

Progress toward the Application of a Localized Particle Filter for Numerical Weather Prediction

JONATHAN POTERJOY

University of Maryland, College Park, College Park, Maryland, and Cooperative Institute for Mesoscale Meteorological Studies, and University of Oklahoma, and NOAA/National Severe Storms Laboratory, Norman, Oklahoma, and NOAA/Atlantic Oceanographic and Meteorological Laboratory, Miami, Florida

LOUIS WICKER

NOAA/National Severe Storms Laboratory, Norman, Oklahoma

MARK BUEHNER

Meteorological Research Division, Environment and Climate Change Canada, Dorval, Quebec, Canada

(Manuscript received 13 November 2017, in final form 11 September 2018)

ABSTRACT

A series of papers published recently by the first author introduce a nonlinear filter that operates effectively as a data assimilation method for large-scale geophysical applications. The method uses sequential Monte Carlo techniques adopted by particle filters, which make no parametric assumptions for the underlying prior and posterior error distributions. The filter also treats the underlying dynamical system as a set of loosely coupled systems to effectively localize the effect observations have on posterior state estimates. This property greatly reduces the number of particles—or ensemble members—required for its implementation. For these reasons, the method is called the local particle filter. The current manuscript summarizes algorithmic advances made to the local particle filter following recent tests performed over a hierarchy of dynamical systems. The revised filter uses modified vector weight calculations and probability mapping techniques from earlier studies, and new strategies for improving filter stability in situations where state variables are observed infrequently with very accurate measurements. Numerical experiments performed on low-dimensional data assimilation problems provide evidence that supports the theoretical benefits of the new improvements. As a proof of concept, the revised particle filter is also tested on a high-dimensional application from a real-time weather forecasting system at the NOAA/National Severe Storms Laboratory (NSSL). The proposed changes have large implications for researchers applying the local particle filter for real applications, such as data assimilation in numerical weather prediction models.

1. Introduction

The success of ensemble Kalman filters (EnKFs) in oceanography, meteorology, and other fields of geoscience has motivated recent efforts to develop more general Monte Carlo filters for data assimilation. This research includes advancing particle filters (PFs) for geophysical models. Strategies for implementing these filters in high-dimensional problems tend to fall into one or more of the following categories: 1) filters that

manipulate the transition density between observation times or use a carefully chosen proposal density to reduce the number of particles having low likelihood (van Leeuwen 2010; Chorin et al. 2010); 2) filters that combine PFs with EnKFs to maintain the benefits of Kalman filters for situations where Gaussian assumptions are appropriate (e.g., Majda et al. 2014; Frei and Künsch 2013; Slivinski et al. 2015; Chustagulprom et al. 2016); and 3) filters that break the data assimilation problem into a set of independent problems via spatial localization or other means (e.g., Bengtsson et al. 2003; Lei and Bickel 2011; Poterjoy 2016, hereafter P16; Poterjoy and Anderson 2016, hereafter PA16; Penny and Miyoshi 2016;

Corresponding author: Dr. Jonathan Poterjoy, poterjoy@umd.edu

DOI: 10.1175/MWR-D-17-0344.1

© 2019 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](#) (www.ametsoc.org/PUBSReuseLicenses).

Lee and Majda 2016). This research presents multiple pathways for incorporating benefits of PFs into pre-existing environmental modeling systems designed for EnKFs. For example, Robert et al. (2018) successfully applied a localized ensemble transform Kalman particle filter for data assimilation in the Consortium for Small-Scale Modeling (COSMO) framework. Recent work by the German Meteorological Office testing a Markov Chain particle filter for numerical weather prediction have also been encouraging (Potthast 2016).

The current study summarizes recent progress developing a localized sequential importance resampling (SIR) PF for geophysical data assimilation. The method, first introduced in P16 as the “local PF,” operates by assimilating observations with independent errors sequentially and combining sampled particles and prior particles for each observation. The local PF satisfies the SIR PF solution for state variables located in close geographical proximity to observations in the sequence, but maintains the prior particles for state variables located far from observations. A smooth correlation function that tapers to zero at a finite user-specified distance controls the spatial influence observations have on posterior estimates, which greatly reduces the number of particles needed for geophysical data assimilation problems. Therefore, the resulting filter falls into the third category of PFs described above. An appealing aspect of the local PF is it transitions into the SIR PF as the localization length scale approaches infinity, which would be done with a very large number of particles. Therefore, the local PF converges to the Bayesian solution as the number of particles and localization length scale increase.

Recent studies by PA16 and Poterjoy et al. (2017, hereafter PSA17), demonstrate that the local PF operates effectively for high-dimensional systems. In PA16, the local PF provides accurate posterior representations of baroclinic Rossby waves over year-long data assimilation experiments performed with a simplified atmospheric general circulation model. Following this work, PSA17 show benefits of the local PF in an idealized (i.e., observing systems simulation experiment) convective-scale ensemble analysis and prediction system when compared to a conventional EnKF system. These two studies find the local PF to operate effectively using 25 and 100 particles, respectively, which motivates recent applications for real data assimilation problems. Experiments performed during the course of this work, as well as ongoing efforts applying the local PF for the analysis and prediction of severe convective storms using the experimental “Warn-on-Forecast” prediction framework

(Wheatley et al. 2015; Jones et al. 2016) at NOAA’s National Severe Storms Laboratory (NSSL), lead to several important improvements relative to the filter formulation outlined in P16. The purpose of this manuscript is to discuss obstacles encountered when applying the local PF for applications of increasing complexity and summarize updates made to the filter formulation outlined in P16.

The manuscript is organized in the following manner. Section 2 introduces the SIR PF and the local PF. Sections 3 and 4 provide several improvements for the P16 formulation and filter stabilization techniques required for situations where localization alone cannot prevent the collapse of particle weights. Section 5 presents numerical experiments performed with a low-order model, which justify the algorithmic improvements in sections 3 and 4. Section 6 shows results from a real weather forecasting example to demonstrate that the revised local PF can operate effectively for a high-dimensional geophysical application. The last section summarizes the main findings of this study and discusses the potential of the revised filter for real numerical weather prediction.

2. Sequential importance resampling and the local particle filter

The data assimilation problem of interest considers the probability of a system state \mathbf{x} , represented as a random vector of length N_x , conditioned on independent noisy observations stored in \mathbf{y} , a vector of length N_y . For simplicity, assume \mathbf{x} and \mathbf{y} are valid at the same time, thus allowing time indices to be ignored. Observations relate to \mathbf{x} through

$$\mathbf{y} = H[\mathbf{x}] + \varepsilon, \quad (1)$$

where H is a function that maps the model state to observation space and ε is the observation error. From Bayes’ theorem, the posterior density is proportional to the product of the prior density and the likelihood of particles:

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{\int p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) d\mathbf{x}}. \quad (2)$$

Given N_e equally likely samples from $p(\mathbf{x})$, the PF approximates this density with a sum of delta functions. Consequentially, the posterior density is approximated as a sum of weighted delta functions:

$$p(\mathbf{x}|\mathbf{y}) \approx \sum_{n=1}^{N_e} w_n \delta(\mathbf{x} - \mathbf{x}_n). \quad (3)$$

Here, the weights are proportional to the likelihood of particles given observations:

$$w_n = \frac{p(\mathbf{y}|\mathbf{x}_n)}{\sum_{m=1}^{N_e} p(\mathbf{y}|\mathbf{x}_m)}. \quad (4)$$

More efficient versions of this filter can be formulated by drawing samples from a proposal distribution conditioned on additional information, such as current observations; see [Doucet et al. \(2001\)](#) for a review. The local PF adopts the formulation of the SIR PF described here, which uses the prior density as a proposal. Future formulations of the local PF may include other choices of proposal density, but the current configuration is sufficient for describing important components of the algorithm.

The weights in (4) can inform how to sample particles from $p(\mathbf{x}|\mathbf{y})$. In its simplest form, the SIR PF duplicates particles with large weights and removes particles with small weights ([Gordon et al. 1993](#)), thus sampling most particles from regions where the posterior probability is highest. After resampling, all weights are set equal to $1/N_e$ before propagating particles to the next time observations are available.

The SIR PF is easy to implement and converges to the Bayesian solution as $N_e \rightarrow \infty$, but the number of particles required for its practical application in high-dimensional systems is computationally prohibitive ([Bickel et al. 2008](#); [Bengtsson et al. 2008](#); [Snyder et al. 2008](#)). Recent studies by [Snyder et al. \(2015\)](#) and [Slivinski and Snyder \(2016\)](#) suggest limitations of the SIR PF may also persist when particles are sampled from the “optimal proposal” to reduce the N_e required for its stability. The optimal proposal refers to the proposal density conditioned on the previous model state during sequential data assimilation cycles and the most recent observations ([Doucet et al. 2000](#)). For this reason, most current efforts for alleviating the dimensionality constraint of PFs focus on spatial localization (see [section 1](#) for examples). In this context, localization refers to a filter construction that limits the calculation of posterior quantities to discrete regions of the state space, based on a priori assumptions of the system dynamics and observing system, and the degree of sampling error introduced by applying Monte Carlo approximations of probabilities. This strategy was first adopted by the atmospheric science community to reduce the ensemble size needed for EnKFs in global numerical weather prediction models ([Houtekamer and Mitchell 2001](#); [Hamill and Whitaker 2001](#)). More recently, [P16](#), [PA16](#), and [PSA17](#) adopt a similar approach to apply the SIR PF for comparable geophysical data assimilation problems.

The filter introduced in [P16](#) performs localization by first expanding each scalar w_n in (4) to vectors $\boldsymbol{\omega}_n$ of length N_x . Provided that observation errors are independent, which is often assumed for data assimilation in geoscience, $p(\mathbf{y}|\mathbf{x}_n)$ factors into the product $\prod_{i=1}^{N_y} p(y_i|\mathbf{x}_n)$. This factorization provides a straightforward means of constructing weights that reflect only the local influence of observations on neighboring state variables. [P16](#) specifies the weights for each state variable in \mathbf{x} (denoted by index j) using the likelihood of particles for each observation in \mathbf{y} and a localization function, $l(y_i, x_j, r_{\text{loc}})$:

$$\boldsymbol{\omega}_{n,j} = \prod_{i=1}^{N_y} \{ [p(y_i|\mathbf{x}_n) - 1] l(y_i, x_j, r_{\text{loc}}) + 1 \}. \quad (5)$$

Each $\boldsymbol{\omega}_n$ is then normalized by $\boldsymbol{\Omega} = \sum_{m=1}^{N_e} \boldsymbol{\omega}_m$ so all elements sum to one. For $l(y_i, x_j, r_{\text{loc}})$, we choose a correlation function introduced by Eq. (4.10) of [Gaspari and Cohn \(1999\)](#), which is a continuous, smooth function that has a maximum value of unity when the Euclidean distance between y_i and x_j is zero, and tapers to zero monotonically with distance at a rate controlled by the coefficient r_{loc} . With this function, the weights for each state variable approach the scalar weights in (4) as $r_{\text{loc}} \rightarrow \infty$, yielding the original SIR PF weights. Using the chosen localization function with a value of r_{loc} much larger than the model grid spacing helps produce continuity in the weighting vectors, which is beneficial for most geophysical applications. This function also has compact support, which allows for efficient parallelization of the algorithm.

[P16](#) presents a two-step process for generating posterior samples that reflect the localized vector weights. The first step resamples particles based on scalar weights proportional to the likelihood of particles for each observation in \mathbf{y} , then merges sampled and prior particles in a manner that satisfies the SIR PF solution near the observation, and the posterior mean and variance given by the localized weight equations everywhere else in model space. The resampling process occurs sequentially over all current observations, much like serial implementations of EnKFs ([Anderson 2001](#); [Whitaker and Hamill 2002](#)). As one would assume, the process of localizing particle updates follows a similar strategy. The second step applies a nonparametric probability mapping algorithm to make higher-order corrections to the particles. We describe the details of both steps in [section 3b](#) and encourage readers to review [P16](#), which provides a derivation for update equations, a schematic illustrating each part of the algorithm, and numerical experiments demonstrating the impact of the probability

mapping step. We also refer readers to section 3b of PSA17, which provides two-dimensional examples demonstrating how particles are updated within the localization region, and how the solution differs from a localized EnKF.

3. A revised local particle filter

Since its introduction in P16, the local PF has undergone testing for data assimilation problems of increasing complexity. These experiments range from a simplified general circulation model in PA16 to current testing using a regional convective-scale weather prediction system at NSSL. The extended use of this method motivates several changes to the original algorithm, which we describe in this section. To supplement our description, we also provide a list of important symbols and definitions, and a pseudocode local PF algorithm in appendixes A and B, respectively.

a. Reformulation of vector weight equations

In this section, we describe circumstances that can cause the local PF to become unstable. These situations typically occur when using small numbers of particles to assimilate measurements that have error variance much smaller than the prior variance – in which case, it is unlikely any particles will yield large likelihoods. We use numerical simulations to demonstrate how the shape of vector weights depends greatly on the sum of particle likelihoods, which becomes problematic when likelihoods are very small. We then introduce a revised weight equation, which avoids this problem by normalizing likelihoods before introducing them in vector weight equations. Like the P16 local PF, the revised filter converges to the SIR PF as r_{loc} becomes large. The modification, however, changes how localization influences the shape of particle weighting vectors.

To understand the effect of accurate observations (i.e., those with low error compared to the prior) on the shape of vector weight calculations, recall that the set of vector weights calculated in (5) are normalized by Ω to guarantee they sum to unity. For a single observation y , the elements of Ω equal N_e outside the localization region, but approach $\sum_{m=1}^{N_e} p(y|\mathbf{x}_m)$ near the observation; that is, setting $N_y = 1$ and taking the sum of (5) gives $\Omega_j = l(y, x_j, r_{\text{loc}}) \sum_{m=1}^{N_e} p(y|\mathbf{x}_m) + [1 - l(y, x_j, r_{\text{loc}})]N_e$. Therefore, both $l(y, x_j, r_{\text{loc}})$ and $\sum_{m=1}^{N_e} p(y|\mathbf{x}_m)$ determine the shape of each ω_n . For an observation that produces large likelihoods for many particles [i.e., large $\sum_{m=1}^{N_e} p(y|\mathbf{x}_m)$], the effect will be to “broaden” the spatial influence of the observation




FIG. 1. P16 formulated weighting vectors (red lines) are compared with the new weighting vectors (black lines) for a single observation of variable 20. The figure shows cases using (a) $\sigma_y = 1$, (b) $\sigma_y = 0.2$, and (c) $\sigma_y = 0.02$ to demonstrate the shape dependence on likelihood sums (denoted by “norm” at the top of plots).

compared to those that yield smaller likelihoods—by decreasing the rate at which vector weights converge from the nonlocalized solution to an equal-weights solution. To demonstrate this property, we perform a set of data assimilation experiments for a system of dimension $N_x = 40$, defined on a uniform 1D grid. The test problem uses $N_e = 20$ particles to assimilate a single observation of the 20th state variable, assuming $\varepsilon \sim N(0, \sigma_y^2)$. To examine how the shape of weighting vectors depends on the sum of likelihoods, we set the observation equal to the prior mean and consider three different values for observation error standard deviation: $\sigma_y = 1$, $\sigma_y = 0.2$, and $\sigma_y = 0.02$. For each experiment, we calculate the localized vectors weights (one for each particle) and examine how their shapes change with different observation error standard deviation. These weights, plotted as red lines in Figs. 1a–c, clearly demonstrate a shape dependence on normalization. As the sum of likelihoods becomes very small (Fig. 1c), the particles closest to the observation obtain vector weights that approach delta functions. The rapid decrease in likelihoods with decreasing σ_y follows from the use of an

exponential function (i.e., a Gaussian) for observation errors, which is common practice for data assimilation in geophysics.

In addition to reducing the control a user has over the influence of localization on particle weights, as demonstrated in Fig. 1, the P16 weight formulation makes the local PF prone to round-off error. More notably, the filter performs suboptimally when many observations are located in the neighborhood of a given grid point, causing the normalization to be a product of several low likelihood values. The small numbers resulting from this product can lead to spurious deviations from the smooth transitions of weights provided by the localization function and cause the local PF to fail for large applications (not shown). Following this realization, we introduce a modification to the original weight calculations that eliminates the negative consequences of assimilating accurate measurements. The new approach localizes the PF in a manner that is more consistent with preexisting serial EnKFs, such as the ensemble square root filter (Whitaker and Hamill 2002) and the ensemble adjustment Kalman filter (Anderson 2001).

For a single observation, serial EnKFs update each element of the posterior mean \bar{x}_j^a by adding a correction to the prior mean \bar{x}_j^f :

$$\bar{x}_j^a = \bar{x}_j^f + l(y, x_j, r_{loc})\Delta x_j. \tag{6}$$

In (6), Δx_j is the correction for the j th state variable coming from the Kalman filter update equations. This correction is damped by a localization coefficient, so that only variables with a nonzero $l(y, x_j, r_{loc})$ receive an update. For the PF calculation of the mean, (6) can be written

$$\bar{x}_j^a = \frac{1}{N_e} \sum_{n=1}^{N_e} x_{n,j} + l(y, x_j, r_{loc}) \left[\sum_{n=1}^{N_e} \hat{w}_n x_{n,j} - \frac{1}{N_e} \sum_{n=1}^{N_e} x_{n,j} \right], \tag{7}$$

where \bar{x}_j^f is replaced by the prior sample mean, Δx_j is expressed as the difference between the nonlocalized PF posterior mean and the prior mean. In (7), the non-localized posterior mean estimate comes directly from (3), which can approximate posterior expected values for any measurable function, $f(\mathbf{x})$:

$$\begin{aligned} \overline{f(\mathbf{x})} &= \int f(\mathbf{x})p(\mathbf{x}|\mathbf{y}) d\mathbf{x}, \\ &\approx \sum_{n=1}^{N_e} \hat{w}_n f(\mathbf{x}_n). \end{aligned} \tag{8}$$

In the above expressions, the weights \hat{w}_n are the normalized likelihoods calculated for the observation:

$$\hat{w}_n = \frac{p(y|\mathbf{x}_n)}{\sum_{m=1}^{N_e} p(y|\mathbf{x}_m)}. \tag{9}$$

To find the vector weights that satisfy this type of localization, we set the right-hand side of (7) to a vector weighted sum of prior particles and solve for these weights. After simplification, the weight equation for each element becomes

$$\omega_{n,j} = \frac{(N_e \hat{w}_n - 1)l(y, x_j, r_{loc}) + 1}{N_e}, \tag{10}$$

which generalizes for multiple observations to

$$\omega_{n,j} = \prod_{i=1}^{N_y} \left[\frac{(N_e \hat{w}_n^{(y_i)} - 1)l(y_i, x_j, r_{loc}) + 1}{N_e} \right], \tag{11}$$

where $\hat{w}_n^{(y_i)} = p(y_i|\mathbf{x}_n) / \left[\sum_{m=1}^{N_e} p(y_i|\mathbf{x}_m) \right]$.

When we repeat the single-observation experiments described above, the new weighting vectors (black lines in Fig. 1) exhibit a broader structure than the previous weight formulation, but maintain the same shape regardless of the likelihoods. From (10) it is clear that the weighting vectors are proportional to the specified function l with the added offset $1/N_e$. The major difference between the two weight formulations in (5) and (11), is (11) normalizes likelihoods before using them to form the vector weights. This step keeps all elements of Ω close to constant, thus forcing the weights to be independent of the magnitude of likelihoods. The revised weight formulation also alleviates computational issues related to round off errors. To further reduce round off errors, (11) can be normalized after each y_i in the product to yield the same result as normalizing the full product, which we find to be necessary when assimilating large networks of dense observations.

b. Reformulation of update equations

The local PF described in P16 uses (5) to derive a set of equations needed to update particles. Therefore, revisions to the particle update equations must follow any modifications to the vector weight equation. For the i th observation in sequence, the local PF considers the original set of prior particles not updated by observations at the current time, $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_e}\}$, and the set of particles updated by all observations in $(y_1 y_2 \dots y_{i-1})^T$ at the current time, which we denote by $\{\mathbf{x}_1^{(y_{1:i-1})}, \dots, \mathbf{x}_{N_e}^{(y_{1:i-1})}\}$. The local PF first samples

particles with replacement from $\{\mathbf{x}_1^{(y_{1:i-1})}, \dots, \mathbf{x}_{N_e}^{(y_{1:i-1})}\}$, using the scalar weights:

$$\tilde{w}_n^{(y_{1:i})} = \frac{p(y_i | \mathbf{x}_n^{(y_{1:i})})}{\sum_{m=1}^{N_e} p(y_i | \mathbf{x}_m^{(y_{1:i})})}. \tag{12}$$

This process follows the same bootstrap resampling step used by standard PF methods, which considers the likelihood of each particle in the set $\{\mathbf{x}_1^{(y_{1:i-1})}, \dots, \mathbf{x}_{N_e}^{(y_{1:i-1})}\}$ to determine the sampling. The sampled particles are then merged with this set to form localized posterior particles:

$$\mathbf{x}_{n,j}^{(y_{1:i})} = \bar{\mathbf{x}}_j^{(y_{1:i})} + r_{1,j} \left(\mathbf{x}_{k_n,j}^{(y_{1:i-1})} - \bar{\mathbf{x}}_j^{(y_{1:i})} \right) + r_{2,j} \left(\mathbf{x}_{n,j}^{(y_{1:i-1})} - \bar{\mathbf{x}}_j^{(y_{1:i})} \right). \tag{13}$$

Here, k_n is the index of each sampled particle, $\bar{\mathbf{x}}_j^{(y_{1:i})}$ is the mean conditioned on $(y_1 y_2 \dots y_i)^T$, and $r_{1,j}$ and $r_{2,j}$ determine how to merge the sampled particles with current particles. The coefficients used for the merging are stored in vectors \mathbf{r}_1 and \mathbf{r}_2 , and specified to satisfy the SIR PF solution near the current observation and the prior solution where localization coefficients are zero. Everywhere else in state space, the local PF satisfies the posterior mean, $\bar{\mathbf{x}}^{(y_{1:i})}$, and the trace of the posterior covariance, $\sigma^{(y_{1:i})2}$, calculated from $\{\omega_1^{(y_{1:i})}, \dots, \omega_{N_e}^{(y_{1:i})}\}$ and $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_e}\}$, where each $\omega_n^{(y_{1:i})}$ reflects the likelihood of the original set of prior particles given $(y_1 y_2 \dots y_i)^T$. Therefore, coefficients derived for the update step reflect: 1) posterior calculations made independently of random sampling performed during serial processing of measurements [i.e., through $\bar{\mathbf{x}}^{(y_{1:i})}$ and $\sigma^{(y_{1:i})2}$]; and 2) the posterior particles resulting from bootstrap resampling near observations in state space. The derivation for \mathbf{r}_1 and \mathbf{r}_2 , which is provided in the appendix of P16, yields the following set of equations for each element of the vectors:

$$r_{1,j} = \sqrt{\frac{\sigma_j^{(y_{1:i})2}}{\frac{1}{N_e - 1} \sum_{n=1}^{N_e} \left[\mathbf{x}_{k_n,j}^{(y_{1:i-1})} - \bar{\mathbf{x}}_j^{(y_{1:i})} + c_j \left(\mathbf{x}_{n,j}^{(y_{1:i-1})} - \bar{\mathbf{x}}_j^{(y_{1:i})} \right) \right]^2}}, \tag{14}$$

$$r_{2,j} = c_j r_{1,j}. \tag{15}$$

Expressions for both coefficients contain the term c_j , which depends on the choice of vector weight formulation used to localize the impact of observations. For the P16 weight formulation given by (5),

$$c_j = \frac{N_e (1 - l[x_j, y_i, r_{1\text{loc}}])}{l[x_j, y_i, r_{1\text{loc}}] \sum_{m=1}^{N_e} p(y_i | \mathbf{x}_m)}. \tag{16}$$

As $l[x_j, y_i, r_{1\text{loc}}] \rightarrow 0$, P16 shows that $r_{1,j} \rightarrow 0$ and $r_{2,j} \rightarrow 1$. Equation (16), however, contains $\sum_{m=1}^{N_e} p(y_i | \mathbf{x}_m)$ in the denominator, which can cause the filter to be unstable when an observation yields very small likelihoods for all particles. This potential source of numerical instability follows from the weight formulation chosen in P16. Modifications to the vector weight equations have a direct effect on c_j alone, which we discuss in the remaining parts of this subsection.

The expression for c_j is derived in P16 after a simplification of its original form, given by

$$c_j = \frac{\bar{\mathbf{x}}_j^{(y_{1:i})} - \tilde{\bar{\mathbf{x}}}_j^{(y_{1:i})}}{\bar{\mathbf{x}}_j^{(y_{1:i-1})} - \tilde{\bar{\mathbf{x}}}_j^{(y_{1:i})}}, \tag{17}$$

where $\tilde{\bar{\mathbf{x}}}_j^{(y_{1:i})}$ is the posterior mean resulting from not localizing the impact of the current observation. As before, we approximate this mean with a weighted sum of particles. In this case, each vector weight is proportional to the product of the normalized likelihood for y_i , and the localized vector weights from previous observations in the sequence:

$$\tilde{\bar{\mathbf{x}}}_j^{(y_{1:i})} \approx \sum_{n=1}^{N_e} \frac{\hat{w}_n^{(y_i)} \omega_{n,j}^{(y_{1:i-1})}}{\hat{\Omega}_j^{(y_{1:i})}} \mathbf{x}_{n,j}. \tag{18}$$

A new expression for c_j is then given by

$$c_j = \frac{\sum_{n=1}^{N_e} \frac{\omega_{n,j}^{(y_{1:i})}}{\Omega_j^{(y_{1:i})}} \mathbf{x}_{n,j} - \sum_{n=1}^{N_e} \frac{\hat{w}_n^{(y_i)} \omega_{n,j}^{(y_{1:i-1})}}{\hat{\Omega}_j^{(y_{1:i})}} \mathbf{x}_{n,j}}{\sum_{n=1}^{N_e} \frac{\omega_{n,j}^{(y_{1:i-1})}}{\Omega_j^{(y_{1:i-1})}} \mathbf{x}_{n,j} - \sum_{n=1}^{N_e} \frac{\omega_{n,j}^{(y_{1:i})}}{\Omega_j^{(y_{1:i})}} \mathbf{x}_{n,j}}, \tag{19}$$

$$= \frac{\sum_{n=1}^{N_e} \left(\frac{\omega_{n,j}^{(y_{1:i})}}{\Omega_j^{(y_{1:i})}} - \frac{\hat{w}_n^{(y_i)} \omega_{n,j}^{(y_{1:i-1})}}{\hat{\Omega}_j^{(y_{1:i})}} \right) \mathbf{x}_{n,j}}{\sum_{n=1}^{N_e} \left(\frac{\omega_{n,j}^{(y_{1:i-1})}}{\Omega_j^{(y_{1:i-1})}} - \frac{\omega_{n,j}^{(y_{1:i})}}{\Omega_j^{(y_{1:i})}} \right) \mathbf{x}_{n,j}},$$

where $\Omega_j^{(y_{1:i})}$ normalizes $\omega_{n,j}^{(y_{1:i})}$ and $\hat{\Omega}_j^{(y_{1:i})}$ normalizes $\hat{w}_n^{(y_i)} \omega_{n,j}^{(y_{1:i-1})}$. An important difference between this formulation and the P16 formulation [see Eq. (A6) of P16] is that (19) expresses each mean using weighted sums of $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_e}\}$ instead of $\{\mathbf{x}_1^{(y_{1:i-1})}, \dots, \mathbf{x}_{N_e}^{(y_{1:i-1})}\}$. The revised calculations of c_j no longer rely on past updates

made to particles when assimilating $(y_1 y_2 \dots y_{i-1})^T$ in the sequence, thus improving the accuracy of this formulation over the one outlined in P16. This form of c_j is feasible because (11) uses normalized scalar weights $\hat{w}_n^{(y_i)}$, instead of the observation likelihoods directly in the vector weight calculations. As a result, the number of variables in (19) can be reduced by writing

$$\omega_{n_j}^{(y_{1:i})} = \omega_{n_j}^{(y_{1:i-1})} \left[\frac{(N_e \hat{w}_n^{(y_i)} - 1)l(y_i, x_j, r_{loc}) + 1}{N_e} \right]. \quad (20)$$

In this form, c_j depends only on weights updated by previous observations in the sequence, the current likelihoods, and the localization coefficients used for y_i . This subtle difference from the P16 formulation also means $\Omega_j^{(y_{1:i})}$ can be written as a function of $\hat{\Omega}_j^{(y_{1:i})}$, which further simplifies the expression. Using (20) and realizing that each term in the denominator of (19) divides each term in the numerator to yield the same constant for all n , c_j reduces to

$$c_j = \frac{1 - l(y_i, x_j, r_{loc})}{\hat{\Omega}_j^{(y_{1:i})} N_e l(y_i, x_j, r_{loc})}, \quad (21)$$

which requires calculating only $\hat{\Omega}_j^{(y_{1:i})}$ to update c_j for each observation during data assimilation. In addition to being more accurate, the new formulation is more stable than (16) because the denominator of (21) is less likely to approach zero in the neighborhood of observations.

c. Probability mapping step

The probability mapping step of the local PF provides additional correction when a mismatch exists between the weights and posterior particles resulting from sampling and merging steps. This procedure helps compensate for the fact that the local PF satisfies the SIR solution near observations, but considers only the first two moments outside this region. We perform this step independently for each state variable by mapping particles into kernel-estimated marginal probability distributions calculated using prior particles and their weights. As shown in P16, this step provides little benefit when prior errors are close to Gaussian, but improves filtering results when large deviations from Gaussianity occur.

The general strategy focuses on matching quantiles of an input distribution $g(x)$ with quantiles of some target distribution $q(x)$. For the local PF, we use particles resulting from the sampling and merging steps as input particles for the probability mapping. For simplicity,

refer to $\{x_1, \dots, x_{N_e}\}$ and $\{\tilde{x}_1, \dots, \tilde{x}_{N_e}\}$ as sets of scalar particles before and after the mapping.

The mapping performs an inverse of the target cumulative density function (cdf) at each particle quantile to sample from the desired distribution:

$$\tilde{x}_n = Q^{-1}[G(x_n)]. \quad (22)$$

Here, $G(x)$ and $Q(x)$ are the initial and target cdfs, respectively. As in McGinnis et al. (2015), we represent $g(x)$ and $q(x)$ nonparametrically with

$$g(x) = \frac{1}{N_e} \sum_{m=1}^{N_e} K(x_m, b_m), \quad (23)$$

$$q(x) = \sum_{m=1}^{N_e} w_m K(\hat{x}_m, b_m), \quad (24)$$

where $K(x_m, b_m)$ is a Gaussian kernel centered on the m th input particle x_m with a standard deviation or “bandwidth” of b_m and $w_m K(\hat{x}_m, b_m)$ is a weighted kernel centered on the n th original prior particle, denoted \hat{x}_m . The kernel-estimated $g(x)$ and $q(x)$ are then integrated numerically over a specified domain to estimate cdfs. In the current study, we adopt a less arduous approach, which estimates quantiles directly as sums of Gaussian cdfs. At the input particle locations, we calculate quantiles using

$$G(x_n) = \frac{1}{2N_e} \sum_{m=1}^{N_e} \left[1 + \operatorname{erf} \left(\frac{x_n - x_m}{\sqrt{2}b_m} \right) \right], \quad (25)$$

which is derived from integrating the sum of Gaussian kernels in (24). Likewise, we estimate the target cdf using

$$Q(x) = \frac{1}{2} \sum_{n=1}^{N_e} w_n \left[1 + \operatorname{erf} \left(\frac{x - \hat{x}_n}{\sqrt{2}b_n} \right) \right], \quad (26)$$

where x is a domain defined by the set of evenly spaced points between x_{\min} and x_{\max} . We then estimate the inverse of (26) at each quantile by interpolating $Q(x)$ to each $G(x_n)$, which provides the adjusted posterior particles, $\{\tilde{x}_1, \dots, \tilde{x}_{N_e}\}$. For these experiments, we set x_{\min} (x_{\max}) to be the minimum (maximum) particle value minus (plus) twice the difference between the maximum and minimum particle values, and use 500 points between x_{\min} and x_{\max} . The domain grid spacing [i.e., $(x_{\max} - x_{\min})/500$] is small enough to ensure no systematic decrease in filter accuracy occurs due to interpolation errors. We arrived at this configuration after performing multiple numerical experiments with increasingly more grid points for the domain (not shown).

The probability mapping preserves relationships across state variables, despite being performed independently for each element of \mathbf{x} . A part of this result comes from multivariate information being passed through the spatially smooth weights used to define the target cdfs during the mapping. We also find it necessary to use a constant bandwidth for all Gaussian kernels used in the cdf estimates to preserve this property. For experiments shown in this study, we set each bandwidth equal to the sample standard deviation of $\{x_1, \dots, x_{N_e}\}$ for all b_n used in the calculation of G and Q .

4. Improved PF inflation methodologies

Monte Carlo filters can grossly underestimate uncertainty when N_e is small and when errors in the dynamical model are quantified inaccurately or ignored (Moritz and Sutera 1981; Buizza et al. 1993; Houtekamer et al. 1996; Anderson and Anderson 1999; Poterjoy et al. 2014). These problems can be reduced, but not fully solved, by localization strategies. For example, EnKFs typically rely on heuristic methods for inflating prior or posterior error variance when ensemble error statistics underestimate root-mean-square errors (RMSEs) averaged over many sequential data assimilation cycles. This inflation can take several forms, such as multiplying ensemble perturbations by a factor greater than unity (Anderson and Anderson 1999), adding noise to samples (Houtekamer and Mitchell 2005), or relaxing a portion of the posterior sample update back to the prior sample (Zhang et al. 2004; Whitaker and Hamill 2012). Inflating the sample variance in a similar manner is nontrivial in a PF framework, since no parametric form is assumed for these errors. These strategies are also not guaranteed to prevent weights from collapsing onto a single particle when assimilating dense networks of accurate observations.

Previous studies (i.e., P16, PA16, and PSA17) enforce a minimum bound on likelihood calculations to make the local PF resilient to weight collapse. To implement this strategy, the localization function is multiplied by a scalar coefficient α , which is less than one and forces particle weights to be more uniform. Because the goal of this coefficient is to prevent spurious reduction in the posterior sample variance after resampling, we call this method “ α inflation.” For a univariate problem, this approach leads to a weight calculation of the following form:

$$w_n = [p(y|x_n) - 1]\alpha + 1. \quad (27)$$

As discussed in PA16, (27) has a minimum value of $1 - \alpha$, which stabilizes the filter when localization alone

is insufficient for preventing weight collapse with small samples. The α inflation strategy, however, is not effective when estimated observation errors are much smaller than prior errors; Lee and Majda (2016) demonstrate this deficiency using experiments with sparse, accurate observations. In this case, the probability of prior particles yielding likelihoods greater than $1 - \alpha$ is low, causing the filter to assign near equal weights to each particle. Therefore, the method begins ignoring observations when their accuracy increases. This deficiency motivates the use of a different filter stabilization strategy for small samples, which we introduce here. Like the previous technique, the new method acts to reduce the impact of observations that collapse particle weights if left untreated. Instead of enforcing a minimum value for likelihood calculations, we increase the assigned observation error variances by a factor β . For this study, we choose β based on a metric that is inversely proportional to the variance in particle weights called “effective sample size,” defined by $N_{\text{eff}} = [\sum w_n^2]^{-1}$ (Liu and Chen 1998). The new method, which we refer to as “ β inflation,” assumes posterior calculations made using small N_{eff} are prone to sampling error, so the benefits of maintaining $N_{\text{eff}} > 1$ outweigh the negative impact of artificially increasing the error variance assigned to potentially accurate measurements. This strategy follows similar approaches that perform a sampling error correction based on N_{eff} (e.g., Penny and Miyoshi 2016).

The inflation scheme also helps stabilize the filter when a mismatch exists between physical processes captured by observations and processes represented by numerical models, or when a persistent bias exists in the model or measurement operators. To cope with this issue, a number of past studies introduce adaptive observation error inflation methods (Geer and Bauer 2011; Okamoto et al. 2014; Zhu et al. 2016; Minamide and Zhang 2017), which emphasize the importance of nonstationary error statistics for observations measuring complex physical processes, such as all-sky radiance measurements. These studies formulate schemes for variational and Kalman filtering methods in a manner that is effective for a Gaussian data assimilation framework. Strategies of this type, however, are not always appropriate for PFs. For example, a single accurate observation can collapse the SIR PF when provided with a relatively small prior sample. Presented with the same data assimilation problem, an EnKF would produce a posterior sample containing a reasonably accurate mean with nonzero variance.

For the SIR PF, drawing samples from an appropriate proposal density is one solution for improving filter stability for small samples (e.g., Chorin et al. 2010;

Morzfeld et al. 2018). Unfortunately, constructing this density can be quite difficult for complex non-Gaussian applications, such as the convective-scale forecasting system presented later in this manuscript. Other approaches for improving filter stability include regularization (e.g., Casarin and Marin 2009), additive noise (e.g., Penny and Miyoshi 2016), and using an EnKF for assimilating observations that would lead to weight collapse (Lee and Majda 2016). In this study, we choose to inflate observation error variance based on N_{eff} to stabilize the filter, which amounts to a smaller local PF correction during data assimilation. The new algorithm operates only when N_{eff} falls below a target value of N_{eff}^t . For example, an observation with error drawn from a Gaussian would yield $w_n \propto \exp\{-[y - H(x_n)]^2/2\hat{\sigma}_y^2\}$, where $\hat{\sigma}_y^2 = \beta\sigma_y^2$. The method requires finding β so that the target N_{eff}^t is reached:

$$N_{\text{eff}}^t = \left(\sum_{n=1}^{N_e} \exp\left\{ \frac{-[y - H(x_n)]^2}{2\beta\sigma_y^2} \right\} \right)^2 \times \left(\sum_{n=1}^{N_e} \exp\left\{ \frac{-[y - H(x_n)]^2}{\beta\sigma_y^2} \right\} \right)^{-1}. \quad (28)$$

This inflation mechanism serves a similar purpose as α inflation; that is, it prevents weights from collapsing to a single particle under the circumstances described above. Unlike α inflation, the new approach does not completely neglect observations when all likelihoods are very small.

a. Univariate demonstration

The relative merits of β inflation and α inflation can be inferred easily from numerical simulations. For this purpose, we perform experiments estimating the posterior mean and variance for a univariate random variable x , conditioned on a noisy measurement y . We first draw samples $\{x_1, \dots, x_{N_e}\}$ and a true state x^t from a prior distribution $N(0, 1)$, and simulate y by adding an error drawn from $N(0, \sigma_y^2)$ to x^t . We choose $\sigma_y = 0.2$ so that the variance in prior particles is much larger than the estimated observation error variance, thus increasing the probability of particle weights collapsing to a single particle for small N_e . Each experiment uses the SIR PF to assimilate y using the two stability methods described above, with α coefficients from 0.9–1.0 and β coefficients from 1.0–16.0. These ranges cover a large enough portion of the parameter space to provide optimal results under a variety of circumstances. We perform these experiments 10^4 times to accumulate samples for calculating mean RMSE and sample standard deviation (spread) statistics for each data assimilation configuration.

We summarize the results in Fig. 2 by plotting average RMSEs and ratios of spread to RMSE (left panels) as a function of α and β . In this figure, the ordinate values reflect increasing impact of inflation (i.e., decreasing α and increasing β), with the origin reflecting no inflation. We also plot the ratio N_{eff}/N_e for each stability parameter (right panels in Fig. 2) to show how α and β impact N_{eff} for a given N_e . For the univariate Gaussian data assimilation problem, we can designate a “target inflation” as the α or β yielding an average spread to RMSE ratio of unity; coefficients producing this result are indicated by \times s in the figure. As demonstrated by the simulations, the two parameters can be tuned so the posterior variance matches the corresponding posterior mean RMSEs, and the target inflation decreases as N_e increases (values are shown for $N_e = 10, 20, 40,$ and 80). Both methods also have the expected property that modifying the likelihood calculations to prevent weight collapse increases the average RMSEs of the data assimilation experiments, in which case, we find β inflation to be a more effective strategy for maintaining filter stability because it provides the lowest RMSEs for target inflation values. Posterior RMSEs also exhibit less sensitivity to β , compared to α , demonstrating that choosing a sub-optimal value for this parameter is less likely to increase errors in the mean.

Repeating the experiments with a smaller observation error of $\sigma_y = 0.02$ shows even larger differences between the two strategies. In this case, α inflation becomes far less effective than β inflation for maintaining variance in the posterior particles without large increases in RMSEs (Fig. 3). The target β in each configuration also corresponds to a N_{eff}/N_e that is comparable to the values found in experiments for $\sigma_y = 0.2$. This result motivates our choice of using N_{eff} to tune the inflation for real applications, where measurements come from diverse observing networks with heterogeneous accuracy and sparsity. In addition, the results in Figs. 2 and 3 show that the relationship between β and N_{eff}/N_e is relatively independent of the sample size used, thus making it easier to specify a threshold value for N_{eff} .

b. Multivariate β inflation

For systems of multiple spatial dimensions and multiple observations, the collapse of particle weights at the location of a model variable can be caused by assimilating one or more distant observations. Provided with a network of observations located at different geographical locations, one strategy is to find the vector of inflation coefficients \mathbf{b} of length N_y , which gives $N_{\text{eff}} \geq N_{\text{eff}}^t$ for weights calculated at each observation

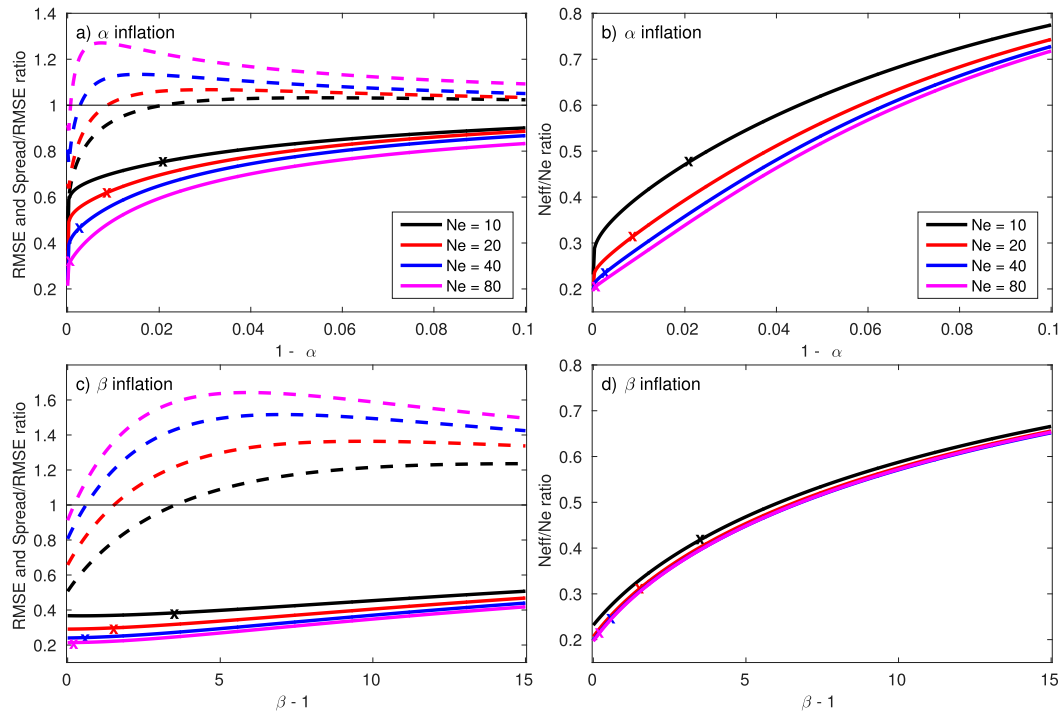


FIG. 2. Two filter stability techniques compared for a univariate application assimilating an observation with error $\sigma_y = 0.2$ for $N_e = 10$ (black), $N_e = 20$ (red), $N_e = 40$ (blue), and $N_e = 80$ (magenta). (a),(c) RMSEs (solid lines) and ratios of sample standard deviation to RMSE as a function of stability coefficient (dashed lines) for (top) α inflation and (bottom) β inflation. (b),(d) The ratio of N_{eff} to N_e as a function of inflation coefficient for each configuration.

location. This problem is trivial to solve when observations reflect independent portions of the state space, but the solution is no longer unique when one or more dependent variables are observed. In the latter case, finding a \mathbf{b} that satisfies this criteria also requires solving numerically for the set of inflation coefficients, since the normalization of weights and inclusion of localization complicates any analytical solutions to this problem. This process would require solving (28) multiple times for each β_i , which is computationally prohibitive for large data assimilation problems.

An effective heuristic approach is to solve (28) once for each observation to obtain the set $\{\tilde{\beta}_1, \dots, \tilde{\beta}_{N_y}\}$, then propagate the inflation values spatially using

$$\beta_i = 1 + \sum_{k=1}^{N_y} (\tilde{\beta}_k - 1) l(y_i, y_k, r_{\text{loc}}), \quad (29)$$

where the localization function $l(y_i, y_k, r_{\text{loc}})$ determines the impact of each $\tilde{\beta}_k$ value in the calculation of β_i . This method constructs each β_i by aggregating coefficients from nearby observations to determine the inflation needed for preventing weight collapse. As a result, the inflation passes the task of stabilizing the filter to

multiple sources, which we find to be an effective strategy for dense networks of accurate observations.

c. Filter degeneracy during resampling

The P16 local PF processes observations serially to produce posterior samples that reflect the localized particle weights. Because a resampling step is necessary for each independent observation in \mathbf{y} , successive sampling steps performed for closely spaced observations can lead to filter degeneracy faster than if they were assimilated at once. One solution is to assimilate these observations in batches, which is a straightforward option when provided with measurements located in close proximity to each other. For example, measurements of reflectivity and radial velocity can be assimilated simultaneously in weather models. This strategy also significantly reduces the computational cost of the local PF for dense observation networks and provides a means of assimilating observations with correlated errors.

The α inflation strategy adopted in previous studies provides another mechanism for reducing issues related to the serial processing of observations. In addition to preventing weight collapse, this method reduces part of the update formed by merging particles together for




FIG. 3. As in Fig. 2, but for $\sigma_y = 0.02$.

each observation, thus allowing more diversity in particles at observation locations. We maintain this strategy in the revised local PF by introducing a mixing parameter similar to what is used in Zhang et al. (2004) for EnKFs, which operates on the update coefficients \mathbf{r}_1 and \mathbf{r}_2 :

$$\mathbf{r}_1 \rightarrow \gamma \mathbf{r}_1, \tag{30}$$

$$\mathbf{r}_2 \rightarrow \gamma(\mathbf{r}_2 - 1) + 1. \tag{31}$$

The coefficient, γ is a scalar between 0 and 1, which forces the local PF to update particles using combinations of the current particles and resampled particles everywhere in state space, including at the location of observations. Following this modification of \mathbf{r}_1 and \mathbf{r}_2 , we center and scale posterior particles resulting from (13) to match $\bar{\mathbf{x}}^{(y_{1:i})}$ and $\sigma^{(y_{1:i})}$, so successive assimilation steps follow assumptions made in deriving the update equations presented in section 3a. For small N_e , setting $\gamma < 1$ is useful for reducing filter degeneracy due to sampling errors in the serial processing of observations, but setting $\gamma = 1$ is more appropriate for moderate to large N_e .

5. 40-variable Lorenz (1996) application

In this section, we evaluate updates made to the local PF algorithm using a set of data assimilation

experiments with the Lorenz (1996) model (hereafter L96). The dynamical system consists of $N_x = 40$ variables evolved in time using the differential equations,

$$\frac{dx_i}{dt} = (x_{i+1} - x_{i-2})x_{i-1} - x_i + F, \tag{32}$$

with $F = 8$, and cyclic boundaries: $x_{i+N_x} = x_i$ and $x_{i-N_x} = x_i$. We integrate (32) forward numerically using the fourth-order Runge–Kutta method with a time step of 0.05 time units. For the initial set of experiments, we create a sparse network of $N_y = 10$ observations, simulated every 0.05 time units on every fourth model grid point, by selecting values from a truth simulation and adding uncorrelated Gaussian errors selected from $N(0, \sigma_y^2 I)$. The test problem is similar to the setup described in sections 4a and 4b of P16, except we use half the number of measurements and space them equally across the domain. From this application, we can form a set of data assimilation problems that present little difficulty for most filtering techniques, like EnKFs, but reproduce known challenges for the P16 formulation of the local PF. We compare the local PF performance with α and β inflation using a fixed sample size of $N_e = 40$, and decreasing values of observation error standard deviations: $\sigma_y = 1$, $\sigma_y = 0.2$, and $\sigma_y = 0.02$. The objective is to test the new inflation and weight formulations for a set of problems that are increasingly more difficult for the previous filter




FIG. 4. Domain mean posterior RMSEs (thick solid lines) and spread (thick dashed lines) from experiments performed with the L96 model using observation errors of (a) $\sigma_y = 1$, (b) $\sigma_y = 0.2$, and (c) $\sigma_y = 0.02$. Values are plotted on a log scale for each data assimilation cycle to compare α inflation (red), β inflation (blue), and β inflation with new weight formulation (black). For reference, the thin dashed line indicates the observation error used for each experiment.

configuration (i.e., sparse, but accurate observation networks). The experiments constructed here are similar to those performed by Lee and Majda (2016), which demonstrate deficiencies in previous formulations of the local PF (see their Fig. 2).

Figure 4 shows posterior mean RMSEs and ensemble standard deviation averaged over all state variables every data assimilation cycle for 1000 cycles using α and β inflation. For these comparisons, we use the same local PF weight formulation as P16 to evaluate the two stability mechanisms, which we label “ α inflation” and “ β inflation” in Fig. 4. We also include a third case using β inflation with the new weight formulation described in section 3a, which is labeled “new weights.” Results shown in this figure reflect configurations of the local PF tuned to provide the lowest posterior RMSEs when averaged over 1000 cycles. In the α inflation case, we use $\alpha = 0.99$ for the $\sigma_y = 1$ experiment, but find the method to be ineffective for $\sigma_y = 0.2$ and $\sigma_y = 0.02$, regardless of the choice for α . For the second two β inflation cases, we use fixed values of $N_{\text{eff}}^l = 8$ and $\gamma = 0.5$ for all three observation errors. We also use the same localization length of $r_{\text{loc}} = 5$ for the two experiments with the P16 weight formulation, but find the new weight formulation to provide the best results when $r_{\text{loc}} = 3.6$. As shown in Fig. 1, the new weight formulation produces a broader update than the P16 formulation when provided with the same localization function, thus demonstrating why a smaller r_{loc} is necessary for this experiment.

The three configurations provide comparable skill when $\sigma_y = 1$ (Fig. 4a). This result occurs because the data assimilation application does not produce a situation that reveals the weaknesses of α inflation and the P16 weight formulation. When we decrease σ_y to 0.2, α inflation no longer provides an effective strategy for preventing filter divergence (Fig. 4b), because the probability of producing one or more particles from regions with high likelihood becomes very small.

The β inflation experiment takes several cycles to begin converging to the truth, but eventually produces steady filtering results with mean posterior RMSEs below the observation error (indicated by the dashed lines in Fig. 4). With the new weight formulation, the local PF begins converging to the truth after the first cycle, likely because the normalization of likelihoods no longer affects the spread of information in the localization region, as described in section 3a. Setting $\sigma_y = 0.02$ further slows convergence for the β inflation case, while having no negative impact on the experiment using the new weight formulation (Fig. 4c).

The results shown in Fig. 4 suggest that the filter can benefit substantially from the new weight formulation and from adaptively broadening the high likelihood region until particle weights maintain a threshold effective sample size. To perform a more rigorous test of this hypothesis, we carried out a larger set of sensitivity experiments using different observation networks for the L96 model. The observation networks consist of $N_y = 40, 20,$ and 10 equally spaced measurements—each collected at periods of 0.05, 0.1, and 0.2 time units (totaling 9 sets of measurements). The observations contain random errors sampled from $N(0, \sigma_y^2 I)$ with $\sigma_y = 0.2$, which is similar to the second case described above. We first perform a thorough tuning of each filter to identify values of localization and inflation yielding the lowest time- and domain-averaged posterior RMSEs for each observation network, then run each configuration for 5200 cycles (observation times), using the first 200 for spin up. Average RMSEs and spread resulting from the last 5000 cycles are summarized in Fig. 5.

The results in Fig. 5 confirm some of the findings discussed in the first set of experiments, where a progressively more accurate observation network yields a sudden decline in filter accuracy for the original P16 filter configuration. These experiments test a similar




FIG. 5. Time-averaged domain mean posterior RMSEs (solid lines) and spread (dashed lines) for L96 experiments as a function of observation period using (a) $N_y = 40$, (b) $N_y = 20$, and (c) $N_y = 10$. Values are plotted on a log scale for α inflation (red), β inflation (blue), and β inflation with new weight formulation (black). The thin dashed line indicates the observation error ($\sigma_y = 0.2$).

range of regimes, where the ratio of domain-average prior errors to measurement errors changes significantly over the nine sets of observation networks. All three configurations provide comparable RMSEs for the dense networks (i.e., $N_y = 40$), but the experiments demonstrate large discrepancies when fewer measurements are available to constrain the solution.

To further explore the behavior of each PF configuration, we carry out a rank histogram verification of the experiments performed with the sparsest observation network; that is, $N_y = 10$ with an observation period of 0.2 time units. The verification tallies the number of instances the true model state lands in bins

formed by ranking posterior particles in ascending order (Anderson 1996; Hamill and Colucci 1996; Talagrand et al. 1997). Figure 6 shows rank histograms calculated for the first and third variables of the L96 model (i.e., x_1 and x_3) over the 5000 cycles to demonstrate differences in posterior statistics for observed and nonobserved variables. The negative effects of α inflation are illustrated well by Fig. 6a. In addition to producing posterior samples of x_1 that are often overdispersed (indicated by an elevated frequency in the histogram center), the method frequently causes the local PF to lose track of the true solution (indicated by large frequency in the histogram edges for both Figs. 6a and 6d). The outcome is consistent with the low RMSEs calculated for this configuration (Fig. 5c). While β inflation reduces some of the deficiencies occurring in the α inflation experiments – regarding overdispersion – the true solution also lands outside the span of the particles too frequently for both x_1 and x_3 (Figs. 6b,e). The experiment performed with β inflation and the new weight configuration, however, yields relatively flat histograms for both variables, thus showing no obvious shortcomings in the probabilistic representation of the two variables. Therefore, the algorithmic changes introduced in section 3 are anticipated to yield benefits for real geophysical data assimilation problems, where dense networks of observations are not always available for constraining the most important dynamical processes.

6. Real weather application

The algorithmic changes introduced in this study are partially motivated by recent tests of the local PF for forecasting convective-scale weather events. In this section, we briefly describe results from a set of experiments demonstrating how each change affects the performance of the local PF for a real application. These tests use NSSL's Experimental Warn-on-forecast System for ensembles (NEWS-e) framework, which is a convective-scale ensemble analysis and prediction system developed to investigate whether the frequent assimilation of measurements for high-resolution weather models can eventually augment severe thunderstorm and tornado warnings (Stensrud et al. 2009, 2013). The NEWS-e system also acts as a test bed for implementing new data assimilation and modeling strategies that may eventually transition into future developments for operational weather forecasting systems, such as the NOAA High Resolution Rapid Refresh (HRRR; Benjamin et al. 2016). For this purpose, scientists at NSSL adopted an EnKF data assimilation system, which provides the required balance between filter

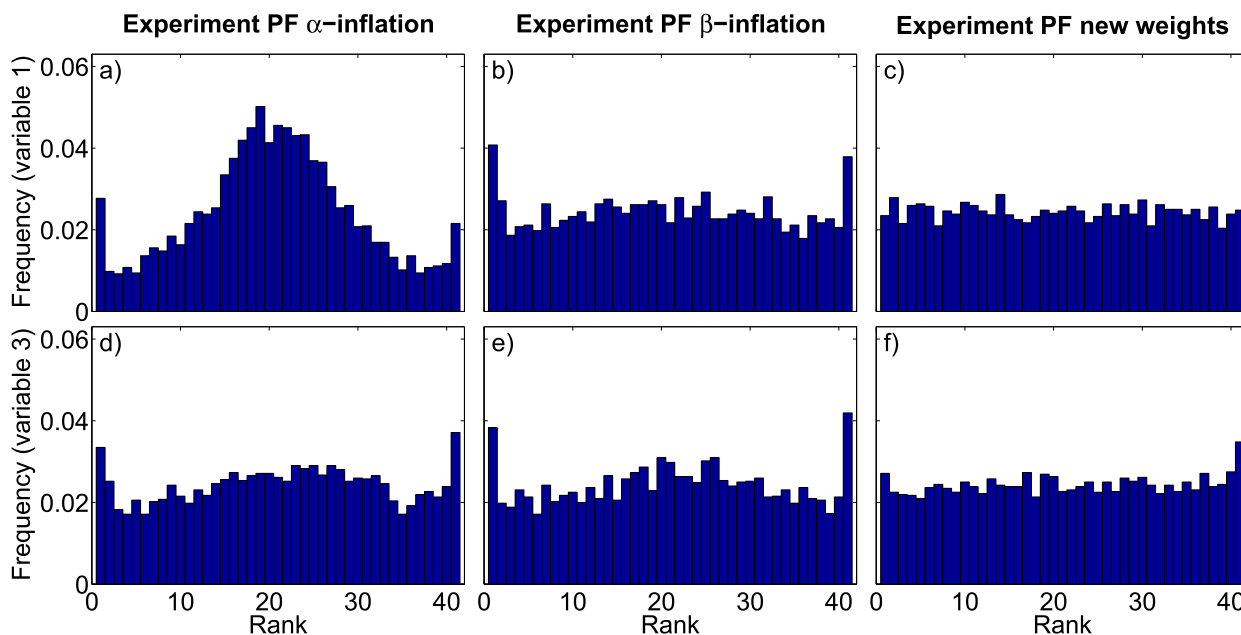


FIG. 6. Rank histograms generated for (a)–(c) variable 1 and (d)–(f) variable 3 for (left) α inflation, (middle) β inflation, and (right) β inflation with new weight formulation, using $N_y = 10$ and an observation period of 0.2.

performance and computational efficiency needed for weather forecasting on short (0–90 min) time scales.

The NEWS-e system has been run experimentally each May since 2015 to provide a real-time rapidly updating ensemble of convective storm forecasts (Wheatley et al. 2015; Jones et al. 2016; Lawson et al. 2018) during NOAA’s Hazardous Weather Testbed experiment (see <https://hwt.nssl.noaa.gov>). This system uses version 3.6.1 of the National Center for Atmospheric Research (NCAR) Weather Research and Forecasting (WRF) Model (Skamarock et al. 2008), and the NCAR Data Assimilation Research Testbed (DART) software package (Anderson et al. 2009) with the ensemble adjustment Kalman filter (EAKF) introduced by Anderson (2001). The ensemble is comprised of 36 members, run at a convection-permitting 3-km horizontal grid spacing with 51 vertical levels and a model top at 15 hPa over a 750 km \times 750 km domain. This domain is relocated each day to regions where severe weather is expected. The NEWS-e EAKF assimilates radar velocity and reflectivity, Oklahoma mesonet observations, and cloud water path retrievals every 15 min, starting from the 1800 UTC experimental NOAA HRRR ensemble, and ending at 0300 UTC the next day. Using this setup, the EAKF posterior initializes 18-member ensemble forecasts every 30 min during select convective weather outbreaks over the western plains of the United States. For additional details regarding the NEWS-e system, including the configuration

of the EAKF, we refer readers to Wheatley et al. (2015) and Jones et al. (2016).

The local PF is included in the suite of data assimilation methods available in DART, which allows for a seamless transition from the EAKF to local PF in the NEWS-e system. In this section, we briefly discuss experiments testing the stability and accuracy of the revised local PF algorithm with NEWS-e. These data assimilation experiments focus on a single high-impact event that produced several tornadoes over western Kansas on 24 May 2016. On this day, multiple convective cells formed along a dryline late in the afternoon, before organizing into a mesoscale convective system (MCS) by early evening (Fig. 7). Figure 7 shows composite reflectivity observations at three times representative of major changes in the MCS evolution, which are plotted over the same domain used for data assimilation and forecasting experiments.

As indicated by black markers in Fig. 7, a radar located in Dodge City, Kansas, captured large portions of the MCS’s upscale development from convective cells. We calculate root-mean-square differences (RMSDs) using the Dodge City radial velocity observations and ensemble members projected into observation space to verify predictions generated during cycling data assimilation experiments. These predictions come from 18-member ensemble forecasts, run every 30 min from 2230 UTC 24 May to 0300 UTC 25 May. The verification period covers the full development of the MCS targeted




FIG. 7. Composite reflectivity every 5 dBZ at three times during the evolution of the mesoscale convective system targeted for data assimilation experiments with the WRF model. The radial coverage of verifying observations from the Dodge City radar is indicated by the black dots.

for the experiments, and provides a suitable sample of model forecasts for testing whether changes made to the local PF algorithm are appropriate for high-dimensional geophysical applications, such as weather forecasting. We take a spatial mean of these statistics, calculated over the verifying region, and average the values over all forecasts to quantify a mean time evolution of errors over the life cycle of the event. The resulting verification summarizes how each data assimilation configuration affects predictions for the flow field in and around the developing MCS at different forecast lead times.

When applied in the NEWS-e system, initial tests of the P16 version of the local PF revealed several of the weaknesses described above. The α inflation mechanism, combined with the original weight formulation of the local PF, produces a filter that is unable to maintain adequate RMSE/spread statistics in the presence of large biases in the model, measurement operators, and observations. The configuration resulted in frequent low likelihood calculations for all particles (see section 4)

and a tendency of the filter to ignore most observations available during the experiment. These factors often produced nonsmooth weights (see section 3), which caused large imbalances in model initial conditions for some particles, and the occasional failure of particle model integrations between cycles. For these reasons, we do not show results with this configuration. Only after adopting the observation error inflation strategy introduced in section 4b (β inflation), were we able to achieve stable filtering results over the entire weather event. For the experiments shown here, we estimate β coefficients adaptively using an $N_{\text{eff}}^t = 0.2N_e$ and set $\gamma = 0.5$. We arrive at these parameters after performing a limited number of sensitivity experiments with the local PF, which are guided by results performed using low-order models.

Figure 8 shows time series plots of radar wind RMSDs and expected errors, which we calculate from 0–90-min ensemble forecasts as described previously in this section. The expected error is taken as the square root of




FIG. 8. Domain mean RMSDs to radar winds (solid lines) and expected observation-space errors (dashed lines), averaged over WRF forecasts initialized during the verification period. Values are plotted for the benchmark NEWS-e EAKF (gold), the local PF with β inflation (red), the local PF with β inflation and new weight formulation (blue), and the local PF after tuning (black).

the sum of ensemble variance and observation error variance at each observation location. In addition to comparing different formulations of the local PF, we provide results from the real-time EAKF system used in NEWS-e. The EAKF has been tuned for severe convective storm applications at NSSL, thus providing an appropriate benchmark for testing whether the local PF provides satisfactory filtering performance. Before comparing RMSDs from each configuration, it is important to note that the EAKF and local PF forecasts provide noticeable differences in expected forecast error. Ensemble forecasts verified in Fig. 8 consistently underestimate the average prediction errors during verification because of unresolved processes in the forecast model and the small sample size used to estimate this uncertainty. The local PF experiments, however, tend to provide larger expected errors than the EAKF experiment, partly due to the use of β inflation. While comparable observation error inflation techniques exist for ensemble Kalman filters, we do not deviate from the EAKF configuration currently used by the real-time NEWS-e system.

After performing the experiments, we find progressive improvements in the local PF forecasts from three different changes in configuration. First, we test the P16 weight formulation with the current NEWS-e localization length scales, which are tuned for the EAKF. This experiment produces the results plotted in red (Fig. 8), which yield slightly larger forecast RMSDs than the EAKF in gold. We then run an experiment with the modified local PF weight equations described in sections 3a and 3b, which produces the RMSDs plotted in blue. This configuration results in forecast errors that are comparable to the benchmark EAKF. Finally, we tune the localization length scales in the local PF to arrive at the black RMSDs. For several observation networks, this tuning resulted in a 75% reduction in the current r_{loc} used by the EAKF. After introducing these updates, the resulting forecasts produce smaller wind RMSDs than all other configurations of the local PF, which further suggests the revised local PF algorithm is appropriate for geophysical data assimilation.

The verification shows the final configuration of the local PF outperforming the benchmark EAKF experiment in these tests. Though not shown here, we repeated the EAKF experiments with the same reduced localization length scales used for the “tuned” configuration of the local PF, but did not find a similar reduction in RMSEs. Nevertheless, the limited number of cases used in this demonstration does not provide enough evidence to conclude the local PF performs better than the EAKF for the given application. The results, however, suggest

that the revised filter operates effectively for complex high-dimensional problems with small ensembles, and can provide comparable results to current techniques used for data assimilation in weather models.

7. Discussion and conclusions

This paper summarizes recent progress toward the development of a Bayesian filter for data assimilation in geophysics. The method discussed in this study is the local particle filter (PF) of Poterjoy (2016), which adopts sequential importance resampling techniques from PFs (Gordon et al. 1993), and localization strategies first used for ensemble Kalman filters (Houtekamer and Mitchell 2001; Hamill and Whitaker 2001) to construct a nonlinear filter that operates effectively for applications with large spatial domains. Since Poterjoy (2016), the local PF has been applied for a hierarchy of dynamical systems, including multiple high-dimensional geophysical models. In particular, Poterjoy and Anderson (2016) compare the local PF with deterministic and perturbed observation ensemble Kalman filters for generating posterior representations of baroclinic Rossby waves in a simplified atmospheric general circulation model. Following this study, Poterjoy et al. (2017) apply the local PF for an idealized squall line in the Weather Research and Forecasting (WRF) Model to examine its potential for convective-scale data assimilation and forecasting. Finally, the analysis and prediction experiments presented here within NSSL’s NEWS-e framework demonstrate strong potential for applying the local PF for severe convective storms in the future.

Numerical experiments performed up to this point motivate several algorithmic improvements to the local PF, including: 1) a new formulation of localized weights and filter update equations; 2) a more efficient probability mapping procedure; and 3) new filter stabilization methods for situations where localization is insufficient for preventing particle weight collapse. We use numerical simulations, ranging from a univariate Gaussian problem to a real weather forecasting application in the WRF model, to justify these improvements. In general, the changes introduced in this study improve the local PF’s stability in situations where sample variance in prior particles is much larger than the observation error variance and when unknown errors exist in the model, observation operators, and estimation of measurement uncertainty. These situations occur frequently in geophysical filtering problems, such as convective-scale data assimilation for weather models, and lead to round off errors in the original localized weight formulation and failure of the filter stabilization

technique used in past studies with the local PF (Poterjoy 2016; Poterjoy and Anderson 2016; Poterjoy et al. 2017). We also emphasize that the algorithmic changes introduced in this study are applicable to other PF methods that may adopt a similar type of localization strategy.

Real data tests performed with an experimental convective-scale forecasting system at NOAA NSSL demonstrate the potential of the updated local PF algorithm for numerical weather prediction. Despite using only 36 particles, forecasts generated from the local PF are about as accurate as forecasts generated from an EAKF system tuned over multiple seasons for this application. Future research will focus on a more thorough analysis of the local PF for convective-scale forecasting and its comparison to the real-time EAKF system run at NOAA NSSL. This research also provides additional incentive to explore large ensemble simulations of severe convective storms using the local PF, where higher-order posterior errors can be examined more faithfully.

Acknowledgments. Funding was provided by NOAA/Office of Oceanic and Atmospheric Research under NOAA–University of Oklahoma Cooperative Agreement NA11OAR4320072, U.S. Department of Commerce. Parts of this research was also performed while the first author held an National Research Council Research Associateship award at the NOAA/Atlantic Oceanographic and Atmospheric Laboratory. The first author thanks Jason Sippel and Altug Aksoy for providing comments that improved the clarity of the manuscript.

APPENDIX A

List of Symbols Used in Manuscript

Symbol	Description
\mathbf{x}	Model state vector
\mathbf{y}	Observation vector
N_x	Length of model state vector
N_y	Length of observation state vector
N_e	Ensemble size
ε	Observation error
σ_y^2	Observation error variance
$\bar{\mathbf{x}}^{(y_{1:i})}$	Posterior mean following assimilation of $[y_1, \dots, y_i]$
$\sigma_{(y_{1:i})}^2$	Trace of model state error covariance following assimilation of $[y_1, \dots, y_i]$
\mathbf{x}_n	n th prior particle
$\mathbf{x}_n^{(y_{1:i})}$	n th particle following assimilation of $[y_1, \dots, y_i]$

$\hat{w}_n^{(y_i)}$	n th scalar weight calculated from $p(y_i \mathbf{x}_n)$
$\tilde{w}_n^{(y_{1:i})}$	n th scalar weight calculated from $p(y_i \mathbf{x}_n^{(y_{1:i})})$
$\boldsymbol{\omega}^{(y_{1:i})}$	$N_e \times N_x$ matrix storing localized particle weighting vectors in each row, following assimilation of $[y_1, \dots, y_i]$
$\boldsymbol{\omega}_n^{(y_{1:i})}$	n th row of $\boldsymbol{\omega}^{(y_{1:i})}$
$\boldsymbol{\Omega}$	Normalization vector for rows of $\boldsymbol{\omega}$
$\hat{\boldsymbol{\Omega}}$	Normalization vector for rows of $\hat{\boldsymbol{\omega}}$
$l(y_i, x_j, r_{loc})$	Localization coefficient for model variable x_j during assimilation of y_i
r_{loc}	Localization length scale parameter
\mathbf{r}_1	Weighting vector for sampled particles
\mathbf{r}_2	Weighting vector for prior particles
$G(x)$	Input cdf for probability mapping
$Q(x)$	Target cdf for probability mapping
b_m	Kernel bandwidth for m th particle
N_{eff}	Effective ensemble size
β	Observation error inflation coefficient
γ	Mixing parameter

APPENDIX B

Local PF Algorithm

This appendix provides a pseudocode description of the revised local PF for the case of Gaussian observation errors. Modifying the algorithm for other forms of observation error distributions requires changing the observation error inflation and weight calculation steps to match the desired error distribution.

Algorithm 1 Local particle filter algorithm
Input: Initial weighting matrix $\boldsymbol{\omega}$ with elements $\omega_{n,j} = 1/N_e$ for $n = 1, \dots, N_e$, $j = 1, \dots, N_x$, prior ensemble $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_e}\}$, observation vector \mathbf{y} , observation error covariance matrix with diagonal elements $[\sigma_{y_1}^2, \dots, \sigma_{y_{N_e}}^2]$, N_{eff} for observation error inflation, localization length scale r_{loc} , mixing parameter γ , CDF domain for probability mapping \mathbf{x}_d , and kernel bandwidth parameters $\{b_1, \dots, b_{N_e}\}$.
Observation error inflation coefficients for observations with Gaussian errors
for $i = 1:N_y$ **do**

$$\hat{\beta}_i = \arg \min_{\beta} \left[N_{eff}^i - \left(\sum_{n=1}^{N_e} \exp \left\{ \frac{-[y_i - H(\mathbf{x}_n)]^2}{2\beta\sigma_{y_i}^2} \right\} \right)^2 \right. \\ \left. \times \left(\sum_{n=1}^{N_e} \exp \left\{ \frac{-[y_i - H(\mathbf{x}_n)]^2}{\beta\sigma_{y_i}^2} \right\} \right)^{-1} \right]$$

for $k = 1:N_y$ **do**

$$\beta_k \leftarrow \beta_k + (\hat{\beta}_i - 1)l(y_i, y_k, r_{\text{loc}})$$

end for

end for

Sequential assimilation loop

for $i = 1:N_y$ **do**

Calculate scalar weights for Gaussian likelihoods

for $n = 1:N_e$ **do**

$$\hat{w}_n^{(y_i)} \leftarrow \exp\left\{\frac{-[y_i - H(\mathbf{x}_n)]^2}{2\beta_i\sigma_{y_i}^2}\right\}$$

$$\tilde{w}_n^{(y_{1:i})} \leftarrow \exp\left\{\frac{-[y_i - H(\mathbf{x}_n^{(y_{1:i-1})})]^2}{2\beta_i\sigma_{y_i}^2}\right\}$$

end for

$$\hat{W}^{(y_i)} \leftarrow \sum_{n=1}^{N_e} \hat{w}_n^{(y_i)}, \quad \tilde{W}^{(y_{1:i})} \leftarrow \sum_{n=1}^{N_e} \tilde{w}_n^{(y_{1:i})}$$

for $n = 1:N_e$ **do**

$$\hat{w}_n^{(y_i)} \leftarrow \frac{\hat{w}_n^{(y_i)}}{\hat{W}^{(y_i)}}$$

$$\tilde{w}_n^{(y_{1:i})} \leftarrow \frac{\tilde{w}_n^{(y_{1:i})}}{\tilde{W}^{(y_{1:i})}}$$

end for

Perform resampling

Draw N_e particles from $\{\mathbf{x}_1^{(y_{1:i-1})}, \dots, \mathbf{x}_{N_e}^{(y_{1:i-1})}\}$ according to the weights $\{\tilde{w}_1^{(y_{1:i})}, \dots, \tilde{w}_{N_e}^{(y_{1:i})}\}$.

Assign indices k_1, \dots, k_{N_e} to sampled particles.

for $j = 1:N_x$ **do**

Calculate normalization needed for update step

$$\hat{\Omega}_j^{(y_{1:i})} \leftarrow \sum_{n=1}^{N_e} \hat{w}_n^{(y_i)} \omega_{n,j}^{(y_{1:i-1})}$$

Calculate localized particle weights

for $n = 1:N_e$ **do**

$$\omega_{n,j}^{(y_{1:i})} \leftarrow \omega_{n,j}^{(y_{1:i-1})} \left[\frac{(N_e \hat{w}_n^{(y_i)} - 1)l(y_i, \mathbf{x}_j, r_{\text{loc}}) + 1}{N_e} \right]$$

end for

$$\Omega_j^{(y_{1:i})} \leftarrow \sum_{n=1}^{N_e} \omega_{n,j}^{(y_{1:i})}$$

for $n = 1:N_e$ **do**

$$\omega_{n,j}^{(y_{1:i})} \leftarrow \frac{\omega_{n,j}^{(y_{1:i})}}{\Omega_j^{(y_{1:i})}}$$

end for

Calculate weighted posterior mean and variance

$$\bar{x}_j^{(y_{1:i})} \leftarrow \sum_{n=1}^{N_e} \omega_{n,j}^{(y_{1:i})} x_{n,j}$$

$$\sigma_j^{(y_{1:i})2} \leftarrow \frac{1}{1 - \sum_{n=1}^{N_e} \omega_{n,j}^{(y_{1:i})2}} \sum_{n=1}^{N_e} \omega_{n,j}^{(y_{1:i})} [x_{n,j} - \bar{x}_j^{(y_{1:i})}]^2$$

Calculate vectors needed for particle update

$$c_j \leftarrow \frac{1 - l(y_i, x_j, r_{\text{loc}})}{\hat{\Omega}_j^{(y_{1:i})} N_e l(y_i, x_j, r_{\text{loc}})}$$

$$r_{1,j} \leftarrow \sqrt{\frac{\sigma_j^{(y_{1:i})2}}{\frac{1}{N_e - 1} \sum_{n=1}^{N_e} [x_{k_n,j}^{(y_{1:i-1})} - \bar{x}_j^{(y_{1:i})}]^2 + c_j (x_{k_n,j}^{(y_{1:i-1})} - \bar{x}_j^{(y_{1:i})})^2}}$$

$$r_{2,j} \leftarrow c_j r_{1,j}$$

Update particles

$$x_{n,j}^{(y_i)} \leftarrow \bar{x}_j^{(y_{1:i})} + \gamma r_{1,j} (x_{k_n,j}^{(y_{1:i-1})} - \bar{x}_j^{(y_{1:i})}) + [\gamma(r_{2,j} - 1) + 1] (x_{n,j}^{(y_{1:i-1})} - \bar{x}_j^{(y_{1:i})})$$

end for

end for

Probability mapping**for** $j = 1:N_x$ **do**

Estimate posterior CDF

$$Q(x) \leftarrow \frac{1}{2} \sum_{n=1}^{N_e} \omega_{n,j}^{(y_{1,j})} \left[1 + \operatorname{erf} \left(\frac{\mathbf{x}_d - x_{n,j}}{\sqrt{2}b_n} \right) \right]$$

for $n = 1:N_e$ **do**

Calculate particle quantiles in kernel-estimated prior CDF

$$G(x_n) \leftarrow \frac{1}{2N_e} \sum_{m=1}^{N_e} \left[1 + \operatorname{erf} \left(\frac{x_{n,j}^{(y_i)} - x_{m,j}^{(y_i)}}{\sqrt{2}b_m} \right) \right]$$

Interpolate to find particle located at $G(x_n)$ in posterior CDF

$$x_{n,j}^{(y_i)} \leftarrow \operatorname{interp}[Q(x), G(x_n)]$$

end for**end for**

REFERENCES

- Anderson, J. L., 1996: A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Climate*, **9**, 1518–1530, [https://doi.org/10.1175/1520-0442\(1996\)009<1518:AMFPAE>2.0.CO;2](https://doi.org/10.1175/1520-0442(1996)009<1518:AMFPAE>2.0.CO;2).
- , 2001: An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.*, **129**, 2884–2903, [https://doi.org/10.1175/1520-0493\(2001\)129<2884:AEAKFF>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<2884:AEAKFF>2.0.CO;2).
- , and S. L. Anderson, 1999: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev.*, **127**, 2741–2758, [https://doi.org/10.1175/1520-0493\(1999\)127<2741:AMCIOT>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<2741:AMCIOT>2.0.CO;2).
- , T. Hoar, K. Raeder, H. Liu, N. Collins, R. Torn, and A. Avellano, 2009: The Data Assimilation Research Testbed: A community facility. *Bull. Amer. Meteor. Soc.*, **90**, 1283–1296, <https://doi.org/10.1175/2009BAMS2618.1>.
- Bengtsson, T., C. Snyder, and D. Nychka, 2003: Toward a nonlinear ensemble filter for high-dimensional systems. *J. Geophys. Res.*, **108**, 8775, <https://doi.org/10.1029/2002JD002900>.
- , P. Bickel, and B. Li, 2008: Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. *Probability and Statistics: Essays in Honor of David A. Freedman*, D. Nolan and T. Speed, Eds., Vol. 2, Institute of Mathematical Statistics, 316–334.
- Benjamin, S. G., and Coauthors, 2016: A North American hourly assimilation and model forecast cycle: The Rapid Refresh. *Mon. Wea. Rev.*, **144**, 1669–1694, <https://doi.org/10.1175/MWR-D-15-0242.1>.
- Bickel, P., B. Li, and T. Bengtsson, 2008: Sharp failure rates for the bootstrap particle filter in high dimensions. *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, B. S. Clarke and S. Ghosal, Eds., Vol. 3, Institute of Mathematical Statistics, 318–329.
- Buizza, R., J. Tribbia, F. Molteni, and T. Palmer, 1993: Computation of optimal unstable structures for a numerical weather prediction model. *Tellus*, **45A**, 388–407, <https://doi.org/10.3402/tellusa.v45i5.14901>.
- Casarin, R., and J.-M. Marin, 2009: Online data processing: Comparison of Bayesian regularized particle filters. *Electron. J. Stat.*, **3**, 239–258, <https://doi.org/10.1214/08-EJS256>.
- Chorin, A., M. Morzfeld, and X. Tu, 2010: Implicit particle filters for data assimilation. *Commun. Appl. Math. Comput. Sci.*, **5**, 221–240, <https://doi.org/10.2140/ca.mcos.2010.5.221>.
- Chustagulprom, N., S. Reich, and M. Reinhardt, 2016: A hybrid ensemble transform particle filter for nonlinear and spatially extended dynamical systems. *SIAM/ASA J. Uncertainty Quantif.*, **4**, 592–608, <https://doi.org/10.1137/15M1040967>.
- Doucet, A., S. Godsill, and C. Andrieu, 2000: Sequential Monte Carlo methods for Bayesian filtering. *Stat. Comput.*, **10**, 197–208, <https://doi.org/10.1023/A:1008935410038>.
- , N. de Freitas, and N. Gordon, Eds., 2001: An introduction to sequential Monte Carlo methods. *Sequential Monte Carlo Methods in Practice*, Springer-Verlag, 2–14.
- Frei, M., and H. R. Künsch, 2013: Bridging the ensemble Kalman and particle filters. *Biometrika*, **100**, 781–800, <https://doi.org/10.1093/biomet/ast020>.
- Gaspari, G., and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, **125**, 723–757, <https://doi.org/10.1002/qj.49712555417>.
- Geer, A. J., and P. Bauer, 2011: Observation errors in all-sky data assimilation. *Quart. J. Roy. Meteor. Soc.*, **137**, 2024–2037, <https://doi.org/10.1002/qj.830>.
- Gordon, N. J., D. J. Salmond, and A. F. M. Smith, 1993: Novel approach to nonlinear/non-Gaussian state estimation. *IEE Proc., F, Radar Signal Process.*, **140**, 107–113, <https://doi.org/10.1049/ip-f-2.1993.0015>.
- Hamill, T. M., and S. J. Colucci, 1996: Random and systematic error in NMC's short-range Eta ensembles. Preprints, *13th Conf. on Probability and Statistics in the Atmospheric Sciences*, San Francisco, CA, Amer. Meteor. Soc., 51–56.
- Hamill, T. H., and J. S. Whitaker, 2001: Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.*, **129**, 2776–2790, [https://doi.org/10.1175/1520-0493\(2001\)129<2776:DDFOBE>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<2776:DDFOBE>2.0.CO;2).
- Houtekamer, P. L., and H. L. Mitchell, 2001: A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, **129**, 123–137, [https://doi.org/10.1175/1520-0493\(2001\)129<0123:ASEKFF>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<0123:ASEKFF>2.0.CO;2).
- , and —, 2005: Ensemble Kalman filtering. *Quart. J. Roy. Meteor. Soc.*, **131**, 3269–3289, <https://doi.org/10.1256/qj.05.135>.
- , L. Lefavre, J. Derome, H. Ritchie, and H. L. Mitchell, 1996: A system simulation approach to ensemble prediction. *Mon. Wea. Rev.*, **124**, 1225–1242, [https://doi.org/10.1175/1520-0493\(1996\)124<1225:ASSATE>2.0.CO;2](https://doi.org/10.1175/1520-0493(1996)124<1225:ASSATE>2.0.CO;2).
- Jones, T. A., K. Knopfmeier, D. Wheatley, G. Creager, P. Minnis, and R. Palikonda, 2016: Storm-scale data assimilation and ensemble forecasting with the NSSL experimental Warn-on-Forecast system. Part II: Combined radar and satellite data experiments. *Wea. Forecasting*, **31**, 297–327, <https://doi.org/10.1175/WAF-D-15-0107.1>.
- Lawson, J. R., J. S. Kain, N. Yussouf, D. C. Dowell, D. M. Wheatley, K. H. Knopfmeier, and T. A. Jones, 2018: Advancing from

- convection-allowing NWP to Warn-on-Forecast: Evidence of progress. *Wea. Forecasting*, **33**, 599–607, <https://doi.org/10.1175/WAF-D-17-0145.1>.
- Lee, Y., and A. J. Majda, 2016: State estimation and prediction using clustered particle filters. *Proc. Natl. Acad. Sci. USA*, **113**, 14 609–14 614, <https://doi.org/10.1073/pnas.1617398113>.
- Lei, J., and P. Bickel, 2011: A moment matching ensemble filter for nonlinear non-Gaussian data assimilation. *Mon. Wea. Rev.*, **139**, 3964–3973, <https://doi.org/10.1175/2011MWR3553.1>.
- Liu, J. S., and R. Chen, 1998: Sequential Monte Carlo methods for dynamic systems. *J. Amer. Stat. Assoc.*, **93**, 1032–1044, <https://doi.org/10.1080/01621459.1998.10473765>.
- Lorenz, E. N., 1996: Predictability: A problem partly solved. *Proc. Seminar on Predictability*, Vol. 1, Reading, United Kingdom, ECMWF, <https://www.ecmwf.int/node/10829>.
- Majda, A. J., D. Qi, and T. P. Sapsis, 2014: Blended particle filters for large-dimensional chaotic dynamical systems. *Proc. Natl. Acad. Sci. USA*, **111**, 7511–7516, <https://doi.org/10.1073/pnas.1405675111>.
- McGinnis, S., D. Nychka, and L. O. Mearns, 2015: A new distribution mapping technique for climate model bias correction. *Machine Learning and Data Mining Approaches to Climate Science*, V. Lakshmanan et al., Eds., Springer, 91–99 https://doi.org/10.1007/978-3-319-17220-0_9.
- Minamide, M., and F. Zhang, 2017: Adaptive observation error inflation for assimilating all-sky satellite radiance. *Mon. Wea. Rev.*, **145**, 1063–1081, <https://doi.org/10.1175/MWR-D-16-0257.1>.
- Moritz, R. E., and A. Sutera, 1981: The predictability problem: Effects of stochastic perturbations in multiequilibrium systems. *Advances in Geophysics*, Vol. 23, Academic Press, 345–383, [https://doi.org/10.1016/S0065-2687\(08\)60333-7](https://doi.org/10.1016/S0065-2687(08)60333-7).
- Morzfeld, M., D. Hodys, and J. Poterjoy, 2018: Variational particle smoothers and their localization. *Quart. J. Roy. Meteor. Soc.*, **144**, 806–825, <https://doi.org/10.1002/qj.3256>.
- Okamoto, K., A. P. McNally, and W. Bell, 2014: Progress towards the assimilation of all-sky infrared radiances: An evaluation of cloud effects. *Quart. J. Roy. Meteor. Soc.*, **140**, 1603–1614, <https://doi.org/10.1002/qj.2242>.
- Penny, S. G., and T. Miyoshi, 2016: A local particle filter for high dimensional geophysical systems. *Nonlinear Processes Geophys.*, **23**, 391–405, <https://doi.org/10.5194/npg-23-391-2016>.
- Poterjoy, J., 2016: A localized particle filter for high-dimensional nonlinear systems. *Mon. Wea. Rev.*, **144**, 59–76, <https://doi.org/10.1175/MWR-D-15-0163.1>.
- , and J. L. Anderson, 2016: Efficient assimilation of simulated observations in a high-dimensional geophysical system using a localized particle filter. *Mon. Wea. Rev.*, **144**, 2007–2020, <https://doi.org/10.1175/MWR-D-15-0322.1>.
- , F. Zhang, and Y. Weng, 2014: The effects of sampling errors on the EnKF assimilation of inner-core hurricane observations. *Mon. Wea. Rev.*, **142**, 1609–1630, <https://doi.org/10.1175/MWR-D-13-00305.1>.
- , R. A. Sobash, and J. L. Anderson, 2017: Convective-scale data assimilation for the Weather Research and Forecasting Model using the local particle filter. *Mon. Wea. Rev.*, **145**, 1897–1918, <https://doi.org/10.1175/MWR-D-16-0298.1>.
- Potthast, R., 2016: On ensemble and particle filters for large-scale data assimilation and inverse problems. *Mathematical and Algorithmic Aspects of Data Assimilation in the Geosciences*, Oberwolfach, Germany, Mathematisches Forschungsinstitut Oberwolfach, 2712–2714, <https://doi.org/10.4171/OWR/2016/47>.
- Robert, S., D. Leuenberger, and H. R. Künsch, 2018: A local ensemble transform Kalman particle filter for convective-scale data assimilation. *Quart. J. Roy. Meteor. Soc.*, **144**, 1279–1296, <https://doi.org/10.1002/qj.3116>.
- Skamarock, W. C., and Coauthors, 2008: A description of the Advanced Research WRF version 3. NCAR Tech. Note NCAR/TN-475+STR, 113 pp., <https://doi.org/10.5065/D68S4MVH>.
- Slivinski, L., and C. Snyder, 2016: Exploring practical estimates of the ensemble size necessary for particle filters. *Mon. Wea. Rev.*, **144**, 861–875, <https://doi.org/10.1175/MWR-D-14-00303.1>.
- , E. Spiller, A. Apte, and B. Sandstede, 2015: A hybrid particle–ensemble Kalman filter for Lagrangian data assimilation. *Mon. Wea. Rev.*, **143**, 195–211, <https://doi.org/10.1175/MWR-D-14-00051.1>.
- Snyder, C., T. Bengtsson, P. Bickel, and J. Anderson, 2008: Obstacles to high-dimensional particle filtering. *Mon. Wea. Rev.*, **136**, 4629–4640, <https://doi.org/10.1175/2008MWR2529.1>.
- , —, and M. Morzfeld, 2015: Performance bounds for particle filters using optimal proposal. *Mon. Wea. Rev.*, **143**, 4750–4761, <https://doi.org/10.1175/MWR-D-15-0144.1>.
- Stensrud, D. J., and Coauthors, 2009: Convective-scale warn-on-forecast system: A vision for 2020. *Bull. Amer. Meteor. Soc.*, **90**, 1487–1500, <https://doi.org/10.1175/2009BAMS2795.1>.
- , and Coauthors, 2013: Progress and challenges with Warn-on-Forecast. *Atmos. Res.*, **123**, 2–16, <https://doi.org/10.1016/j.atmosres.2012.04.004>.
- Talagrand, O., R. Vautard, and B. Strauss, 1997: Evaluation of probabilistic prediction systems. *Proc. ECMWF Workshop on Predictability*, Reading, United Kingdom, ECMWF, 1–25, <https://www.ecmwf.int/node/12555>.
- van Leeuwen, P. J., 2010: Nonlinear data assimilation in geosciences: An extremely efficient particle filter. *Quart. J. Roy. Meteor. Soc.*, **136**, 1991–1999, <https://doi.org/10.1002/qj.699>.
- Wheatley, D. M., K. H. Knopfmeier, T. A. Jones, and G. J. Creager, 2015: Storm-scale data assimilation and ensemble forecasting with the NSSL experimental Warn-on-Forecast system. Part I: Radar data experiments. *Wea. Forecasting*, **30**, 1795–1817, <https://doi.org/10.1175/WAF-D-15-0043.1>.
- Whitaker, J. S., and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924, [https://doi.org/10.1175/1520-0493\(2002\)130<1913:EDAWPO>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<1913:EDAWPO>2.0.CO;2).
- , and —, 2012: Evaluating methods to account for system errors in ensemble data assimilation. *Mon. Wea. Rev.*, **140**, 3078–3089, <https://doi.org/10.1175/MWR-D-11-00276.1>.
- Zhang, F., C. Snyder, and J. Sun, 2004: Impacts of initial estimate and observation availability on convective-scale data assimilation with an ensemble Kalman filter. *Mon. Wea. Rev.*, **132**, 1238–1253, [https://doi.org/10.1175/1520-0493\(2004\)132<1238:IOIEAO>2.0.CO;2](https://doi.org/10.1175/1520-0493(2004)132<1238:IOIEAO>2.0.CO;2).
- Zhu, Y., and Coauthors, 2016: All-sky microwave radiance assimilation in NCEP’s GSI analysis system. *Mon. Wea. Rev.*, **144**, 4709–4735, <https://doi.org/10.1175/MWR-D-15-0445.1>.