Key Points:
- A way to implement model-space localization in a serial EnKF is presented
- Model-space localization outperforms observation-space localization for radiance observations in the global atmospheric data assimilation
- Implementing model-space localization in an EnKF allows to reach errors similar to those of EnVar with pure ensemble covariances

Correspondence to:
A. Shlyaeva,
anna.v.shlyaeva@noaa.gov

Model-Space Localization in Serial Ensemble Filters

Anna Shlyaeva[1] [ID], **Jeffrey S. Whitaker[2]**, and **Chris Snyder[3]**

[1] Physical Sciences Division, Cooperative Institute for Research in Environmental Sciences at the NOAA Earth System Research Laboratory, University of Colorado Boulder, Boulder, CO, USA, [2] Physical Sciences Division, NOAA Earth System Research Laboratory, Boulder, CO, USA, [3] National Center for Atmospheric Research, Boulder, CO, USA[†]

Abstract Ensemble-based data assimilation systems typically use covariance localization to dampen spurious correlations associated with sampling error while increasing the rank of the covariance estimate. Variational methods use model-space localization, in which localization is applied to ensemble estimates of covariances between model variables and is based on distances between those variables, while ensemble filters apply observation-space localization to estimates of model-observation covariances, based on distances between model variables and observations. It has been shown that for nonlocal observations, such as satellite radiances, model-space localization can be superior. This paper demonstrates a new method for performing model-space localization in serial ensemble filters using the linearized observation operators (or Jacobians). Results of radiance-only assimilation in a global forecast system show the benefit of using model-space localization relative to observation-space localization. The serial ensemble square root filter with vertical model-space localization gives results similar to those of the Ensemble Variational system (without outer loops or extra balance constraints) while increasing the runtime compared to the filter with observation-space localization by a factor between 2 and 8, depending on how sparse the Jacobian matrices are. The results are also similar to another approach to model-space localization in ensemble filters: ensemble Kalman filter with modulated ensembles.

1. Introduction

Ensembles of short-range forecasts can bring valuable flow-dependent information about background error covariances to a data assimilation system. The drawback is that the ensemble sizes that are affordable are small compared to the state sizes, leading to sampling errors and severe rank deficiency of ensemble background error covariances. To mitigate this, ensemble covariance localization is universally used. Localization typically means dampening of ensemble error covariances between variables that are far from each other in some dimension (physical space, time, and type of variable).

Different ways of solving the ensemble data assimilation problem may use different types of localization. Ensemble Variational (EnVar) systems like Buehner (2005) and Lorenc (2003) localize ensemble background error covariances by multiplying them elementwise by a correlation matrix, thereby damping covariances between distant model state variables. The localization in EnVar, thus, depends only on the distances between model state variables and is independent of the observations being assimilated. We will call this type of localization *model-space localization* through the rest of the paper.

On the other hand, most Ensemble Kalman filters (EnKFs) localize ensemble error covariances based on distances between model variables and observations, by using one of the two common EnKF approaches to localization. Serial ensemble filters, like the Ensemble Square Root Filter (EnSRF) (Whitaker & Hamill, 2002), Ensemble Adjustment Kalman Filter (Anderson, 2001), and the perturbed observation EnKF (Houtekamer & Mitchell, 1998), assimilate observations sequentially and localize ensemble model-observation prior error covariances and ensemble observation-prior error covariances, based on the distance between model state variables and observations and between observations. The Local Ensemble Transform Kalman Filter (LETKF; Hunt et al., 2007) uses "observation error localization" (Greybush et al., 2011) which increases observation error variances based on the distance between the observation and the model state variable being updated. In both of the approaches to localization in the EnKFs, the localization depends on the distance between model state variables and observations. We will call it *observation-space localization* through the rest of the paper

Observation-space localization can be problematic when the observation does not have a clearly defined location (e.g., remotely sensed observations which observe some sort of volume integral of the state), since the concept of distance between the nonlocal observation and a given state variable is not well defined. For example, one of the main sources of data for the current global atmospheric data assimilation systems is satellite radiances for which vertical location is not well defined. EnKF systems usually use the peak of the weighting function in the radiative transfer equation as the observation "location," but it has been shown (Campbell et al., 2010) that this approximation may hurt EnKF performance. Lei and Whitaker (2015) showed that observation-space localization can actually perform better for satellite radiances in certain special circumstances. Despite this caveat and consistent with Campbell et al. (2010), we will show here that model-space localization slightly outperforms observation-space localization when satellite radiances are the primary source of observations.

It is possible to use model-space localization in the context of an EnKF. Bishop and Hodyss (2009) introduced the concept of a "modulated ensemble" to incorporate model-space localization in EnKF systems. Bishop et al. (2017) demonstrated its use in a simple model, and Lei et al. (2018) implemented it for the vertical localization in LETKF for global Numerical Weather Prediction (NWP). In this approach, the background ensemble is reconstructed so that it approximates the square root of the localized background-error covariance matrix. This is achieved by replacing the original ensemble with an expanded ensemble formed by the "modulation product" between each ensemble member and the scaled eigenvectors of the localization matrix. The EnKF can then be used to compute analysis increments without applying any additional localization (since it is already implied by the structure of the background ensemble perturbations). Another path to achieve model-space localization is direct solution of full matrix equations that include localization, as in EnVar. Steward et al. (2018) report progress in this direction for the square-root formulation of the EnKF.

In this paper we show that it is also possible to implement model-space localization in a serial EnKF, by localizing background error covariances directly, provided the linearized observation operator is available.

The general formulation of model-space localization EnKF systems is in section 2. In section 3, implementation of a vertical model-space localization in a serial EnSRF is presented. The results of radiance-only assimilation experiments comparing (i) model-space localization in a serial EnSRF, (ii) modulated-ensemble LETKF, and (iii) EnVar to observation-space localization in a serial EnSRF and LETKF are presented in section 4. Conclusions are in section 5.

2. A New Approach to Model-Space Localization in the EnKF

The Kalman filter update is

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{K}(\mathbf{y} - \mathbf{H}\mathbf{x}^b), \tag{1}$$

where $\mathbf{x}^b$ is the background state vector of dimension N_x, $\mathbf{y}$ is the observation vector of dimension N_y, $\mathbf{H}$ is the linear observation operator, $\mathbf{x}^a$ is the analysis vector. The Kalman gain $\mathbf{K}$ is

$$\mathbf{K} = \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1}, \tag{2}$$

with the background error covariance matrix $\mathbf{B}$ and the observation error covariance matrix $\mathbf{R}$.

2.1. Localization in the EnKF

In EnKFs, ensembles of state vectors are used to approximate error covariance matrices. Let $\mathbf{x}_i$, $i = 1, \ldots, \mathbf{x}_{N_e}$ be the ensemble of N_e background vectors and $\overline{\mathbf{x}} = 1/N_e \sum_{i=1}^{N_e} x_i$ be the ensemble mean. Let $\mathbf{X}$ be the $N_x \times N_e$ matrix of the normalized deviations of the ensemble members from the mean: $\mathbf{X} = (N_e - 1)^{-1/2}[\mathbf{x}_1 - \overline{\mathbf{x}}, \ldots, \mathbf{x}_{N_e} - \overline{\mathbf{x}}]$. Then the background error covariance matrix $\mathbf{B}$ may be approximated by $\mathbf{X}\mathbf{X}^T$, the sample covariance matrix.

Sample covariances may be used in an approximation of $\mathbf{K}$ in either of two mathematically equivalent forms,

$$\mathbf{K} \approx (\mathbf{X}\mathbf{X}^T)\mathbf{H}^T\left[\mathbf{H}(\mathbf{X}\mathbf{X}^T)\mathbf{H}^T + \mathbf{R}\right]^{-1} = \tag{3}$$

$$\mathbf{X}\mathbf{Y}^T(\mathbf{Y}\mathbf{Y}^T + \mathbf{R})^{-1}, \tag{4}$$

where the latter equality follows by defining a matrix of prior observation-space deviations, analogous to $\mathbf{X}$, by

$$\mathbf{Y} = \mathbf{HX}. \tag{5}$$

The direct use of (3) for data assimilation is known to perform poorly in realistic applications where $N_e \ll N_x, N_y$. As discussed in section 1, a crucial element of effective EnKFs is covariance localization, in which the sample covariances between state variables separated by large distances (or between state variables that otherwise might be assumed to have small correlation) are reduced or set to zero.

The two forms of ensemble approximation of $\mathbf{K}$ in (3) lead to two different approaches to covariance localization. First, the localization may be applied separately to the model-observation covariances ($\mathbf{XY}^T$) and observation-observation covariances ($\mathbf{YY}^T$), as suggested by the second form of $\mathbf{K}$ in (3):

$$\mathbf{K} = \mathbf{C} \circ \mathbf{XY}^T \left(\mathbf{D} \circ \mathbf{YY}^T + \mathbf{R} \right)^{-1}. \tag{6}$$

Here $\mathbf{C}$ is the background-observation prior error covariance localization matrix (size $N_x \times N_y$), $\mathbf{D}$ is the observation-prior error covariance localization matrix (size $N_y \times N_y$), and $\circ$ is the Schur, or element-by-element, matrix product. We will term this approach observation-space localization.

Second, localization can be applied directly to the model covariance ($\mathbf{XX}^T$), as suggested by the first form of $\mathbf{K}$ in (3):

$$\mathbf{K} = \left(\mathbf{L} \circ \mathbf{XX}^T \right) \mathbf{H}^T \left[\mathbf{H} \left(\mathbf{L} \circ \mathbf{XX}^T \right) \mathbf{H}^T + \mathbf{R} \right]^{-1}, \tag{7}$$

where $\mathbf{L}$ is the background error covariance localization matrix (size $N_x \times N_x$). We will term this approach model-space localization.

2.2. Model-Space Localization in a Serial Ensemble Filter

In the traditional serial ensemble filter, observations are processed one at a time (under the assumption that observation errors are uncorrelated), and the updated state after assimilating $k-1$ observations is the background for assimilating the kth observation. This is the case we treat below, for simplicity. This implementation of model-space localization also generalizes naturally to serial processing of batches of observations, analogous to the approach of Houtekamer and Mitchell (2001) with observation-space localization.

The update of the ith ensemble-mean state variable with the kth observation is

$$\bar{x}_i^{(k)} = \bar{x}_i^{(k-1)} + K_{ik} \left(y_k - \mathbf{H}_k \bar{\mathbf{x}}^{(k-1)} \right), \tag{8}$$

where subscripts for scalar (nonbold) variables indicate indices in vector or matrix quantities (i.e., $\bar{x}_i$ is the ith element of $\bar{\mathbf{x}}$ and K_{ik} is the (i,k)th element of $\mathbf{K}$), parenthesized superscripts denote indices in the serial update (i.e., $\bar{\mathbf{x}}^{(k-1)}$ is the state estimate after processing $y_1, \ldots, y_{k-1}$), and matrix quantities with subscripts indicate the corresponding row of the matrix (i.e., $\mathbf{H}_k$ is the kth row of $\mathbf{H}$).

The Kalman gain with the model-space localization as in (7) is then

$$K_{ik} = \frac{\left(\mathbf{L} \circ \mathbf{XX}^T \right)_i \mathbf{H}_k^T}{\mathbf{H}_k \left(\mathbf{L} \circ \mathbf{XX}^T \right) \mathbf{H}_k^T + \sigma_k^2}. \tag{9}$$

Here σ_k^2 is the observation error variance for the kth observation, and $\mathbf{X}$ is computed from the ensemble resulting from the assimilation of the first $k-1$ observations. In addition to the ensemble-mean update (8), the ensemble deviations must also be updated (e.g., as in Whitaker & Hamill, 2002); model-space localization appears in the gain for the deviation update as in (9).

The same approach can be applied to a serial, perturbed-observation EnKF. An analog of (8) can be used for all ensemble members, with $\bar{\mathbf{x}}$ replaced by the jth member and y_k by the perturbed observation for the jth member. The Kalman gain is the same as in (9).

For comparison, the Kalman gain for a serial EnKF with the traditional observation space localization as in (6) is

$$K_{ik} = \frac{C_{ik} \mathbf{X}_i \mathbf{Y}_k^T}{\mathbf{Y}_k \mathbf{Y}_k^T + \sigma_k^2} = \frac{C_{ik} \left(\mathbf{XX}^T \mathbf{H}_k^T \right)_i}{\mathbf{H}_k \mathbf{XX}^T \mathbf{H}_k^T + \sigma_k^2}. \tag{10}$$

2.3. Operation Count for Model-Space Localization in a Serial Filter

We next examine the increased number of operations when computing the Kalman gain from (9) rather than from (10).

First, consider the computations required when observation-space localization is used, as in (10). For each observation assimilated, one needs to compute $\mathbf{Y}_k$:

$$\mathbf{Y}_k = \mathbf{H}_k \mathbf{X} = \sum_{j=1}^{N_x} H_{kj} \mathbf{X}_j = \sum_{m=1}^{N_{nz}} H_{kj(m)} \mathbf{X}_{j(m)}, \tag{11}$$

where $N_{nz} = N_{nz}(k)$ is the number of nonzero elements of $\{j : H_{kj} \neq 0\}$, and $j(m)$ specifies precisely those elements for $m = 1, \ldots, N_{nz}$. Computing $\mathbf{Y}_k$ then requires $O(N_{nz} \times N_e)$ operations. This computation only needs to happen once per observation. The observation-observation covariances in the denominator of the Kalman gain $(\mathbf{Y}_k \mathbf{Y}_k^T)$ require $O(N_e)$ operations, also only computed once for kth observation. Finally, $O(N_e)$ operations are required to compute $C_{ik} \mathbf{X}_i \mathbf{Y}_k^T$, the localized state-observation covariances for the ith updated state variable and the kth observation. Let $N_{loc}(k)$ be the number of nonzero elements of the kth column of $\mathbf{C}$ (i.e., the number of nonzero C_{ik} for a given k). Thus, when using observation-space localization, updating the state given the kth observation uses $O(N_e N_{loc}(k) + N_e N_{nz} + N_e) = O(N_e(N_{loc}(k) + N_{nz} + 1))$ operations.

When model-space localization is used, as in (9), the localized state-observation covariances are

$$(\mathbf{L} \circ \mathbf{X}\mathbf{X}^T)_i \mathbf{H}_k^T = \sum_{j=1}^{N_x} L_{ij} \mathbf{X}_i \mathbf{X}_j^T H_{kj} = \mathbf{X}_i \sum_{m=1}^{N_{nz}} L_{ij(m)} H_{kj(m)} \mathbf{X}_{j(m)}^T. \tag{12}$$

Computing (12) requires $O(N_{nz} \times N_e)$ operations. Similarly, the cost of computing localized observation-observation covariances in the denominator of the Kalman gain (9) is $O(N_{nz} \times N_{nz} \times N_e)$ operations. As in observation-space localization, (12) only needs be computed when it has a nonzero value, namely, for i such that $L_{ij(m)} H_{kj(m)}$ is nonzero for at least some $j(m)$. Abusing notation, let $N_{loc}(k)$ also denote the number of those i. For model-space localization, the update given the kth observation then uses $O(N_e N_{nz} N_{loc}(k) + N_e N_{nz} N_{nz}) = O(N_e N_{nz}(N_{loc}(k) + N_{nz}))$ operations.

Note that $N_{loc}(k)$ for the observation-space and model-space localization will not necessarily be the same but should be similar when similar localization length scales are used.

Comparing these two estimates shows that the operation count when using model-space localization is increased by a factor of N_{nz} relative to observation-space localization. This increase will be the smallest when the observation operators are relatively local, so that the $\mathbf{H}_k$ have few nonzero elements. Of course, these operation counts do not translate directly into execution times, and the actual cost increase of the model-space localization will depend on software optimization, optimal (or suboptimal) use of memory, and the details of the computational platform, among other things.

If a parallel algorithm is used and the observations and model state variables are distributed across the processors, model-space localization may also result in increase in data transfers between the processors, compared to observation-space localization. Whereas for observation-space localization, $\mathbf{Y}_k$ (a vector size N_e) has to be used on all the processors that are calculating the state update; for model-space localization, all of those processors need to have access to the $h_{kj(m)} \mathbf{X}_{j(m)}^T, m = 1, \ldots, N_{nz}$—matrices size $N_e \times N_{nz}$. Thus, for model-space localization, it would be important to reduce the number of data transfers and use the parallel algorithm that distributes observations and model variables from local areas to the same processor, for example, as in Houtekamer et al. (2014).

2.4. Nonlinear Observation Operators

The foregoing discussion has covered only the case of linear observation operators. When the observations are related to the state by a nonlinear function $h(\mathbf{x})$, two additional questions emerge. The first is how to compute the observation-space quantities needed for the gain $\mathbf{K}$. With observation-space localization, the natural approach is to form the observation-prior ensemble $h(\mathbf{x}_i), i = 1, \ldots, N_e$, using the nonlinear observation operator and as in Houtekamer and Mitchell (1998) define

$$\mathbf{Y} = (N_e - 1)^{-1/2} \left[h(\mathbf{x}_1) - \overline{h(\mathbf{x})}, \ldots, h(\mathbf{x}_{N_e}) - \overline{h(\mathbf{x})} \right], \tag{13}$$

where $\overline{h(\mathbf{x})} = N_e^{-1} \sum_{i=1}^{N_e} h(\mathbf{x}_i)$. The necessary sample covariances in (6) and (9) can then be computed directly as $\mathbf{XY}^T$ and $\mathbf{YY}^T$. With model-space localization, one should instead follow Shlyaeva and Whitaker (2018) and linearize $h(\mathbf{x})$ about $\overline{\mathbf{x}}$, defining $\mathbf{H} = (\partial h/\partial \mathbf{x})|_{\mathbf{x}=\overline{\mathbf{x}}}$ and proceed as for a linear observation operator in (5). These two approaches are mathematically equivalent for linear observation operators, since $\mathbf{Y}$ from (13) is then equal to $\mathbf{HX}$.

The second question is how to generalize the innovation, $\mathbf{d} = \mathbf{y} - \mathbf{Hx}^b$, in (1) to the nonlinear case. As for the ensemble perturbations in the observation space, there are two possibilities. The first, corresponding to the use (13) for $\mathbf{Y}$ and following Houtekamer and Mitchell (2001), is to set $\mathbf{d} = \mathbf{y} - \overline{h(\mathbf{x})}$. The second, corresponding to employing $\mathbf{H} = (\partial h/\partial \mathbf{x})|_{\overline{\mathbf{x}}}$ to compute $\mathbf{Y}$, is to set $\mathbf{d} = \mathbf{y} - h(\overline{\mathbf{x}})$.

It is also possible, and may be computationally advantageous, to employ the linearized observation operator $\mathbf{H}$ in the algorithm with observation-space localization (Shlyaeva & Whitaker, 2018). The innovation is then computed as $\mathbf{d} = \mathbf{y} - h(\overline{\mathbf{x}})$ for both observation-space and model-space localization. The experiments presented in section 4 follow that implementation.

If $h(\mathbf{x})$ becomes more nonlinear, these aspects of the observation-space and model-space localization have the potential to produce different results. In this regime, the linearization $\mathbf{H} = (\partial h/\partial \mathbf{x})|_{\overline{\mathbf{x}}}$ may yield a poor approximation to the full $h(\mathbf{x})$. But use of the linearized $\mathbf{H}$ also provides the possibility of iterating by relinearizing about the updated estimate, as in the middle loop of variational assimilation schemes. Which approach is to be preferred as h becomes increasingly nonlinear remains an open question. Indeed, for sufficiently nonlinear h, both approaches will be suboptimal, and all approximations in (8) relative to full Bayesian update must be reconsidered.

3. Implementation of Vertical Model-Space Localization in EnSRF

In this section we outline the details of the implementation of vertical model-space localization in the serial EnSRF that is used for operational weather prediction at the National Oceanic and Atmospheric Administration (NOAA). An option to use the linearized observation operator (as in (5)) instead of the full observation operator (as in (13)) has been recently added to that EnSRF (Shlyaeva & Whitaker, 2018), which allows for the implementation of model-space localization as described in section 2.

As discussed in section 1, model-space localization can perform better than observation-space localization if observation operators are nonlocal (Campbell et al., 2010). The primary nonlocal observations for global weather prediction are satellite radiances, which, owing to the underlying radiative-transfer calculations, typically depend significantly on the model state across several model levels in the vertical. At the same time, observation operators for radiances are local in the horizontal. Because the cost of model-space localization for a given observation scales directly with the number of nonzero entries in $\mathbf{H}_k$ in (12), it is preferable to apply the model-space localization only in the vertical.

Horizontal and vertical localization can be applied separately if observation operator is separable horizontally and in vertical: $\mathbf{H} = \mathbf{H}^v \mathbf{H}^h$ where $\mathbf{H}^h$ is the horizontal observation operator and $\mathbf{H}^v$ is the vertical observation operator. The Kalman gain then becomes

$$K_{ik} = \frac{C_{ik}^h \left(\mathbf{L}^v \circ \mathbf{X}(\mathbf{H}^h\mathbf{X})^T\right)_i (\mathbf{H}_k^v)^T}{\mathbf{H}_k^v \left(\mathbf{L}^v \circ \mathbf{H}^h\mathbf{X}(\mathbf{H}^h\mathbf{X})^T\right)(\mathbf{H}_k^v)^T + \sigma_k^2}, \tag{14}$$

where C_{ik}^h is the horizontal covariance localization between the kth observation and the ith updated state variable, and $\mathbf{L}^v$ is the vertical background error covariance localization matrix.

The EnSRF with observation-space localization from Shlyaeva and Whitaker (2018) and the EnSRF with model-space localization implemented here also differ in the details of their respective update steps.

In both implementations, the background ensemble mean in observation space $\overline{\mathbf{y}}^{(0)} = h\left(\overline{\mathbf{x}}^{(0)}\right)$ is precomputed before any observations are assimilated.

The implementation of the EnSRF with observation-space localization also precomputes the background ensemble perturbations in observation space $\mathbf{Y}^{(0)}$ as in (5). An extended-state update as in Anderson and Collins (2007) is then applied to the observation-space quantities $\overline{\mathbf{y}}^{(k)}$ and $\mathbf{Y}^{(k)}$ at kth iteration of the sequential update, together with the state mean $\overline{\mathbf{x}}^{(k)}$ and state ensemble perturbations $\mathbf{X}^{(k)}$. For example, the

ensemble mean in the observation space $\overline{\mathbf{y}}^{(k)}$ is updated at each iteration of sequential update as

$$\overline{\mathbf{y}}^{(k)} = \overline{\mathbf{y}}^{(k-1)} + \tilde{K}_k \left(y_k - \mathbf{H}_k \overline{\mathbf{x}}^{(k-1)} \right), \tag{15}$$

where $\tilde{K}_k$ is a Kalman gain for update in the observation space. $\mathbf{Y}^{(k)}$ is updated in a similar way.

The model-space localization implementation presented here does not use an extended-state update and instead makes use of the linearized $\mathbf{H}$ to compute observation-space quantities from the updated state as needed. For example, the ensemble mean in the observation space $\overline{\mathbf{y}}^{(k)}$ is updated based on the following linearization:

$$\overline{\mathbf{y}}^{(k)} = \overline{\mathbf{y}}^{(0)} + \mathbf{H} \left(\overline{\mathbf{y}}^{(k)} - \overline{\mathbf{y}}^{(0)} \right). \tag{16}$$

$\mathbf{Y}^{(k)}$ is never explicitly computed in the model-space localization implementation, and Kalman gain is computed as in (9).

Because both implementations employ the linearized observation operator $\mathbf{H}$, they will be mathematically equivalent in the absence of localization.

In the implementation presented in this paper, $h\left(\overline{\mathbf{x}}^{(0)}\right)$ and $\mathbf{H} = (\partial h/\partial \mathbf{x})|_{\mathbf{x}=\overline{\mathbf{x}}^{(0)}}$ linearized about the background ensemble mean $\overline{\mathbf{x}}^{(0)}$ are precomputed before assimilation. It is possible that being able to recompute the full forward operator and its linearization based on the latest update $\overline{\mathbf{x}}^{(k-1)}$ could improve the results presented in the following section.

The parallel algorithm is similar to the one used in the observation-space localization EnSRF implementation: The observations and model grid points are distributed randomly across the processors. As stated in section 2, in the model-space localization case, this leads to increased data transfers between all processors for each updated observation when using model-space localization. We expect that if the observations and model grid points are distributed so that nearby observations and grid points are on the same processor, the computational cost of model-space localization can be significantly reduced compared to our implementation.

4. Experiments

In the experiments presented in this section, different approaches to model-space localization are compared to different approaches to observation-space localization, in the global atmospheric data assimilation setup.

The model-space localization algorithms used in the experiments are

- serial EnSRF with model-space localization, described in section 3;
- LETKF using the modulated ensemble approach (Lei et al., 2018); and
- pure-ensemble (nonhybrid) EnVar Global Statistical Interpolation System (Kleist & Ide, 2015). To highlight the effect of localization and make a clean comparison with the above two algorithms, the EnVar experiments did not use "middle loops" (the observation operators were not relinearized about the updated state), and the incremental tangent-linear normal-mode balance constraint (Kleist et al., 2009) was turned off.

The observation-space localization algorithms used in the experiments are

- serial EnSRF with observation-space localization (Shlyaeva & Whitaker, 2018), using the extended state update and linearized observation operator to precompute ensemble perturbations in observation space; and
- LETKF with observation-space localization (Shlyaeva & Whitaker, 2018). As in the above algorithm, linearized observation operator is used to precompute ensemble perturbations in observation space.

4.1. Single Observation Analysis Increments

First, single radiance observation assimilation experiments were run to compare analysis increments when using the different localization algorithms. A single Advanced Microwave Sounding Unit-A channel 7 observation was assimilated using all of the algorithms described above. All experiments used the same localization length scales, with the cutoff of 1.5 scale heights in the vertical and 1,500 km in the horizontal.
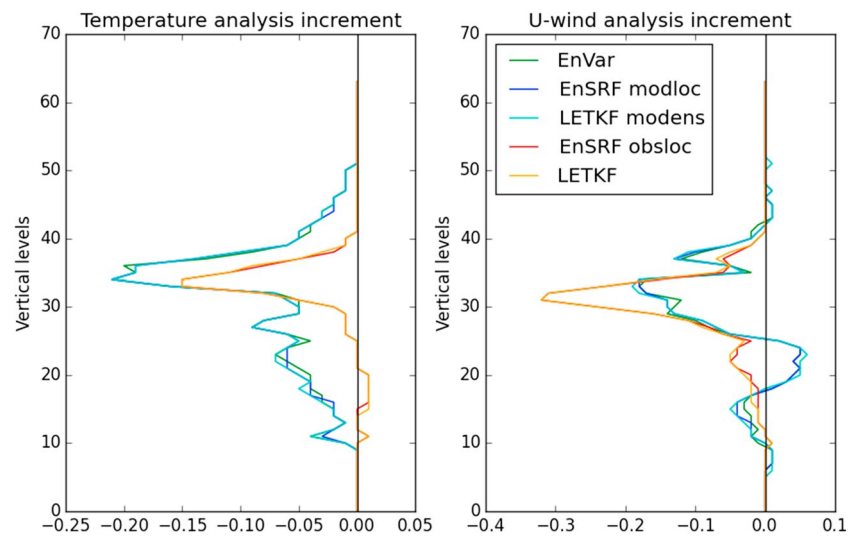
Figure 1. Vertical profiles of analysis increments at observation location for temperature (left) and u-wind component (right) for EnVar (green), serial EnSRF with model-space localization (blue), LETKF with modulated ensemble (cyan), serial EnSRF with observation-space localization (red), and LETKF with observation-space localization (orange). Single Advanced Microwave Sounding Unit-A channel 7 observation was assimilated. EnVar = Ensemble Variational; EnSRF = Ensemble Square Root Filter; LETKF = Local Ensemble Transform Kalman Filter.

The operational background ensemble from 06UTC 1 January 2016 was used in all the experiments. Vertical profiles of analysis increments for temperature and the horizontal components of the wind at the horizontal location of the observation are presented in Figure 1.

Figure 1 shows that all the algorithms with model-space localization produce similar analysis increments. Algorithms with observation-space localization (EnSRF and LETKF) produce an increment that looks very different from model-space localization algorithms. As in Buehner et al. (2010), if the localization length scale is increased, the differences between analysis increments obtained with the model-space and observation-space localization become smaller.

4.2. Data Assimilation Cycling Experiments

Cycled data assimilation experiments with a low-resolution (60 km) version of NOAA Global Forecast System were performed in order to compare different localization strategies. To emphasize the effect of different types of vertical localization with nonlocal observations, only satellite radiances were assimilated in all of the experiments. The following satellite radiances were assimilated: Advanced Microwave Sounding Unit-A from NOAA-15, NOAA-18, NOAA-19, Aqua, and Metop-A and Metop-B satellites; Atmospheric Infrared Sounder from Aqua satellite; Advanced Technology Microwave Sounder and Cross-track Infrared Sounder from Suomi NPP; Infrared Atmospheric Sounding Interferometer from Metop-A and Metop-B satellites; and Microwave Humidity Sounder from NOAA-18, NOAA-19, and Metop-A and Metop-B satellites. On average, about 3.3 million observations were assimilated in each 6 hr assimilation cycle.

In the EnVar cycling data assimilation experiment, the EnSRF with observation-space localization was used to update ensemble perturbations, while the EnVar was used to update the ensemble mean.

Bias correction coefficients for the satellite radiances from the operational high-resolution run that assimilated all available observations were used. All of the experiments used a single time-level background (6-hr forecast) and produced analysis at a single time level. The experiments were initialized from the operational ensemble on from 00UTC 1 January 2016 and run through 18UTC 1 February 2016. The localization length scales with the cutoff of 1.5 scale heights in the vertical and 1,500 km in the horizontal were used in all experiments.

Figure 2 shows the root-mean-square differences between the ensemble-mean background 6-hr forecast and in situ wind and temperature observations (which were not assimilated), averaged over the duration of the experiments (the first four spin-up days were discarded). Figure 3 presents the same information as the right
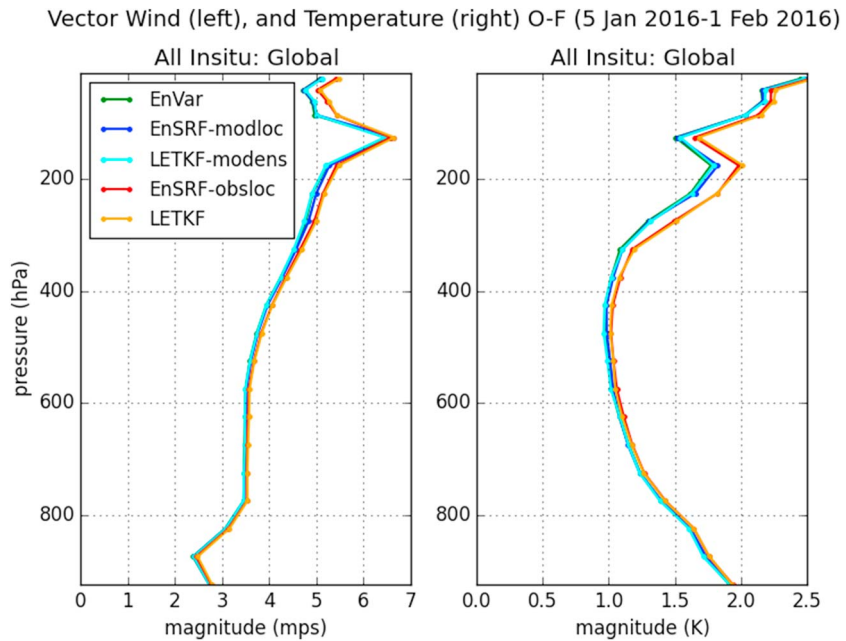
Figure 2. Root-mean-square error of 6-hr mean ensemble forecast wind speed (left) and temperature (right) compared with in situ wind and temperature observations, averaged from 5 January 2016 to 1 February 2016 for EnVar (green), serial EnSRF with model-space localization (blue), LETKF with modulated ensemble (cyan), serial EnSRF with observation-space localization (red), and LETKF with observation-space localization (orange). EnVar = Ensemble Variational; EnSRF = Ensemble Square Root Filter; LETKF = Local Ensemble Transform Kalman Filter; O-F = observation minus forecast.
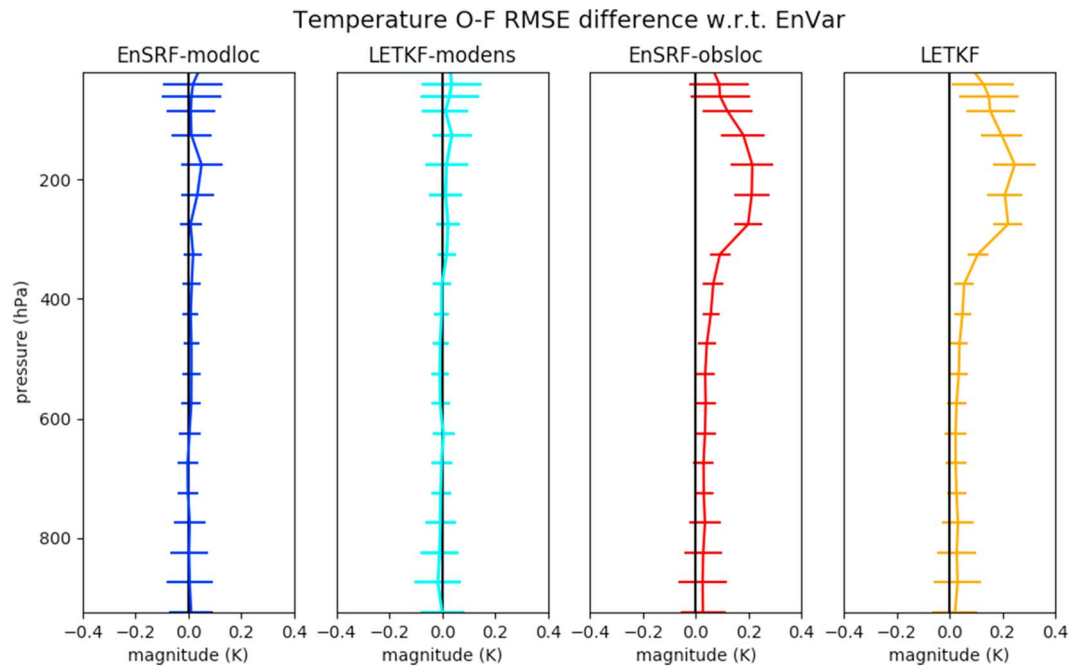
Figure 3. Relative differences in the RMSE of 6-hr mean ensemble forecast temperature compared with EnVar experiment, for serial EnSRF with model-space localization (blue), LETKF with modulated ensemble (cyan), serial EnSRF with observation-space localization (red), and LETKF with observation-space localization (orange). The bars indicate 95% confidence intervals (using t test for independent samples). EnVar = Ensemble Variational; EnSRF = Ensemble Square Root Filter; LETKF = Local Ensemble Transform Kalman Filter; RMSE = root-mean-square errors; O-F = observation minus forecast.

panel of Figure 2, showing relative differences in temperature root-mean-square errors with respect to the EnVar experiment. The bars on the plots indicate 95% confidence intervals.

The ensemble-mean background forecasts are closer to the withheld in situ observations in the upper-troposphere when model-space localization is used. The EnSRF with model-space localization performs comparably to EnVar and to the LETKF using modulated ensembles to simulate model-space localization. The results from the two data assimilation approaches using observation-space localization (LETKF and EnSRF) are similar to each other and inferior compared to the model-space localization approaches, consistent with the conclusion of Campbell et al. (2010) that model-space localization is preferred when nonlocal observations are assimilated.

Cycled data assimilation experiments with assimilating only vertically local (radiosonde) observations were also performed, the results of model-space and observation-space localization are nearly identical due to locality of the observation operator.

For the experiments in this section, EnSRF with model-space localization was on average about 8 times slower than EnSRF with observation-space localization. The experiments with assimilating radiosonde observations were about 1.5 times slower for EnSRF with model-space localization.

5. Conclusion

In this paper we presented a new approach to applying model-space localization to EnKFs. The approach is applicable to serial ensemble filters and requires the availability of the linearized observation operator so that the observation operator can be applied after direct localization of background error covariances.

This approach was implemented and tested in the NOAA serial EnSRF for global atmospheric data assimilation. Since observation-space and model-space localizations behave differently only for nonlocal observation operators and all observation operators used currently at NOAA are local in the horizontal space, we implemented model-space localization only in vertical, leaving the horizontal localization in the observation space.

The application of model-space localization in vertical allowed for better assimilation of satellite radiances that use vertically nonlocal observation operators. Cycled data assimilation experiments with the NOAA global atmospheric forecast system assimilating only radiance observations were performed to compare results of EnSRF with observation-space localization, LETKF; EnSRF with model-space localization, LETKF using modulated ensemble to emulate model-space localization and EnVar. The application of model-space localization in EnSRF or in LETKF (via modulated ensemble approach) allowed to reach the same level of accuracy for the background ensemble mean compared to the in situ observations, as in EnVar data assimilation (that uses model-space localization), whereas EnSRF with observation-space localization and LETKF with observation-space localization showed worse results in the upper troposphere. This shows that one of the reasons of EnVar systems usually outperforming EnKF systems is the differences in localization approach. It should be noted that for the purpose of isolating the effect of localization, in these experiments, EnVar used pure ensemble (not hybrid) covariances, no relinearizations of observation operator were performed, and the incremental tangent-linear normal-mode balance constraint was turned off. In the experiments not presented in this paper, EnVar results improve over the results presented here when those options are used as in the operational data assimilation setup, with the main benefit coming from using hybrid background error covariances.

The relative increase in computation cost of the model-space localization in the serial EnKF depends on the locality of the observation operator and the chosen localization distance. For the experiments with the global atmospheric model presented in this paper, the change of vertical localization to the model-space localization resulted in an 8 times increase in runtime compared to EnSRF with observation-space localization when satellite radiances were assimilated and 1.5 times increase when radiosondes were assimilated, with the same number of processors used. The implementation of model-space localization used for the experiments in this paper was devised to be the proof of concept, and it could be optimized. The extra computations can be easily parallelized. We also expect the computational time increase be smaller if in the parallel algorithm data distribution is so that observations and model state variables that are in the same local area are on the same processor, which would avoid excessive data communications in our current implementation. The approach to model-space localization presented in this paper becomes more expensive when less

severe localization in used. On the contrary, for modulated ensemble approach, a more severe localization is more expensive. It might be beneficial to use the approach presented in this paper for tighter localization and modulated ensemble approach for wider localization. Using observation-space localization and increasing the localization length scale remains a viable option in the absence of resources for using model-space localization.

References

Anderson, J. L. (2001). An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, *129*(12), 2884–2903. https://doi.org/10.1175/1520-0493(2001)129<2884:AEAKFF>2.0.CO;2

Anderson, J. L., & Collins, N. (2007). Scalable implementations of ensemble filter algorithms for data assimilation. *Journal of Atmospheric and Oceanic Technology*, *24*(8), 1452–1463. https://doi.org/10.1175/JTECH2049.1

Bishop, C. H., & Hodyss, D. (2009). Ensemble covariances adaptively localized with ECO-RAP. Part 2: A strategy for the atmosphere. *Tellus A*, *61*(1), 97–111. https://doi.org/10.1111/j.1600-0870.2008.00372.x

Bishop, C. H., Whitaker, J. S., & Lei, L. (2017). Gain form of the ensemble transform Kalman filter and its relevance to satellite data assimilation with model space ensemble covariance localization. *Monthly Weather Review*, *145*(11), 4575–4592. https://doi.org/10.1175/MWR-D-17-0102.1

Buehner, M. (2005). Ensemble-derived stationary and flow-dependent background-error covariances: Evaluation in a quasi-operational NWP setting. *Quarterly Journal of the Royal Meteorological Society*, *131*(607), 1013–1043. https://doi.org/10.1256/qj.04.15

Buehner, M., Houtekamer, P. L., Charette, C., Mitchell, H. L., & He, B. (2010). Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part I: Description and single-observation experiments. *Monthly Weather Review*, *138*(5), 1550–1566. https://doi.org/10.1175/2009MWR3157.1

Campbell, W. F., Bishop, C. H., & Hodyss, D. (2010). Vertical covariance localization for satellite radiances in ensemble Kalman filters. *Monthly Weather Review*, *138*(1), 282–290. https://doi.org/10.1175/2009MWR3017.1

Greybush, S. J., Kalnay, E., Miyoshi, T., Ide, K., & Hunt, B. R. (2011). Balance and ensemble Kalman filter localization techniques. *Monthly Weather Review*, *139*(2), 511–522. https://doi.org/10.1175/2010MWR3328.1

Houtekamer, P. L., He, B., & Mitchell, H. L. (2014). Parallel implementation of an ensemble Kalman filter. *Monthly Weather Review*, *142*(3), 1163–1182. https://doi.org/10.1175/MWR-D-13-00011.1

Houtekamer, P. L., & Mitchell, H. L. (1998). Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, *126*(3), 796–811. https://doi.org/10.1175/1520-0493(1998)126<0796:DAUAEK>2.0.CO;2

Houtekamer, P. L., & Mitchell, H. L. (2001). A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, *129*(1), 123–137. https://doi.org/10.1175/1520-0493(2001)129<0123:ASEKFF>2.0.CO;2

Hunt, B. R., Kostelich, E. J., & Szunyogh, I. (2007). Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, *230*(1), 112–126. https://doi.org/https://doi.org/10.1016/j.physd.2006.11.008

Kleist, D. T., & Ide, K. (2015). An OSSE-based evaluation of hybrid variational ensemble data assimilation for the NCEP GFS. Part I: System description and 3D-hybrid results. *Monthly Weather Review*, *143*(2), 433–451. https://doi.org/10.1175/MWR-D-13-00351.1

Kleist, D. T., Parrish, D. F., Derber, J. C., Treadon, R., Errico, R. M., & Yang, R. (2009). Improving incremental balance in the GSI 3DVAR analysis system. *Monthly Weather Review*, *137*(3), 1046–1060. https://doi.org/10.1175/2008MWR2623.1

Lei, L., & Whitaker, J. S. (2015). Model space localization is not always better than observation space localization for assimilation of satellite radiances. *Monthly Weather Review*, *143*(10), 3948–3955. https://doi.org/10.1175/MWR-D-14-00413.1

Lei, L., Whitaker, J. S., & Bishop, C. (2018). Improving assimilation of radiance observations by implementing model space localization in an ensemble Kalman filter. JAMES.

Lorenc, A. C. (2003). The potential of the ensemble Kalman filter for NWP—A comparison with 4D-Var. *Quarterly Journal of the Royal Meteorological Society*, *129*(595), 3183–3203. https://doi.org/10.1256/qj.02.132

Shlyaeva, A., & Whitaker, J. S. (2018). Using the linearized observation operator to calculate observation-space ensemble perturbations in ensemble filters. *Journal of Advances in Modeling Earth Systems*, *10*, 1414–1420. https://doi.org/10.1029/2018MS001309

Steward, J. L., Roman, J. E., Davina, A. L., & Aksoy, A. (2018). Parallel direct solution of the covariance-localized ensemble square root Kalman filter equations with matrix functions. *Monthly Weather Review*, *146*(9), 2819–2836. https://doi.org/10.1175/MWR-D-18-0022.1

Whitaker, J. S., & Hamill, T. M. (2002). Ensemble data assimilation without perturbed observations. *Monthly Weather Review*, *130*(7), 1913–1924. https://doi.org/10.1175/1520-0493(2002)130<1913:EDAWPO>2.0.CO;2