

Abstract

A local particle filter (LPF) is introduced that outperforms traditional ensemble Kalman filters in highly nonlinear/non-Gaussian scenarios, both in accuracy and computational cost. The standard Sampling Importance Resampling (SIR) particle filter is augmented with an observation-space localization approach, for which an independent analysis is computed locally at each gridpoint. The deterministic resampling approach of Kitagawa is adapted for application locally and combined with interpolation of the analysis weights to smooth the transition between neighboring points. Gaussian noise is applied with magnitude equal to the local analysis spread to prevent particle degeneracy while maintaining the estimate of the growing dynamical instabilities. The approach is validated against the Local Ensemble Transform Kalman Filter (LETKF) using the 40-variable Lorenz-96 model. The results show that: (1) the accuracy of LPF surpasses LETKF as the forecast length increases (thus increasing the degree of non-linearity), (2) the cost of LPF is significantly lower than LETKF as the ensemble size increases, and (3) LPF prevents filter divergence experienced by LETKF in cases with non-Gaussian observation error distributions.

1 Introduction

The Particle Filter (PF) has been explored in the data assimilation community since the introduction of its Gaussian linear variant, the Ensemble Kalman Filter (EnKF) in the mid-1990's (Evensen, 1994). While general PFs have been intractable for high dimensional systems, the EnKF has experienced great success in numerical weather prediction (NWP) (e.g., Kleist, 2012; Hamrud et al., 2014) and ocean data assimilation (e.g., Penny et al., 2015). However, at least two limitations are on the horizon for EnKFs. Perhaps counter-intuitively, these limitations arise due to *increased* computational resources, and have already become challenges at the RIKEN Advanced Institute for Computational Science (AICS, e.g., Miyamoto et al., 2013; Miyoshi et al., 2014, 2015).

NPGD

2, 1631–1658, 2015

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



First, global models will be pushed to higher resolutions in which they begin to resolve highly nonlinear processes. To maintain the Gaussian linear assumption required for the EnKF, much smaller timesteps are needed. For example, the standard 6 h analysis cycles used for the atmosphere may need to be decreased to 5 min or even 30 s.

Second, large ensembles (e.g., with ensemble size $k > 10\,000$ members) will become feasible for lower-resolution models. While at first this may seem an advantage rather than a limitation, the computational cost of the local ensemble transform Kalman filter (LETKF) (Hunt et al., 2007), for example, increases at a rate $O(k^3)$ with the ensemble size k . Thus as the ensemble size increases, the cost of computing the analysis increases at a much greater rate. Alternative EnKFs feasible for large geophysical systems scale in computational cost with the observation dimension, which is typically multiple orders of magnitude larger than the ensemble dimension.

The PF is generally applicable to nonlinear non-Gaussian systems, including cases with multi-modal distributions or nonlinear observation operators. With little difficulty, PFs can explicitly include representation of model error, nonlinear observation operators (Nakano, 2007; Lei and Bickel, 2011), non-diagonal observation error covariance matrices, and non-Gaussian likelihood functions. For example, observed variables such as precipitation are inherently non-Gaussian and cannot be effectively assimilated by standard EnKF techniques (e.g., Lien et al., 2013, 2015). In the expansion to sea-ice and land data assimilation applications, the non-Gaussian quantities such as ice concentration, ice thickness, snow cover, and soil moisture outnumber those that can be modeled with Gaussian error. Bocquet et al. (2010) further review the difficulties using observations with non-Gaussian error distributions. All of these problem-specific variations can create great difficulties for standard methods, such as the EnKF or variational approaches (3-D-Var/4-D-Var), as used in current operational systems.

Sampling Importance Resampling (SIR) (also known as the bootstrap filter, Gordon et al., 1993) is a commonly used enhancement to the basic Sequential Importance Sampling (SIS) particle filter. However, even with resampling the number of ensemble members required by the SIR particle filter to capture the high probability region of the

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



posterior in high-dimensional geophysical applications is too large to make SIR usable (Ades and van Leeuwen, 2013). Snyder et al. (2008) found that the number of required ensemble members scales exponentially with the size of the system, giving the example that a 200 dimensional system would require 10^{11} members. However, Snyder et al. note that clever choices of the proposal distribution could overcome the need for these exponentially large ensemble sizes in high-dimensional systems, which has been more recently explored by Snyder et al. (2015). Applying such an approach, van Leeuwen (2003) considers a model for the Agulhas Current with dimension of roughly 2×10^5 . Further, Beskos et al. (2012) discuss recursive methods for estimating the proposal densities, similar to the Running-in-Place algorithm (Yang et al., 2012a, b; Penny et al., 2013) that has been used with LETKF in meteorological and oceanographic data assimilation.

Techniques such as localization and inflation are typically applied as modifications to make the EnKF operationally feasible. Inspired by this practice, we introduce a local particle filter (LPF) designed for geophysical systems that is scalable to high dimensions and has computational cost $O(k)$. Spatial localization is typically justified by the fact that long distance correlations are either spurious or weak in comparison to nearby correlations, particularly when the ensemble is under-sampled. We use this same approach to reduce the required ensemble size for the LPF.

2 Methodology

Localization is used in most operational NWP data assimilation systems, either through a direct scaling of the background error covariance matrix (e.g., Whitaker and Hamill, 2002) or by a scaling of the observation error covariance matrix (Hunt et al., 2007). Because the computation of a background error covariance matrix is not needed for the PF, the latter approach is applied here to develop an efficient PF for high-dimensional geophysical systems. Localization reduces the dimensionality of the solution space,

thus requiring fewer ensemble members to sample the phase space. Gaussian noise is applied as an additive inflation to prevent particle degeneracy.

2.1 The standard SIR particle filter

There are many variations of the PF (Stewart and McCarty, 1992; Gordon et al., 1993; Kitagawa, 1996; Hurzeler and Kunsch, 1998; Liu and Chen, 1998). In essence it is simply a Monte Carlo estimation of Bayes Theorem, reformulated as a recursion (Doucet et al., 2001),

$$\rho(\mathbf{x}_t | \mathbf{y}_{1:t}) = \frac{\rho(\mathbf{y}_t | \mathbf{x}_t) \rho(\mathbf{x}_t | \mathbf{y}_{1:t-1})}{\rho(\mathbf{y}_t | \mathbf{y}_{1:t-1})} \quad (1)$$

where $\rho(\mathbf{x}_t | \mathbf{y}_{1:t})$ is the probability of the state \mathbf{x} at time t , given all observations \mathbf{y} up to time t . We consider the model domain to be described by vectors \mathbf{x} with dimension m , the observation domain to be described by vector \mathbf{y} with dimension l , and an ensemble of size k . The term in the numerator can be expressed using the Chapman-Kolmogorov equation as,

$$\rho(\mathbf{x}_t | \mathbf{y}_{1:t-1}) = \int \rho(\mathbf{x}_t | \mathbf{x}_{t-1}) \rho(\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1}) d\mathbf{x}_{t-1} \quad (2)$$

and similarly the term in the denominator can be expressed as,

$$\rho(\mathbf{y}_t | \mathbf{y}_{1:t-1}) = \int \rho(\mathbf{y}_t | \mathbf{x}_t) \rho(\mathbf{x}_t | \mathbf{y}_{1:t-1}) d\mathbf{x}_t. \quad (3)$$

The two factors in the numerator of Eq. (1) are sampled using a numerical model f ,

$$\rho(\mathbf{x}_t | \mathbf{y}_{1:t-1}) \approx \frac{1}{K} \sum_{i=1}^k \delta(\mathbf{x}_t - f(\mathbf{x}_{t-1}^i)) \quad (4)$$

$$\rho(\mathbf{y}_t | \mathbf{x}_t) = g(\mathbf{y}_t | \mathbf{x}_t). \quad (5)$$

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page	
Abstract	Introduction
Conclusions	References
Tables	Figures
◀	▶
◀	▶
Back	Close
Full Screen / Esc	
Printer-friendly Version	
Interactive Discussion	



The term in Eq. (5) is typically called the likelihood, because the probability of \mathbf{y} given \mathbf{x} is equivalent to the likelihood of \mathbf{x} given \mathbf{y} , i.e., $\rho(\mathbf{y}|\mathbf{x}) \equiv \ell(\mathbf{x}|\mathbf{y})$. The function g is general and can represent any distribution for the observations.

For the experiments here we generate two experiment cases, each with a different likelihood function. First we use a Gaussian likelihood function corresponding to that used for EnKFs,

$$g(\mathbf{y}_t|\mathbf{x}_t) = \frac{1}{\sqrt{(2\pi)^m |\mathbf{R}|}} \exp \left[-\frac{1}{2} (\mathbf{y}_t - h(\mathbf{x}_t))^T \mathbf{R}^{-1} (\mathbf{y}_t - h(\mathbf{x}_t)) \right] \quad (6)$$

where the function h is a general, possibly nonlinear, observation operator mapping from the model state space to the observation space. For the LPF, it is straightforward to generalize to arbitrary non-Gaussian likelihood functions. As an example, we also apply a multivariate Gaussian mixture model (GM₂) following Fowler and van Leeuwen (2013) with pdf,

$$\rho(\mathbf{y}|\mathbf{x}) \propto v_w \exp \left[(\mathbf{y} + v_1 \mathbf{1} - h(\mathbf{x}))^T \mathbf{R}^{-1} (\mathbf{y} + v_1 \mathbf{1} - h(\mathbf{x})) \right] + (1 - v_w) \exp \left[(\mathbf{y} + v_2 \mathbf{1} - h(\mathbf{x}))^T \mathbf{R}^{-1} (\mathbf{y} + v_2 \mathbf{1} - h(\mathbf{x})) \right]. \quad (7)$$

Let each ensemble member be identified with an index, i . Normalized weights are evaluated for each member,

$$w_t^i = \frac{\rho(\mathbf{y}_t|\mathbf{x}_t^i)}{\sum_{j=1}^k \rho(\mathbf{y}_t|\mathbf{x}_t^j)}. \quad (8)$$

Then the posterior is,

$$\rho(\mathbf{x}_t|\mathbf{y}_{1:t}) \approx \sum_{i=1}^k w_t^i \delta(\mathbf{x}_t - f(\mathbf{x}_{t-1}^i)). \quad (9)$$

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Based on Liouville's theorem, the evolution of a probability measure in a dynamical system satisfies the property that “the probability of finding trajectories inside the time-variant volume $W(t)$ is constant during the evolution of the dynamical system” (Property 2, http://www.ulb.ac.be/di/map/gbonte/ftp/bontempi_fpde.pdf). If the solution manifold expands in some directions, so will the pdf represented by the particles. Thus, the fidelity of the distribution will quickly become insufficient to sample a solution manifold around the true trajectory. A resampling procedure is used to refocus the particles on the densest areas of the distribution at each analysis step. For the experiments here, we use a resampling procedure that resembles resampling with replacement. After resampling we have,

$$\rho(\mathbf{x}_t | \mathbf{y}_{1:t}) \approx \frac{1}{k} \sum_{i=1}^k \delta(\mathbf{x}_t - \mathbf{x}_t^i). \quad (10)$$

2.2 The transform interpretation

The PF can be interpreted similarly to the Ensemble Transform Kalman Filter (ETKF) of Bishop et al. (2001). Namely, we define the PF solution as a transformation of the background ensemble to the analysis ensemble,

$$\mathbf{X}^a = \mathbf{X}^b \mathbf{T} \quad (11)$$

where each column of \mathbf{X}^b is a background ensemble member, and each column of \mathbf{X}^a is an analysis ensemble member.

Let \mathbf{b} be the vector of background particle indices and \mathbf{a} be the vector of analysis particle indices,

$$\mathbf{b} = \left\{ \mathbf{z} \in \mathbb{Z}^k \mid \mathbf{z} = (1, 2, 3, \dots, k) \right\} \quad (12)$$

$$\mathbf{a} = \left\{ \mathbf{z} \in \mathbb{Z}^k \mid \mathbf{z} = (a_1, a_2, a_3, \dots, a_k), a_j \in [1, k] \right\} \quad (13)$$

$$\mathbf{e}_j = \left\{ \mathbf{z} \in \{0, 1\}^k \mid \mathbf{z} = (0, \dots, 1_j, \dots, 0) \right\}. \quad (14)$$

If \mathbf{e}_j are the canonical basis vectors then we can define,

$$\mathbf{E}_{k \times k} = \left[\mathbf{e}_{a_1} \mathbf{e}_{a_2} \cdots \mathbf{e}_{a_k} \right]. \quad (15)$$

For the standard PF, the indicator matrix \mathbf{E} is made up of k (not necessarily unique) standard basis vectors \mathbf{e}_j , with entries 0 and 1 that we will interpret as weights. Thus the analysis ensemble for the PF is defined simply by the transform,

$$\mathbf{X}_{m \times k}^a = \mathbf{X}_{m \times k}^b \mathbf{E}_{k \times k}. \quad (16)$$

We note that by using this approach, each new analysis member, with index i , maintains the continuity properties of its associated background member, a_j .

For reference in the next section, the components of the analysis matrix \mathbf{X}^a will have the form,

$$\mathbf{X}^a = \begin{bmatrix} x_{1,i} \mathbf{e}_{i,1} & x_{1,i} \mathbf{e}_{i,2} & \cdots & x_{1,i} \mathbf{e}_{i,k} \\ x_{2,i} \mathbf{e}_{i,1} & x_{2,i} \mathbf{e}_{i,2} & \cdots & x_{2,i} \mathbf{e}_{i,k} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m,i} \mathbf{e}_{i,1} & x_{m,i} \mathbf{e}_{i,2} & \cdots & x_{m,i} \mathbf{e}_{i,k} \end{bmatrix}. \quad (17)$$

Here we have used the Einstein tensor notation for the elements, in which $x_{1,i} \mathbf{e}_{i,1}$ represents a summation over the index i (i.e., the inner product of row 1 of \mathbf{X}^b and column 1 of \mathbf{E}). While the summation index could be represented generically by any symbol, we reuse the symbol “ i ” due to its correspondence with the background particle indices as defined above.

2.3 The localization approach

Snyder et al. (2008) note that when either the model dimension or observation count is large, the PF requires significantly more particles to give an adequate representation

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page	
Abstract	Introduction
Conclusions	References
Tables	Figures
◀	▶
◀	▶
Back	Close
Full Screen / Esc	
Printer-friendly Version	
Interactive Discussion	



of the system. Localization reduces both the model and observation dimensions by dividing the problem into a series of sub-domains, thus reducing the required number of particles for accurate filtering. The LPF uses the approach of Hunt et al. (2007) to select nearby observations for independent analysis at each grid point. Nearby grid points thus assimilate nearly identical sets of observations to derive their analyses.

We use the deterministic resampling of Kitagawa (1996), with complexity $O(k)$, adapted for local use as described next. A uniform partition of the interval $[0,1]$ with width $1/k$ is first generated globally, with an offset applied from a uniform distribution over the interval $[0,1/k]$. The same partition is used locally for resampling at each grid point. Cumulative sums of the normalized weights (Eq. 8),

$$\check{w}_t^j = \sum_{i=1}^j w_t^i \quad (18)$$

are compared with the elements of the partition. Traversing from $j = 1, \dots, k$, all unassigned particles with index j having a corresponding cumulative sum with index j that surpasses the next element of the partition (ordered monotonically increasing) are assigned as particles of the resampled (analysis) ensemble. For a given grid point, when the cumulative sums of the particle weights are near one of the partition values, there may be sensitivity in neighboring grid points that lead to discontinuities between local analyses. The analysis ensemble at this grid point consists of a subset of background particle indices (1 through k) with repetitions. To eliminate the discontinuities with neighboring grid points, with the particle indices we associate weights of a local transform function T , nominally either 1.0 (full weight) or 0.0 (no weight). This is partially inspired by the “weight interpolation” of Bowler (2006), applied to LETKF by Yang et al. (2009), who found that interpolation of weights was more robust than interpolation of state values. At a single grid point, there are k pieces of background information about the possible system state at that point. In the standard PF, only 1 out of these k pieces of information is retained for each analysis ensemble member, based on the overall agreement with observations. In the LPF we use anywhere from 1 to k mem-

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

[Title Page](#)[Abstract](#)[Introduction](#)[Conclusions](#)[References](#)[Tables](#)[Figures](#)[Back](#)[Close](#)[Full Screen / Esc](#)[Printer-friendly Version](#)[Interactive Discussion](#)

bers to construct each analysis member based on the agreement with observations within a local radius.

For the LPF, a new transform is defined for each point in the model domain to generate a set of m indicator matrices, $\mathbf{E}_{k \times k}^{(j)}$, so that for each point (x_j , for $j = 1 \dots m$),

$$\mathbf{X}_{m \times k}^a{}^j = \mathbf{X}_{m \times k}^b \mathbf{E}_{k \times k}^{(j)}. \quad (19)$$

Using the summation tensor notation described in the previous section, the analysis ensemble can be written,

$$\mathbf{X}^a = \begin{bmatrix} x_{1,i} e_{i,1}^{(1)} & x_{1,i} e_{i,2}^{(1)} & \cdots & x_{1,i} e_{i,k}^{(1)} \\ x_{2,i} e_{i,1}^{(2)} & x_{2,i} e_{i,2}^{(2)} & \cdots & x_{2,i} e_{i,k}^{(2)} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m,i} e_{i,1}^{(m)} & x_{m,i} e_{i,2}^{(m)} & \cdots & x_{m,i} e_{i,k}^{(m)} \end{bmatrix}. \quad (20)$$

The transform matrix may have any degree of sophistication. We apply a smoothing operator by modifying the weights \mathbf{e}_{a_j} , associated with each analysis particle index $a_{(i)}$ from a binary value to a continuous value between 0 and 1, while maintaining all column sums equal to one, and call this new transform matrix \mathbf{W} . This smoothing is performed only in the ensemble space; no explicit interpolation is applied in the model space.

We define the concept of a “neighbor point” abstractly as a point near the analyzed grid point based on a specified distance metric. If there are N neighbor points, then there will be at most $\min(N + 1, k)$ collocated pieces of background information that can be utilized to construct each analysis ensemble member at this point. An example is given in Fig. 1. In our case, with a sufficiently large set of observations the indices for these neighbor points are calculated from nearly identical observation innovations. Therefore, when there is a sufficiently large ensemble size (k) the difference between the states associated with different particle indices will be small. The transform function T is applied across all background indices (i.e., for particles $1, \dots, k$) at this grid point to compute the analysis.

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



2.4 Particle degeneracy

The particle selection process of the PF reduces the rank of the ensemble. For a linear deterministic system this leads to a rapid collapse of the ensemble and divergence of the filter. For a sufficiently stochastic nonlinear system the full rank is recovered after a single forecast step. If the nonlinear system is not sufficiently stochastic, then we must address the ensemble initialization problem at every analysis cycle. Pazo et al. (2010) discuss the desirable properties in an initial ensemble, namely the members: (1) should be well-embedded in the attractor, (2) should be statistically equivalent but have enough diversity to represent a significant portion of the phase space, (3) should adequately represent the error between the analysis and true state, and (4) should sample the fastest growing directions in phase space. We wish to avoid particle degeneracy while also engendering some of these qualities. Therefore we employ a simple approach: at each cycle we add Gaussian noise with variance scaled locally to a magnitude matching the analysis error variance and apply this to each analysis member prior to the subsequent ensemble forecast.

2.5 Computational complexity

A data assimilation system is comprised of many components. We simplify the cost analysis in order to gain an approximate relative measure of the algorithms presented here. Let m be the model dimension, l be the observation dimension, and \bar{l}_i be the average local observation dimension. The total cost (C_T) of an analysis cycle is equal to the overhead (C_H) of the assimilation system plus m times the average local cost (C_L) of the assimilation method plus k times the cost of one model forecast (C_M) of duration τ ,

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



$$C_T(k, l, m) = C_H(k, l, m) + m \cdot C_L(k, \bar{l}_i) + k \cdot C_M(\tau, m). \quad (21)$$

We will assume that between the two methods the overhead and model costs are approximately equal. The primary difference in cost between the two systems is then the average local cost,

$$C_L^{\text{LPF}} = O(k\bar{l}_i) \quad (22)$$

$$C_L^{\text{LETKF}} = O(k^2\bar{l}_i + k^3). \quad (23)$$

If as is typically the case, the system size m is large and the ensemble size k is small, then

$$C_T(k, l, m) \approx C_H(k, l, m) + O(m), \quad (24)$$

and the difference in cost between LETKF and LPF is small. However for large k , we see that the average local cost of LETKF,

$$C_T^{\text{LETKF}}(k, l, m) = C_H(k, l, m) + O(mk^3) \quad (25)$$

exceeds that of the LPF,

$$C_T^{\text{LPF}}(k, l, m) = C_H(k, l, m) + O(mk). \quad (26)$$

Subtracting the overhead costs, in this case the LPF is a factor of k^2 cheaper than LETKF.

2.6 Data assimilation methods

We enumerate the benefits of the LPF vs. the benchmark LETKF, an ensemble square root filter that performs its analysis in the ensemble space at each grid point using

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



of the system (Fig. 2). On the contrary, LETKF performs well even with few ensemble members and few observations per cycle ($k = 20$, $l = 10$). Localization is consistent between each method, using $r = 2$ gridpoints. For a given ensemble size, increasing the localization radius degraded the accuracy of both methods. To explore the relative advantages of each approach, we will describe a series of cases in which the LETKF outperforms the LPF, and in which the LPF outperforms LETKF.

3.1 Case 1: typical forecast lengths ($dt = 0.05$, or 6 h)

Lorenz (1996) introduced the $dt = 0.05$ timescale as being comparable to the error doubling taking place over 6 h in the operational forecasting systems of the early 1990's. In this relatively linear timescale of the L96 system, LETKF clearly outperforms the LPF at a given ensemble size. This is expected as EnKFs take advantage of the Gaussian/linear assumption. When the experiment parameters match such assumptions (even loosely), LETKF performs quite well. However, using localization, the LPF can perform adequately (i.e., avoid filter divergence) in a similar parameter regime (Fig. 3). Thus for this case, we find that LETKF attains higher accuracy than the LPF, but the LPF improves upon the accuracy and stability of the standard SIR PF for a given ensemble size.

3.2 Case 2: long forecast lengths ($dt = 0.50$, or 60 h)

To increase the degree of nonlinearity in a data assimilation system using L96, it is typical to increase the analysis cycle length (e.g., Lei and Bickel, 2011). The LPF has superior performance for more nonlinear regimes of the L96 system (e.g., $dt = 0.5$) provided there are many ensemble members, e.g., $O(100)$. Using 80 observations per cycle and 100 ensemble members, for example, LETKF produces occasional errors that propagate eastward (along the positive x direction). The LPF does not produce such effects, and the errors are generally lower than with LETKF (Fig. 4). We consider

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



this a relevant scenario because the majority of observations in operational weather forecasting are discarded (Ochatta et al., 2005).

Exploring a more complete parameter space, we examine the forecast error for LETKF over a range of observation coverage ($l = 2, \dots, 80$ per analysis cycle) and ensemble sizes ($k = 10, \dots, 400$), and compare the relative difference to LPF. Figure 5 shows the average absolute error over 600 analysis cycles of length $dt = 0.5$ for 1600 different parameter combinations of observation coverage (l) and ensemble size (k). The LPF is more accurate than LETKF when using many observations (e.g., $l > 20$) and large ensemble sizes (e.g., $k > 100$ – 200)

Further, when examining the computational cost of the LPF vs. LETKF, the relative costs reflect the analytical assessment given above in Sect. 2.5. Namely, the elapsed time of the LETKF experiments grows exponentially with ensemble size, while the elapsed time of the LPF is significantly lower at large ensemble sizes (Fig. 6).

3.3 Case 3: non-Gaussian observation error

The previous section examined the impacts of nonlinearity and non-Gaussianity on the forecast. We now examine the impacts of non-Gaussian observation error. Using a multivariate Gaussian mixture model (GM_2) following Fowler and van Leeuwen (2013), we apply a corresponding random error to each observation and compare the impacts on LETKF and LPF. We use the LETKF without modification, but modify the likelihood function of LPF to the definition of GM_2 as in Sect. 2.1, Eq. (7). Using $\nu_w = 0.1$, $\nu_1 = -1$, $\nu_2 = 1$, we create a bimodal distribution biased toward the second Gaussian mode. The analysis cycle is $dt = 0.05$ (6 h) as in experiment case 1, Sect. 3.1. Figure 7 compares LETKF and LPF using $k = 100$ ensemble members and $l = 80$ observations. An additional result is given for LPF with $l = 20$ observations. The introduction of a strong non-Gaussianity in the observation error distribution disrupts LETKF and eventually creates errors that propagate throughout the entire domain. Using the same ensemble size and observation count, the LPF gains significant advantage in its ability to explic-

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Doucet, A., De Freitas, N., and Gordon, N. J.: An introduction to Sequential Monte Carlo Methods, in: SMC in Practice, available at: http://www.stats.ox.ac.uk/~doucet/smc_resources.html (last access: 25 November 2015), 2001.

Evenson, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *J. Geophys. Res.*, 99, 10143–10162, 1994.

Fowler, A. and van Leeuwen, P. J.: Observation impact in data assimilation: the effect of non-Gaussian observation error, *Tellus A*, 65, 20035, doi:10.3402/tellusa.v65i0.20035, 2013.

Gordon, N. J., Salmond, D., and Smith, A. F. M.: Novel approach to nonlinear/non-Gaussian Bayesian state estimation, *IEE Proc.-F*, 140, 107, 1993.

Hoffman, R. N. and Kalnay, E.: Lagged Average Forecasting: an alternative to Monte Carlo Forecasting, *Tellus*, 35, 100–118, 1983.

Hou, D., Toth, Z., and Zhu, Y.: A Stochastic Parameterization Scheme Within NCEP Global Ensemble Forecast System. *Am. Met. Soc.*, 18 Conference on Probability and Statistics in the Atmospheric Sciences, available at: <http://ams.confex.com/ams/pdfpapers/101401.pdf> (last access: 25 November 2015), 2006.

Hou, D., Toth, Z., Zhu, Y., Yang, W., and Wobus, R.: A Stochastic Total Tendency Perturbation Scheme Representing Model-Related Uncertainties in the NCEP Global Ensemble Forecast System, NOAA/NCEP/EMC, available at: http://www.emc.ncep.noaa.gov/gmb/yzhu/gif/pub/Manuscript_STTP_Tellus_A_HOU-1.pdf (last access: 25 November 2015), 2010.

Hurzeler, M. and Kunsch, H.: Monte Carlo approximations for general state-space models, *J. Comput. Grap. Stat.*, 7, 175–193, doi:10.2307/1390812, 1998.

Kitagawa, G.: Monte Carlo filter and smoother for non-Gaussian non-linear state space models, *J. Comput. Graph. Stat.*, 5, 1–25, 1996.

Kolczynski, W., Pegion, P., Hamill, T., Whitaker, J. S., Hou, D., Zhu, Y., and Zhou, X.: Investigating a New Stochastic Physics Suite for Use in the NCEP Global Ensemble. *Am. Met. Soc.*, 27 Conference On Weather Analysis And Forecasting/23 Conference On Numerical Weather Prediction, recorded presentation available at: <https://ams.confex.com/ams/27WAF23NWP/webprogram/Paper273838.html>, last access: 25 November 2015.

Lei, J. and Bickel, P.: A moment matching ensemble filter for nonlinear non-Gaussian data assimilation, *Mon. Weather Rev.*, 139, 3964–3973, 2011.

Lien, G.-Y., Kalnay, E., and Miyoshi, T.: Effective assimilation of global precipitation: simulation experiments, *Tellus A*, 65, 19915, doi:10.3402/tellusa.v65i0.19915, 2013.

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



- Lien, G.-Y., Miyoshi, T., and Kalnay, E.: Assimilation of TRMM Multisatellite Precipitation Analysis with a low-resolution NCEP Global Forecasting System, *Mon. Weather Rev.*, doi:10.1175/MWR-D-15-0149.1, in press, 2015.
- Liu, J. S. and Chen, R.: Sequential Monte Carlo methods for dynamic systems, *JAS A.*, 93, 1032–1044, doi:10.2307/2669847, 1998.
- Miyamoto, Y., Kajikawa, Y., Yoshida, R., Yamaura, T., Yashiro, H., and Tomita, H.: Deep moist atmospheric convection in a subkilometer global simulation, *Geophys. Res. Lett.*, 40, 4922–4926, doi:10.1002/grl.50944, 2013.
- Miyoshi, T., Kondo, K., and Imamura, T.: The 10240-member ensemble Kalman filtering with an intermediate AGC M, *Geophys. Res. Lett.*, 41, 5264–5271, doi:10.1002/2014GL060863, 2014.
- Miyoshi, T., Kondo, K., and Terasaki, K.: Big ensemble data assimilation in numerical weather prediction, *Computer*, 48, 15–21, doi:10.1109/MC.2015.332, 2015.
- Nakano, S., Ueno, G., and Higuchi, T.: Merging particle filter for sequential data assimilation, *Nonlin. Processes Geophys.*, 14, 395–408, doi:10.5194/npg-14-395-2007, 2007.
- Nerger, L.: On serial observation processing in Localized Ensemble Kalman Filters, *Mon. Weather Rev.*, 143, 1554–1567, doi:10.1175/MWR-D-14-00182.1, 2015.
- Ochatta, T., Gebhardt, C., Saupe, D., and Wergen, W.: Adaptive thinning of atmospheric observations in data assimilation with vector quantization and filtering methods, *Q. J. Roy. Meteorol. Soc.*, 131, 3427–3437, 2005.
- Ott, E., Hunt, B. R., Szunyogh, I., Zimin, A. V., Kostelich, E. J., Corazza, M., Kalnay, E., Patil, D. J., and Yorke, J. A.: A Local Ensemble Kalman Filter for atmospheric data assimilation, *Tellus A*, 56, 415–428, doi:10.1111/j.1600-0870.2004.00076.x, 2004.
- Pazo, D., Rodriguez, M. A., and Lopez, J. M.: Spatio-temporal evolution of perturbations in ensembles initialized by bred, Lyapunov, and singular vectors, *Tellus A*, 62, 10–23, 2010.
- Penny, S. G.: The Hybrid Local Ensemble Transform Kalman Filter, *Mon. Weather Rev.*, 142, 2139–2149, doi:10.1175/MWR-D-13-00131.1, 2014.
- Penny, S. G., Kalnay, E., Carton, J. A., Hunt, B. R., Ide, K., Miyoshi, T., and Chepurin, G. A.: The local ensemble transform Kalman filter and the running-in-place algorithm applied to a global ocean general circulation model, *Nonlin. Processes Geophys.*, 20, 1031–1046, doi:10.5194/npg-20-1031-2013, 2013.

Yang, S.-C., Kalnay, E., and Hunt, B. R.: Handling nonlinearity in Ensemble Kalman Filter: experiments with the three-variable Lorenz model, *Mon. Weather Rev.*, 140, 2628–2646, doi:10.1175/MWR-D-11-00313.1, 2012a.

5 Yang, S.-C., Kalnay, E., and Miyoshi, T.: Improving EnKF spin-up for typhoon assimilation and prediction, *Weather Forecast.*, 27, 878–897, 2012b.

NPGD

2, 1631–1658, 2015

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

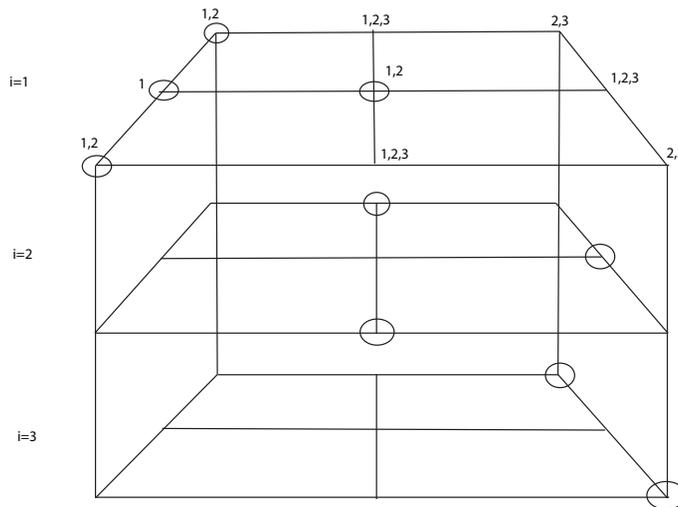


Figure 1. A hypothetical example depicting the construction of a single analysis member. Each level represents a different background ensemble member (particle), with a model space composed of a 3×3 grid. The nodes of the grid are circled if the member is chosen for the construction of analysis member 1 by the LPF. The numerals indicate the ids for the background members that will be averaged at the corresponding node, in this case based on the immediately adjacent neighbor points of that node.

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

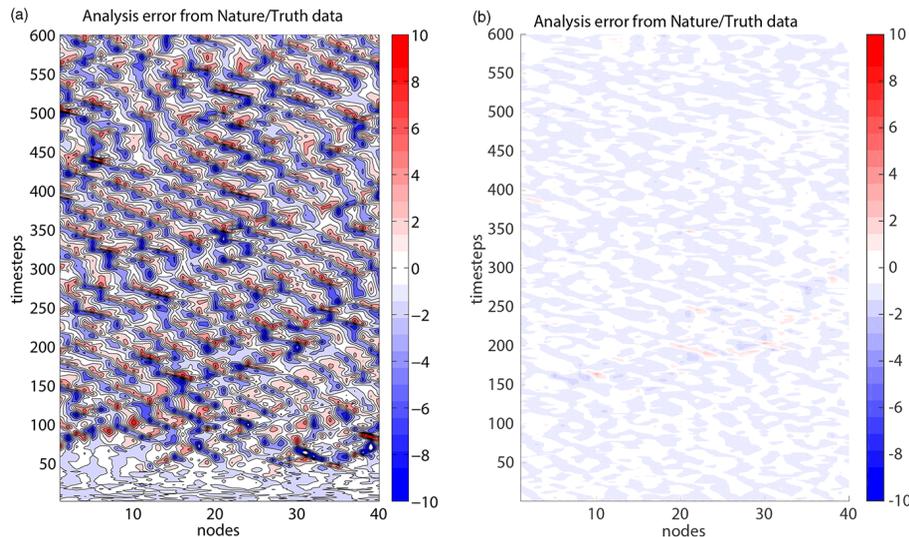


Figure 2. Analysis error using an analysis cycle window length $dt = 0.05$ (6 h) for **(a)** the standard SIR PF using $k = 1500$ particles with $l = 20$ observations per analysis cycle, and **(b)** LETKF with localization radius $r = 2$, $k = 20$ ensemble members, and $l = 10$ observations per analysis cycle, sampled randomly on the domain.

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

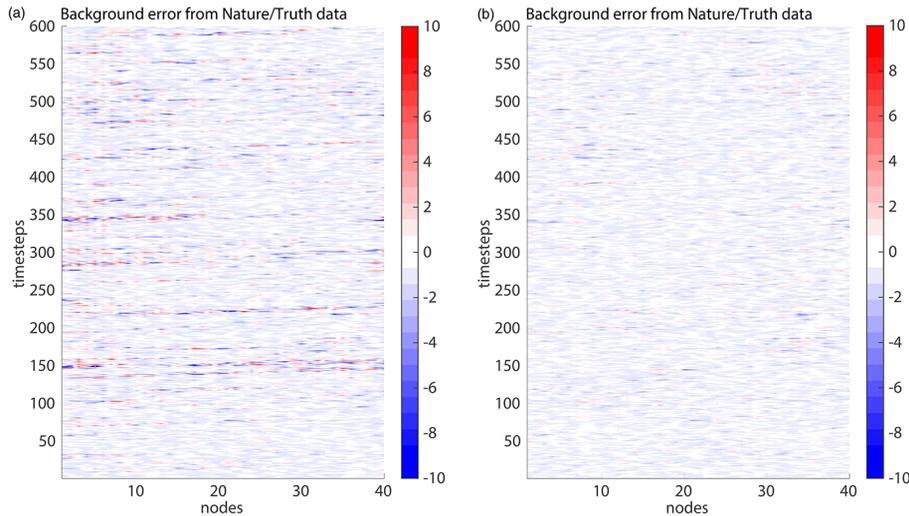


Figure 4. Forecast error for **(a)** LETKF and **(b)** LPF, using an analysis cycle window length $dt = 0.5$ (60 h), localization radius $r = 2$ grid points, $k = 100$ ensemble members, and $l = 80$ observations sampled randomly on the domain (observation locations and errors are identical for both methods).

[Title Page](#)[Abstract](#)[Introduction](#)[Conclusions](#)[References](#)[Tables](#)[Figures](#)[Back](#)[Close](#)[Full Screen / Esc](#)[Printer-friendly Version](#)[Interactive Discussion](#)

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

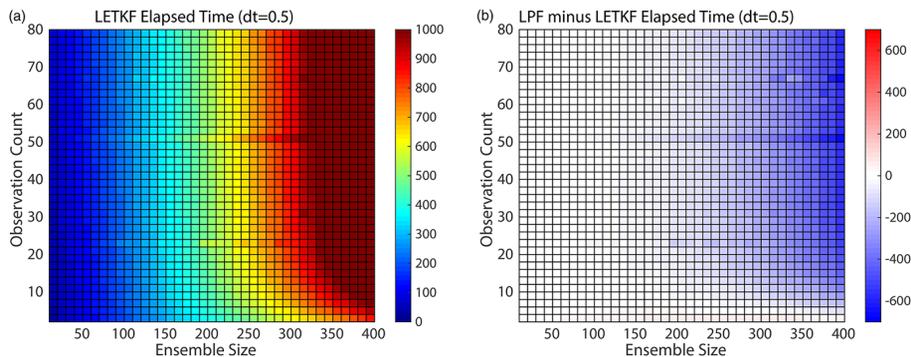


Figure 6. Elapsed time in seconds for (a) LETKF and (b) LPF minus LETKF.

A local particle filter for high dimensional geophysical systems

S. G. Penny and
T. Miyoshi

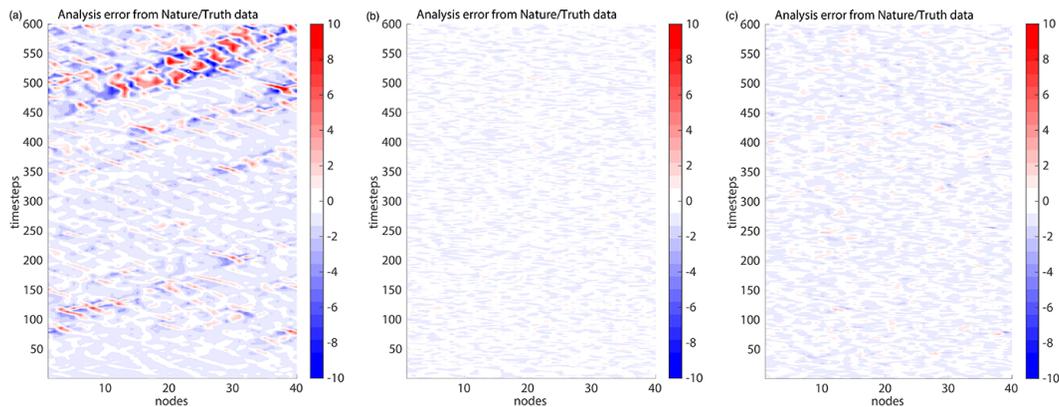


Figure 7. Analysis error for **(a)** LETKF and **(b)** LPF, using $l = 80$ observations and $k = 100$ ensemble members. The observations used between **(a)** and **(b)** are identical. In **(c)**, the number of observations for LPF is reduced to $l = 20$, but improvement in accuracy vs. LETKF remains.

[Title Page](#)[Abstract](#)[Introduction](#)[Conclusions](#)[References](#)[Tables](#)[Figures](#)[Back](#)[Close](#)[Full Screen / Esc](#)[Printer-friendly Version](#)[Interactive Discussion](#)