



A Bayesian approach to multivariate adaptive localization in ensemble-based data assimilation with time-dependent extensions

Andrey A. Popov and Adrian Sandu

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

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

Received: 5 October 2018 – Discussion started: 11 October 2018

Revised: 12 May 2019 – Accepted: 21 May 2019 – Published: 14 June 2019

Abstract. Ever since its inception, the ensemble Kalman filter (EnKF) has elicited many heuristic approaches that sought to improve it. One such method is covariance localization, which alleviates spurious correlations due to finite ensemble sizes by using relevant spatial correlation information. Adaptive localization techniques account for how correlations change in time and space, in order to obtain improved covariance estimates. This work develops a Bayesian approach to adaptive Schur-product localization for the deterministic ensemble Kalman filter (DEnKF) and extends it to support multiple radii of influence. We test the proposed adaptive localization using the toy Lorenz'96 problem and a more realistic 1.5-layer quasi-geostrophic model. Results with the toy problem show that the multivariate approach informs us that strongly observed variables can tolerate larger localization radii. The univariate approach leads to markedly improved filter performance for the realistic geophysical model, with a reduction in error by as much as 33 %.

1 Introduction

Data assimilation (Asch et al., 2016; Law et al., 2015; Evensen, 2009; Reich and Cotter, 2015) fuses information from the model forecast states (Eq. 1) and observations (Eq. 3) in order to obtain an improved estimation of the truth at any given point in time. Data assimilation approaches include the ensemble Kalman filters (EnKFs) (Evensen, 1994, 2009; Constantinescu et al., 2007b) that rely on Gaussian assumptions, particle filters for non-Gaussian distributions (Van Leeuwen et al., 2015; Attia et al., 2017; Attia and Sandu, 2015), and variational approaches, rooted in control

theory (Le Dimet and Talagrand, 1986; Sandu et al., 2005; Carmichael et al., 2008).

EnKF is an important family of data assimilation techniques that propagate both the mean and covariance of the state uncertainty through the model using a Monte Carlo approach. While large dynamical systems of interest have a large number of modes along which errors can grow, the number of ensemble members used to characterize uncertainty remains relatively small due to computational costs. As a result, inaccurate correlation estimates obtained through Monte Carlo sampling can profoundly affect the filter results. Techniques such as covariance localization and inflation have been developed to alleviate these problems (Petrie and Dance, 2010).

Localization techniques take advantage of the fundamental property of geophysical systems that correlations between variables decrease with spatial distance (Kalnay, 2003; Asch et al., 2016). This prior knowledge is used to scale ensemble-estimated covariances between distant variables so as to reduce inaccurate, spurious correlations. The prototypical approach to localization is a Schur-product-based tapering of the covariance matrix (Bickel and Levina, 2008); theoretical results ensure that the covariance matrices estimated using small ensembles sampled from a multivariate normal probability distribution, upon tapering, approach quickly the true covariance matrix. A formal theory of localization (Flowerdew, 2015) has been attempted, though a true multivariate theory is still out of our grasp. Practical implementations of localization rely on restricting the information flow, either in state space or in observation space, to take place within a given “radius of influence” (Hunt et al., 2007). Some variants of EnKF like the maximum likelihood ensemble filter (MLEF) (Zupanski, 2005) reduce the need for localization,

while others use localization in order to efficiently parallelize the analysis cycles in space (Nino-Ruiz et al., 2015).

The performance of the EnKF algorithms critically depends on the correct choice of localization radii (also known as the decorrelation distances), since values that are too large fail to correct for spurious correlations, while values that are too small throw away important correlation information. However, the physical values of the spatial decorrelation scales are not known a priori, and they change with the temporal and spatial location. At the very least the decorrelation scales depend on the current atmospheric flow. In atmospheric chemistry systems, because of the drastic difference in reactivity, each chemical species has its own individual localization radius (Constantinescu et al., 2007a). Multi-scale schemes (Buehner and Shlyayeva, 2015) for localization are an immediate necessity. Clearly, the widely used approach of estimating decorrelation distances from historical ensembles of simulations and then using a constant averaged value throughout the space and time domain leads to a suboptimal performance of ensemble Kalman filters.

Adaptive localization schemes seek to estimate decorrelation distances from the data, so as to optimize the filter performance according to some criteria. One approach to adaptive localization utilizes an ensemble of ensembles to detect and mitigate spurious correlations (Anderson, 2007). Relying purely on model dynamics and foregoing reliance on spatial properties of the model, the method is very effective for small-scale systems, but its applicability to large-scale geophysical problems is unclear. There is, however, evidence that optimal localization depends more on ensemble size and observation properties than on model dynamics (Kirchgeßner et al., 2014) and that adaptive approaches whose correlation functions follow these dynamics do not show a significant improvement over conventional static localization (Bishop and Hodyss, 2011). Methods such as sampling error correction (Anderson, 2012) take advantage of these properties to build correction factors and apply them as an ordinary localization scheme. A different approach uses the inherent properties of correlation matrices to construct Smoothed ENsemble CORrelations Raised to a Power (SEN-CORP) (Bishop and Hodyss, 2007) matrices that in the limiting case remove all spurious correlations. This method relies purely on the statistical properties of correlation matrices and ignores the model dynamics and the spatial and temporal dependencies it defines. A more theoretically sound approach (Ménétrier et al., 2015) would try to create a localization matrix such that the end result would better approximate the asymptotic infinite-ensemble covariance matrix. A recent approach considers machine learning algorithms to capture hidden properties in the propagation model that affect the localization parameters (Moosavi et al., 2018). Still another recent approach considers localization (and inflation) in terms of optimal experimental design (Attia and Constantinescu, 2018).

This work develops a Bayesian framework to dynamically learn the parameters of the Schur-product-based localization from the ensemble of model states and the observations during the data assimilation in geophysical systems. Specifically, the localization radii are considered random variables described by parameterized distributions and are retrieved as part of the assimilation step together with the analysis states. One of the primary goals of this paper is to develop ways in which such an approach could be extended to both multivariate and time-dependent 4-D-esque cases. We prove the approach's empirical validity through a type of idealized variance minimization that has access to the true solution (which we call an oracle). We then show that the approach provides a more stable result with a much larger initial radius guess. The exploration of the idea is done through the use of several test problems such as that of the Lorenz'96 problem, a multivariate variant of which we introduce specifically for this paper, and a more realistic quasi-geostrophic model to showcase the applicability of the method to scenarios more in line with operational ones.

The paper is organized as follows. Section 2 reviews background material for EnKF and Schur-product localization. Section 3 provides a framework for naturally extending univariate localization to the multivariate case. Section 4 describes the proposed theoretical framework for localization and the resulting optimization problems for maximum likelihood solutions. Numerical results with three test problems reported in Sect. 5 provide empirical validation of the proposed approach.

2 Background

We consider a computational model that approximates the evolution of a physical dynamical system such as the atmosphere:

$$\mathbf{x}_{i+1} = \mathcal{M}_{i,i+1}(\mathbf{x}_i) + \boldsymbol{\xi}_i. \quad (1)$$

The (finite-dimensional) state of the model $\mathbf{x}_i \in \mathbb{R}^n$ at time t_i approximates the (finite-dimensional projection of the) physical true state $\mathbf{x}_i^t \in \mathbb{R}^n$. The computational model (Eq. 1) is inexact, and we assume the model error to be additive and normally distributed, $\boldsymbol{\xi}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_i)$.

The initial state of the model is also not precisely known, and to model this uncertainty we consider that it is a random variable drawn from a specific probability distribution:

$$\mathbf{x}_0 \sim \mathcal{N}(\mathbf{x}_0^t, \mathbf{P}_0^b). \quad (2)$$

Consider also observations of the truth,

$$\mathbf{y}_{i+1} = \mathcal{H}(\mathbf{x}_{i+1}^t) + \boldsymbol{\eta}_{i+1}, \quad (3)$$

that are corrupted by normal observation errors $\boldsymbol{\eta}_{i+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_{i+1})$. We consider here the case with a linear observation operator, $\mathcal{H} := \mathbf{H}$.

Consider an ensemble of N states $\mathbf{x} \in \mathbb{R}^{n \times N}$ sampling the probability density that describes the uncertainty in the state at a given time moment (the time subscripts are omitted for simplicity of notation). The ensemble mean, ensemble anomalies, and ensemble covariance are

$$\bar{\mathbf{x}} = \frac{1}{N} \mathbf{x} \mathbf{1}_N, \quad \mathbf{X} = \mathbf{x} - \bar{\mathbf{x}} \mathbf{1}_N^\top, \quad \mathbf{P} = \frac{1}{N-1} \mathbf{X} \mathbf{X}^\top, \quad (4)$$

respectively. Typically N is smaller than the number of positive Lyapunov exponents in our dynamical system, and much smaller than the number of states (Bergemann and Reich, 2010). Consequently, the ensemble statistics (Eq. 4) are marred by considerable sampling errors. The elimination of sampling errors, manifested as spurious correlations in the covariance matrix (Evensen, 2009), leads to the need for localization.

2.1 Kalman analysis

The mean and the covariance are propagated first through the forecast step. Specifically, each ensemble member is advanced to the current time using the model (1) to obtain the ensemble forecast $\mathbf{x}^f$ (with mean $\bar{\mathbf{x}}^f$ and covariance $\mathbf{P}^f$) at the current time. A covariance inflation step $\mathbf{X}^f \leftarrow \alpha \mathbf{X}^f$, $\alpha > 1$ can be applied to prevent filter divergence (Anderson, 2001).

The mean and covariance are then propagated through the analysis step, which fuses information from the forecast mean and covariance and from observations (Eq. 3) to provide an analysis ensemble $\mathbf{x}^a$ (with mean $\bar{\mathbf{x}}^a$ and covariance $\mathbf{P}^a$) at the same time. Here we consider the deterministic EnKF (DEnKF) (Sakov and Oke, 2008), which obtains the analysis as follows:

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{K} \mathbf{d}, \quad (5a)$$

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^\top (\mathbf{H} \mathbf{P}^f \mathbf{H}^\top + \mathbf{R})^{-1}, \quad (5b)$$

$$\mathbf{d} = \mathbf{y} - \mathbf{H} \bar{\mathbf{x}}^f, \quad (5c)$$

$$\mathbf{X}^a = \mathbf{X}^f - \frac{1}{2} \mathbf{K} \mathbf{H} \mathbf{X}^f, \quad (5d)$$

$$\mathbf{x}^a = \bar{\mathbf{x}}^a \mathbf{1}_N^\top + \mathbf{X}^a, \quad (5e)$$

where $\mathbf{K}$ is the Kalman gain matrix and $\mathbf{d}$ the vector of innovations. DEnKF is chosen for simplicity of implementation and ease of amending to support Schur-product-based localization. However, the adaptive localization techniques discussed herein are general – they do not depend on this choice and are equally applicable to any EnKF algorithm.

2.2 Schur-product localization

Covariance localization involves the Schur (element-wise) product between a symmetric positive semi-definite matrix $\boldsymbol{\rho}$ and the ensemble forecast covariance:

$$\mathbf{P}^f \leftarrow \boldsymbol{\rho} \circ \mathbf{P}^f, \quad \mathbf{P}_{i,j}^f \leftarrow \rho_{i,j} \mathbf{P}_{i,j}^f. \quad (6)$$

By the Schur-product theorem (Schur, 1911), if $\boldsymbol{\rho}$ and $\mathbf{P}^f$ are positive semi-definite, then their Schur product is positive semi-definite. The matrix $\boldsymbol{\rho}$ is chosen such that it reduces the sampling errors and brings the ensemble covariance closer to the true covariance matrix. We note that one can apply the localization in ways that mitigate the problem of storing full covariance matrices (Houtekamer and Mitchell, 2001). Efficient implementation aspects are not discussed further as they do not impact the adaptive localization approaches developed herein.

We seek to generate the entries of the localization matrix $\boldsymbol{\rho}$ from a localization function $\ell : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$, used to represent the weights applied to each individual covariance:

$$\boldsymbol{\rho} = [\ell(d(i, j)/r)]_{1 \leq i, j \leq n}. \quad (7)$$

The function ℓ could be thought of as a regulator which brings the ensemble correlations in line with the physical correlations, which are often compactly supported functions (Gneiting, 2002). The metric d quantifies the physical distance between model variables, such that $d(i, j)$ represents the spatial distance between the state-space variables x_i and x_j . The “radius of influence” parameter r represents the correlation spatial scale; the smaller r is, the faster variables decorrelate with increasing distance.

If the spatial discretization is time-invariant, and ℓ and r are fixed, then the matrix $\boldsymbol{\rho}$ is also time-invariant. The goal of the adaptive localization approach is to learn the best value of r dynamically from the ensemble.

A common localization function used in production software is due to Gaspari and Cohn (Gaspari and Cohn, 1999; Gneiting, 2002; Petrie, 2008). Here we use the Gaussian function

$$\ell(u) = \exp(-u^2/2) \quad (8)$$

to illustrate the adaptive localization strategy. However, our approach is general and can be used with any suitable localization function.

3 Extension to multivariate localization functions

It is intuitively clear that different physical effects propagate spatially at different rates, leading to different correlation distances. Consequently, different state-space variables should be analyzed using different radii of influence. This raises the additional question of how to localize the covariance of two variables when each of them is characterized by a different radius of influence. One approach (Roh et al., 2015) is to use a multivariate compactly supported function (Askey, 1973; Porcu et al., 2013) to calculate the modified error statistics. We however wish to take advantage of already developed univariate compactly supported functions.

We define the mapping operator $\tau : \mathbb{N}_n \rightarrow \mathbb{N}_g$ that assigns each state-space component to a group. All state variables

assigned to the same group are analyzed using the same localization radius. Assuming a fixed mapping, the localization parameters are the radii values for each group $\mathbf{v} \in \mathbb{R}^g$. These values can be tuned independently during the adaptive algorithm. The corresponding prolongation operator $\mathbf{p} : \mathbb{R}^g \rightarrow \mathbb{R}^n$ assigns one of the g radii values to each of the n state-space components:

$$\mathbf{r} = \mathbf{p}(\mathbf{v}). \quad (9)$$

Setting $g = 1$ and $\mathbf{p}(\mathbf{v}) = v\mathbf{1}_n$ recovers the univariate localization approach.

Petrie (2008) showed that true square-root filters such as the LETKF (Hunt et al., 2007) are not amenable to Schur-product-based localization, and therefore they need to rely on sequentially assimilating every observation space variable. Here we wish to combine both the ease-of-use of Schur-product-based localization and the utility of multivariate localization techniques.

To this end, we introduce a commutative, idempotent, binary operation, $m : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, that computes a “mean localization function value” so as to harmonize the effects of different values of the localization radii. More explicitly, given $0 \leq a \leq b$, m should have the properties that $m(a, a) = a$, $m(b, b) = b$, and $m(a, b) = m(b, a)$. We also impose the additional common sense property that $a \leq m(a, b) \leq b$. We consider here several simple mean functions m as follows:

$$m_{\min}(\lambda_i, \lambda_j) = \min\{\lambda_i, \lambda_j\}, \quad (10a)$$

$$m_{\max}(\lambda_i, \lambda_j) = \max\{\lambda_i, \lambda_j\}, \quad (10b)$$

$$m_{\text{mean}}(\lambda_i, \lambda_j) = (\lambda_i + \lambda_j) / 2, \quad (10c)$$

$$m_{\text{sqr}}(\lambda_i, \lambda_j) = \sqrt{\lambda_i \lambda_j}, \quad (10d)$$

$$m_{\text{rms}}(\lambda_i, \lambda_j) = \sqrt{\lambda_i^2 + \lambda_j^2} / \sqrt{2}, \quad (10e)$$

$$m_{\text{harm}}(\lambda_i, \lambda_j) = \frac{2\lambda_i \lambda_j}{\lambda_i + \lambda_j}. \quad (10f)$$

Assume that the variables x_i and x_j have the localization radii r_i and r_j , respectively. We extend the definition of the localization matrix $\boldsymbol{\rho}$ to account for multiple localization radii associated with different variables as follows:

$$\boldsymbol{\rho} = \left[m\left(\ell(d(i, j)/r_i), \ell(d(j, i)/r_j)\right) \right]_{1 \leq i, j \leq n}. \quad (11)$$

The localization sub-matrix of state-space variables x_i and x_j

$$\begin{bmatrix} \ell(0) & \rho_{i,j} \\ \rho_{j,i} & \ell(0) \end{bmatrix} \quad (12)$$

is a symmetric matrix for any mean function (10a). When ℓ is a univariate compactly supported function, our approach implicitly defines its m -based multivariate counterpart.

One of the common criticisms of a distance assumption about the correlation of geophysical systems is that two variables in close proximity to each other might have very weak

correlations. For example, in a model that takes into account the temperature and concentration of stationary cars at any given location, the two distinct types of information might not at all be correlated with each other. The physical distance between the two, however, is 0, and thus any single correlation function will take the value 1 and does not remove any spurious correlations. One can mitigate this problem by considering univariate localization functions for each pair of components $\ell_{i,j} = \ell_{j,i}$ and extending the localization matrix as follows:

$$\boldsymbol{\rho} = \left[m\left(\ell_{i,j}(d(i, j)/r_i), \ell_{j,i}(d(j, i)/r_j)\right) \right]_{1 \leq i, j \leq n}. \quad (13)$$

We will not analyze multiple localization functions.

4 Bayesian approach

We denote the analysis step of the filter by

$$\mathbf{x}^a = \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}), \quad (14)$$

where $\mathbf{v}$ are the variable localization and inflation parameters. In this paper, we look solely at varying localization and will keep inflation constant.

In the Bayesian framework, we consider the localization parameters to be random variables with an associated prior probability distribution. Specifically, we assume that each of the radii $v^{(j)}$ is distributed according to a univariate gamma distribution $v^{(j)} \sim \Gamma(\alpha^{(j)}, \beta^{(j)})$. The gamma probability density

$$f_{\Gamma}(v; \alpha, \beta) = \frac{\beta^{\alpha} v^{\alpha-1} e^{-\beta v}}{\Gamma(\alpha)}, \quad v, \alpha, \beta > 0 \quad (15)$$

has mean $\bar{v} = \alpha/\beta$ and variance $\text{Var}(v) = \alpha/\beta^2$. We have chosen the gamma distribution as it is the maximum entropy probability distribution for a fixed mean (Singh et al., 1986), i.e., is the best distribution that one can choose without additional information. It is supported on the interval $(0, \infty)$ and has exactly two free parameters (allowing us to control both the mean and variance).

The assimilation algorithm computes a posterior (analysis) probability density over the state space considering the probabilities of observations and parameters. We start with Bayes' identity:

$$\begin{aligned} \pi(\mathbf{x}, \mathbf{v} | \mathbf{y}) &\propto \pi(\mathbf{y} | \mathbf{x}, \mathbf{v}) \pi(\mathbf{x} | \mathbf{v}) \pi(\mathbf{v}) \\ &= \pi(\mathbf{y} | \mathbf{x}, \mathbf{v}) \pi(\mathbf{x} | \mathbf{v}, \mathbf{y}) \pi(\mathbf{v}) \pi(\mathbf{y}). \end{aligned} \quad (16)$$

Note that $\pi(\mathbf{y})$ is a constant scaling factor. Here $\pi(\mathbf{v})$ represents the (prior) uncertainty in the parameters, $\pi(\mathbf{x} | \mathbf{y}, \mathbf{v})$ represents the uncertainty in the state for a specific value of the parameters, and $\pi(\mathbf{y} | \mathbf{x}, \mathbf{v})$ represents the likelihood of observations with respect to both state and parameters. Under

Gaussian assumptions about the state and observation errors, Eq. (16) can be explicitly written out as

$$\pi(\mathbf{x}, \mathbf{v} | \mathbf{y}) \propto \exp \left(-\frac{1}{2} \left\| \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) - \mathbf{x}^f \right\|_{\mathbf{P}_{(v)}^f}^2 \right) \exp \left(-\frac{1}{2} \left\| \mathbf{y} - \mathbf{H} \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) \right\|_{\mathbf{R}^{-1}}^2 \right) f_{\mathcal{P}}(\mathbf{v}), \quad (17)$$

with $f_{\mathcal{P}}(\mathbf{v})$ given by Eq. (15). Note that, once the analysis scheme (14) has been chosen, the analysis state becomes a function of $\mathbf{v}$, and the conditional probability in Eq. (17) represents the posterior probability of $\mathbf{v}$.

The negative log likelihood of the posterior probability (Eq. 17) is

$$\begin{aligned} \mathcal{J}(\mathbf{v}) &= -\log \left[\exp \left(-\frac{1}{2} \left\| \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) - \mathbf{x}^f \right\|_{\mathbf{P}_{(v)}^f}^2 \right) \right] \\ &= -\log \left[\exp \left(-\frac{1}{2} \left\| \mathbf{y} - \mathbf{H} \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) \right\|_{\mathbf{R}^{-1}}^2 \right) \right] \\ &\quad \prod_{j=1}^g v^{(j), \alpha-1} \exp \left(-\beta^{(j)} v^{(j)} \right) \\ &= \frac{1}{2} \left\| \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) - \mathbf{x}^f \right\|_{\mathbf{P}_{(v)}^f}^2 \\ &= +\frac{1}{2} \left\| \mathbf{y} - \mathbf{H} \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) \right\|_{\mathbf{R}^{-1}}^2 \\ &= +\sum_{j=1}^g \left(\beta^{(j)} v^{(j)} - (\alpha^{(j)} - 1) \log(v^{(j)}) \right). \end{aligned} \quad (18)$$

Through liberal use of the chain rule and properties of symmetric semi-positive definite matrices, the gradient of the cost function with respect to individual localization parameters reads as

$$\begin{aligned} \frac{\partial \mathcal{J}}{\partial v^{(j)}} &= \frac{1}{2} \left\| \mathbf{P}_{(v)}^f \right\|^{-1} \left(\mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) - \mathbf{x}^f \right) \left\| \frac{\partial \mathbf{P}_{(v)}^f}{\partial v^{(j)}} \right\| \\ &= +\frac{\partial \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v})}{\partial v^{(j)}} \left\| \mathbf{P}_{(v)}^f \right\|^{-1} \left(\mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) - \mathbf{x}^f \right) \\ &= -\left(\mathbf{H} \frac{\partial \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v})}{\partial v^{(j)}} \right)^{\top} \mathbf{R}^{-1} \left(\mathbf{y} - \mathbf{H} \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) \right) \\ &= +\beta^{(j)} - (\alpha^{(j)} - 1) \frac{1}{v^{(j)}}, \end{aligned} \quad (19)$$

where we took advantage of the properties of symmetric matrices. Note that without the assumption that parameters are independent, the gradient would involve higher-order tensors.

4.1 Solving the optimization problem

Under the assumptions that the analysis function (14) is based on DENKF (Eq. 5a) and (the not-quite-correct assumption)

that all ensemble members are i.i.d., we obtain

$$\begin{aligned} \left\| \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) - \mathbf{x}^f \right\|_{\mathbf{P}_{(v)}^f}^2 &= \sum_{e=1}^N \left\| \mathbf{K}_{(v)} \mathbf{z}^{(:,e)} \right\|_{\mathbf{P}_{(v)}^{f-1}}^2 \\ &= \sum_{e=1}^N \left\| \mathbf{S}_{(v)}^{-1} \mathbf{z}^{(:,e)} \right\|_{\mathbf{H} \mathbf{P}_{(v)}^f \mathbf{H}^{\top}}^2, \\ \mathbf{z} &= \mathbf{d} \mathbf{1}_N^{\top} - \frac{1}{2} \mathbf{H} \mathbf{X}^f, \end{aligned} \quad (20)$$

$$\begin{aligned} \left\| \mathbf{y} - \mathbf{H} \mathcal{A}(\mathbf{x}^f, \mathbf{y}, \mathbf{v}) \right\|_{\mathbf{R}^{-1}}^2 &= \sum_{e=1}^N \left\| \mathbf{g}_{(v)}^{(:,e)} \right\|_{\mathbf{R}^{-1}}^2, \\ \mathbf{g}_{(v)} &= (\mathbf{I} - \mathbf{H} \mathbf{K}_{(v)}) \mathbf{d} \mathbf{1}_N^{\top} \\ &= -\mathbf{H} \mathbf{X}^f \\ &= +\frac{1}{2} \mathbf{H} \mathbf{K}_{(v)} \mathbf{H} \mathbf{X}^f, \\ \mathbf{K}_{(v)} &= \mathbf{P}_{(v)}^f \mathbf{H}^{\top} \mathbf{S}_{(v)}^{-1}, \\ \mathbf{S}_{(v)} &= \mathbf{H} \mathbf{P}_{(v)}^f \mathbf{H}^{\top} + \mathbf{R}. \end{aligned} \quad (21)$$

Combining Eqs. (18), (20), and (21) leads to the cost function form:

$$\begin{aligned} \mathcal{J}(\mathbf{v}) &= \sum_{e=1}^N \left[\frac{1}{2} \left\| \mathbf{S}_{(v)}^{-1} \mathbf{z}^{(:,e)} \right\|_{\mathbf{H} \mathbf{P}_{(v)}^f \mathbf{H}^{\top}}^2 + \frac{1}{2} \left\| \mathbf{g}_{(v)}^{(:,e)} \right\|_{\mathbf{R}^{-1}}^2 \right] \\ &= +\sum_{j=1}^g \left(\beta^{(j)} v^{(j)} - (\alpha^{(j)} - 1) \log(v^{(j)}) \right), \end{aligned} \quad (22)$$

which only requires the computation of the background covariance matrix $\mathbf{H} \mathbf{P}_{(v)}^f \mathbf{H}^{\top}$ in observation space, thereby reducing the amount of computation and storage. The equivalent manipulations of the gradient lead to

$$\begin{aligned} \frac{\partial \mathcal{J}}{\partial v^{(j)}} &= \sum_{e=1}^N \left[\frac{1}{2} \left\| \mathbf{S}_{(v)}^{-1} \mathbf{z}^{(:,e)} \right\|_{\mathbf{H} \frac{\partial \mathbf{P}_{(v)}^f}{\partial v^{(j)}} \mathbf{H}^{\top}}^2 \right. \\ &= \sum \left[-\mathbf{z}^{(:,e), \top} \mathbf{H} \frac{\partial \mathbf{P}_{(v)}^f}{\partial v^{(j)}} \mathbf{H}^{\top} \mathbf{S}_{(v)}^{-2} \mathbf{H} \mathbf{P}_{(v)}^f \mathbf{H}^{\top} \mathbf{S}_{(v)}^{-1} \mathbf{z}^{(:,e)} \right. \\ &= \sum \left[-\mathbf{H} \frac{\partial \mathbf{K}_{(v)}}{\partial v^{(j)}} \mathbf{z}^{(:,e)} \right] \\ &= +\beta^{(j)} - (\alpha^{(j)} - 1) \frac{1}{v^{(j)}}, \end{aligned} \quad (23)$$

$$\begin{aligned} \frac{\partial \mathbf{K}_{(v)}}{\partial v^{(j)}} &= \frac{\partial \mathbf{P}_{(v)}^f}{\partial v^{(j)}} \mathbf{H}^{\top} \mathbf{S}_{(v)}^{-1} \\ &= -\mathbf{P}_{(v)}^f \mathbf{H}^{\top} \mathbf{H} \frac{\partial \mathbf{P}_{(v)}^f}{\partial v^{(j)}} \mathbf{H}^{\top} \mathbf{S}_{(v)}^{-2}, \end{aligned} \quad (24)$$

where $\frac{\partial \mathbf{P}_{(v)}^f}{\partial v^{(j)}} = \frac{\partial \rho_{(v)}}{\partial v^{(j)}} \circ \mathbf{P}^f$. Calculation of the gradient (23) requires computations only in observation space.

One will note that the form of our cost function (22) is similar to that of other hybrid approaches (Bannister, 2017) to data assimilation such as 3DENVar (Hamill and Snyder,

2000). As a lot of similar computation is done, this technique could potentially be used as a preprocessing step in a larger hybrid data assimilation scheme.

The choice of using the DEnKF is a bit arbitrary. From the above, however, it is evident that a method that decouples the anomaly updates from the mean updates would most likely be more advantageous. A perturbed observation EnKF does not have this property and thus would incur significantly more computational effort in optimizing the cost function. Extending this idea to a square-root filter, like the ETKF, would require significant algebraic manipulation and heuristics which are outside the scope of this paper.

4.2 Adaptive localization via a time-distributed cost function

We now seek to extend the 3D-Var-like cost function (18) to a time-dependent 4D-Var-like version. As the ensemble is essentially a reduced-order model, we do not expect the accuracy of the forecast to hold over the long term as we might in traditional 4D-Var. We thereby propose a limited extension of the cost function to include a small number K of additional future times. Assuming that we are currently at the i th time moment, the extended cost function reads as

$$\begin{aligned} \mathcal{J}_i(\mathbf{v}) &= \sum_{e=1}^N \left[\frac{1}{2} \left\| \mathbf{S}_{(\mathbf{v}),i}^{-1} \mathbf{z}_i^{(:,e)} \right\|_{\mathbf{H}\mathbf{P}_{(\mathbf{v}),i}^f \mathbf{H}^\top}^2 \right. \\ &= \sum \left[+ \frac{1}{2} \left\| \mathbf{g}_{(\mathbf{v}),i}^{(:,e)} \right\|_{\mathbf{R}_i^{-1}}^2 \right. \\ &= \sum \left[+ \sum_{k=1}^K \frac{1}{2} \left\| \mathbf{y}_{i+k} - \mathbf{H}\mathbf{x}_{(\mathbf{v}),i+k}^{f,(:,e)} \right\|_{\mathbf{R}_{i+k}^{-1}}^2 \right] \\ &= + \sum_{j=1}^g \left(\beta_i^{(j)} v^{(j)} - \left(\alpha_i^{(j)} - 1 \right) \log(v^{(j)}) \right). \end{aligned} \quad (25)$$

One will notice that in the 4-D part of the cost function the future forecasts are now dependent on $\mathbf{v}$ as they are obtained by running the model from $\mathbf{x}_{i,(v)}^a$. Similar to some 4-D ensemble approaches, the gradient computations can be approximated by the tangent linear model with the adjoint not being required. In fact all that is required is matrix vector products which can be approximated with finite differences (Tranquilli et al., 2017).

Various 4-D-type approximation strategies are also applicable to this cost function extension, though they are outside of the scope of this paper.

4.3 Algorithm

In practice, instead of dealing with the Gamma distribution parameters of α and β , we use the parameters $\bar{v}$ and $\text{Var}(v)$, such that $\alpha = \frac{\bar{v}^2}{\text{Var}(v)}$ and $\beta = \frac{\bar{v}}{\text{Var}(v)}$. For the sake of simplicity we assume that all $\mathbf{v}$ are identically distributed, but this is not required for the algorithm to function. The initial guess

for our minimization procedure is the vector of means. After minimizing the cost function, the radii for different components will be different. These radii, along with the corresponding localization and m functions, are used to build the localization matrix ρ . An outline is presented in Algorithm 1.

Algorithm 1 DEnKF Adaptive Localization Algorithm

Require: $[\bar{v}^{(j)}]_j, [\text{Var}(v^{(j)})]_j$

- 1: $\alpha^{(j)} \leftarrow \frac{\bar{v}^{(j),2}}{\text{Var}(v^{(j)})}$
- 2: $\beta^{(j)} \leftarrow \frac{\bar{v}^{(j)}}{\text{Var}(v^{(j)})}$
- 3: $\mathbf{v}_0^* \leftarrow [\bar{v}^{(j)}]_j$
- 4: $\mathbf{v}^* \leftarrow \text{argmin}_{\mathbf{v}} \mathcal{J}(\mathbf{v})$ {Eq. (22)}
- 5: $\mathbf{r}^* \leftarrow \mathbf{p}(\mathbf{v}^*)$ {Eq. (9)}
- 6: $\rho \leftarrow [m(l_i, l_j)]_{1 \leq i, j \leq n}$ {Eq. (13)}
- 7: **with**
- 8: $l_i \leftarrow \ell(d(i, j)/r^{(i),*})$ {Possibly Eq. (8)}
- 9: $l_j \leftarrow \ell(d(i, j)/r^{(j),*})$
- 10: **return** ρ

5 Numerical experiments and results

In order to validate our methodology, we carry out twin experiments under the assumption of identical perfect dynamical systems for both the truth and the model. The analysis accuracy is measured by the spatio-temporally averaged root mean square error:

$$\text{RMSE} = \sqrt{\frac{1}{n \cdot n_t} \sum_{i=1}^{n_t} \sum_{j=1}^n ([x_i^t]_j - [\bar{x}_i^a]_j)^2}, \quad (26)$$

where n_t is the number of data assimilation cycles (the number of analysis steps).

For each of the models we repeat the experiments with different values of the inflation constant α and report the best (minimal RMSE) results.

All initial value problems used were independently implemented (Roberts et al., 2019).

5.1 Oracles

We will make use of oracles to empirically evaluate the performance of the multivariate approach to Schur-product localization. An oracle is an idealized procedure that produces close to optimal results by making use of all the available information, some of which is unknown to the data assimilation system. In our case the oracle minimizes cost functions involving the true model state. Specifically, in an ideal filtering scenario one seeks to minimize the error of the analysis with respect to the truth, i.e., the cost function,

$$\mathcal{J}(\mathbf{x}^a) = \text{RMSE}(\mathbf{x}^a - \mathbf{x}^t). \quad (27)$$

Our oracle will pick the best parameters, in this case the radii, that minimize an ideal cost function $\mathcal{J}(\mathbf{v}) = \text{RMSE}(\bar{\mathbf{x}}_{(v)}^a -$

$\mathbf{x}^t$). This can be viewed as an ideal variance minimization of the state space in parameter space.

5.2 Lorenz'96

The 40-variable Lorenz model (Lorenz, 1996; Lorenz and Emanuel, 1998) is the first standard test employed. This problem is widely used in the testing of data assimilation algorithms.

5.2.1 Model setup

The Lorenz'96 model equations,

$$\frac{dx_i}{dt} = (x_{i+1} - x_{i-2})x_{i-1} - x_i + F, \quad i = 1, \dots, n, \quad (28)$$

are obtained through a coarse discretization of a forced Burger equation with Newtonian damping (Reich and Cotter, 2015). We impose $x_0 = x_n$, $x_{-1} = x_{n-1}$, and $x_{n+1} = x_1$, where $n = 40$. We take the canonical value for the external forcing factor, $F = 8$. Using known techniques for dynamical systems (Parker and Chua, 2012), one can approximately compute that this particular system has 13 positive Lyapunov exponents (Strogatz, 2014) and 1 zero exponent, with a Kaplan–Yorke dimension of approximately 27.1.

The initial conditions used for the experiments are obtained by starting with

$$[x_0]_i = \begin{cases} 8 & i \neq 20 \\ 8.008 & i = 20 \end{cases} \quad (29)$$

and integrating Eq. (28) forward for 1 time unit in order to reach the attractor.

The physical distance between x_i and x_j is the shortest cyclic distance between any two state variables:

$$d(i, j) = \min\{\|i - j\|, \|n + i - j\|, \|n + j - i\|\}, \quad (30)$$

where the distance between two neighbors is 1.

For the numerical experiments, we consider a perfect model and noisy observations. We take a 6 h assimilation window (corresponding to $\Delta t_{\text{obs}} = 0.05$ model time units) and calculate the RMSE on the assimilation cycle interval [100, 1100] in the case of oracle testing or [5000, 55000] in the case of adaptive localization testing. Lorenz'96 is an ergodic system (Fatkullin and Vanden-Eijnden, 2004), and therefore its behavior in time for most initial conditions is the same as the behavior of all its possible phase spaces on any orbit around a strange attractor at any point in time, meaning that a long enough time averaged run should be the same as a collection of shorter space-averaged runs. We use 10 ensemble members (just under the number of positive Lyapunov exponents) and observe the 30 model variables (2, 4, ..., 18, 20, 21, ..., 39, 40), with an observation error variance of 1 for each observed state.




Figure 1. Generalized mean oracles for Lorenz'96. Comparison of the various m -based multivariate localization techniques with that of standard univariate localization and constant “best radius” localization. Results were obtained using the Lorenz'96 model over the assimilation cycles [100, 1100] and were run for various values of the inflation parameter.

5.2.2 Oracle results

Figure 1 shows a visualization of multivariate oracle runs for the Lorenz'96 test problem using the best constant radius, a univariate oracle, and a multivariate oracle utilizing the m functions from Eq. (10a). The constant “best” radius was chosen on an individual basis for every value of the inflation parameter, while the univariate and multivariate cases were allowed to minimize Eq. (27) for a single global radius and for 40 local radii, respectively. This means that in the multivariate case each of the 40 variables is given an independent radius. As one can see, the univariate oracle performs no better than a constant radius, while several of the multivariate approaches provide much better results. The schemes (Eq. 10a) that closely mirror an unbiased mean, namely m_{mean} , m_{sqrt} , m_{rms} , and m_{harm} , yield the best results, while the conservative scheme m_{min} performs no better than a univariate approach. This indicates that the problem is better suited for multivariate localization. Most m functions perform similarly, and in further experiments we will only consider m_{mean} as a representative and easy-to-implement option. We note that for Lorenz'96 with DEnKF and our experimental setup, no 3-D univariate adaptive localization scheme which aims for the analysis estimate to optimally approach the truth outperforms the best constant localization radius.

We also test both the validity of arbitrarily grouping the radii and the validity of using a time-distributed cost function (Eq. 25). Figure 2 presents results for arbitrary radii




Figure 2. Time-dependent 4-D oracle for Lorenz'96. Comparison of the RMSE for a radius oracle that is both multivariate and time-dependent. The y axis represents the number of independent radii values (groups of the model state components). The assimilation cycle interval is [100, 1100]; a fixed inflation value of $\alpha = 1.02$ and the function m_{mean} are used.

groupings and for a limited run 4-D approach. The results were calculated with arbitrary choices of radii groupings such that each group would contain an equal number of variables, and with the fixed inflation value of $\alpha = 1.02$ with the function m_{mean} . There is significant benefit in using more radii groupings, but marginal benefit from the 4-D approach.

5.2.3 Adaptive localization results

Adaptive localization results for Lorenz'96 are shown in Fig. 3. The optimal constant localization radius was found and a search of possible input mean and variance values was performed around it for the adaptive case. The value of $\bar{v}$ was varied by additive factors of one of $-1, -0.5, +0, +0.5, \text{ or } +1$ with respect to the best univariate radius with $\text{Var}(v)$ chosen to be one of $1/8, 1/4, 1/2, 1, \text{ or } 2$. As accurately predicted by the oracle results, a univariate adaptive approach for this model cannot do better than the best univariate radius (Fig. 1), as no meaningful reduction in error was detected.

5.3 Multivariate Lorenz'96

The canonical Lorenz'96 model is ill suited for multivariate adaptive localization as each variable in the problem behaves identically to all the others. This means that for any univariate localization scheme a constant radius is close to optimal.

5.3.1 Model setup

We modify the problem in such a way that the average behavior remains very similar to that of the original model, but that




Figure 3. Lorenz'96 adaptive localization results. Comparison of the best univariate localization radius results with their corresponding adaptive localization counterparts. $\bar{v}$ was varied by additive factors of one of $-1, -0.5, +0, +0.5, \text{ or } +1$ with respect to the best univariate radius, with $\text{Var}(v)$ chosen to be one of $1/8, 1/4, 1/2, 1, \text{ or } 2$.

instantaneous behavior requires different localization radii. In order to accomplish this we use the time-dependent forcing function that is different for each variable:

$$[F(t)]_i = 8 + 4 \cos \left(\omega \left(t + \frac{(i-1) \bmod q}{q} \right) \right), \quad (31)$$

where $\omega = 2\pi$ (in the context of Lorenz'96 the equivalent period is 5 d), i is the variable index, and q is an integer factor of n . Here we set $q = 4$. The different forcing values will create differences between the dynamics of different variables.

For each individual variable the forcing value cycles between 4 and 12, with an average value of 8, just like in the canonical Lorenz'96 formulation. If taken to be a constant, the forcing factor of 4 will make the equation lightly chaotic with only one positive Lyapunov exponent, whilst a constant value of 12 will make the dynamics have about 15 positive Lyapunov exponents. Our modified system still has the same average behavior with 13 positive Lyapunov exponents. The mean doubling times of the two problems are also extremely similar at around 0.36. This is the ideal desired behavior. Figure 4 shows a comparison of numerically computed covariance matrices for this modified problem. There is considerable difference between the size of non-diagonal entries for different state variables over the course of a single step, but this difference disappears after averaging. This indicates that for this problem the best constant univariate radius is the same as for the canonical Lorenz'96 model but that instantaneous adaptive radii are different. This shows that an




Figure 4. Calculated optimal single-time and time-averaged covariance matrices for the multivariate Lorenz'96 model. Comparison of ensemble covariance matrices for the multivariate Lorenz'96 equations for a single time step (a) with that for a time-averaged run (b) for 100 000 and 10 000 ensemble members, respectively.

adaptive approach to the instantaneous covariance could be advantageous.

Error measurements will be carried out over the interval [500, 5500], with the omitted steps acting as spinup. The rest of the problem setup is identical to that of the canonical Lorenz'96.

5.3.2 Adaptive localization results

Figure 5 shows results with the multivariate Lorenz'96 model, with $g = q = 4$ radii groups and the 4-D parameter set to $K = 1$. The inflation is kept constant for each experiment, with multiple experiments spanning $\alpha \in [1.02, 1.1]$. The localization radius is varied in increments of 0.5 over the range [0.5, 16] for the constant case. The corresponding minimal-error radii are used as mean inputs into our adaptive algorithm. For the adaptive case we choose four arbitrary groupings of radii ($g = 4$) with the mean function m_{mean} . The parameters $\bar{v}$ take one of three possible values, the optimal constant radius -1 , $+0$, and $+1$. The parameters $\text{Var}(v)$ take one value from $\{1/4, 1, 4\}$. The 4-D variable is set to $K = 1$ to look at an additional observation in the future.

As before, the results for the constant radius were their optimal value for each given inflation value, while the adaptive results were obtained through a search for possible means and variances around that value. The largest reduction in error is only about 8 %; however, this is a significant improvement over the behavior of the univariate Lorenz'96. In the canonical Lorenz'96 there is no meaningful choice of grouping other than arbitrary, but in this case, the groups were chosen such that all related variables have the same forcing from Eq. (31). The results show a significant improvement over the univariate case, especially for low inflation values. We note that the filter spinup takes significantly longer for the adaptive localization case than for the constant univariate case.




Figure 5. Time-dependent multivariate Lorenz'96 4-D adaptive localization. The constant radius case shows the minimal error when the localization radius is varied between set predefined values. The adaptive localization case has four radii groupings, $g = q = 4$, with the time-dependence parameter, $K = 1$. The parameters $\bar{v}$ were varied by -1 , $+0$, and $+1$ from the best constant radius. The parameters $\text{Var}(v)$ take one value from $\{1/4, 1, 4\}$.

Consequently, the assimilation cycles are chosen in the time interval [500, 5500] units. An idea to mitigate this might be to run the filter with a constant radius for a few assimilation cycles before switching to the adaptive localization strategy, so as to allow the filter to quickly catch up with the shadow attractor trajectory.




Figure 6. ML96 sample radii. Radii for the configuration in Fig. 5 with $\alpha = 1.02$, $\bar{v} = 6$, and $\text{Var}(v) = 1$. Each unique marker type represents a unique radius grouping with a red $\cdot$ representing the first group, a green $\cdot$ representing the second, a blue $\cdot$ representing the third, and a black $\cdot$ representing the fourth. Groups 2 and 4 contain variables that were observed fully, and groups 1 and 3 contain variables that were observed partially. The graph has been truncated to only show radii from 3 to 9 in order to better capture the general tendencies.

Figure 6 demonstrates a sampling of the radii obtained through the adaptive 4-D multivariate method. Taking $g = q = 4$, $\alpha = 1.02$, $\bar{v} = 6$, $\text{Var}(v) = 1$, and $K = 1$, we look at the radii generated by our method. For the post-spinup case, the means (5.8449, 5.8669, 5.8441, 5.8602) of the radii for each grouping are fairly similar; however, the variances (0.0300, 0.0731, 0.0295, 0.0821) differ wildly. The groupings with less observations, groups 1 and 3, have much more conservative variances, while groups 2 and 4, which are observed fully, have much greater variances.

This gives us insight into a potential way of choosing multivariate localization groups. Based on some measure of the observability of any given state-space variable, similarly “observable” state-space variables should have similar radii.

Tightly coupled models like the multivariate Lorenz’96 have rapidly diverging solutions, and constraining them requires more information about the underlying dynamics. Incorporating future observations and adding degrees of freedom to the cost function increase the performance of our analysis. In the limiting case of one radius per variable and general information from the future one approaches a variant of 4DenVar, which is in principle superior to any pure filtering method.

5.4 Quasi-geostrophic model (QGSO)

The 1.5-layer quasi-geostrophic model of Sakov and Oke (Sakov and Oke, 2008), obtained by non-dimensionalizing potential vorticity, is given by the equations

$$\begin{aligned} q_t &= -\psi_x - \varepsilon J(\psi, q) - A \Delta^3 \psi + 2\pi \sin(2\pi y), \\ \Delta \psi - F \psi &= q, \quad J(\psi, q) \equiv \psi_x q_y - \psi_y q_x. \end{aligned} \quad (32)$$

The variable ψ can be thought of as the stream function, and the spatial domain is the square $(x, y) \in [0, 1]^2$. The constants are $F = 1600$, $\varepsilon = 10^{-5}$, and $A = 2 \times 10^{-11}$. We use homogeneous Dirichlet boundary conditions.

A second-order central finite difference spatial discretization of the Laplacian operator Δ is performed over the interior of a 129×129 grid, leading to a model dimension $n = 127^2 = 16\,129$. The time discretization is the canonical fourth-order explicit Runge–Kutta method with a time step of 1 time unit. The Helmholtz equation on the right-hand side of Eq. (32) is solved by an offline pivoted sparse Cholesky decomposition. J is calculated via the canonical Arakawa approximation (Arakawa, 1966; Ferguson, 2008). The Δ^3 operator is implemented by repeated application of our discrete Laplacian.

The time between consecutive observations is 5 time units, and the model is run for 3300 such cycles. The first 300 cycles, corresponding to the filter spinup, are discarded, and therefore the assimilation interval is $[300, 3300]$ time units. Observations are performed with a standard 300-component observation operator as shown in Fig. 7. An observation error variance of 4 is taken for each component. The physical distance between two components is defined as

$$d(i, j) = \sqrt{(i_x - j_x)^2 + (i_y - j_y)^2}, \quad (33)$$

with (i_x, i_y) and (j_x, j_y) the spatial grid coordinates of the state-space variables i and j , respectively.

Our rough estimate of the number of positive Lyapunov exponents of this model is 1451, with a Kaplan–Yorke dimension estimate of 6573.4; thus, we will take a conservative 25 ensemble members whose initial states are derived from a random sampling of a long run of the model.

This model has been tested extensively with both the DEnKF and various localization techniques (Sakov and Bertino, 2011; Bergemann and Reich, 2010; Moosavi et al., 2018).

Adaptive localization results

The adaptive localization results for the quasi-geostrophic problem are shown in Fig. 8. As before, a constant best univariate localization radius was calculated (iteratively taken in increments of 5 over the range $[5, 45]$) for each inflation value ($\alpha \in [1.02, 1.16]$), and was used as a seed for varying the results in the adaptive case. For the adaptive case we vary




Figure 7. Quasi-geostrophic model. A typical model state of the 1.5-layer quasi-geostrophic model. **(a)** shows the stream function values, with the red dots representing the locations of the variables that are observed. **(b)** shows the corresponding vorticity.

the parameters $\bar{v}$ by taking one of three possible values, differing from the optimal constant radius by -5 , 0 , and $+5$. The adaptive $\text{Var}(v)$ takes one of the values $1/2$, 1 , 2 , or 4 . An improvement of RMSE of up to about 33 % can be seen for certain inflation values, meaning that the adaptive localization procedure results in significant reductions in analysis error, while for other values no significant benefits are observed.

A better representation of how well the adaptive localization scheme works is by showing its consistency. The empirical utility of the adaptive localization technique is further analyzed in Fig. 9, which compares the error results of a sub-optimal constant radius with that of an adaptive run with the mean parameter set to the same values as the constant ones. The adaptive results are – except in a few cases of filter divergence – always as good as or better than their constant localization counterparts, with a reduction in error as large as 33 % with the same mean radius as the constant radius, and as much as 50 % with different radii. Even in the case where the adaptive filter diverged, the constant localization filter diverged as well. This indicates that our localization scheme is significantly better than a corresponding sub-optimal constant scheme with the same parameters, as is typically the case in real-world production codes. This opens up the possibility of adapting existing systems that use a conservative suboptimal constant localization to an adaptive localization scheme.

Figure 10 shows a sample of the radii obtained by the adaptive algorithm. One will notice that during the spinup time the algorithm is much more conservative, with the radius of influence signaling that there should be an over-reliance on the observations instead of the model prior. After the spinup time however, the algorithm tends to select radii greater than the mean, signifying a greater confidence in the observations.




Figure 8. Quasi-geostrophic model adaptive localization. The inflation factor is kept constant, and α values from 1.02 to 1.16 are represented on the x axis. The constant localization radius varies in increments of 5 over the range $[5, 45]$. Only the best results are plotted, and they are used as the mean seeds for the adaptive algorithm. For the adaptive case we vary the parameters $\bar{v}$ by taking one of three possible values, differing from the optimal constant radius by -5 , 0 , and $+5$. The adaptive $\text{Var}(v)$ takes one of the values $1/2$, 1 , 2 , and 4 .

6 Conclusions

This paper proposes a novel Bayesian approach to adaptive Schur-product-based localization. A multivariate approach is developed, where multiple radii corresponding to different types of variables are taken into account. The Bayesian problem is solved by constructing 3-D and 4-D cost functions and




Figure 9. Quasi-geostrophic model adaptive localization raw RMSE. The green line represents the same optimal constant localization radius as in Fig. 8. The red line represents the error for the radius -5 units below, and the blue line, if it had not experienced filter divergence, would have represented $+5$. The correspondingly colored areas represent the ranges of error of the adaptive localization scheme obtained by fixing the mean, $\bar{v}$, to be that of the constant scheme and ranging over the variance.




Figure 10. QGSO sample radii. For the configuration in Fig. 8 with $\alpha = 1.08$, $\bar{v} = 25$, and $\text{Var}(v) = 4$. Each dot represents a radius, with the line representing the mean.

minimizing them to obtain the maximum a posteriori estimates of the localization radii. We show that in the case of the DEnKF these cost functions and their gradients are computationally inexpensive to evaluate and can be relatively easily implemented within existing frameworks. We provide a new approach for assessing the performance of adaptive localization approaches through the use of restricted cost function oracles.

The adaptive localization approach is tested using the Lorenz'96 and quasi-geostrophic models. Somewhat surprisingly, the adaptivity produces better results for the larger quasi-geostrophic problem. This may be due to the ensemble analysis anomaly independence assumption made in Sect. 4.1, an assumption that holds better for a large system with sparse observations than for a small tightly coupled system with dense observations. The performance of the adaptive approach on the small, coupled Lorenz'96 system is increased by using multivariate and 4-D extensions of the cost function. The approach yields modest gains of about an 8 % reduction in error, but more importantly shows empirical evidence for a quite intuitive observation that variables that are more strongly observed can accept a greater variance of localization radii and can be less conservative with their radii choices.

We believe that the algorithm presented herein has a strong potential to improve existing geophysical data assimilation systems that use ensemble-based filters such as the DEnKF. In order to avoid filter divergence in the long term, these systems often use a conservative localization radius and a liberal inflation factor. The QG model results indicate that, in such cases, our adaptive method outperforms the approach based on a constant localization. The approach leads to a reduction of as much as 33 % in error. For a severely undersampled ensemble, the approach appears to improve the quality of the analysis substantially, potentially because the need for localization is significantly greater than for a small toy problem like L96. The new adaptive methodology can replace the existing approach with a relatively modest implementation effort.

Future work to extend the methodology includes finding good approximations of the probability distribution of the localization parameters, perhaps through a machine learning approach, and reducing the need for assumption that the ensemble members are independent identically distributed random variables. A future direction of interest is applying this methodology to a larger operational model, e.g., the Weather Research and Forecasting Model (WRF) (Skamarock et al., 2008), which is computationally feasible in the short term.

Data availability. No data sets were used in this article.

Author contributions. All the authors contributed equally to this work.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors thank the Computational Science Laboratory at Virginia Tech. The authors also thank the anonymous referees and the editor for their constructive input that significantly helped with the quality of the paper.

Financial support. This research has been supported by the Air Force Office of Scientific Research (grant no. DDDAS FA9550-17-1-0015), the National Science Foundation, Division of Computing and Communication Foundations (grant no. CCF-1613905), and the National Science Foundation, Division of Advanced Cyberinfrastructure (grant no. ACI-17097276).

Review statement. This paper was edited by Zoltan Toth and reviewed by two anonymous referees.

References

- Anderson, J. L.: An ensemble adjustment Kalman filter for data assimilation, *Mon. Weather Rev.*, 129, 2884–2903, 2001.
- Anderson, J. L.: Exploring the need for localization in ensemble data assimilation using a hierarchical ensemble filter, *Physica D*, 230, 99–111, 2007.
- Anderson, J. L.: Localization and sampling error correction in ensemble Kalman filter data assimilation, *Mon. Weather Rev.*, 140, 2359–2371, 2012.
- Arakawa, A.: Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. Part I, *J. Comput. Phys.*, 1, 119–143, 1966.
- Asch, M., Bocquet, M., and Nodet, M.: Data assimilation: methods, algorithms, and applications, SIAM, Philadelphia, PA, USA, 2016.
- Askey, R.: Radial characteristics functions, Tech. rep., Wisconsin University Madison Mathematics Research Center, Madison, WI, USA, 1973.
- Attia, A. and Constantinescu, E.: An Optimal Experimental Design Framework for Adaptive Inflation and Covariance Localization for Ensemble Filters, arXiv preprint arXiv:1806.10655, 2018.
- Attia, A. and Sandu, A.: A hybrid Monte Carlo sampling filter for non-gaussian data assimilation, *AIMS Geosciences*, 1, 41–78, 2015.
- Attia, A., Ștefănescu, R., and Sandu, A.: The reduced-order hybrid Monte Carlo sampling smoother, *Int. J. Numer. Meth. Fl.*, 83, 28–51, 2017.
- Bannister, R.: A review of operational methods of variational and ensemble-variational data assimilation, *Q. J. Roy. Meteor. Soc.*, 143, 607–633, 2017.
- Bergemann, K. and Reich, S.: A localization technique for ensemble Kalman filters, *Q. J. Roy. Meteor. Soc.*, 136, 701–707, 2010.
- Bickel, P. J. and Levina, E.: Regularized estimation of large covariance matrices, *Ann. Stat.*, 36, 199–227, 2008.
- Bishop, C. H. and Hodyss, D.: Flow-adaptive moderation of spurious ensemble correlations and its use in ensemble-based data assimilation, *Q. J. Roy. Meteor. Soc.*, 133, 2029–2044, 2007.
- Bishop, C. H. and Hodyss, D.: Adaptive ensemble covariance localization in ensemble 4D-VAR state estimation, *Mon. Weather Rev.*, 139, 1241–1255, 2011.
- Buehner, M. and Shlyueva, A.: Scale-dependent background-error covariance localisation, *Tellus A*, 67, 28027, <https://doi.org/10.3402/tellusa.v67.28027>, 2015.
- Carmichael, G., Chai, T., Sandu, A., Constantinescu, E., and Daescu, D.: Predicting air quality: improvements through advanced methods to integrate models and measurements, *J. Comput. Phys.*, 227, 3540–3571, <https://doi.org/10.1016/j.jcp.2007.02.024>, 2008.
- Constantinescu, E., Sandu, A., Chai, T., and Carmichael, G.: Ensemble-based chemical data assimilation. II: Covariance localization, *Q. J. Roy. Meteor. Soc.*, 133, 1245–1256, <https://doi.org/10.1002/qj.77>, 2007a.
- Constantinescu, E., Sandu, A., Chai, T., and Carmichael, G.: Ensemble-based chemical data assimilation. I: General approach, *Q. J. Roy. Meteor. Soc.*, 133, 1229–1243, <https://doi.org/10.1002/qj.76>, 2007b.
- Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *J. Geophys. Res.-Oceans*, 99, 10143–10162, 1994.
- Evensen, G.: Data assimilation: the ensemble Kalman filter, Springer Science & Business Media, 2009.
- Fatkullin, I. and Vanden-Eijnden, E.: A computational strategy for multiscale systems with applications to Lorenz 96 model, *J. Comput. Phys.*, 200, 605–638, 2004.
- Ferguson, J.: A numerical solution for the barotropic vorticity equation forced by an equatorially trapped wave, Master's thesis, University of Victoria, 2008.
- Flowerdew, J.: Towards a theory of optimal localisation, *Tellus A*, 67, 25257, <https://doi.org/10.3402/tellusa.v67.25257>, 2015.
- Gaspari, G. and Cohn, S. E.: Construction of correlation functions in two and three dimensions, *Q. J. Roy. Meteor. Soc.*, 125, 723–757, 1999.
- Gneiting, T.: Compactly supported correlation functions, *J. Multivariate Anal.*, 83, 493–508, 2002.
- Hamill, T. M. and Snyder, C.: A hybrid ensemble Kalman filter–3D variational analysis scheme, *Mon. Weather Rev.*, 128, 2905–2919, 2000.
- Houtekamer, P. L. and Mitchell, H. L.: A sequential ensemble Kalman filter for atmospheric data assimilation, *Mon. Weather Rev.*, 129, 123–137, 2001.
- Hunt, B. R., Kostelich, E. J., and Szunyogh, I.: Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter, *Physica D*, 230, 112–126, 2007.
- Kalnay, E.: Atmospheric modeling, data assimilation and predictability, Cambridge University Press, Cambridge, UK, 2003.
- Kirchgessner, P., Nerger, L., and Bunse-Gerstner, A.: On the choice of an optimal localization radius in ensemble Kalman filter methods, *Mon. Weather Rev.*, 142, 2165–2175, 2014.
- Law, K., Stuart, A., and Zygalakis, K.: Data assimilation: a mathematical introduction, vol. 62, Springer, Cham, Switzerland, 2015.
- Le Dimet, F.-X. and Talagrand, O.: Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects, *Tellus A*, 38, 97–110, 1986.

- Lorenz, E. N.: Predictability: A problem partly solved, in: Proc. Seminar on predictability, vol. 1, 1996.
- Lorenz, E. N. and Emanuel, K. A.: Optimal sites for supplementary weather observations: Simulation with a small model, *J. Atmos. Sci.*, 55, 399–414, 1998.
- Ménétrier, B., Montmerle, T., Michel, Y., and Berre, L.: Linear filtering of sample covariances for ensemble-based data assimilation. Part I: Optimality criteria and application to variance filtering and covariance localization, *Mon. Weather Rev.*, 143, 1622–1643, 2015.
- Moosavi, A., Attia, A., and Sandu, A.: A Machine Learning Approach to Adaptive Covariance Localization, *ArXiv e-prints*, 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*, p. 4, ACM, Austin, Texas, 2015.
- Parker, T. S. and Chua, L.: *Practical numerical algorithms for chaotic systems*, Springer Science & Business Media, New York, NY, USA, 2012.
- Petrie, R.: Localization in the ensemble Kalman filter, *MSc Atmosphere, Ocean and Climate University of Reading*, Reading, USA, 2008.
- Petrie, R. E. and Dance, S. L.: Ensemble-based data assimilation and the localisation problem, *Weather*, 65, 65–69, 2010.
- Porcu, E., Daley, D. J., Buhmann, M., and Bevilacqua, M.: Radial basis functions with compact support for multivariate geostatistics, *Stoch. Env. Res. Risk A.*, 27, 909–922, 2013.
- Reich, S. and Cotter, C.: *Probabilistic forecasting and Bayesian data assimilation*, Cambridge University Press, Cambridge, UK, 2015.
- Roberts, S., Popov, A. A., and Sandu, A.: ODE Test Problems: a MATLAB suite of initial value problems, *arXiv preprint arXiv:1901.04098*, 2019.
- Roh, S., Jun, M., Szunyogh, I., and Genton, M. G.: Multivariate localization methods for ensemble Kalman filtering, *Nonlin. Processes Geophys.*, 22, 723–735, <https://doi.org/10.5194/npg-22-723-2015>, 2015.
- Sakov, P. and Bertino, L.: Relation between two common localisation methods for the EnKF, *Comput. Geosci.*, 15, 225–237, 2011.
- Sakov, P. and Oke, P. R.: A deterministic formulation of the ensemble Kalman filter: an alternative to ensemble square root filters, *Tellus A*, 60, 361–371, 2008.
- Sandu, A., Daescu, D., Carmichael, G., and Chai, T.: Adjoint sensitivity analysis of regional air quality models, *J. Comput. Phys.*, 204, 222–252, <https://doi.org/10.1016/j.jcp.2004.10.011>, 2005.
- Schur, J.: Bemerkungen zur Theorie der beschränkten Bilinearformen mit unendlich vielen Veränderlichen., *Journal für die reine und angewandte Mathematik*, 140, 1–28, 1911.
- Singh, V. P., Rajagopal, A., and Singh, K.: Derivation of some frequency distributions using the principle of maximum entropy (POME), *Adv. Water Resour.*, 9, 91–106, 1986.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Barker, D. M., Duda, M. G., Huang, X., Wang, W., and Powers, J. G.: A description of the advanced research WRF version 3, *Tech. rep.*, National Center For Atmospheric Research, 2008.
- Strogatz, S. H.: *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*, Westview Press, Boulder, CO, 2014.
- Tranquilli, P., Glandon, S. R., Sarshar, A., and Sandu, A.: Analytical Jacobian-vector products for the matrix-free time integration of partial differential equations, *J. Comput. Appl. Math.*, 310, 213–223, 2017.
- Van Leeuwen, P. J., Cheng, Y., and Reich, S.: *Nonlinear data assimilation*, vol. 2, Springer, Cham, Switzerland, 2015.
- Zupanski, M.: Maximum likelihood ensemble filter: Theoretical aspects, *Mon. Weather Rev.*, 133, 1710–1726, 2005.