigriding [47], then 150 distinct spatial points (using an observation operator similar to [41]) from the non-linear observation operator,

$$\mathcal{H}(\psi) = \sqrt{\psi_x^2 + \psi_y^2}, \quad (4.6)$$

representing zonal wind magnitude, are taken. The observation error is unbiased, with covariance $\mathbf{R} = 4\mathbf{I}_{150}$. The number of synthetic ensemble members is fixed at a constant $M = 100$, as to be more than the number of full model run ensemble members, but significantly less than the rank of the covariance. Observations are taken $\Delta t = 0.010886$ time units (representing one day in model space) apart. We run a total of 350 assimilation steps, taking the first 50 as spinup. Results are averaged over 5 model runs (with the same nature run, but different initializations of the dynamic ensemble), with diverging runs treated as de-facto infinite error.

4.4. QG assimilation results

Figure 3 reports the results with the QG model. Comparing our methodology to the LETKF with an optimally tuned Gaspari-Cohn (GC) [15] localization (such that both error and stability are prioritized), we see that GC significantly decreases the error for larger values of N and α , but is not stable for more operational under-sampled dynamic ensemble sizes and low inflation factors, as opposed to our shrinkage method. Possible sources of error are both the nonlinear observations and the coarse approximation to the covariance estimate.

These results lend additional support to the argument that shrinkage alone is not enough. Localization is still required in operational settings, and combining both might yield a positive result.

The quasi-geostrophic results indicate that our methodology holds promise to be of use for practical data assimilation systems, and that the methodology can handle observations that are non-linear transformations of the state representation. However, the methodology needs to be refined with more optimal shrinkage factors for operational undersampled empirical covariances.

An operational implementation of the LETKF requires $m \times N$ linear solves and m matrix square roots, while our stochastic shrinkage algorithm requires $N + M$ linear solves and one matrix square root. Thus as the number of observations grows, the stochastic shrinkage methodology becomes a lot more compelling.

4.5. Shallow water on a sphere (SWS)

The last round of experiments aims at validating the Localized Shrinkage ETKF on a different geophysical problem of interest. To that end we employ the shallow water equations [14, 29] on the sphere, which represent an approximation of the atmospheric dynamic over Earth. We use a modification of the Cartesian shallow water equations,

$$h_t = -\nabla \cdot (h \cdot \mathbf{u}), \quad (4.7)$$

$$\mathbf{u}_t = -(\mathbf{u} \cdot \nabla)\mathbf{u} - f(\mathbf{p} \times \mathbf{u}) - g\nabla h, \quad (4.8)$$

under the constraint that the flow is confined to a spherical approximation to the Earth; the radius of the sphere is one spatial unit. Here, f is the Coriolis force, g is gravity, h is the height of the water, and $\mathbf{p}$ and $\mathbf{u}$ are the x , y , and z positions and velocities, respectively. We follow the radial basis function formulation in [14] for the spatial discretization with 100 points, for a total state




FIG. 4: Left panel: initial condition of the water height with blue represented lower than average and yellow representing higher than average, and observation locations (red points). Right panel: analysis RMSE for the localized shrinkage ETKF, and the localized ETKF with the Gaspari-Cohn decorrelation function, with the error bars representing two standard deviations.

space dimension of $n = 400$. We take the order three Buhmann function [6] with a Cartesian radius of $r = 2$ on the unit sphere (representing full coverage). We use a third order adaptive strong stability preserving method [27] for time integration of this system.

We observe the height at ten locations over the domain; the velocities are unobserved. Observations are taken every $\Delta t = 1$ day over the assimilation window. The observation covariance is $\mathbf{R} = 100 \mathbf{I}_{10}$, to simulate a noisy observation scenario.

We compare the localized variant of the stochastic shrinkage approach (see section 3.2), which we term the L-Shr-ETKF, against LETKF. We select a synthetic ensemble size $M = 250$. For localization, we use a great circle radius of $r = \pi/5$ spatial units, with the Gaspari-Cohn decorrelation function, as this was found to be approximately optimal for the LETKF by manual tuning. The best inflation factors obtained by manual tuning are used, as follows: $\alpha = 1.001$, and $\alpha = 1.05$ for LETKF.

We run a total of three months of observations for January, February and March (90 days), discarding the first 31 days of January as spinup, and observing the analysis RMSE for a range of physical ensemble sizes, $N \in [4, 40]$. Twelve total independent runs are taken to account for possible spurious results.

4.6. SWS Results

The left panel of figure 4 show the initial conditions, and observations for the shallow water equations. The initial conditions were chosen to be quasi-stable so that they would slowly diverge from a cyclic solution. The right panel shows the results comparing L-Shr-ETKF against the state-of-the-art LETKF. It can be clearly seen that even for large ensemble sizes up to 40, the LETKF error is higher than the observation error of 100. The LETKF also suffers from large variability in the error from various different initial ensembles. L-Shr-ETKF, on the other hand, matches the observation error for an ensemble size of $N = 4$ dynamical members, and has lower error than that of the observations for all larger dynamical ensemble sizes. The variability of the L-Shr-ETKF error is also substantially smaller than that of LETKF.

The results clearly demonstrate that, in a small-ensemble high-observation-error regime where LETKF performs relatively poorly, the proposed L-Shr-ETKF algorithm provides robust analyses.

5. DISCUSSION

Shrinkage covariance matrix estimators were shown to greatly improve the performance of the EnKF [30]. This work extends the the idea of covariance shrinkage to the ensemble transform Kalman filter. Instead of enhancing the covariance estimate, we propose enhancing the ensemble with a synthetic ensemble derived from the target matrix of the shrinkage approach. By applying the ETKF formulas to this enhanced ensemble, we develop the Shr-ETKF, whose internal representation of the Kalman gain is approximately based on the shrinkage estimate of the covariance.

We compare Shr-ETKF to the current state-of-the-art LETKF algorithm on several test problems. Lorenz '96 model results indicate that the new filter performs worse in the under-sampled regime than the best 'static' shrinkage method, and performs better (in terms of less variance in the error) than an optimal dynamic shrinkage method for the sufficiently sampled case. Results with QG model indicate that our method could potentially be used to augment operational LETKF implementations, but not in the low-observation-error regime. Results with the shallow-water equation on a sphere model show that a localized stochastic covariance shrinkage ETKF can perform significantly better than the LETKF in a high-observation-error regime.

These results indicate that L-Shr-ETKF can be potentially utilized in an operational framework to improve the performance of LETKF while keeping the dynamical ensemble size (the number of forecast model runs) small. Additional work is needed to devise better heuristic estimates of the

shrinkage factor γ .

ACKNOWLEDGMENTS

The first two authors would like to thank Traian Iliescu and Changhong Mou for their in-depth help with understanding of the Quasi-Geostrophic model, Steven Roberts, the primary maintainer of the ODE Test Problems package, Amit N Subrahmanya for his help with the shallow water on a sphere problem, and the rest of the members of the Computational Science Laboratory at Virginia Tech for their continuous support.

The authors would like to thank the four previous referees for their insightful feedback that lead to an improved paper.

The first two authors were supported, in part, by the National Science Foundation through awards NSF ACI-1709727, NSF CCF-1613905, NSF CDS&E-MSS 1953113, AFOSR through the award AFOSR DDDAS 15RT1037, and by DOE ASCR through the award DE-SC0021313.

The last author was supported by the Research Council of Norway and the companies AkerBP, Wintershall-DEA, Vår Energy, Petrobras, Equinor, Lundin and Neptune Energy, through the Petromaks-2 project (280473) DIGIRES (<http://digires.no>)

REFERENCES

-
- [1] Anderson, J. L. (2001). An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, 129(12):2884–2903.
 - [2] Anderson, J. L. (2012). Localization and sampling error correction in ensemble Kalman filter data assimilation. *Monthly Weather Review*, 140(7):2359–2371.
 - [3] Arakawa, A. (1966). Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. Part I. *Journal of Computational Physics*, 1(1):119–143.
 - [4] Asch, M., Bocquet, M., and Nodet, M. (2016). *Data assimilation: methods, algorithms, and applications*. SIAM.
 - [5] Bishop, C. H., Etherton, B. J., and Majumdar, S. J. (2001). Adaptive sampling with the ensemble transform kalman filter. part i: Theoretical aspects. *Monthly weather review*, 129(3):420–436.
 - [6] Buhmann, M. D. (1998). Radial functions on compact support. *Proceedings of the Edinburgh Mathematical Society*, 41(1):33–46.

- [7] Burgers, G., Jan van Leeuwen, P., and Evensen, G. (1998). Analysis scheme in the ensemble kalman filter. *Monthly weather review*, 126(6):1719–1724.
- [8] Chen, Y., Wiesel, A., Eldar, Y. C., and Hero, A. O. (2010). Shrinkage algorithms for mmse covariance estimation. *IEEE Transactions on Signal Processing*, 58(10):5016–5029.
- [9] Chen, Y., Wiesel, A., and Hero, A. O. (2009). Shrinkage estimation of high dimensional covariance matrices. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2937–2940. IEEE.
- [10] Chen, Y., Wiesel, A., and Hero, A. O. (2011). Robust shrinkage estimation of high-dimensional covariance matrices. *IEEE Transactions on Signal Processing*, 59(9):4097–4107.
- [11] Computational Science Laboratory (2020). ODE test problems.
- [12] Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*, 99(C5):10143–10162.
- [13] Evensen, G. (2009). *Data assimilation: the ensemble Kalman filter*. Springer Science & Business Media.
- [14] Flyer, N. and Wright, G. B. (2009). A radial basis function method for the shallow water equations on a sphere. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2106):1949–1976.
- [15] Gaspari, G. and Cohn, S. E. (1999). Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(554):723–757.
- [16] Hairer, E., Nørsett, S. P., and Wanner, G. (1991). *Solving ordinary differential equations. 1, Nonstiff problems*. Springer-Vlg.
- [17] Hamill, T. M. (2001). Interpretation of rank histograms for verifying ensemble forecasts. *Monthly Weather Review*, 129(3):550–560.
- [18] Hamill, T. M. and Snyder, C. (2000). A hybrid ensemble Kalman filter–3d variational analysis scheme. *Monthly Weather Review*, 128(8):2905–2919.
- [19] Hunt, B. R., Kostelich, E. J., and Szunyogh, I. (2007). Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, 230(1-2):112–126.
- [20] Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge university press.
- [21] Jespersen, D. C. (1974). Arakawa’s method is a finite-element method. *Journal of computational physics*, 16(4):383–390.
- [22] Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45.
- [23] Kalnay, E. (2003). *Atmospheric modeling, data assimilation and predictability*. Cambridge university press.
- [24] Law, K., Stuart, A., and Zygalakis, K. (2015). *Data assimilation: a mathematical introduction*, volume 62. Springer.
- [25] Ledoit, O. and Wolf, M. (2004). A well conditioned estimator for large dimensional covariance matrices.

- Journal of Multivariate Analysis*, 88(2):365–411.
- [26] Lorenz, E. N. (1996). Predictability: A problem partly solved. In *Proc. Seminar on predictability*, volume 1.
 - [27] Macdonald, C. B. (2003). *Constructing high-order Runge-Kutta methods with embedded strong-stability-preserving pairs*. PhD thesis, Theses (Dept. of Mathematics)/Simon Fraser University.
 - [28] Mou, C., Liu, H., Wells, D. R., and Iliescu, T. (2019). Data-driven correction reduced order models for the quasi-geostrophic equations: A numerical investigation. *arXiv preprint*, <http://arxiv.org/abs/1908.05297>.
 - [29] Neta, B., Giraldo, F., and Navon, I. (1997). Analysis of the turkel–zwas scheme for the two-dimensional shallow water equations in spherical coordinates. *Journal of Computational Physics*, 133(1):102–112.
 - [30] Nino-Ruiz, E. and Sandu, A. (2015). Ensemble Kalman filter implementations based on shrinkage covariance matrix estimation. *Ocean Dynamics*, 65(11):1423–1439.
 - [31] Nino-Ruiz, E. and Sandu, A. (2017). An ensemble Kalman filter implementation based on modified Cholesky decomposition for inverse covariance matrix estimation. *SIAM Journal on Scientific Computing*, submitted.
 - [32] Nino-Ruiz, E. and Sandu, A. (2018). Efficient parallel implementation of DDDAS inference using an ensemble Kalman filter with shrinkage covariance matrix estimation. *Cluster Computing*, pages 1–11.
 - [33] Nino-Ruiz, E. D., Sandu, A., and Deng, X. (2015). 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*, volume Supercomputing 2015 of *ScalA '15*, Austin, Texas.
 - [34] Petersen, K. B., Pedersen, M. S., et al. (2008). The matrix cookbook. *Technical University of Denmark*, 7(15):510.
 - [35] Petrie, R. (2008). Localization in the ensemble Kalman filter. *MSc Atmosphere, Ocean and Climate University of Reading*.
 - [36] Popov, A. A. and Sandu, A. (2019). A bayesian approach to multivariate adaptive localization in ensemble-based data assimilation with time-dependent extensions. *Nonlinear Processes in Geophysics*, 26(2):109–122.
 - [37] Popov, A. A. and Sandu, A. (2020). An explicit probabilistic derivation of inflation in a scalar ensemble Kalman filter for finite step, finite ensemble convergence.
 - [38] Reich, S. and Cotter, C. (2015). *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press.
 - [39] Roberts, S., Popov, A. A., and Sandu, A. (2019). ODE test problems: a MATLAB suite of initial value problems. *CoRR*, abs/1901.04098.
 - [40] Ruiz, E. N., Sandu, A., and Anderson, J. (2014). An efficient implementation of the ensemble Kalman filter based on an iterative Sherman–Morrison formula. *Statistics and Computing*, pages 1–17.
 - [41] Sakov, P. and Oke, P. R. (2008a). A deterministic formulation of the ensemble Kalman filter: an alternative to ensemble square root filters. *Tellus A*, 60(2):361–371.

- [42] Sakov, P. and Oke, P. R. (2008b). Implications of the form of the ensemble transformation in the ensemble square root filters. *Monthly Weather Review*, 136(3):1042–1053.
- [43] San, O. and Iliescu, T. (2015). A stabilized proper orthogonal decomposition reduced-order model for large scale quasigeostrophic ocean circulation. *Advances in Computational Mathematics*, 41(5):1289–1319.
- [44] Stoica, P., Li, J., Zhu, X., and Guerri, J. R. (2008). On using a priori knowledge in space-time adaptive processing. *IEEE Transactions on Signal Processing*, 56(6):2598–2602.
- [45] Zhang, Y., Liu, N., and Oliver, D. S. (2010). Ensemble filter methods with perturbed observations applied to nonlinear problems. *Computational Geosciences*, 14(2):249–261.
- [46] Zhu, X., Li, J., and Stoica, P. (2011). Knowledge-Aided Space-Time Adaptive Processing. *IEEE Transaction on Aerospace And Electronic Systems*, 47(2):1325–1336.
- [47] Zubair, H. B. (2009). *Efficient Multigrid Methods based on Improved Coarse Grid Correction Techniques*. PhD thesis, Delft University of Technology, Netherlands.