A comparison of nonlinear extensions to the ensemble Kalman filter

Gaussian anamorphosis and two-step ensemble filters

Ian Grooms¹

Received: 15 July 2021 / Accepted: 9 February 2022 / Published online: 5 March 2022
© The Author(s), under exclusive licence to Springer Nature Switzerland AG 2022

Abstract

Ensemble Kalman filters are based on a Gaussian assumption, which can limit their performance in some non-Gaussian settings. This paper reviews two nonlinear, non-Gaussian extensions of the Ensemble Kalman Filter: Gaussian anamorphosis (GA) methods and two-step updates, of which the rank histogram filter (RHF) is a prototypical example. GA-EnKF methods apply univariate transforms to the state and observation variables to make their distribution more Gaussian before applying an EnKF. The two-step methods use a scalar Bayesian update for the first step, followed by linear regression for the second step. The connection of the two-step framework to the full Bayesian problem is made, which opens the door to more advanced two-step methods in the full Bayesian setting. A new method for the first part of the two-step framework is proposed, with a similar form to the RHF but a different motivation, called the ‘improved RHF’ (iRHF). A suite of experiments with the Lorenz-‘96 model demonstrate situations where the GA-EnKF methods are similar to EnKF, and where they outperform EnKF. The experiments also strongly support the accuracy of the RHF and iRHF filters for nonlinear and non-Gaussian observations; these methods uniformly beat the EnKF and GA-EnKF methods in the experiments reported here. The new iRHF method is only more accurate than RHF at small ensemble sizes in the experiments reported here.

Keywords Data assimilation · Ensemble Kalman filter · Gaussian anamorphosis · Rank histogram filter

Mathematics Subject Classification (2010) MSC 60G35 · 62M20

1 Introduction

Data assimilation (DA) is the process of combining observational data and a dynamical model to estimate and predict the state of a dynamical system. It encompasses a variety of methods with varying degrees of complexity, efficiency, and robustness [34, 58, 62]. DA is widely used in a variety of applications including weather forecasting [1], atmospheric and oceanic sciences [97], oil & gas reservoir modeling [2, 36], and epidemiology [93], to name just a few. Bayesian estimation theory provides a mathematical framework for DA (e.g. [107]). Bayes’ theorem relates the posterior probability distribution describing the state of the system given observations to the probability of the observations given the system state (the ‘likelihood’) and the prior probability of

the state of the system without knowing the observations. The Bayesian *filtering* problem is to find the posterior on the current system state given all past and current observations; by contrast, *smoothing* problems allows future observations to impact the posterior. The focus here is on large-scale applications as occur in the geosciences. For example, a global ocean model with modest resolution and no biogeochemistry has in the vicinity of 20 million degrees of freedom, while coupled climate models at higher resolution can have on the order of billions of degrees of freedom.

When the dynamics and the observations are linear and all distributions are Gaussian, the Bayesian filtering problem is exactly solved by the Kalman filter [57]. Ensemble Kalman filtering, which comprises an entire class of methods, is an approach to solving the filtering problem based on ensemble approximations to the Kalman filter formulas [34]. It is known that the mean and covariance generated by ensemble Kalman filters (EnKFs) converge to those of the Kalman filter in the limit of infinite ensemble size for linear systems with Gaussian noise [60, 73]. EnKFs are routinely used in systems with nonlinear dynamics and

✉ Ian Grooms
ian.grooms@colorado.edu

¹ Department of Applied Mathematics, University of Colorado, Boulder, CO, USA

non-Gaussian distributions but there is no guarantee of convergence for these systems. Indeed, EnKFs are expected *not* to converge to the correct Bayesian posterior for nonlinear, non-Gaussian systems. Nevertheless, EnKFs are relatively easy to implement compared to alternatives like 4D-Var [102] and yield surprisingly good results in many nonlinear problems.

There is a group of methods that aims to extend EnKF to improve performance in problems with non-Gaussian distributions and nonlinear observations while retaining algorithmic similarity to EnKFs. This group includes Gaussian mixture methods [12, 13, 17, 38, 49, 70, 79, 101], methods based on gamma/inverse-gamma distributions [19, 83], methods that target higher moments of the posterior [47, 48], methods based on rank statistics [9–11, 75], and ‘Gaussian Anamorphosis’ methods [18, 22, 23, 95, 108]. Gaussian anamorphosis (GA) methods were originally motivated by the desire to keep certain state variables - like concentration, mass, or volume - positive, a constraint not respected by EnKFs. GA methods operate by applying changes of variable to the prior to make it closer to Gaussian, then applying the EnKF, then transforming back. The rank-histogram filter (RHF) methods of [9–11, 75] are another example of methods that aim to relax Gaussian assumptions while retaining algorithmic similarity to EnKFs; the RHF and related methods operate within a two-step assimilation framework developed by Anderson [8] that unites Bayesian estimation and linear regression.

The focus here is on nonlinear and non-Gaussian extensions of ensemble Kalman filters. The two main categories of competing methods, in the context of large-scale geoscience applications, are variational methods and particle filters. Though these are outside the scope of the review, a brief overview is included here for interested readers.

Variational methods [102] construct an objective function - usually proportional to the negative log of the posterior - with two terms: one that penalizes deviations from the prior mean and one that penalizes deviations from observations. Optimization methods are then used to seek a global minimizer, which is the Maximum A Posteriori (MAP) estimator of the true state, though in practice the methods might only find a local minimizer. Variational methods can be crudely divided into two broad categories as follows: In 4D-Var there are observations of the state at future times, whereas in 3D-Var there are not. 4D-Var handles nonlinearity in both the dynamics and observations more accurately than basic EnKF-type methods since it computes the MAP based on the true posterior. To reduce cost, 4D-Var is often applied to a sequence of shorter time windows instead of to a single long time window. One main drawback of classical 4D-Var is that the penalty on the initial condition does not vary from one time window to the next; for

example, the forecast error model is the same on a rainy day as on a dry one. The ability of 4D-Var to provide excellent state estimates has resulted in widespread use in operational weather and ocean DA [35, 41, 53, 76, 88]. One way to improve the performance of 4D-Var is to use an ensemble forecast to codify the forecast uncertainty, which can then be used to define the term penalizing deviations from the forecast mean in the 4D-Var objective function. There are many ways to do this [14]; in ‘pure’ ensemble-variational (EnVar) methods the ensemble forecast covariance matrix is used directly, while in ‘hybrid’ EnVar the ensemble forecast covariance is combined with a time-independent background. These EnVar methods are the current state of the art for numerical weather prediction [1, 105].

Particle filters are a class of sequential Monte Carlo methods [32] that solve the Bayesian smoothing and filtering problems via sequential importance sampling. The classical bootstrap or ‘SIR’ particle filter of [42] generates an empirical measure that provably converges weakly to the correct posterior distribution in the limit of an infinite ensemble size [30, 62], but the convergence is prohibitively slow in high-dimensional systems [16, 98, 99]. In view of this limitation there is a large body of research aiming to design alternatives to the classical particle filter for applications with extremely high dimension, including [3, 4, 26–28, 52, 80, 82, 84, 85, 87, 89–91, 103]. Particle filters have also been hybridized with EnKFs [29, 37, 45] and variational methods [77].

The goal of the present paper is to compare GA methods and the RHF as nonlinear extensions of the EnKF. We begin with a review of GA methods and of the two-step framework on which both the GIGG-EnKF (Gamma, Inverse-Gamma, and Gaussian; [19, 83]) and the RHF are built. New versions of both the GA and RHF filters are presented. Then, GA and RHF methods are tested in the Lorenz-‘96 model with nonlinear observations, with an EnKF for comparison. In these tests the RHF and its newly-proposed extension are the most accurate, and do not incur significant extra computational cost.

2 Gaussian anamorphosis in the ensemble Kalman filter

2.1 Ensemble Kalman filters

Let the vector of variables that we are trying to estimate be denoted by $\mathbf{x}$, and let our prior uncertainty about $\mathbf{x}$ be encoded in the distribution of a random vector $\mathbf{X}$. Assume that there is another random vector $\mathbf{Y}$ jointly distributed with $\mathbf{X}$, and we have a single ‘observation’ vector $\mathbf{y}$ which is a realization of $\mathbf{Y}$. For example, one might have $\mathbf{Y} = \mathbf{H}(\mathbf{X}) + \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon}$ is a random variable independent of

X that represents observation error, and H some nonlinear function. This type of relationship between X and Y is widely assumed to hold in the DA literature, but is not by any means required, and will not be used in some of the numerical experiments reported below. The prior is the marginal distribution on X , the likelihood is the probability density function (pdf) of $Y|X$ evaluated at $Y = y$ and then viewed as a function of the values of $X = x$, and the posterior distribution that we seek is the conditional distribution $X|Y = y$.

The usual derivation of the Kalman Filter (KF) formulas that form the foundation for the EnKF assumes that X has a Gaussian distribution with mean x_b and covariance B , that $Y = HX + \epsilon$ where H is a constant matrix (independent of X and ϵ) and ϵ is a centered Gaussian random vector independent of X with covariance R . An equivalent and particularly useful alternative derivation of the KF formulas makes a Gaussian approximation of the joint distribution of X and Y . It is well known that if a joint distribution is Gaussian then the conditional is also Gaussian, with mean and covariance

$$\mathbb{E}[X|Y = y] = x_b + \text{Cov}[X, Y]\text{Cov}[Y]^{-1}(y - \bar{y}) \quad (1)$$

and

$$\text{Cov}[X|Y = y] = \text{Cov}[X] - \text{Cov}[X, Y]\text{Cov}[Y]^{-1}\text{Cov}[Y, X]. \quad (2)$$

Under the assumptions stated above we find

$$\bar{y} = Hx_b, \quad (3)$$

$$\text{Cov}[X] = B, \quad (4)$$

$$\text{Cov}[X, Y] = BH^T, \quad (5)$$

$$\text{Cov}[Y] = HBH^T + R. \quad (6)$$

Inserting these into (1) and (2) yields the familiar Kalman update formulas

$$x_a = x_b + K(y - Hx_b), \quad (7)$$

$$C = (I - KH)B, \quad (8)$$

$$K = BH^T(HBH^T + R)^{-1}, \quad (9)$$

where x_a is the posterior mean, also called the analysis mean, C is the posterior covariance, and K is the Kalman gain matrix.

The first EnKF was proposed by Evensen [33]. An equally-weighted ensemble was used to obtain an approximation of the background covariance matrix B , and then the Kalman update formula (7) with x_b replaced by x_i was applied to each ensemble member, where x_i with $i = 1, \dots, N$ denotes an ensemble of N state vectors. The original EnKF contained a flaw in that while the analysis ensemble mean was consistent with (7), the analysis

ensemble covariance was not consistent with (8). This was simultaneously fixed in [25, 50] using the update formula

$$x_i^+ = x_i + K(y - Hx_i - \epsilon_i), \quad (10)$$

where ϵ_i , $i = 1, \dots, N$ are independent samples from the distribution of ϵ , and the superscript $+$ denotes the analysis ensemble members. This version of the EnKF has come to be known as ‘the’ EnKF, or the ‘stochastic’ EnKF, or the ‘perturbed-observations’ EnKF. Several alternative EnKFs that do not require an ensemble of observation errors were proposed in later years [7, 20, 81, 106]; these can be called deterministic EnKFs. See [51] for a review of the literature on EnKFs for atmospheric data assimilation.

2.2 Gaussian anamorphosis

As noted above, the EnKF uses a Gaussian approximation of the joint distribution of X and Y , which is patently erroneous when the scalar marginal distributions, i.e. the distributions of each element of X and Y , are far from normal. In such cases the Gaussian approximation can be relaxed by transforming to variables with Gaussian marginals, then applying the EnKF, then transforming back. The idea of using univariate transformations of state variables in the EnKF was first proposed in [18], with the main goal being to force the EnKF update to respect bounds on the variables, e.g. to prevent the EnKF from producing negative mass. There are other methods for dealing with linear inequality constraints in the EnKF, e.g. by using truncated normals [61] or using constrained optimization [5, 56, 66]. The point here is not to deal with inequality constraints per se, but to introduce transforms as a means of generalizing the EnKF to non-Gaussian distributions.

Some methods within the literature explicitly aim to construct the univariate transformations so that the marginals are all Gaussian, while others merely design their transformations so that the state variables are guaranteed to respect desired boundedness properties like positivity after the EnKF update. The former are known as ‘Gaussian anamorphosis’ methods [18] and ‘normal score’ methods [65, 108] in the ensemble DA literature. The correct probability-integral transformation of a real-valued scalar random variable X to a standard normal random variable $\hat{X}$ is $\hat{X} = \Phi^{-1}(F_X(X))$ where F_X is the cumulative distribution function (cdf) of X and Φ is the standard normal cdf. This only applies when X has a non-atomic density; it is not possible to transform X to a standard normal when the F_X is discontinuous. The ad hoc transformations appearing in the literature involve poor (but still consistent) estimators of F_X using samples of X , mainly based around piecewise-linear non-parametric approximations (e.g. [95]). Application of these methods has been fairly limited, comprising ocean ecosystem models [15, 31, 39, 95], groundwater models [65,

[108], and some limited use in meteorological applications [59, 67–69].

Using GA transformations in the EnKF relaxes the EnKF approximation that the joint distribution of $\mathbf{X}$ and $\mathbf{Y}$ is Gaussian to an approximation that the copula is Gaussian (see the appendix for a definition of the copula of jointly-distributed random variables). As emphasized in [6], this updated approximation is not necessarily an improvement, and can in fact lead to *worse* results than a standard EnKF. On the other hand, practical experience [15, 31, 39, 59, 65, 67–69, 95, 108] and many of the examples in [6] show that GA usually improves the performance significantly. Furthermore, GA is a legitimate non-Gaussian extension of the EnKF: GA methods should converge towards sampling from the true posterior for a wider range of problems than the EnKF, namely problems with Gaussian copulas. A proof of this conjecture is beyond the scope of the current article. Multivariate transformations based on optimal transport [74] could potentially be used to relax even the assumption of a Gaussian copula, and there are some promising preliminary results using these transformations in the context of data assimilation [100].

We close this review of GA methods by proposing a perturbed-observations GA-EnKF:

1. Start with an ensemble $\mathbf{x}_i$, $i = 1, \dots, N$ and use it to generate an ensemble $\mathbf{y}_i$, as in the stochastic EnKF, by generating $\mathbf{y}_i$ as a realization of $\mathbf{Y}|\mathbf{X} = \mathbf{x}_i$.
2. Use the ensembles to find approximations to the marginal cdfs of the elements of $\mathbf{X}$ and $\mathbf{Y}$.
3. Apply univariate transformations of the form $\Phi^{-1}(F(\cdot))$ to each scalar state variable in the ensembles (a different transform for each variable) to get transformed ensembles $\hat{\mathbf{x}}_i$, $\hat{\mathbf{y}}_i$ and a transformed observation vector $\hat{\mathbf{y}}$.
4. Apply the perturbed-observation EnKF in the transformed space using the formula

$$\hat{\mathbf{x}}_i^+ = \hat{\mathbf{x}}_i + \text{Cov}[\hat{\mathbf{X}}, \hat{\mathbf{Y}}] \text{Cov}[\hat{\mathbf{Y}}]^{-1} (\hat{\mathbf{y}} - \hat{\mathbf{y}}_i). \quad (11)$$

In this update step, use ensemble approximations to $\text{Cov}[\hat{\mathbf{X}}, \hat{\mathbf{Y}}]$ and $\text{Cov}[\hat{\mathbf{Y}}]$.

5. Apply inverse transformations of the form $F^{-1}(\Phi(\cdot))$ to the updated ensemble members.

A simple yet crucial aspect of the proposed algorithm is the use of the conditional Gaussian formula (1) rather than the KF formulas (7) and (9). There has been some confusion in the literature surrounding the question of how to use the KF update formulas in the transformed space, since the KF formulas need $\mathbf{H}$ and $\mathbf{R}$, which are not clearly defined in the transformed space. Note that the above GA-EnKF algorithm does not specify details for estimating the marginal cdfs F ; two specific methods for approximating the cdfs F and their inverses are detailed in Section 5.5.

Step 4 uses ensemble approximations of $\text{Cov}[\hat{\mathbf{X}}, \hat{\mathbf{Y}}]$ and $\text{Cov}[\hat{\mathbf{Y}}]$. There is, of course, no need to form these matrices, which are too big to fit in memory for many applications; their factored form in terms of scaled ensemble perturbation matrices can be stored and used instead. Let $\hat{\mathbf{A}}_X$ be a matrix whose columns are

$$\frac{1}{\sqrt{N-1}} (\hat{\mathbf{x}}_i - \hat{\bar{\mathbf{x}}}) \quad (12)$$

and similarly for $\hat{\mathbf{A}}_Y$. Then the update formula becomes

$$\hat{\mathbf{x}}_i^+ = \hat{\mathbf{x}}_i + \hat{\mathbf{A}}_X \hat{\mathbf{A}}_Y^T (\hat{\mathbf{A}}_Y \hat{\mathbf{A}}_Y^T)^{-1} (\hat{\mathbf{y}} - \hat{\mathbf{y}}_i). \quad (13)$$

If the observation errors are independent, serial scalar assimilation can be used in the transformed space. If simultaneous assimilation is desired or required and the ensemble size N is small enough that $\hat{\mathbf{A}}_Y \hat{\mathbf{A}}_Y^T$ is not invertible, then Schur-product localization (or some other regularization, e.g. [86]) can be used to generate an invertible approximation to $\text{Cov}[\hat{\mathbf{Y}}]$.

Note that the transformations used in GA are univariate, and thus do not change the spatial locations of the variables, i.e. the spatial location of the j^{th} element of $\mathbf{x}$ is the same as the spatial location of the j^{th} element of $\hat{\mathbf{x}}$. It is thus reasonable to suppose that the optimal length scales for localization in the transformed space may be similar to those in the original space. Similarly, the fact that the transformations are univariate implies that observation errors that are independent before transformation remain independent after transformation, which is important if one intends to use serial assimilation in the transformed space.

Deterministic EnKFs can also be designed based on the conditional Gaussian formulas (1) and (2). For example, an ETKF update [20, 104] of the scaled transformed ensemble perturbation matrix $\hat{\mathbf{A}}_X$ could be accomplished as follows. To exactly satisfy (2), the updated scaled transformed ensemble perturbation matrix $\hat{\mathbf{A}}_{X,+}$ should satisfy

$$\hat{\mathbf{A}}_{X,+} \hat{\mathbf{A}}_{X,+}^T = \hat{\mathbf{A}}_X \left(\mathbf{I} - \hat{\mathbf{A}}_Y^T (\hat{\mathbf{A}}_Y \hat{\mathbf{A}}_Y^T)^{-1} \hat{\mathbf{A}}_Y \right) \hat{\mathbf{A}}_X^T. \quad (14)$$

If $\mathbf{T}$ is a symmetric square root

$$\mathbf{T} = \left(\mathbf{I} - \hat{\mathbf{A}}_Y^T (\hat{\mathbf{A}}_Y \hat{\mathbf{A}}_Y^T)^{-1} \hat{\mathbf{A}}_Y \right)^{1/2}, \quad (15)$$

then

$$\hat{\mathbf{A}}_{X,+} = \hat{\mathbf{A}}_X \mathbf{T} \quad (16)$$

will exactly satisfy (2). This ETKF update in the transformed space could be localized using either a local analysis [24, 54] or the gain form of the ETKF [21].

3 Two-step Bayesian updating

In a much-cited paper [8], Anderson developed a widely-used two-step framework for ensemble DA based on a combination of Bayesian estimation and least-squares regression. The connection of the two-step framework to the full Bayesian estimation problem is developed here. The two-step framework of [8] is based on serial assimilation of independent scalar observations (though it can be extended to batches of observations), and thus deals with scalar observations y . It is straightforward to generalize to vector observations, but the presentation here is limited to the scalar case for clarity.

First, introduce a scalar random variable Z with the property that Y conditioned on Z is independent of X . In the traditional case where $Y = H(X) + \epsilon$ one could let $Z = H(X)$. If the observation Y is a function of a single element of X , then one could let Z be the observed element of X . A general approach to choosing Z is discussed in Section 5.4.

Next note that Bayes' rule for this problem can be written as follows

$$\begin{aligned} [x|y] &= \int [x, z|y] dz \\ &\propto \int [y|x, z][x, z] dz \\ &= \int ([y|z][z]) [x|z] dz \\ &\propto \int [z|y][x|z] dz \end{aligned} \quad (17)$$

where the notation $[\cdot]$ represents the pdf of a random variable and the proportionality constants are required to normalize the expressions into pdfs. Sampling from the posterior can thus be performed in a two-step process: First draw samples z_i^+ , $i = 1, \dots, N$ from the posterior distribution $Z|Y = y$, then draw x_i^+ from the conditional distribution $X|Z = z_i^+$. No assumptions of Gaussianity have been made in deriving this two-step update.

The first step involves sampling from a Bayesian posterior on a scalar random variable, which can be accomplished by any of a variety of methods, including the ensemble adjustment Kalman filter (EAKF; [7, 43]) as originally proposed by Anderson [8]. Any one of a variety of alternatives could be used, including particle filters, Gaussian mixture models, and GA-EnKFs, though no publications have developed these options in the two-step filtering context. The GIGG-EnKF [19, 83] and the rank histogram filter (RHF; [9]) are specifically developed as options for the first step of the two-step framework. These are all examples of methods that use deterministic transformations to transform a sample from the prior on Z to a sample from the posterior $Z|Y = y$. The two-step

framework in principle also enables the assimilation of categorical observations, i.e. situations where the range of Y is discrete, as long as a method can be constructed for sampling from the scalar posterior $Z|Y = y$.

The second step in the two-step framework was performed using linear regression by [8]. To understand the connection, note that the linear model

$$X = \beta_0 + \beta_1 Z + \eta \quad (18)$$

where η is a centered Gaussian random vector independent of Z and with covariance Σ implies that $X|Z = z \sim \mathcal{N}(\beta_0 + \beta_1 z, \Sigma)$. Anderson [8] finds estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ of the regression coefficients from the prior ensemble x_i, z_i using ordinary least squares (the hat serves to distinguish the true coefficients from their estimates), and then computes the residuals

$$\eta_i = x_i - \hat{\beta}_0 - \hat{\beta}_1 z_i. \quad (19)$$

Approximate samples from the posterior (the distribution is not exactly the posterior unless $\hat{\beta}_j = \beta_j$) are generated as

$$x_i^+ = \hat{\beta}_0 + \hat{\beta}_1 z_i^+ + \eta_i. \quad (20)$$

Re-use of the residuals in this manner avoids the need to estimate the covariance Σ , and essentially amounts to a deterministic method of transforming an ensemble sampled from $X|Z = z_i$ to samples from $X|Z = z_i^+$.

The above discussion clarifies the relationship between regression in the second step of the two-step framework of [8] and the full Bayesian update, and shows the way to enabling EnKF-like algorithms for a wide range of distributions. For example one might use a general linear model of the form

$$X = \eta + \sum_{j=1}^J \beta_j \phi_j(z) = \eta + \mu(z) \quad (21)$$

where the coefficients β_j are estimated using ordinary least squares. Polynomial regression is a simple example where $\phi_j(z) = z^j$, though polynomial regression is likely to prove unstable whenever any of the z_i^+ lies outside the range of the prior ensemble z_i . General linear models correspond to a conditional distribution of the form $X|Z = z \sim \mathcal{N}(\mu(z), \Sigma)$. The model could be updated by incorporating heteroscedasticity, i.e. by allowing the variance of η to be a function of z , which would update the conditional distribution to $X|Z = z \sim \mathcal{N}(\mu(z), \Sigma(z))$. In a model incorporating heteroscedasticity, the residuals η_i would not be simply re-used in the update, but would instead require element-wise transformations of η_i . A range of non-Gaussian conditional distributions could be accommodated by generalized linear models of the form

$$g(X) = \eta + \sum_{j=1}^J \beta_j \phi_j(z) \quad (22)$$

where g is a vector link function, which must be invertible. The link function could be used, for example, to maintain bounds on the state variables. Extensions to heteroscedasticity and generalized linear models would require techniques more sophisticated than ordinary least squares to estimate the coefficients β_j , but the overall two-step framework remains the same:

1. Estimate the unknown model coefficients β_j using the data in the prior ensemble
2. Obtain samples z_i^+ from $Z|Y = y$
3. Obtain samples from the posterior using

$$x_i^+ = g^{-1} \left(\eta_i^+ + \sum_{j=1}^J \beta_j \phi_j(z_i^+) \right) \quad (23)$$

where, for a homoscedastic model, $\eta_i^+ = \eta_i$.

The methods of [10, 11, 75] are aimed to replace linear regression in the second step with nonlinear transformations. Novel models for the second step of the two-step process are not explored further here.

4 An improved rank histogram filter

This section reviews a specific algorithm for the first step of the two-step process above, namely the Rank Histogram Filter (RHF; [9]), and then proposes a similar algorithm that is expected to be more accurate at small ensemble sizes.

4.1 The rank histogram filter

The RHF [9] transforms samples from the prior on Z into samples from the posterior $Z|Y = y$ using an approximation of

$$z_i^+ = F_+^{-1}(F_Z(z_i)) \quad (24)$$

where F_Z is the prior cdf of Z and F_+ is the posterior cdf of $Z|Y = y$. Both cdfs are estimated from piecewise approximations of the prior and the likelihood. Let $z_{k_1} < \dots < z_{k_N}$ be a sorted permutation of the samples $z_1, \dots, z_N$. (Note that for non-atomic distributions the samples z_i can be sorted without ties with probability 1.) To represent the prior, the range of Z is broken into $N + 1$ subintervals where the left and right boundaries are the boundaries of the domain, and the interior boundaries are z_{k_i} . Each subinterval is assumed to account for total probability $1/(N + 1)$. The prior pdf is represented as piecewise-constant on each of the interior subintervals:

$$p(z) = \frac{z_{k_{i+1}} - z_{k_i}}{N + 1} \text{ for } z \in [z_{k_i}, z_{k_{i+1}}). \quad (25)$$

In the tails $p(z)$ is represented as a Gaussian pdf using the empirical mean and standard deviation of the sample, with

each tail rescaled to have total probability $1/(N + 1)$. (When the range of Z is not $\mathbb{R}$ one can instead set one or both tails to piecewise-constant, as needed.) From this representation it is straightforward to obtain the prior cdf F_Z ; it is only necessary to evaluate this cdf at z_{k_i} , where it takes values $i/(N + 1)$.

The RHF approximates the likelihood $[y|z]$ as a piecewise function defined on the same subintervals as the prior. On the interior subintervals the likelihood is simply the piecewise-linear interpolant of the likelihood evaluated at z_{k_i} . In the original presentation of [9] the likelihood was modeled as Gaussian in the tails, though the simpler form of approximating it as a constant in the tails was adopted more recently by Anderson in [10], and this simpler form is used here. Specifically

$$\ell(z) = \begin{cases} [y|z_{k_1}] & \text{for } z \leq z_{k_1}, \text{ and} \\ [y|z_{k_N}] & \text{for } z_{k_N} \leq z. \end{cases} \quad (26)$$

Note that the notation $\ell(z)$ is used here to denote the approximation to the likelihood, rather than the likelihood itself, which is simply denoted $[y|z]$. With these approximations to the prior and likelihood, the approximation to the posterior becomes a piecewise function with Gaussian tails and a piecewise-linear (but discontinuous) interior:

$$[z|y] \approx \frac{1}{P_0} p(z) \ell(z) \quad (27)$$

where

$$P_0 = \int p(z) \ell(z) dz \quad (28)$$

and the integral is over the range of Z . The integral is available in closed form. The RHF approximation to the posterior cdf is

$$F_+(z) \approx \frac{1}{P_0} \int_{-\infty}^z p(s) \ell(s) ds \quad (29)$$

where s is a dummy integration variable. This approximation is an error function in the tails, and piecewise-quadratic in the interior, and can be easily inverted.

4.2 Improving the rank histogram filter

The RHF described in the preceding section is extremely convenient and quite general, though no proof of its convergence as $N \rightarrow \infty$ has yet been provided. This subsection proposes a different algorithm - an improved RHF (iRHF) - motivated by the form and ease of use of the RHF, but making use of more traditional approximations.

Prior PDF In the first place, the RHF's approximation of the prior pdf is quite novel. A more traditional approach to

scalar density estimation would use kernel density estimation

$$p(z) = \frac{1}{N} \sum_{i=1}^N h_i^{-1} K\left(\frac{z - z_i}{h_i}\right) \quad (30)$$

where the kernel is K and the bandwidth of the i^{th} kernel is h_i . In order to preserve the simplicity and speed of the RHF, it is here proposed to use the most basic kernel: an indicator function (also known as a ‘top hat’ or ‘boxcar’ function)

$$K(s) = \begin{cases} 0 & \text{for } s \leq -\frac{1}{2} \text{ or } s \geq \frac{1}{2}, \text{ and} \\ 1 & \text{for } -\frac{1}{2} < s < \frac{1}{2}. \end{cases} \quad (31)$$

The bandwidth is estimated using the standard rule-of-thumb for the Gaussian [94], with the prefactor modified to fit the top-hat kernel

$$\bar{h} = 3.13 \min \left\{ \hat{\sigma}, \frac{\text{IQR}}{1.34} \right\} N^{-1/5} \quad (32)$$

where $\hat{\sigma}$ denotes the empirical standard deviation of the prior sample, and IQR denotes the inter-quartile range. It is possible with this bandwidth for $p(z)$ to have zero probability density over some subintervals; this happens if there is any i for which $z_{k_{i+1}} - z_{k_i} > \bar{h}$. This would be undesirable, because it would imply that the approximation to the posterior cdf F_+ would no longer be invertible. To avoid this situation, the kernel bandwidth is allowed to vary as follows. For $i = 2, \dots, N - 1$

$$h_i = \frac{1}{2} \max \{ z_{k_{i+1}} - z_{k_i}, z_{k_i} - z_{k_{i-1}}, 2\bar{h} \} \quad (33)$$

and

$$h_1 = \frac{1}{2} \max \{ z_{k_2} - z_{k_1}, 2\bar{h} \}, \quad (34)$$

$$h_N = \frac{1}{2} \max \{ z_{k_N} - z_{k_{N-1}}, 2\bar{h} \}. \quad (35)$$

This approximation, like that of the RHF, is piecewise-constant on the interior subintervals, though this new approximation has $2N - 1$ interior subintervals rather than $N - 1$. The discontinuities occur at z_m , $m = 1 \dots, 2N$, and the locations can be found simply by sorting $z_{k_i} \pm h_i$. The prior cdf F_Z is piecewise linear, and can be evaluated easily, though not quite so easily as for the RHF. This approximation of the prior cdf is used to evaluate $F_Z(z_i)$ in the update $z_i^+ = F_+^{-1}(F_Z(z_i))$. The approximation used here could be updated using, e.g., a triangular or quadratic kernel. Such an approximation would be piecewise-polynomial and therefore easy to handle numerically, but the resulting differences in performance might be small.

Posterior CDF The foregoing approximation to the prior pdf has no tails, i.e. the probability of z is zero outside some bounded interval. This would be problematic whenever the

likelihood is large outside the support of $p(z)$, which is not unusual. The approximation of the prior pdf is therefore updated before using it to generate an approximation of the posterior pdf. Specifically, the new approximation to the prior is

$$\hat{p}(z) = \frac{1}{P_1} \left(\text{tails} + \frac{1}{N} \sum_{i=1}^N h_i^{-1} K\left(\frac{z - z_i}{h_i}\right) \right) \quad (36)$$

where the ‘tails’ are those of a Gaussian pdf using the empirical mean and variance of the prior sample. Unlike in the RHF, these tails are not each normalized to have total probability $1/(N + 1)$. The normalization constant P_1 is used to force total probability 1, and is available in closed form. (When the range of Z is not $\mathbb{R}$, standard kernel density estimation techniques like transforms and reflections rather than Gaussian tails can be used to deal with boundaries [94].)

With this new approximation of the prior pdf, we simply need a piecewise approximation of the likelihood. It is here proposed to use a shape-preserving piecewise-cubic interpolant on the interior subintervals (in the numerical experiments this is provided by Matlab’s `pchip` interpolator), and the same constant approximation as used by the RHF in the tails. Constructing this approximation requires evaluating the likelihood at $2N$ points as compared to the RHF’s N evaluations of the likelihood, which might lead to cost considerations, depending on the application. A shape-preserving interpolant is used to avoid producing negative values of the likelihood in between the interpolation points, which would lead to a non-invertible approximation of the posterior cdf.

With these approximations of the prior and likelihood, the posterior pdf is approximated as a piecewise polynomial with Gaussian tails. The posterior cdf is approximated as

$$F_+(z) \approx \frac{1}{P_2} \int_{-\infty}^z \hat{p}(s) \ell(s) ds \quad (37)$$

where the normalization constant $P_2 = \int_{\mathbb{R}} \hat{p}(z) \ell(z) dz$ is available in closed form, as is the value of $F_+(z)$ on any subinterval. This approximation is an error function in the tails, and piecewise-quartic in the interior.

Inverting the Posterior CDF The piecewise-quartic posterior cdf of the iRHF is not as easy to invert as the piecewise-quadratic posterior cdf of the RHF. A good first guess for a rootfinding algorithm would be to linearly interpolate the posterior cdf at the subinterval boundaries, since it is trivial to invert a piecewise-linear monotonic function. We propose to take this a step further, by simply using this first guess without further rootfinding. (Rootfinding is still carried out for values in the Gaussian tails.) In our numerical experiments we find that this time-saving approximation has no discernible impact on the results.

4.3 Comparing the RHF and iRHF in simple examples

This section compares the RHF and iRHF for Gaussian problems. The RHF and iRHF approximations to a standard normal pdf are first compared by drawing an ensemble of $N = 20$ samples from $\mathcal{N}(0, 1)$. These same samples are used to construct the RHF and iRHF approximations to the prior pdf, which are plotted in Fig. 1, together with the standard normal pdf. Although the RHF and iRHF both use crude piecewise-constant approximations, the iRHF approximation is clearly superior. Because RHF assigns equal probability to regions between ensemble members, whenever two ensemble members are close the pdf has a large spike. For example, there is a large spike in the RHF approximation to the PDF near $z = -1$ in Fig. 1 that extends to a height of approximately 39.

Next consider a scalar problem where the prior on Z is Gaussian with mean μ_z and variance σ_z^2 , and the observation y is a draw from $Y = Z + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \gamma^2)$. Without loss of generality, let $\mu_z = 0$ and $\sigma_z^2 = 1$. In this case, the map from samples of the prior to samples of the posterior is affine and is given by the formula

$$z_i^+ = \frac{y}{\gamma^2 + 1} + \frac{\gamma}{\sqrt{1 + \gamma^2}} z_i = F_+^{-1}(F_Z(z_i)). \quad (38)$$

To explore the performance of the RHF and the iRHF in approximating this affine map we perform a suite of experiments. The observation value y and the observation error γ are both ranged from 0 to 2. For each parameter value we run 100 independent experiments at an ensemble

Fig. 1 The standard normal pdf, together with the approximations produced by the RHF and iRHF. The ensemble used to construct the RHF and iRHF approximations is marked along the horizontal axis

size $N = 20$, and another 100 independent experiments at $N = 80$. For each experiment the maximum absolute error in z_i^+ is recorded; the median over all 100 trials is shown in Fig. 2 as a function of y and γ for both filters and for $N = 20$ and 80. Results using the root mean squared error in z_i^+ are qualitatively similar (not shown).

The most basic conclusion is that for these examples the iRHF performs more accurately than RHF, and that performance improves from $N = 20$ to $N = 80$, though the RHF at $N = 80$ is still not as accurate as the iRHF at $N = 20$. Both filters perform well when the observation error γ is small and the observation y is within one standard deviation of the prior mean. Both filters' performance decreases when either γ increases or when y is far from the prior mean; both of these correspond to situations where the likelihood has large values in the shoulders or tails of the prior distribution. This is not surprising; since the RHF and iRHF approximations are based on samples from the prior, the RHF and iRHF filters will both have worse representations of the transformation in regions where there are few samples. Probably the crudest approximation made in the RHF and iRHF is to treat the likelihood as constant in the tails. Generally speaking, one expects to know the exact functional form of the likelihood function for every observation assimilated, and this information could be used to improve the approximation. That said, it is not clear if there is a general strategy that works better than the constant approximation for a wide range of likelihoods, without incurring too much extra cost.

5 Numerical experiments: configuration

This section describes the configuration of a suite of data assimilation experiments and their results, with the goal of comparing a range of nonlinear extensions of the EnKF. All code used in the numerical experiments is available at [44].

5.1 Model configuration

The dynamical model is the Lorenz-'96 model [71]

$$\frac{dx_k}{dt} = -x_{k-1}(x_{k-2} - x_{k+1}) - x_k + F \quad (39)$$

with $F = 8$, $k = 1, \dots, 40$, and the periodic convention $X_{k \pm 40} = X_k$. A reference solution is initialized from a standard normal distribution and integrated forward by 9 nondimensional Model Time Units (MTUs). According to the classical accounting based on comparing the error doubling time for this model to that for midlatitude synoptic atmospheric dynamics 0.2 MTUs correspond to one 'day' [72]. Starting at $t = 9$ MTU, a reference solution is saved every 0.05 MTUs for a total of 5,500 time points. Time

Fig. 2 Maximum absolute error between the correct value of z_i^+ and the values reported by the RHF (left) and iRHF (right) filters, reported as a median over 100 independent trials, as a function of y (the observation value) and γ (the observation error standard deviation), for both $N = 20$ (upper) and $N = 80$ (lower)

integration is carried out using Matlab's `ode45`, which internally uses a fourth-order Runge-Kutta method with an adaptive step size. Although the degree of nonlinearity in a forecast depends on both the initial distribution and the length of the time window, 0.05 MTUs is fairly short for this model - not long enough to generate a strongly non-Gaussian forecast.

5.2 Observing systems

Three observing systems are configured. The first is direct observations of all model variables

$$\mathbf{y} = \mathbf{x} + \boldsymbol{\epsilon} \quad (40)$$

where $\mathbf{x}$ is the reference solution and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. This configuration is referred to as 'linear observations' and is provided simply for reference, as a means to validate the performance of all methods in a well-studied and nearly-Gaussian problem.

The second observing system takes the form

$$\mathbf{y} = \frac{1}{1 + \exp\{0.5 \times (\mathbf{x} - 2.5) + \boldsymbol{\epsilon}\}} \quad (41)$$

where the observation error is again standard normal $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. These observations are bounded between 0 and 1, and the expression $0.5 \times (\mathbf{x} - 2.5)$ was chosen so that observations will cover nearly the whole range of values from 0 to 1. The error distribution can be considered 'logit-normal' in the sense that

$$\text{logit}(\mathbf{Y}) = \log\left(\frac{\mathbf{Y}}{1 - \mathbf{Y}}\right) \quad (42)$$

is Gaussian for fixed $\mathbf{x}$. This configuration is referred to as 'logit-normal observations.'

Although the observation operator is clearly nonlinear, the likelihood function is actually Gaussian. Nevertheless, the EnKF does not 'approximate the likelihood as Gaussian' - in fact it makes no use of the likelihood function at all. Instead, as noted in Section 2.1, the EnKF approximates the joint distribution of $\mathbf{X}$ and $\mathbf{Y}$ as Gaussian, which is far from true for this observation operator. Whenever $\mathbf{X}$ is normally distributed, the joint distribution of $\mathbf{Y}$ and $\mathbf{X}$ has a Gaussian copula, so GA methods should converge towards the true distribution.

The third observing system takes the form

$$\mathbf{y} = \exp\{0.5 \times |\mathbf{x} - 2.5| + \boldsymbol{\epsilon}\} \quad (43)$$

where the observation error is again standard normal $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The observations are positive, and the error distribution is log-normal in the sense that

$$\log(\mathbf{Y}) = 0.5 \times |\mathbf{x} - 2.5| + \boldsymbol{\epsilon} \quad (44)$$

is Gaussian for fixed $\mathbf{x}$. The likelihood is bimodal because of the absolute value, and the joint distribution of $\mathbf{X}$ and $\mathbf{Y}$ does not have a Gaussian copula when $\mathbf{X}$ is Gaussian. The expression $0.5 \times |\mathbf{x} - 2.5|$ ensures that bimodality primarily occurs for values of $\mathbf{x}$ near 2.5, which is close to the time-mean value of $\mathbf{x}$, so that bimodality occurs frequently.

5.3 Performance metrics

The methods compared are described in the following subsections. Each method generates an ensemble forecast

and an ensemble analysis. The first performance metric is the root mean squared error (RMSE) in the ensemble mean, defined as

$$\left[\frac{1}{40} \sum_{k=1}^{40} (x_k - \bar{x}_k)^2 \right]^{1/2}, \quad \bar{x} = \frac{1}{N} \sum_{i=1}^N x_i \quad (45)$$

where x_k denotes the k^{th} element of the true solution vector $\mathbf{x}$. The forecast and analysis RMSE are stored at each of the 5,500 assimilation cycles, and the median values over the last 5,000 cycles are reported as a summary. The forecast and analysis spread are also tracked; the spread is defined as

$$\left[\frac{1}{40} \sum_{k=1}^{40} \text{Var}_i[x_k] \right]^{1/2} \quad (46)$$

where $\text{Var}_i[x_k]$ denotes the ensemble variance of the variable x_k . Generally speaking, one expects the spread to match the RMSE. Finally, the continuous ranked probability score (CRPS; [40, 46]) is tracked for each variable x_k at each assimilation cycle. The CRPS measures the quality of a probabilistic estimate, with 0 being the best possible score (though it bears noting that the true Bayesian posterior does not have a CRPS of 0).

5.4 Overview of methods

The methods compared here are

- A perturbed-observation EnKF with simultaneous assimilation of all observations. This method is labeled ‘EnKF.’ The analysis update is based on the conditional-Gaussian formula (1) rather than the EnKF formulas (7) & (9). While this is not the best EnKF for the linear-obs setup, it is crucial for the logit-normal and log-normal configurations where the observing system does not take the traditional form $\mathbf{y} = \mathbf{H}(\mathbf{x}) + \epsilon$. Also, in some nonlinear and non-Gaussian cases the perturbed-observation EnKF can be more robust than deterministic versions [64].
- Two perturbed-observation GA-EnKF methods using the algorithm from Section 2.2, which differ in the way they construct the univariate anamorphosis transformations. One method uses a piecewise-linear anamorphosis transformation and is labeled ‘GA-PL’ while the other uses kernel density estimation and is labeled ‘GA-KDE.’ Further details are given in Section 5.5.
- Two two-step methods: The RHF with linear regression and the iRHF with linear regression; these are labeled ‘RHF’ and ‘iRHF’, respectively. These methods use serial assimilation, and the scalar auxiliary variable Z is always chosen to be equal to x_k , where x_k is the variable being observed. This choice is made so that as much

non-Gaussianity as possible can be dealt with in the first part of the two-step cycle, leaving as much Gaussianity as possible for the linear regression step.

For all of these methods the ensemble is initialized by adding draws from a standard normal distribution to the reference solution.

There is a huge range of nonlinear data assimilation methods that could be compared. Since the focus here is on nonlinear and non-Gaussian extensions of the EnKF, particle filters and variational methods are not implemented. Other notable classes of methods that are omitted here include Gaussian mixture methods, e.g. [12, 17, 38, 70], the method of [75], and the Marginal Adjustment RHF [11]. The GIGG-EnKF of [19, 83] is not used because the observations do not have Gamma or Inverse-Gamma distributions.

Polynomial regression was implemented for the second step of the RHF and iRHF filters, but initial tests indicated that it was highly unstable due to the gigantic increments produced whenever the observation lies outside the ensemble spread. For the specific configurations here, where the relationship between Z and X is not strongly nonlinear, regression models more advanced than linear, including the rank regression method of [10], are unlikely to lead to significant improvements in performance.

For small-scale problems such as this one, it is valuable to compute, when possible, a consistent approximation of the Bayesian posterior, e.g. using Markov Chain Monte Carlo (MCMC) [55, 63]. A sequential importance resampling (SIR) particle filter with an ensemble size of $N = 3.072 \times 10^5$ with particle rejuvenation was run as an attempt to obtain a consistent approximation of the true posterior. Even with this extremely large ensemble size the SIR particle filter provided very poor estimates, much worse than the competing methods.

5.5 GA method details

As noted in Section 2.2, the correct transformation of a scalar random variable X to a Gaussian random variable $\hat{X}$ is $\hat{X} = \Phi^{-1}(F_X(X))$ where F_X is the cdf of X and Φ is the cdf of the Gaussian variable $\hat{X}$. (This assumes that F_X is continuous.) A straightforward approach to approximate this transformation is to use kernel density estimation to obtain an approximation to F_X , and then to let Φ be the cdf of a standard normal. The method GA-KDE tested here uses Matlab’s `ksdensity` to estimate $F_X(X)$ for each element of $\mathbf{x}$ and $\mathbf{y}$. The `ksdensity` method uses a Gaussian kernel with the standard rule-of-thumb bandwidth to approximate the prior pdf. The cdf is then approximated as a sum of error functions, and the inverse is obtained using a rootfinding algorithm. This approach would be overly

costly for large-scale applications, but is useful here as a benchmark for careful and accurate approximation of both F_X and its inverse.

The cdf must be estimated for each element of Y , which have bounded support in the logit-normal and log-normal configurations. The default behavior of `ksdensity` for these situations is to transform the data to unbounded variables before applying the standard kernel density estimation procedure. For the logit-normal data `ksdensity` transforms the data using the logit function, and for the log-normal data it transforms the data using the log function.

The straightforward approach of the GA-KDE method has not been previously used in the GA literature. Simon & Bertino [95] use a complicated transformation whose main features are emulated here using a method labeled ‘GA-PL’ where PL stands for piecewise-linear. Simon & Bertino use a very similar estimate of F_X to what is used in the RHF. Letting $x_{k_1} < \dots < x_{k_N}$ be the sorted prior ensemble, and noting that the transformation from X to $\hat{X}$ only needs to be applied to the prior ensemble x_i , simply let

$$F_X(x_{k_i}) = \frac{i}{N+1}. \quad (47)$$

The full transform becomes

$$\hat{x}_{k_i} = \Phi^{-1}\left(\frac{i}{N+1}\right) \quad (48)$$

where Φ is the standard normal cdf. Note that the output of this transformation is independent of the actual values of x_i .

In the case of the observations, one has both an ensemble of perturbed observations y_i , $i = 1, \dots, N$, and the observation itself y . The above formula is used to transform the ensemble of perturbed observations, and the observation y is transformed to $\hat{y}$ by linearly interpolating the map (48). In the case where y falls outside the range of the ensemble y_i direct linear interpolation is not possible. Simon & Bertino use exponential tails to extrapolate, whereas GA-PL simply adds extra points and then linearly interpolates. Specifically, for linear obs GA-PL adds the points $-10 = \Phi^{-1}(F_Y(\bar{y} - 10\sigma_y))$ and $10 = \Phi^{-1}(F_Y(\bar{y} + 10\sigma_y))$ where σ_y is the standard deviation of the ensemble of perturbed observations. For logit-normal observations the values 0 and 1 are mapped to -20 and 20 , respectively. For log-normal observations the value 0 is mapped to -20 , and the value $\bar{y} + 4\sigma_y$ is mapped to 4. (Preliminary testing indicated better performance from mapping $\bar{y} + 4\sigma_y$ to 4 than from mapping $\bar{y} + 10\sigma_y$ to 10.)

The foregoing discussion explained that the map from X to $\hat{X}$ (and from Y to $\hat{Y}$) is piecewise-linear for GA-PL. The inverse map, which is used to transform the analysis ensemble $\hat{X}$ to the analysis ensemble X , is easily obtained since it is also piecewise-linear, and is implemented using linear interpolation.

5.6 Inflation and localization

Inflation and localization are crucial for the performance of EnKF-type methods. For the EnKF, RHF, and iRHF methods inflation is implemented by multiplying the forecast ensemble perturbations by a positive inflation factor r before assimilation

$$x_i \leftarrow \bar{x} + r(x_i - \bar{x}). \quad (49)$$

where $\bar{x}$ is defined in (45). For the GA methods inflation is implemented in the transformed space

$$\hat{x}_i \leftarrow \hat{\bar{x}} + r(\hat{x}_i - \hat{\bar{x}}). \quad (50)$$

All methods use Schur-product localization. The localization matrix is populated using an exponential localization function

$$(\mathbf{L})_{ij} = \exp\left\{-\frac{1}{2}\left(\frac{d_{ij}}{d}\right)^2\right\} \quad (51)$$

where $d_{ij} = \min\{|i - j|, 40 - |i - j|\}$ and d is the localization radius. In the EnKF method the ensemble approximations to both $\text{Cov}[X, Y]$ and $\text{Cov}[Y]$ are multiplied by $\mathbf{L}$ elementwise. In the GA methods the ensemble approximations to both $\text{Cov}[\hat{X}, \hat{Y}]$ and $\text{Cov}[\hat{Y}]$ are multiplied by $\mathbf{L}$ elementwise. In the RHF and iRHF methods, the analysis increments produced by linear regression are localized based on the distance between x_k and z .

The localization matrix $\mathbf{L}$ is not guaranteed to be positive definite using the distance d_{ij} . This distance was replaced by embedding the 40 Lorenz-‘96 variables on a circle in $\mathbb{R}^2$ before computing distances, and tested using the EnKF method in the linear problem. No significant differences were observed.

For all methods, a range of inflation factors and localization radii are tested. Inflation factors and localization radii are tested in the range

$$r = 1 + \frac{k}{20}, k = 0, \dots, 8 \quad (52)$$

$$d = 0.5, 2k - 1 \text{ for } k = 1, \dots, 8, \text{ and } \infty \quad (53)$$

A wide range was required because the optimal parameter values differed significantly between methods.

6 Numerical experiments: results

We begin by reporting results on the linear-obs problem, primarily as a baseline to verify that the methods are working as expected. Results are only reported for an ensemble size of $N = 120$. Some interesting results arise, which are briefly discussed before turning to the main topic of the nonlinear, non-Gaussian problems.

Table 1 Results for all methods in the linear observation problem with ensemble size $N = 120$

Method	Loc. Radius	Inflation	F. RMSE	F. Spread	F. CRPS	A. RMSE	A. Spread	A. CRPS
EnKF	3	1.05	0.28	0.25	0.11	0.26	0.23	0.10
GA-PL	3	1.05	0.29	0.26	0.11	0.26	0.24	0.10
GA-KDE	5	1.10	0.30	0.26	0.11	0.28	0.23	0.10
RHF	15	1	0.19	0.18	0.07	0.17	0.17	0.06
iRHF	∞	1	0.19	0.20	0.07	0.17	0.18	0.06

For each method results are given at the parameter values that produce the optimal analysis RMSE. Forecast RMSE is abbreviated F. RMSE, and Analysis RMSE is abbreviated A. RMSE, etc

6.1 Linear observations

Table 1 presents the results for all methods in the problem with linear observations at the highest ensemble size $N = 120$. All methods perform well, with RMS errors well below the observation error level (which is 1), but the RHF and iRHF perform significantly better than the EnKF and GA methods. This problem is only weakly nonlinear, i.e. the distributions are only weakly non-Gaussian, so the nonlinear extensions to the EnKF are not expected to produce substantial improvements in performance. This suggests that the improvements seen in the RHF and iRHF methods are primarily the result of better stability to sampling errors, which is supported by the smaller optimal inflation factors and larger optimal localization radii. It is worth recalling that the EnKF implemented here is a perturbed-observation version of the EnKF, which can be more susceptible to sampling errors than deterministic EnKFs, which further supports the idea that the difference between the EnKF and RHF is mainly due to the different sensitivity to sampling errors. The GA-PL and GA-KDE methods have the same perturbed-observation EnKF inside the transformed space, so they are also more susceptible to sampling errors. Thus, the three perturbed-observation methods perform similarly, and the two RHF methods perform similarly. The EnKF, GA-PL, and GA-KDE methods have slightly under-dispersed ensembles, while the RHF and iRHF methods have well-spread ensembles. Since the focus

is on non-Gaussian DA, the results of the linear observation setup are not analyzed further.

6.2 Logit-normal observations

Table 2 presents the results for all methods in the problem with logit-normal observations at ensemble size $N = 120$. The EnKF does quite well despite the nonlinear, non-Gaussian observations, with small RMS errors and a slightly under-dispersed ensemble. Adding Gaussian Anamorphosis to the EnKF via the piecewise-linear transform in the GA-PL method degrades the analysis RMSE by 11%, while the GA-KDE method leads to a 5% improvement in analysis RMSE.

As in the linear observation problem, the RHF and iRHF methods are significantly better than the EnKF and GA methods, and they also require less localization and less inflation. They both produce well-balanced ensembles.

The patterns evident at $N = 120$ persist as the ensemble size is varied. Figure 3 shows the time-median analysis RMSE for all methods as a function of ensemble size N . Of particular note is the fact that the RHF and iRHF methods at $N = 20$ still perform significantly better than the other methods at N up to 200. Also, the difference between RHF and iRHF increases as the ensemble size is lowered, with RHF having 17% worse analysis RMSE than iRHF at $N = 20$. Finally, the RHF and iRHF results appear to have converged by $N = 120$, whereas the other three methods

Table 2 Results for all methods in the logit-normal observation problem with ensemble size $N = 120$

Method	Loc. Radius	Inflation	F. RMSE	F. Spread	F. CRPS	A. RMSE	A. Spread	A. CRPS
EnKF	3	1.05	0.60	0.55	0.23	0.55	0.50	0.21
GA-PL	3	1.05	0.67	0.56	0.25	0.61	0.51	0.23
GA-KDE	3	1.05	0.57	0.49	0.22	0.52	0.45	0.20
RHF	9	1	0.42	0.40	0.16	0.39	0.36	0.14
iRHF	15	1	0.41	0.43	0.15	0.38	0.39	0.14

For each method results are given at the parameter values that produce the optimal analysis RMSE. Forecast RMSE is abbreviated F. RMSE, and Analysis RMSE is abbreviated A. RMSE, etc

Fig. 3 Analysis RMSE in the problem with logit-normal observations as a function of ensemble size N for all methods. The GA-KDE method was unstable at $N = 20$ and failed to produce results at any of the parameter values tested

require $N > 120$ to converge, which further underscores that RHF and iRHF are less sensitive to sampling errors.

6.3 Lognormal observations

Table 3 presents the results for all methods in the problem with log-normal observations at ensemble size $N = 120$. The EnKF is highly unstable with only a few parameter values avoiding blowup and returning results. In those cases the spread is close to machine precision, and the RMS errors are enormous: The filter completely ignores the observations. Adding Gaussian Anamorphosis to the EnKF makes a huge difference here. The GA-KDE method does about 13% better than the GA-PL method, though both perform reasonably well. The GA-KDE method has a slightly over-dispersed ensemble while the GA-PL method has a significantly under-dispersed ensemble, though these differences do not end up leading to large differences in the CRPS. It is possible that the optimal inflation parameter for both methods lies between 1.05 and 1.10.

As in the linear and logit-normal observation problems, the RHF and iRHF methods are significantly better than the EnKF and GA methods, and they also require less localization and less inflation. They both produce well-balanced ensembles.

As in the logit-normal problem, the patterns evident at $N = 120$ persist as the ensemble size is varied. Figure 4

Fig. 4 Analysis RMSE in the problem with log-normal observations as a function of ensemble size N for all methods except EnKF, whose RMSE is off the chart. The GA-KDE method was unstable and failed to produce results at any of the parameter values tested for $N < 120$

shows the analysis RMSE for all methods as a function of ensemble size N . Once again the RHF and iRHF methods at $N = 20$ perform significantly better than the other methods at N up to 200: RHF at $N = 20$ is comparable to GA-PL, and GA-KDE at $N = 200$, but iRHF at $N = 20$ remains significantly better than those methods at $N = 200$. The difference between RHF and iRHF increases as the ensemble size is lowered, with RHF having 33% worse analysis RMSE than iRHF at $N = 20$. Finally, the RHF and iRHF results appear to have converged by $N = 80$, whereas the other three methods require $N > 120$ to converge, which again underscores that RHF and iRHF are less sensitive to sampling errors than EnKF and the GA methods.

GA-KDE was particularly unstable on this problem at low ensemble sizes, and diverged to machine infinity before completing 5,500 assimilation cycles for all parameters tested and all $N < 120$. By contrast, GA-PL remained stable and produced better results than EnKF, despite having very poor performance at small ensemble sizes.

7 Conclusions

This paper reviews two extensions of the EnKF to nonlinear and non-Gaussian problems. The first class of methods, Gaussian anamorphosis (GA) methods, use univariate

Table 3 Results for all methods in the log-normal observation problem with ensemble size $N = 120$

Method	Loc. Radius	Inflation	F. RMSE	F. Spread	F. CRPS	A. RMSE	A. Spread	A. CRPS
EnKF	7	1	5.20	0	3.64	5.20	0	3.64
GA-PL	3	1.05	0.91	0.64	0.33	0.83	0.59	0.31
GA-KDE	3	1.10	0.78	0.87	0.31	0.72	0.80	0.29
RHF	11	1	0.44	0.42	0.16	0.41	0.39	0.15
iRHF	11	1	0.45	0.47	0.17	0.41	0.43	0.15

For each method results are given at the parameter values that produce the optimal analysis RMSE. Forecast RMSE is abbreviated F. RMSE, and Analysis RMSE is abbreviated A. RMSE, etc

transformations to make the univariate marginal distributions of the state and observation vectors Gaussian before applying an EnKF, followed by the inverse transformation back to the original variables. GA-EnKF methods should presumably converge in the large-ensemble limit for problems with a Gaussian copula. The second class of methods is based on a two-step framework introduced by Anderson [8]; the first step is a scalar Bayesian update and the second step is a regression problem. The equivalence of the two-step process to the Bayesian update was previously known for Gaussian problems; this review clarifies the connection of the two-step process to the original Bayesian problem with no assumption of Gaussianity. The new perspective on the two-step process enables a class of EnKF-like data assimilation algorithms that can reasonably be expected to converge in the large-ensemble limit for a wide class of non-Gaussian problems. No analysis of that limit is performed in this paper, which is focused on understanding the methods, clarifying the situations when they are expected to work, and providing examples showing their success in nonlinear, non-Gaussian problems.

Two specific GA-EnKF methods are described and implemented, which differ in the empirical approximations used for their univariate transforms. GA-PL uses a piecewise-linear transformation modeled on the method from [95], while GA-KDE uses standard kernel density estimation tools to construct the transformations. Both of the GA-EnKF methods use perturbed-observation EnKFs, and a matching perturbed-observation EnKF is implemented without the transforms for comparison. Although deterministic EnKFs are sometimes more accurate than perturbed-observation EnKFs, the latter are sometimes more well-behaved in nonlinear and non-Gaussian problems [64]. In addition to the RHF version of the two-step method from [9], a new first step is developed, called iRHF for ‘improved’ RHF. In the experiments reported here, both the RHF and iRHF methods use simple linear regression for the second step of the two-step process. Three suites of numerical experiments are described using the Lorenz-‘96 model [71]: one using linear observations with Gaussian errors, one using logit-normal observations, and one using log-normal observations with a bimodal likelihood. The logit-normal observation configuration has a Gaussian copula, so the GA-EnKF methods should perform well, while the log-normal observation configuration does not have a Gaussian copula.

In the linear observation experiments the EnKF, GA-PL, and GA-KDE methods perform similarly. The RHF and iRHF methods perform similarly to each other, and are both more accurate than the other three methods. This indicates that RHF and iRHF are less sensitive to sampling errors than the EnKF and GA-EnKF methods, though the differences could also be associated with some weak non-Gaussianity from the nonlinear forecast.

In the logit-normal observation experiments, EnKF does quite well. GA-PL is actually worse than EnKF, though this is evidently because of the low quality of the transformations since GA-KDE performs better than EnKF. The RHF and iRHF methods are very similar to each other, and are both significantly more accurate than the other three methods.

In the log-normal observation experiment the EnKF is completely unable to produce a meaningful result; it either diverges to machine infinity, or ignores the observations completely. The GA-PL and GA-KDE methods do quite well, with GA-KDE performing slightly better than GA-PL due to the more-accurate transformations. This is an example of a problem where GA methods should not converge to the correct posterior, but where they still lead to significant gains over the EnKF. By contrast, the RHF and iRHF methods perform similarly and far more accurately than the GA-EnKF methods.

The foregoing discussion notes that RHF and iRHF perform similarly to each other in all tests *at large ensemble sizes*. For small ensembles on the order of $N = 20$ iRHF is significantly more accurate than RHF. This is because the approximations used in the RHF and iRHF are both accurate for large ensemble sizes, but the iRHF uses approximations that remain accurate at small ensemble sizes. The ensemble size where differences between RHF and iRHF arise depends on the details of the problem at hand. The iRHF results at $N = 20$ were better than the EnKF and GA-EnKF results at $N = 200$ in the experiments with logit-normal and log-normal observations.

Future work in this area could be directed towards different regression models for the second step of the two-step process, and towards rigorously proving convergence of GA and two-step methods for classes of non-Gaussian distributions. Some preliminary results were computed using cubic polynomial regression for the second step of the RHF and iRHF filters. These were uniformly more unstable than linear regression due to the large increments produced by a polynomial when extrapolating outside the range of the data, which happens occasionally during a long string of DA cycles. The simple linear regression model used here was quite accurate in all experiments, so the development of more advanced regression models is left to a future work. A simple example where regression methods more advanced than simple linear models would be needed is when one or more of the state variables has a bounded domain, e.g. where concentrations must remain positive.

Acknowledgements The author is grateful to J. L. Anderson for discussions on the two-step framework and the RHF, and to M. El Gharamti for discussions of Gaussian anamorphosis methods. Two anonymous reviewers offered suggestions for clarification and improvement of the presentation. This work utilized resources from the University of Colorado Boulder Research Computing Group, which is

supported by the National Science Foundation (awards ACI-1532235 and ACI-1532236), the University of Colorado Boulder, and Colorado State University.

Declarations

Conflict of Interests The author declares that he has no conflict of interest.

Appendix : Copulas

This appendix provides a brief, incomplete background on copulas. For further details, see [78, 92]. Sklar's theorem [96] relates the joint cdf $H(x_1, \dots, x_d)$ of a vector of random variables to the marginal cdfs $F_k(x_k)$ and a 'copula' C :

$$P[X_1 < x_1, \dots, X_d < x_d] = H(x_1, \dots, x_d) \\ = C(F_1(x_1), \dots, F_d(x_d)). \quad (54)$$

The copula C is thus a function from the d -dimensional hypercube to $[0, 1]$; it is uniquely defined when the marginal cdfs are continuous. The copula encodes the dependence relationships between the variables, in the sense that strictly-increasing transformations of the random variables X_i do not change the copula [78, Theorem 2.3.4]. To give an example relevant to the discussion of GA methods, let $X_1, \dots, X_d$ be jointly Gaussian random variables with zero mean, unit variance, and covariance matrix $\mathbf{C}$. Let $G_1, \dots, G_d$ be arbitrary invertible cdfs, and suppose that we apply a probability integral transform to each X_k , resulting in the new set of random variables $\hat{X}_k = G_k^{-1}(\Phi(X_k))$, where Φ is the standard normal cdf. This new set of random variables $\hat{X}_1, \dots, \hat{X}_d$ is clearly not jointly normal (the marginal pdf of $\hat{X}_k$ is $G'_k(x)$, which is not Gaussian), but the dependence structure between the variables is completely described by a Gaussian copula. There is a huge menagerie of families of copulas describing different kinds of dependence structure between variables. By transforming all the marginal distributions to Gaussian distributions and then applying the EnKF, which assumes that the joint distribution is Gaussian, GA methods are evidently assuming that the distribution of the original variables has a Gaussian copula.

References

- PART II: DATA ASSIMILATION, chap. 2. Integrated Forecasting System Documentation. ECMWF. <https://www.ecmwf.int/node/19306> (2019)
- Aanonsen, S.I., Nævdal, G., Oliver, D.S., Reynolds, A.C., Vallè, s., B., et al.: The ensemble Kalman filter in reservoir engineering—a review. *Spe J.* **14**(03), 393–412 (2009)
- Ades, M., Van Leeuwen, P.J.: An exploration of the equivalent weights particle filter. *Q. J. Roy. Meteor. Soc.* **139**(672), 820–840 (2013)
- Ades, M., Van Leeuwen, P.J.: The equivalent-weights particle filter in a high-dimensional system. *Q. J. Roy. Meteor. Soc.* **141**(687), 484–503 (2015)
- Albers, D.J., Blancquart, P.A., Levine, M.E., Seylabi, E.E., Stuart, A.: Ensemble Kalman methods with constraints. *Inverse Problems* **35**(9), 095007 (2019)
- Amezcu, J., Van Leeuwen, P.J.: Gaussian anamorphosis in the analysis step of the enKF: a joint state-variable/observation approach. *Tellus A* **66**(1), 23493 (2014)
- Anderson, J.L.: An ensemble adjustment Kalman filter for data assimilation. *Mon. Weather Rev.* **129**(12), 2884–2903 (2001)
- Anderson, J.L.: A local least squares framework for ensemble filtering. *Mon. Weather Rev.* **131**(4), 634–642 (2003)
- Anderson, J.L.: A non-Gaussian ensemble filter update for data assimilation. *Mon. Weather Rev.* **138**(11), 4186–4198 (2010)
- Anderson, J.L.: A nonlinear rank regression method for ensemble Kalman filter data assimilation. *Mon. Weather Rev.* **147**(8), 2847–2860 (2019)
- Anderson, J.L.: A marginal adjustment rank histogram filter for non-Gaussian ensemble data assimilation. *Mon. Weather Rev.* **148**(8), 3361–3378 (2020)
- Anderson, J.L., Anderson, S.L.: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Weather Rev.* **127**(12), 2741–2758 (1999)
- Attia, A., Moosavi, A., Sandu, A.: Cluster sampling filters for non-gaussian data assimilation. *Atmosphere* **9**(6), 213 (2018)
- Bannister, R.: A review of operational methods of variational and ensemble-variational data assimilation. *Q. J. Roy. Meteor. Soc.* **143**(703), 607–633 (2017)
- Beal, D., Brasseur, P., Brankart, J.M., Ourmieres, Y., Verron, J.: Characterization of mixing errors in a coupled physical biogeochemical model of the North atlantic: Implications for nonlinear estimation using Gaussian anamorphosis. *Ocean Sci.* **6**, 247–262 (2010)
- Bengtsson, T., Bickel, P., Li, B.: Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems, *Collections*, vol. Volume 2, pp. 316–334. Institute of Mathematical Statistics, Beachwood, Ohio USA. <https://doi.org/10.1214/193940307000000518> (2008)
- Bengtsson, T., Snyder, C., Nychka, D.: Toward a nonlinear ensemble filter for high-dimensional systems. *J. Geophys. Res.-Atmos* **108**(D24) (2003)
- Bertino, L., Evensen, G., Wackernagel, H.: Sequential data assimilation techniques in oceanography. *Int. Stat. Rev.* **71**(2), 223–241 (2003)
- Bishop, C.H.: The GIGG-enKF: ensemble Kalman filtering for highly skewed non-negative uncertainty distributions. *Q. J. Roy. Meteor. Soc.* **142**(696), 1395–1412 (2016)
- Bishop, C.H., Etherton, B.J., Majumdar, S.J.: Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Weather Rev.* **129**(3), 420–436 (2001)
- Bishop, C.H., Whitaker, J.S., Lei, L.: Gain form of the ensemble transform Kalman filter and its relevance to satellite data assimilation with model space ensemble covariance localization. *Mon. Weather Rev.* **145**(11), 4575–4592 (2017)
- Bocquet, M., Pires, C.A., Wu, L.: Beyond Gaussian statistical modeling in geophysical data assimilation. *Mon. Weather Rev.* **138**(8), 2997–3023 (2010)
- Brankart, J.M., Testut, C.E., Béal, D., Doron, M., Fontana, C., Meinville, M., Brasseur, P., Verron, J.: Towards an improved description of ocean uncertainties: Effect of local anamorphic

- transformations on spatial correlations. *Ocean Sci.* **8**(2), 121 (2012)
24. Brusdal, K., Brankart, J.M., Halberstadt, G., Evensen, G., Brasseur, P., van Leeuwen, P.J., Dombrowsky, E., Verron, J.: A demonstration of ensemble-based assimilation methods with a layered OGCM from the perspective of operational ocean forecasting systems. *J. Marine Sys.* **40**, 253–289 (2003)
 25. Burgers, G., van Leeuwen, P., Evensen, G.: Analysis scheme in the ensemble Kalman filter. *Mon. Weather Rev.* **126**(6), 1719–1724 (1998)
 26. Chorin, A., Morzfeld, M., Tu, X.: Implicit particle filters for data assimilation. *Comm App Math Com Sc* **5**(2), 221–240 (2010)
 27. Chorin, A.J., Tu, X.: Implicit sampling for particle filters. *Proc. Natl. Acad. Sci. USA* **106**(41), 17249–17254 (2009)
 28. Chorin, A.J., Tu, X.: An iterative implementation of the implicit nonlinear filter. *ESAIM-Math Model Num.* **46**(3), 535–543 (2012)
 29. Chustagulprom, N., Reich, S., Reinhardt, M.: A hybrid ensemble transform particle filter for nonlinear and spatially extended dynamical systems. *SIAM-ASA. J. Uncertainty Quant.* **4**(1), 592–608 (2016)
 30. Crisan, D., Doucet, A.: A survey of convergence results on particle filtering methods for practitioners. *IEEE T signal proces* **50**(3), 736–746 (2002)
 31. Doron, M., Brasseur, P., Brankart, J.M., Losa, S.N., Melet, A.: Stochastic estimation of biogeochemical parameters from Globcolour ocean colour satellite data in a North Atlantic 3D ocean coupled physical-biogeochemical model. *J. Marine Systems* **117**, 81–95 (2013)
 32. Doucet, A., De Freitas, N., Gordon, N.: An Introduction to Sequential Monte Carlo Methods. In: *Sequential Monte Carlo Methods in Practice*, pp. 3–14. Springer (2001)
 33. Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.-Oceans* **99**(C5), 10143–10162 (1994)
 34. Evensen, G.: Data assimilation: the ensemble kalman filter springer (2009)
 35. Fischer, C., Montmerle, T., Berre, L., Auger, L., ȘtefĂnescu, S.E.: An overview of the variational assimilation in the ALADIN/france numerical weather-prediction system. *Q. J. Roy. Meteor. Soc.* **131**(613), 3477–3492 (2005)
 36. Fossum, K., Mannseth, T., Stordal, A.S.: Assessment of multilevel ensemble-based data assimilation for reservoir history matching. *Computat. Geosci.* **24**(1), 217–239 (2020)
 37. Frei, M., Künsch, H.R.: Bridging the ensemble Kalman and particle filters. *Biometrika* **100**(4), 781–800 (2013)
 38. Frei, M., Künsch, H.R.: Mixture ensemble Kalman filters. *Comput. Stat. Data An.* **58**, 127–138 (2013)
 39. Gharamti, M., Samuelson, A., Bertino, L., Simon, E., Korosov, A., Daewel, U.: Online tuning of ocean biogeochemical model parameters using ensemble estimation techniques: Application to a one-dimensional model in the North Atlantic. *J. Marine Systems* **168**, 1–16 (2017)
 40. Gneiting, T., Raftery, A.E.: Strictly proper scoring rules, prediction, and estimation. *J. Am Stat Assoc* **102**, 359–378 (2007)
 41. Gopalakrishnan, G., Cornuelle, B.D., Hoteit, I., Rudnick, D.L., Owens, W.B.: State estimates and forecasts of the loop current in the Gulf of Mexico using the MITgcm and its adjoint. *J. Geophys. Res.-Oceans* **118**(7), 3292–3314 (2013)
 42. Gordon, N.J., Salmond, D.J., Smith, A.F.: Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In: *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, pp. 107–113. IET (1993)
 43. Grooms, I.: A note on the formulation of the Ensemble Adjustment Kalman Filter. *arXiv:2006.02941* (2020)
 44. Grooms, I.: *iangrooms/Nonlinear-EnKF: Nonlinear EnKF Methods in Lorenz-96*. <https://doi.org/10.5281/zenodo.5889403>. <https://github.com/iangrooms/Nonlinear-EnKF> (2022)
 45. Grooms, I., Robinson, G.: A hybrid particle-ensemble Kalman filter for problems with medium nonlinearity. *PLoS ONE* **16**, e0248266 (2021)
 46. Hersbach, H.: Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather Forecast* **15**, 559–570 (2000)
 47. Hodyss, D.: Accounting for skewness in ensemble data assimilation. *Mon. Weather Rev.* **140**(7), 2346–2358 (2012)
 48. Hodyss, D., Anderson, J.L., Collins, N., Campbell, W.F., Reinecke, P.A.: Quadratic polynomial regression using serial observation processing: Implementation within DART. *Mon. Weather Rev.* **145**(11), 4467–4479 (2017)
 49. Hoteit, I., Pham, D.T., Triantafyllou, G., Korres, G.: A new approximate solution of the optimal nonlinear filter for data assimilation in meteorology and oceanography. *Mon. Weather Rev.* **136**(1), 317–334 (2008)
 50. Houtekamer, P., Mitchell, H.: Data assimilation using an ensemble Kalman filter technique. *Mon. Weather Rev.* **126**(3), 796–811 (1998)
 51. Houtekamer, P.L., Zhang, F.: Review of the ensemble Kalman filter for atmospheric data assimilation. *Mon. Weather Rev.* **144**(12), 4489–4532 (2016)
 52. Hu, C.C., van Leeuwen, P.J.: A particle flow filter for high-dimensional system applications. *Q. J. Roy. Meteor. Soc.* **147**, 2352–2374 (2021)
 53. Huang, X.Y., Xiao, Q., Barker, D.M., Zhang, X., Michalakes, J., Huang, W., Henderson, T., Bray, J., Chen, Y., Ma, Z., et al.: Four-dimensional variational data assimilation for WRF: formulation and preliminary results. *Mon. Weather Rev.* **137**(1), 299–314 (2009)
 54. Hunt, B.R., Kostelich, E.J., Szunyogh, I.: Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D* **230**(1–2), 112–126 (2007)
 55. Iglesias, M.A., Law, K.J., Stuart, A.M.: Evaluation of gaussian approximations for data assimilation in reservoir models. *Computat. Geosci.* **17**(5), 851–885 (2013)
 56. Janjić, T., McLaughlin, D., Cohn, S.E., Verlaan, M.: Conservation of mass and preservation of positivity with ensemble-type Kalman filter algorithms. *Mon. Weather Rev.* **142**(2), 755–773 (2014)
 57. Jazwinski, A.H.: *Stochastic processes and filtering theory* academic press (1970)
 58. Kalnay, E.: *Atmospheric modeling, data assimilation, and predictability*. Cambridge University Press (2002)
 59. Kotsuki, S., Miyoshi, T., Terasaki, K., Lien, G.Y., Kalnay, E.: Assimilating the global satellite mapping of precipitation data with the Nonhydrostatic Icosahedral Atmospheric Model (NICAM). *J. Geophys. Res. Atmos.* **122**(2), 631–650 (2017)
 60. Kwiatkowski, E., Mandel, J.: Convergence of the square root ensemble Kalman filter in the large ensemble limit. *SIAM-ASA. J. Uncertainty Quant.* **3**(1), 1–17 (2015)
 61. Lauvernet, C., Brankart, J.M., Castruccio, F., Broquet, G., Brasseur, P., Verron, J.: A truncated Gaussian filter for data assimilation with inequality constraints: Application to the hydrostatic stability condition in ocean models. *Ocean Model.* **27**(1–2), 1–17 (2009)
 62. Law, K., Stuart, A., Zygalakis, K.: *Data assimilation a mathematical introduction* springer heidelberg (2015)
 63. Law, K.J., Stuart, A.M.: Evaluating data assimilation algorithms. *Mon. Weather Rev.* **140**(11), 3757–3782 (2012)
 64. Lei, J., Bickel, P., Snyder, C.: Comparison of ensemble Kalman filters under non-Gaussianity. *Mon. Weather Rev.* **138**(4), 1293–1306 (2010)

65. Li, L., Zhou, H., Hendricks Franssen, H.J., Gómez-Hernández, J.J.: Groundwater flow inverse modeling in non-MultiGaussian media: performance assessment of the normal-score Ensemble Kalman Filter. *Hydrol. Earth Syst. Sci.* **16**(2), 573–590 (2012)
66. Li, R., Jan, N.M., Huang, B., Prasad, V.: Constrained ensemble Kalman filter based on Kullback–Leibler divergence. *J. Process Control* **81**, 150–161 (2019)
67. Lien, G.Y., Kalnay, E., Miyoshi, T.: Effective assimilation of global precipitation: Simulation experiments. *Tellus A* **65**(1), 19915 (2013)
68. Lien, G.Y., Kalnay, E., Miyoshi, T., Huffman, G.J.: Statistical properties of global precipitation in the NCEP GFS model and TMPA observations for data assimilation. *Mon. Weather Rev.* **144**(2), 663–679 (2016)
69. Lien, G.Y., Miyoshi, T., Kalnay, E.: Assimilation of TRMM multisatellite precipitation analysis with a low-resolution NCEP global forecast system. *Mon. Weather Rev.* **144**(2), 643–661 (2016)
70. Liu, B., Ait-El-Fquih, B., Hoteit, I.: Efficient kernel-based ensemble Gaussian mixture filtering. *Mon. Weather Rev.* **144**(2), 781–800 (2016)
71. Lorenz, E.: Predictability: a problem partly solved. In: *Proceedings of Seminar on Predictability*, vol. 1, pp. 1–18. ECMWF, Reading, UK (1996)
72. Lorenz, E.N., Emanuel, K.A.: Optimal sites for supplementary weather observations: Simulation with a small model. *J. Atmos. Sci.* **55**(3), 399–414 (1998)
73. Mandel, J., Cobb, L., Beezley, J.D.: On the convergence of the ensemble Kalman filter. *Appl. Math.* **56**(6), 533–541 (2011)
74. Marzouk, Y., Moselhy, T., Parno, M., Spantini, A.: *Sampling via Measure Transport: An Introduction*, pp. 1–41. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-11259-6_23-1
75. Metref, S., Cosme, E., Snyder, C., Brasseur, P.: A non-Gaussian analysis scheme using rank histograms for ensemble data assimilation. *Nonlinear Proc. Geoph.* **21**(4), 869–885 (2014)
76. Moore, A.M., Arango, H.G., Broquet, G., Powell, B.S., Weaver, A.T., Zavala-Garay, J.: The Regional Ocean Modeling System (ROMS) 4-dimensional variational data assimilation systems: Part I—system overview and formulation. *Prog. Oceanogr.* **91**(1), 34–49 (2011)
77. Morzfeld, M., Hodyss, D., Poterjoy, J.: Variational particle smoothers and their localization. *Q. J. Roy. Meteor. Soc.* **144**(712), 806–825 (2018)
78. Nelson, R.B.: *An introduction to copulas*, second edn. Springer (2006)
79. Nino-Ruiz, E.D., Cheng, H., Beltran, R.: A robust non-Gaussian data assimilation method for highly non-linear models. *Atmosphere* **9**(4), 126 (2018)
80. Penny, S.G., Miyoshi, T.: A local particle filter for high-dimensional geophysical systems. *Nonlinear Proc. Geoph.* **23**(6), 391–405 (2016)
81. Pham, D.T., Verron, J., Gourdeau, L.: Singular evolutive Kalman filters for data assimilation in oceanography. *Comptes Rendus de l'Academie des Sciences Series IIA Earth and Planetary Science* **4**(326), 255–260 (1998)
82. Pinheiro, F.R., van Leeuwen, P.J., Geppert, G.: Efficient nonlinear data assimilation using synchronization in a particle filter. *Q. J. Roy. Meteor. Soc.* **145**(723), 2510–2523 (2019)
83. Posselt, D.J., Bishop, C.H.: Nonlinear data assimilation for clouds and precipitation using a gamma inverse-gamma ensemble filter. *Q. J. Roy. Meteor. Soc.* **144**(716), 2331–2349 (2018)
84. Poterjoy, J.: A localized particle filter for high-dimensional nonlinear systems. *Mon. Weather Rev.* **144**(1), 59–76 (2016)
85. Poterjoy, J., Wicker, L., Buehner, M.: Progress toward the application of a localized particle filter for numerical weather prediction. *Mon. Weather Rev.* **147**(4), 1107–1126 (2019)
86. Pourahmadi, M.: Covariance estimation: the GLM and regularization perspectives. *Stat. Sci.* pp. 369–387 (2011)
87. Pulido, M., van Leeuwen, P.J.: Sequential Monte Carlo with kernel embedded mappings: The mapping particle filter. *J. Comput. Phys.* (2019)
88. Rawlins, F., Ballard, S., Bovis, K., Clayton, A., Li, D., Inverarity, G., Lorenc, A., Payne, T.: The Met Office global four-dimensional variational data assimilation scheme. *Q. J. Roy. Meteor. Soc.* **133**(623), 347–362 (2007)
89. Rebeschini, P., Van Handel, R.: Can local particle filters beat the curse of dimensionality? *Ann Appl Probab* **25**(5), 2809–2866 (2015)
90. Robert, S., Künsch, H.R.: Localizing the ensemble Kalman particle filter. *Tellus A* **69**(1), 1282016 (2017)
91. Robinson, G., Grooms, I., Kleiber, W.: Improving particle filter performance by smoothing observations. *Mon. Weather Rev.* **146**(8), 2433–2446 (2018)
92. Schölzel, C., Friederichs, P.: Multivariate non-normally distributed random variables in climate research – introduction to the copula approach. *Nonlinear Proc. Geoph.* **15**(5), 761–772 (2008)
93. Sesterhenn, J.L.: Adjoint-based Data Assimilation of an Epidemiology Model for the Covid-19 Pandemic in 2020. *arXiv e-prints* (2020)
94. Silverman, B.: *Density estimation for statistics and data analysis* CRC Press (1998)
95. Simon, E., Bertino, L.: Gaussian anamorphosis extension of the DEnKF for combined state parameter estimation: Application to a 1D ocean ecosystem model. *J. Marine Systems* **89**(1), 1–18 (2012)
96. Sklar, A.: Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris* **8**, 229–231 (1959)
97. Slivinski, L.C., Compo, G.P., Whitaker, J.S., Sardeshmukh, P.D., Giese, B.S., McColl, C., Allan, R., Yin, X., Vose, R., Titchner, H., et al.: Towards a more reliable historical reanalysis: Improvements for version 3 of the Twentieth Century Reanalysis system. *Q. J. Roy. Meteor. Soc.* **145**(724), 2876–2908 (2019)
98. Snyder, C., Bengtsson, T., Bickel, P., Anderson, J.: Obstacles to high-dimensional particle filtering. *Mon. Weather Rev.* **136**(12), 4629–4640 (2008)
99. Snyder, C., Bengtsson, T., Morzfeld, M.: Performance bounds for particle filters using the optimal proposal. *Mon. Weather Rev.* **143**(11), 4750–4761 (2015)
100. Spantini, A., Baptista, R., Marzouk Y.: Coupling techniques for nonlinear ensemble filtering (2019)
101. Stordal, A.S., Karlsen, H.A., Nævdal, G., Skaug, H.J., Vallè, s., B.: Bridging the ensemble Kalman filter and particle filters: the adaptive Gaussian mixture filter. *Computat. Geosci.* **15**(2), 293–305 (2011)
102. Talagrand, O.: Variational Assimilation. In: Lahoz, W., Khattatov, B., Menard, R. (eds.) *Data Assimilation Making Sense of Observations*, pp. 41–67. Springer (2010)
103. Van Leeuwen, P.J.: Particle filtering in geophysical systems. *Mon. Weather Rev.* **137**(12), 4089–4114 (2009)
104. Wang, X., Bishop, C.H., Julier, S.J.: Which is better, an ensemble of positive-negative pairs or a centered spherical simplex ensemble? *Mon. Weather Rev.* **132**(7), 1590–1605 (2004)
105. Wang, X., Lei, T.: GSI-Based four-dimensional ensemble-variational (4DEnsvar) data assimilation: Formulation and single-resolution experiments with real data for NCEP Global Forecast System. *Mon. Weather Rev.* **142**(9), 3303–3325 (2014)

106. Whitaker, J.S., Hamill, T.M.: Ensemble data assimilation without perturbed observations. *Mon. Weather Rev.* **130**(7), 1913–1924 (2002)
107. Wikle, C.K., Berliner, L.M.: A Bayesian tutorial for data assimilation. *Physica D* **230**(1), 1–16 (2007)
108. Zhou, H., Gomez-Hernandez, J.J., Franssen, H.J.H., Li, L.: An approach to handling non-Gaussianity of parameters and state variables in ensemble Kalman filtering. *Adv. Water Resour.* **34**(7), 844–864 (2011)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.