

REVIEW

Open Access



The Ensemble Kalman filter: a signal processing perspective

Michael Roth* , Gustaf Hendeby, Carsten Fritsche and Fredrik Gustafsson

Abstract

The ensemble Kalman filter (EnKF) is a Monte Carlo-based implementation of the Kalman filter (KF) for extremely high-dimensional, possibly nonlinear, and non-Gaussian state estimation problems. Its ability to handle state dimensions in the order of millions has made the EnKF a popular algorithm in different geoscientific disciplines. Despite a similarly vital need for scalable algorithms in signal processing, e.g., to make sense of the ever increasing amount of sensor data, the EnKF is hardly discussed in our field.

This self-contained review is aimed at signal processing researchers and provides all the knowledge to get started with the EnKF. The algorithm is derived in a KF framework, without the often encountered geoscientific terminology. Algorithmic challenges and required extensions of the EnKF are provided, as well as relations to sigma point KF and particle filters. The relevant EnKF literature is summarized in an extensive survey and unique simulation examples, including popular benchmark problems, complement the theory with practical insights. The signal processing perspective highlights new directions of research and facilitates the exchange of potentially beneficial ideas, both for the EnKF and high-dimensional nonlinear and non-Gaussian filtering in general.

1 Introduction

Numerical weather prediction [1] is an extremely high-dimensional geoscientific state estimation problem. The state x comprises physical quantities (temperature, wind speed, air pressure, etc.) at many spatially distributed grid points, which often yields a state dimension n in the order of millions. Consequently, the Kalman filter (KF) [2, 3] or its nonlinear extensions [4, 5] that require the storage and processing of $n \times n$ covariance matrices cannot be applied directly. It is well-known that the application of particle filters [6, 7] is not feasible either. In contrast, the ensemble Kalman filter [8, 9] (EnKF) was specifically developed as algorithm for high-dimensional n .

The EnKF

- is a random-sampling implementation of the KF;
- reduces the computational complexity of the KF by propagating an ensemble of $N < n$ state realizations;
- can be applied to nonlinear state-space models without the need to compute Jacobian matrices;

- can be applied to continuous-time as well as discrete-time state transition functions;
- can be applied to non-Gaussian noise densities;
- is simple to implement;
- does not converge to the Bayesian filtering solution for $N \rightarrow \infty$ in general;
- often requires extra measures to work in practice.

Also in the field of stochastic signal processing (SP) and Bayesian state estimation, high-dimensional problems become more and more relevant. Examples include SLAM [10] where x contains an increasing number of landmark positions, or extended target tracking [11, 12] where x can contain many parameters to describe the shape of the target. Furthermore, scalable SP algorithms are required to make sense of the ever increasing amount of data from sensors in everyday devices.

EnKF approaches hardly appear in the relevant SP journals, though. In contrast, vivid theoretical development is documented in geoscientific journals under the umbrella term data assimilation (DA) [1]. Hence, a relevant SP problem is being addressed with only little participation from the SP community. Conversely, much of the DA literature makes little reference to relevant SP contributions. It is our intention to bridge this interesting gap.

*Correspondence: michael.roth@liu.se
Dept. of Electrical Engineering, Linköping University, Linköping, Sweden

There is further overlap that motivates for a closer investigation of the EnKF. First, the basic EnKF [9] can be applied to nonlinear and non-Gaussian state-space models because it is entirely sampling based. In fact, the state evolution in geoscientific applications is typically governed by large nonlinear black box prediction models derived from partial differential equations. Furthermore, satellite measurements in weather applications are nonlinearly related to the states [1]. Hence, the EnKF has long been investigated as nonlinear filter. Second, the EnKF literature contains so called localization methods [13, 14] to systematically approach high-dimensional problems by only acting on a part of the state vector in each measurement update. These ideas can be directly transferred to sigma point filters [5]. Third, the EnKF offers several interesting opportunities to apply SP techniques, e.g., via the application of bootstrap or regularization methods in the EnKF gain computation.

The contributions of this paper aim at making the EnKF more accessible to SP researchers. We provide a concise derivation of the EnKF based on the KF. A literature review highlights important EnKF papers with their respective contributions and facilitates easier access to the extensive and rapidly developing DA literature on the EnKF. Moreover, we put the EnKF in context with popular SP algorithms such as sigma point filters [4, 5] and the particle filter [6, 7]. Our presentation forms a solid basis for further developments and the transfer of beneficial ideas and techniques between the fields of SP and DA.

The structure of the paper is as follows. After an extensive literature review in Section 2, the EnKF is developed from the KF in Section 3. Algorithmic properties and challenges of the EnKF and the available approaches to face them are discussed in Sections 4 and 5, respectively. Relations to other filtering algorithms are discussed in Section 6. The theoretical considerations are followed by numerical simulations in Section 7 and some concluding remarks in Section 8.

2 Review

The following literature review provides important landmarks for the EnKF novice.

State-space models and the filtering problem have been investigated since the 1960s. Early results include the Kalman filter (KF) [2] as algorithm for linear systems and the Bayesian filtering equations [15] as theoretical solution for nonlinear and non-Gaussian systems. Because the latter approach cannot be implemented in general, approximate filtering algorithms are required. With a leap in computing capacity, the 1990s saw major developments. The sampling-based sigma point Kalman filters [4, 5] started to appear. Furthermore, particle filters [6, 7] were developed to approximately implement

the Bayesian filtering equations through sequential importance sampling.

The first EnKF [8] was proposed in a geoscientific journal in 1994 and introduced the idea of propagating ensembles to mimic the KF. A systematic error that resulted in an underestimated uncertainty was later corrected by processing “perturbed measurements.” This randomization is well motivated in [9] but also used in [13].

Interestingly, [8] remains the most cited EnKF paper¹, followed by the overview article [16] and the monograph [17] by the same author. Other insightful overviews from a geoscientific perspective are [18, 19]. Many practical aspects of operational EnKF for weather prediction and re-analysis are described in [19–21]. Whereas the aforementioned papers were mostly published in geoscientific outlets, a special issue of the *IEEE Control Systems Magazine* appeared with review articles [22–24] and an EnKF case study [25]. Still, the above material was written by EnKF researchers with a geoscientific focus and in the application-specific terminology. Furthermore, references to the recent SP literature and other nonlinear KF variants [5] are scarce.

Some attention has been devoted to the EnKF also beyond the geosciences. Convergence properties for $N \rightarrow \infty$ have been established using different theoretical analyses of the EnKF [26–28]. Statistical perspectives are provided in the thesis [29] and the review [30]. A recommended SP view that connects the EnKF with Bayesian filtering and particle methods, including convergence results for nonlinear systems, is [31]. Examples of the EnKF as tool for tomographic imaging and target tracking are described in [32] and [33], respectively. Brief introductory papers that connect the EnKF with more established SP algorithms include [34] and [35]. The latter also served as basis for this article.

The majority of EnKF advances are still documented in geoscientific publications. Notable contributions include deterministic EnKF that avoid the randomization of [9] and propagate an ensemble of deviations from the ensemble mean [16, 36–38]. Their common basis as square root EnKF and the relation to square root KF [3] is discussed in [39]. The computational advantages in high-dimensional EnKF with small ensembles ($N \ll n$) come at the price of adverse effects, including the risk of filter divergence. The often encountered underestimation of uncertainty can be counteracted with covariance inflation [40]. A scheme with two EnKF in parallel that provide each other with gain matrices to reduce unwanted “inbreeding” has been suggested in [13]. The benefit of such a double EnKF is, however, debated [38, 41]. The low-rank approximation of covariance matrices can yield spurious correlations between supposedly

uncorrelated state components and measurements. Localization techniques such as local measurement updates [13, 16, 42] or covariance tapering [14, 43] let the measurement only affect a part of the state vector. In other words, localization effectively reduces the dimension of each measurement update. Inflation and localization are essential components of operational EnKF [19]. Smoothing algorithms based on the EnKF are discussed in [17] and, more recently, [44]. Approaches that combine variational DA techniques [1] with the EnKF include [45, 46]. A list of further advances in the geoscientific literature is provided in the appendix of [17].

An interesting development for SP researchers is the reconsideration of particle filters (PF) for high-dimensional geoscientific problems, with seemingly little reference to SP literature. An early example is [47]. The well-known challenges, mostly related to the problem of importance sampling in high dimensions, are reviewed in [48, 49]. Several recent approaches [50–52] were successfully tested on a popular EnKF benchmark problem [53] that is also investigated in Section 7.4. Combinations of the EnKF with the deterministic sampling of sigma point filters [5] are given in [54] and [55]. However, the benefit of the unscented transformation [5, 56] in [55] is debated in [57]. Ideas to combine the EnKF with Gaussian mixture approaches are given in [58–60].

3 A signal processing introduction to the ensemble Kalman filter

The underlying framework of our filter presentation are discrete-time state-space models [3, 15]. The Kalman filter and many EnKF variants are built upon the linear model

$$x_{k+1} = Fx_k + Gv_k, \quad (1a)$$

$$y_k = Hx_k + e_k, \quad (1b)$$

with the n -dimensional state x and the m -dimensional measurement y . The initial state x_0 , the process noise v_k , and the measurement noise e_k are assumed to be independent and described by $E(x_0) = \hat{x}_0$, $E(v_k) = 0$, $E(e_k) = 0$, $\text{cov}(x_0) = P_0$, $\text{cov}(v_k) = Q$, and $\text{cov}(e_k) = R$. In the Gaussian case, these moments completely characterize the distributions of x_0 , v_k , and e_k .

Nonlinear relations in the state evolution and measurement equations can be described by a more general model

$$x_{k+1} = f(x_k, v_k), \quad (2a)$$

$$y_k = h(x_k, e_k). \quad (2b)$$

More general noise and initial state distributions can, for example, be characterized by probability density functions $p(x_0)$, $p(v_k)$, and $p(e_k)$.

Both (1) and (2) can be time-varying but the time indices for functions and matrices are omitted for convenience.

3.1 A brief Kalman filter review

The KF is an optimal linear filter [3] for (1) that propagates state estimates $\hat{x}_{k|k}$ and covariance matrices $P_{k|k}$.

The KF time update or prediction is given by

$$\hat{x}_{k+1|k} = F\hat{x}_{k|k}, \quad (3a)$$

$$P_{k+1|k} = FP_{k|k}F^T + GQG^T. \quad (3b)$$

The above parameters can be used to predict the output of (1) and its uncertainty via

$$\hat{y}_{k|k-1} = H\hat{x}_{k|k-1}, \quad (4a)$$

$$S_k = HP_{k|k-1}H^T + R. \quad (4b)$$

The measurement update adjusts the prediction results according to

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k(y_k - \hat{y}_{k|k-1}) \quad (5a)$$

$$= (I - K_kH)\hat{x}_{k|k-1} + K_k y_k, \quad (5b)$$

$$P_{k|k} = (I - K_kH)P_{k|k-1}(I - K_kH)^T + K_k R K_k^T, \quad (5c)$$

with a gain matrix K_k . Here, (5b) resembles a deterministic observer and only requires all eigenvalues of $(I - K_kH)$ inside the unit circle to obtain a stable filter.

The optimal K_k in the minimum variance sense is given by

$$K_k = P_{k|k-1}H^T S_k^{-1} = M_k S_k^{-1}, \quad (6)$$

where M_k is the cross-covariance between the state and output predictions. Alternatives to the covariance update (5c) exist, but the shown Joseph form [3] will simplify the derivation of the EnKF. Furthermore, it is valid for all gain matrices K_k beyond (6) and numerically well-behaved. It should be noted that it is numerically advisable to obtain K_k by solving $K_k S_k = M_k$ rather than explicitly computing S_k^{-1} [61].

3.2 The ensemble idea

The central idea of the EnKF is to propagate an ensemble of $N < n$ (often $N \ll n$) state realizations $\{x_k^{(i)}\}_{i=1}^N$ instead of the n -dimensional estimate $\hat{x}_{k|k}$ and the $n \times n$ covariance $P_{k|k}$ of the KF. The ensemble is processed such that

$$\bar{x}_{k|k} = \frac{1}{N} \sum_{i=1}^N x_k^{(i)} \approx \hat{x}_{k|k}, \quad (7a)$$

$$\bar{P}_{k|k} = \frac{1}{N-1} \sum_{i=1}^N \left(x_k^{(i)} - \bar{x}_{k|k} \right) \left(x_k^{(i)} - \bar{x}_{k|k} \right)^T \approx P_{k|k}. \quad (7b)$$

Reduced computational complexity is achieved because the explicit computation of $\bar{P}_{k|k}$ is avoided in the EnKF recursion. Of course, this reduction comes at the price of a low-rank approximation in (7b) that entails some negative effects and requires extra measures.

For our development, it is convenient to treat the ensemble as an $n \times N$ matrix $X_{k|k}$ with columns $x_k^{(i)}$. This

allows for the compact notation of the ensemble mean and covariance

$$\bar{x}_{k|k} = \frac{1}{N} X_{k|k} \mathbb{1}, \quad (8a)$$

$$\bar{P}_{k|k} = \frac{1}{N-1} \tilde{X}_{k|k} \tilde{X}_{k|k}^T, \quad (8b)$$

where $\mathbb{1} = [1, \dots, 1]^T$ is an N -dimensional vector and

$$\tilde{X}_{k|k} = X_{k|k} - \bar{x}_{k|k} \mathbb{1}^T = X_{k|k} \left(I_N - \frac{1}{N} \mathbb{1} \mathbb{1}^T \right) \quad (9)$$

is an ensemble of deviations from $\bar{x}_{k|k}$, sometimes called ensemble anomalies [17]. The matrix multiplication in (9) provides a compact way to write the anomalies but is not the most efficient way to compute them.

3.3 The EnKF time update

The EnKF time update is referred to as forecast in the geoscientific literature. In analogy to (3), a prediction ensemble $X_{k+1|k}$ is computed that carries the information in $\hat{x}_{k+1|k}$ and $P_{k+1|k}$. An ensemble of N independent process noise realizations $\left\{ v_k^{(i)} \right\}_{i=1}^N$ with zero mean and covariance Q , stored as matrix V_k , is used in

$$X_{k+1|k} = F X_{k|k} + G V_k. \quad (10)$$

An extension to nonlinear state transitions (2a) is given by

$$X_{k+1|k} = f(X_{k|k}, V_k), \quad (11)$$

where we generalized f to act on the columns of its input matrices. Apparently, the EnKF time update amounts to a one-step-ahead simulation of $X_{k|k}$. Consequently, also continuous-time dynamics can be considered by, for example, numerically solving partial differential equations to obtain $X_{k+1|k}$. Also non-Gaussian initial state and process noise distributions with arbitrary densities $p(x_0)$ and $p(v_k)$ can be employed as long as they allow sampling. Perhaps because of this flexibility, the time update is often omitted in the EnKF literature [9, 13].

3.4 The EnKF measurement update

The EnKF measurement update is referred to as analysis in the geoscientific literature. A prediction or forecast ensemble $X_{k|k-1}$ is processed to obtain the filtering ensemble $X_{k|k}$ that encodes the KF mean and covariance. We assume that a gain $\bar{K}_k = K_k$ is given and postpone its computation to the next section.

With \bar{K}_k available, the KF update (5b) can be applied to each ensemble member as follows [8]

$$X_{k|k} = (I - \bar{K}_k H) X_{k|k-1} + \bar{K}_k y_k \mathbb{1}^T. \quad (12)$$

The resulting ensemble average (8a) is the KF mean $\hat{x}_{k|k}$ of (5b). However, with $y_k \mathbb{1}^T$ known, the sample covariance (8b) of $X_{k|k}$ gives only the first term of (5c). Hence, $X_{k|k}$ fails to carry the information in $P_{k|k}$.

A solution [9] is to account for the missing term $\bar{K}_k R \bar{K}_k^T$ by adding artificial zero-mean measurement noise realizations $\left\{ e_k^{(i)} \right\}_{i=1}^N$ with covariance R , stored as matrix E_k , according to

$$X_{k|k} = (I - \bar{K}_k H) X_{k|k-1} + \bar{K}_k y_k \mathbb{1}^T - \bar{K}_k E_k. \quad (13)$$

Then, $X_{k|k}$ has the correct ensemble mean and covariance, $\hat{x}_{k|k}$ and $P_{k|k}$ of (5), respectively. The model (1) is implicit in (13) because the matrix H appears. If we, in analogy to (4), define a predicted output ensemble

$$Y_{k|k-1} = H X_{k|k-1} + E_k \quad (14)$$

that encodes $\hat{y}_{k|k-1}$ and S_k , we can reformulate (13) to an update that resembles (5a):

$$X_{k|k} = X_{k|k-1} + \bar{K}_k \left(y_k \mathbb{1}^T - Y_{k|k-1} \right). \quad (15)$$

In contrast to (13), the update (15) is entirely sampling based. As a consequence, we can extend the algorithm to nonlinear measurement models (2b) by replacing (14) with

$$Y_{k|k-1} = h(X_{k|k-1}, E_k), \quad (16)$$

where we generalized h to accept matrix inputs similar to (11).

In the EnKF literature, the prevailing view of inserting artificial noise is that perturbed measurements $y_k \mathbb{1}^T - E_k$ are processed. This might appear unusual from an SP perspective since it suggests that information is distorted before processing. The introduction of output ensembles $Y_{k|k-1}$, in contrast, yields a direct connection to (4) and highlights the similarities between (15) and (5a).

An interesting point [60] is that the measurement y_k enters linearly in (13) and (15) and merely shifts the ensemble locations. This highlights the EnKF roots in the linear KF in which $P_{k|k}$ also remains unchanged by y_k .

3.5 The EnKF gain

The optimal gain (6) in the KF is computed from the covariance matrices of the predicted state and output. In the EnKF, the required M_k and S_k are not available but must be approximated from the prediction ensembles (10) or (11), and (14) or (16).

A straightforward way to compute the EnKF gain \bar{K}_k is to first compute the deviations or anomalies

$$\tilde{X}_{k|k-1} = X_{k|k-1} \left(I_N - \frac{1}{N} \mathbb{1} \mathbb{1}^T \right), \quad (17a)$$

$$\tilde{Y}_{k|k-1} = Y_{k|k-1} \left(I_N - \frac{1}{N} \mathbb{1} \mathbb{1}^T \right), \quad (17b)$$

and second the sample covariance matrices

$$\bar{M}_k = \frac{1}{N-1} \tilde{X}_{k|k-1} \tilde{Y}_{k|k-1}^T, \quad (17c)$$

$$\bar{S}_k = \frac{1}{N-1} \tilde{Y}_{k|k-1} \tilde{Y}_{k|k-1}^T. \quad (17d)$$

The computations (17) are entirely sampling based, which is useful for the nonlinear case but introduces extra sampling errors. An obvious improvement for additive measurement noise e_k with covariance R is given in Section 5.2, together with the square root EnKF that avoid the insertion of E_k altogether.

Similar to the KF, the gain \bar{K}_k should be obtained from the solution of a linear system of equations

$$\bar{K}_k \tilde{Y}_{k|k-1} \tilde{Y}_{k|k-1}^T = \tilde{X}_{k|k-1} \tilde{Y}_{k|k-1}^T. \quad (18)$$

4 Some properties and challenges of the EnKF

After a brief review of convergence results and the computational complexity of the EnKF, we discuss adverse effects that can occur in EnKF with finite ensemble size N .

4.1 Asymptotic convergence results

In linear Gaussian systems, the EnKF mean and covariance (7) converge to the KF results (5) as $N \rightarrow \infty$. This result has been established from different theoretical perspectives [26–28, 31].

For nonlinear systems, the convergence is not as tangible. An investigation of the EnKF as particle system is given in [31], with the outcome that the EnKF does not give the Bayesian filtering solution except for the linear Gaussian case. An illustration of this property is given in the example of Section 7.2.

4.2 Computational complexity

For the complexity analysis, we assume that we are only interested in the filtering results and that $n > N > m$, that is, the number of measurements is less than the ensemble size and state dimension.

The KF propagates the n -dimensional mean vector $\hat{x}_{k|k}$ and the $n \times n$ covariance matrix $P_{k|k}$ with $n(n+1)/2$ unique entries. These storage requirements of $\mathcal{O}(n^2/2)$ dominate for large $n > m$. The EnKF requires the storage of only nN values. The space required to store the Kalman gain and other intermediate results is similar for the KF and EnKF. A reduction via sequential processing of measurements, as explained in Section 5.1, is possible for both.

For large n , the computational bottleneck of the KF is the covariance time update (3b). Without considering any potential structure in F , slightly less than $\mathcal{O}(n^3)$ floating point operations (flops) are required. Contemporary matrix multiplication routines [61] achieve a reduction to roughly $\mathcal{O}(n^{2.4})$. The EnKF time update requires the propagation of N realizations. If each propagation costs $\mathcal{O}(n^2)$ flops, then time update is achieved in $\mathcal{O}(n^2N)$ flops.

The computation of the KF gain requires $\mathcal{O}(n^2m)$ flops for the computation of M_k and S_k . The solution of (6) for K_k amounts to $\mathcal{O}(m^3)$. The actual measurement update (5) adds further $\mathcal{O}(n^2m)$ flops. For large n , the total cost is $\mathcal{O}(n^2m)$. In contrast, the EnKF parameters \bar{M}_k and

\bar{S}_k can be computed in $\mathcal{O}(nmN)$ flops which, again, dominates the total cost of the measurement update for large n . So, the EnKF flop count scales a factor $\frac{N}{n}$ better.

4.3 Sampling and coupling effects for finite ensemble size

A serious issue in the EnKF is a commonly noted tendency to underestimate the state uncertainty when using $N < n$ ensemble members [13, 18, 19]. In other words, the EnKF becomes over-confident and is likely to diverge [3] for too small N . A number of causes and related effects can be noted.

First, an ensemble $X_{k|k-1}$ with too few members might not cover the relevant regions of the state-space well enough after the time update (10). The underestimated spread persists in the measurement update (13) or (15) and also $X_{k|k}$ shows too little spread.

Second, the ensemble can only transport limited information and provide a sampling covariance $\bar{P}_{k|k}$, (7b) or (8b), of at most rank $N - 1$. Consequently, identically zero entries of $P_{k|k}$ are difficult to reproduce and unwanted spurious correlations show up in $\bar{P}_{k|k}$. An example would be an unreasonably large correlation between the temperature at two distant locations on the globe. Of course, these correlations also affect \bar{M}_k and \bar{S}_k , and thus the EnKF gain \bar{K}_k in (18). As a result, state components that are actually uncorrelated to y_k are erroneously updated in (13) or (15). Again, this leads to a reduction in ensemble spread.

Third, the ensemble members are nonlinearly coupled because the gain (18) is computed from the ensemble. This “inbreeding” [13] increases with each measurement update. An interesting side effect is that the ensemble is not independent and Gaussian, even for linear Gaussian problems. To illustrate this, we combine (18) and (15) to obtain

$$X_{k|k} = X_{k|k-1} + \left(\tilde{X}_{k|k-1} \tilde{Y}_{k|k-1}^T \right) \left(\tilde{Y}_{k|k-1} \tilde{Y}_{k|k-1}^T \right)^{-1} \left(y_k \mathbb{1}^T - Y_{k|k-1} \right) \quad (19)$$

and consider a linear model (1) with $n = 1, H = 1$, and a zero-mean $X_{k|k-1}$. Then, one member of $X_{k|k}$ is given by

$$x_{k|k}^{(i)} = x_{k|k-1}^{(i)} + \frac{\sum_{j=1}^N \left(x_{k|k-1}^{(j)} \right)^2}{\sum_{j=1}^N \left(x_{k|k-1}^{(j)} + e_k^{(j)} \right)^2} \left(y_k - x_{k|k-1}^{(i)} - e_{k|k-1}^{(i)} \right), \quad (20)$$

which clearly shows the nonlinear dependencies that impede Gaussianity of $x_{k|k}^{(i)}$. Although similar conclusions hold for the general case, concise effects on the ensemble spread are difficult to analyze. Some special cases ($n = 1$ and $n = m, H = I, R \propto I$) are investigated in [26] and shown to produce an underestimated $\bar{P}_{k|k}$.

Finally, the random sampling in the measurement update by inserting measurement noise in (14) or (16) adds to the EnKF error budget. The inherent sampling errors can be reduced by using the square root EnKF of Section 5.2.

Experiments suggest that there is a threshold for N above which the EnKF works. A good example is given in [42]. Section 5 discusses methods such as inflation and localization that can reduce this minimum N .

5 Important extensions to the EnKF

The previous section highlighted some of the challenges of the EnKF. Here, we summarize the important extensions that are often essential to achieve a working EnKF with only few ensemble members.

5.1 Sequential updates

For the KF, it is algebraically equivalent to carry out m measurement updates (5) with the scalar components of y_k instead of a batch update with the m -dimensional y_k , if the measurement noise covariance R is diagonal [3]. Although often treated as a side note only, this technique is very useful. It yields a more flexible algorithm with regard to the availability of measurements at each time step k and reduces the computational complexity. After all, the Kalman gain (6) merely requires a scalar division for each component of y_k . An extension to block-diagonal R is imminent.

Motivated by the large number of measurements in geoscientific problems, sequential updates have also been suggested for the EnKF [14]. Because of the randomness inherent to the EnKF, there is no algebraic equivalence between sequential and batch updates. Hence, the order in which measurements are processed has an effect on the filtering results.

Furthermore, an unusual alternative interpretation of sequential updates can be found in the EnKF literature. Namely, measurement updates are carried out “grid point by grid point” [13, 16, 42], that is, an iteration is carried out over state rather than measurement components. We will return to this aspect in Section 5.4.

5.2 Model knowledge in the EnKF and square-root filters

The sampling based derivation of the EnKF in Eqs. (10) through (18) facilitates a compact presentation. However, the randomization through E_k in (14) or (16) adds Monte Carlo sampling errors to the EnKF budget. This section discusses how these errors can be reduced for linear systems (1). Similar results for nonlinear systems with additive noise follow easily. The interpretation of ensembles as (rectangular) matrix square roots is a common theme in the following approaches. In (8b), for instance, $\frac{1}{\sqrt{N-1}}\tilde{X}_{k|k}$ can be seen as an $n \times N$ square root of $\bar{P}_{k|k}$.

A first thing to note is that the cross covariance M_k in the KF and its ensemble equivalent \bar{M}_k should not be influenced by additive measurement noise e_k . Therefore, it is reasonable to replace $\tilde{Y}_{k|k-1}$ with

$$\tilde{Z}_{k|k-1} = H\tilde{X}_{k|k-1} \quad (21a)$$

so as to reduce the Monte Carlo variance of (17) using

$$\bar{M}_k = \frac{1}{N-1}\tilde{X}_{k|k-1}\tilde{Z}_{k|k-1}^T, \quad (21b)$$

$$\bar{S}_k = \frac{1}{N-1}\tilde{Z}_{k|k-1}\tilde{Z}_{k|k-1}^T + R. \quad (21c)$$

The Kalman gain \bar{K}_k is then computed as in the KF (6). Alternatively, a matrix square-root $R^{\frac{1}{2}}$ with $R^{\frac{1}{2}}R^{\frac{1}{2}} = R$ can be used to factorize

$$\bar{S}_k = \left[\frac{1}{\sqrt{N-1}}\tilde{Z}_{k|k-1} \quad R^{\frac{1}{2}} \right] \begin{bmatrix} \frac{1}{\sqrt{N-1}}\tilde{Z}_{k|k-1}^T \\ R^{\frac{1}{2}} \end{bmatrix}. \quad (22)$$

A QR decomposition [61] of the right matrix then yields a triangular $m \times m$ square root of \bar{S}_k , and the computation of \bar{K}_k simplifies to forward and backward substitution. Such ideas have their origin in sigma point KF variants [62].

The KF permits offline computation of the covariance matrices $P_{k|k}$ for all k because they do not depend on the measurements. In an EnKF for a linear system (1), we can mimic this behavior by propagating zero-mean ensembles $\tilde{X}_{k|k}$ that only carry the information of $P_{k|k}$. This is the central idea of different square root EnKF [39] which were suggested in [36] (ensemble adjustment filter, EAKF) or [37, 38] (ensemble transform filter, ETKF). The name square root EnKF stems from a relation to square root KF [3] which propagate $n \times n$ matrix square roots $P_{k|k}^{\frac{1}{2}}$ with $P_{k|k}^{\frac{1}{2}}P_{k|k}^{\frac{1}{2}} = P_{k|k}$. Most importantly, the artificial measurement noise and the inherent sampling error can be avoided.

The following derivation [39] rewrites an alternative expression for (5c) using a square root $P_{k|k-1}^{\frac{1}{2}}$ and its ensemble approximation $\frac{1}{N-1}\tilde{X}_{k|k-1}$:

$$\begin{aligned} P_{k|k} &= (I - K_k H)P_{k|k-1} \\ &= P_{k|k-1}^{\frac{1}{2}} \left(I - P_{k|k-1}^{\frac{1}{2}} H^T S_k^{-1} H P_{k|k-1}^{\frac{1}{2}} \right) P_{k|k-1}^{\frac{1}{2}} \\ &\approx \frac{1}{N-1}\tilde{X}_{k|k-1} \left(I - \frac{1}{N-1}\tilde{Z}_{k|k-1}^T \bar{S}_k^{-1} \tilde{Z}_{k|k-1} \right) \tilde{X}_{k|k-1}^T, \end{aligned} \quad (23a)$$

where (21a) was used. The next step is to factorize

$$\left(I - \frac{1}{N-1}\tilde{Z}_{k|k-1}^T \bar{S}_k^{-1} \tilde{Z}_{k|k-1} \right) = \Pi_k^{\frac{1}{2}} \Pi_k^{\frac{1}{2}}, \quad (23b)$$

which requires the left hand side to be positive definite. This property is easily established for the positive definite \bar{S}_k of (21c) after realizing that the left hand side of (23b) is a Schur complement [61] of a positive definite matrix.

Finally, the $N \times N$ matrix $\Pi_k^{\frac{1}{2}}$ can be used to create a deviation ensemble

$$\tilde{X}_{k|k} = \tilde{X}_{k|k-1} \Pi_k^{\frac{1}{2}} \quad (24)$$

that correctly encodes $P_{k|k}$ without using any random perturbations. Numerically efficient schemes to reduce the computational complexity of ETKF that work on $N \times N$ transform matrices can be found in the literature [39]. Other variants update the deviation ensemble via a multiplication from the left [36], which is more costly for large n . Some more conditions on $\Pi_k^{\frac{1}{2}}$ must be met for $\tilde{X}_{k|k}$ to remain zero-mean [63, 64].

The actual filtering is achieved by updating a single estimate according to

$$\bar{x}_{k|k} = (I - \bar{K}_k H) F \bar{x}_{k-1|k-1} + \bar{K}_k y_k, \quad (25)$$

where \bar{K}_k is computed from the deviation ensembles.

There are indications that in nonlinear and non-Gaussian systems the sampling based EnKF variants should be preferable over their square root counterparts: A low-dimensional example is studied in [65]; the impression is confirmed for a high-dimensional problem in [66].

5.3 Covariance inflation

Covariance inflation is a measure to counteract the tendency of the EnKF to underestimate the state uncertainty for small N and an important ingredient in operational EnKF [18]. The spread of the prediction ensemble $X_{k|k-1}$ is increased according to

$$X_{k|k-1} = \bar{x}_{k|k-1} \mathbb{1}^T + c \tilde{X}_{k|k-1} \quad (26)$$

with a factor $c > 1$. In the EnKF context, this heuristic has been proposed in [40]. Related concepts are dithering in the PF [7] and the “fudge factor” to increase $P_{k|k-1}$ in the KF [67]. Extensions to adaptive inflation, where c is adjusted online, are discussed in [23].

5.4 Localization

Localization is a technique to address the issue of spurious correlations in the EnKF, and a crucial feature of operational EnKF [18, 19]. The underlying idea applies equally well to the EnKF and the KF, and can be used to systematically update only a part of the state vector with each measurement.

In order to explain the concept, we regard the KF measurement update for a linear system (1) with a low-dimensional² measurement y_k . Let $x = x_{k|k-1}$ and $P = P_{k|k-1}$ for notational convenience. It is possible to permute the state components such that

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}, \quad H = [H_1 \ 0 \ 0], \quad P = \begin{bmatrix} P_1 & P_{12} & 0 \\ P_{12}^T & P_2 & P_{23}^T \\ 0 & P_{23}^T & P_3 \end{bmatrix}. \quad (27)$$

Only the part x_1 appears in the measurement Eq. (1b) $y_k = H_1 x_1 + e_k$. While x_2 is correlated to x_1 , there is zero correlation between x_1 and x_3 . As a consequence, many submatrices of P vanish in the computation of

$$PH^T = [H_1 P_1 \ H_1 P_{12} \ 0]^T, \quad HPH^T = H_1 P_1 H_1^T, \quad (28a)$$

and do not contribute to the Kalman gain (6)

$$K_k = \begin{bmatrix} P_1 H_1^T \\ P_{12}^T H_1^T \\ 0 \end{bmatrix} \left(H_1 P_1 H_1^T + R \right)^{-1}. \quad (28b)$$

A KF measurement update (5) with the above K_k does not affect the x_3 estimate or covariance. Hence, there is a lower dimensional measurement update that only alters the statistics of x_1 and x_2 .

Localization in the EnKF enforces the above structure using two prevailing techniques, local updates [13, 16, 42] and covariance tapering [14, 43]. Both rely on prior knowledge of the covariance structure. For example, the state components are often connected to geographic locations in geoscientific applications. From the underlying physics, it is reasonable to assume zero correlation between distant states. Unfortunately, this viewpoint is not transferable to high-dimensional problems in general.

Local updates were introduced for the sampling-based EnKF in [13] and for different square root EnKF in [16, 42]. Nonlinear measurement functions (2b) are linearized in the latter two. All of the above references update the state vector “grid point by grid point,” which appears unusual from a KF perspective [3]. In an iteration, local state vectors of small dimension ($< N$) are chosen and updated with a subset of supposedly relevant measurements. These “full rank” updates avoid some of the problems associated with small N and large n . However, discontinuities between state components are introduced [68]. Some heuristics to combine the local ensembles and further implementation details can be found in [42, 69].

Under the assumption of the structure in (27), a local analysis would amount to an EnKF update of the x_1 - and x_2 -components only, to avoid errors in x_3 .

Covariance tapering was introduced in [13]. It contradicts the EnKF idea in the sense that the ensemble covariance $\bar{P}_{k|k-1}$ of $X_{k|k-1}$ is processed. However, it will become clear that not all entries of $\bar{P}_{k|k-1}$ must be computed. Prior knowledge of a covariance structure as in (27) is used to create an $n \times n$ matrix ρ with entries in $[0, 1]$, and a tapered covariance ($\rho \circ \bar{P}_{k|k-1}$) is computed. Here, \circ denotes the element-wise Hadamard or Schur product [61]. A typical ρ has ones on the diagonal and decays smoothly to zero for unwanted off-diagonal elements. The

standard choice uses a compactly supported correlation function from [70] and is discussed in [14, 43, 68]. Subsequently, the Kalman gain is computed as in the KF (6) using

$$\bar{M}_k = (\rho \circ \bar{P}_{k|k-1})H^T, \quad (29a)$$

$$\bar{S}_k = H(\rho \circ \bar{P}_{k|k-1})H^T + R, \quad (29b)$$

where we assumed a linear measurement relation (1b).

There are some technicalities associated with the tapering operation. Only positive semi-definite ρ guarantee that $(\rho \circ \bar{P}_{k|k-1})$ is a valid covariance [26]. Full rank ρ yield an increased rank in $(\rho \circ \bar{P}_{k|k-1})$ [14]. However, low rank ρ do not necessarily decrease the rank of $(\rho \circ \bar{P}_{k|k-1})$. A closely related problem to finding valid (positive semi-definite or definite) ρ is the creation of covariance functions and kernels in Gaussian processes [71]. Here, a methodology to create more complicated kernels from simpler ones could be used to create ρ .

Unfortunately, the Hadamard product cannot be formulated as an operation on the ensembles in general. Still, the computational requirements can be limited by only working with the non-zero elements of $(\rho \circ \bar{P}_{k|k-1})$. Furthermore, it is common to avoid the computation of $\bar{P}_{k|k-1}$ using

$$\bar{M}_k = \rho_M \circ \bar{M}_k, \quad (30)$$

instead of (29a) and to skip the tapering in S_k altogether [43]. After all, for low-dimensional y_k (small m) \bar{M}_k has the strongest influence on the gain \bar{K}_k . Also, the matrix ρ_M is constructed from prior knowledge about the correlation. In the geoscientific context, where the state components and measurements are associated with geographic locations, this is easy. In general, however, it might not be possible to devise an appropriate ρ_M . Other variants [14, 26, 68] with tapering for \bar{S}_k exist and have in common that they are only identical to (29) for $H = I$.

Some relations between local updates and covariance tapering are discussed in [68]. For the structure in (27), we can suggest a rank-1 taper ρ that establishes a correspondence between the two concepts. Let r_1 and r_2 be vectors of the same dimensions as x_1 and x_2 , respectively, that contain all ones. Let r_3 be a zero vector of the same dimension as x_3 and $r^T = [r_1^T, r_2^T, r_3^T]$. Then, $\rho = rr^T$ removes all entries from $\bar{P}_{k|k-1}$ that would disappear in (28) anyhow. Furthermore, the Hadamard product for the rank-1 ρ can be written as an operation on the ensemble $\tilde{X}_{k|k-1}$ using

$$\begin{aligned} (rr^T) \circ \bar{P}_{k|k-1} &= \text{diag}(r)\bar{P}_{k|k-1}\text{diag}(r) \\ &= \frac{1}{N-1} (\text{diag}(r)\tilde{X}_{k|k-1}) (\text{diag}(r)\tilde{X}_{k|k-1})^T. \end{aligned} \quad (31)$$

The multiplication with $\text{diag}(r)$ merely removes the rows corresponding to x_3 , which establishes an equivalence

between local updates and covariance tapering. By picking a smoothly decaying r , we can furthermore avoid the discontinuities associated with local updates.

5.5 The EnKF gain and least squares

A parallel to least squares problems can be disclosed by closer inspection of the Eq. (18) that is used to compute the EnKF gain \bar{K}_k . Perhaps more apparent in the transpose of (18), in

$$\tilde{Y}_{k|k-1}\tilde{Y}_{k|k-1}^T\bar{K}_k^T = \tilde{Y}_{k|k-1}\tilde{X}_{k|k-1}^T, \quad (32a)$$

appear the normal equations of the least squares problems

$$\tilde{Y}_{k|k-1}^T\bar{K}_k^T = \tilde{X}_{k|k-1}^T, \quad (32b)$$

that are to be solved for each row of \bar{K}_k and $\tilde{X}_{k|k-1}$.

Hence, the EnKF iteration can be carried out without explicitly computing any sample covariance matrices if instead efficient solutions to the problem (32b) are employed. Furthermore, the problem (32b) could be modified using regularization [72] to enforce sparsity in \bar{K}_k . This would be an alternative approach to the localization methods discussed earlier. Related ideas to improve the Kalman gain using bootstrap methods [72] for computing \bar{M}_k and \bar{S}_k in (17) are discussed in [73, 74].

6 Relations to other algorithms

The EnKF for nonlinear systems (2) differs from other sampling-based nonlinear filters such as sigma point KF [5] or particle filters (PF) [7]. One reason for this is that the EnKF approximates the KF algorithm (with the side effect that it can be applied to (2)) rather than trying to solve the nonlinear filtering problem directly.

The biggest difference between the EnKF and sigma point filters [5] such as the unscented KF [4, 56] or divided difference KF [62] is the measurement update. Whereas the EnKF updates its ensembles, the latter carry out the KF measurement update (5) using approximately computed mean values and covariance matrices. That is, the samples or sigma points are condensed into a filtering estimate $\hat{x}_{k|k}$ and its covariance $P_{k|k}$, which entails a loss of information and can be seen as an inherent Gaussian assumption on the filtering density $p(x_k|y_{1:k})$. In contrast, the EnKF can preserve more information and deviations from Gaussianity in the ensemble. Similarities appear in the gain computations of the EnKF and sigma point KF. In both, the Kalman gain appears as a function of the sampling covariance matrices, although with the deterministic sigma points and weights in the latter. With their origin in the KF, both sigma point filters and the EnKF can be expected to share difficulties with multimodal posterior distributions.

Similar to the EnKF, the PF propagates N state realizations that are called particles. For the bootstrap particle filter [6], the prediction step corresponds to the EnKF time

update (11). Apart from that, however, the differences dominate. First, the PF is designed as an approximate solution of the Bayesian filtering equations [15] using sequential importance sampling [7]. For $N \rightarrow \infty$, the PF solution recovers the true filtering density. Second, the samples in basic PF variants are generated from a proposal distribution only once every time instance and then left untouched. The measurement update amounts to updating the particle weights, which leads to a degeneracy problem for large n . In the EnKF, in contrast, the ensemble members are influenced by the time and the measurement update. Third, the PF relies on a crucial resampling step that is not present in the EnKF. An attempt to use the EnKF as proposal density in PF is described in [75]. A unifying interpretation of the EnKF and PF as ensemble transform filters can be found [76].

Still, the EnKF appears as a distinct algorithm besides sigma point KF and PF. Its properties and potential for nonlinear problems remain to be fully investigated. Existing results that the EnKF does not converge to the Bayesian filtering recursion [31] remain to be interpreted in a constructive manner.

7 Instructive simulation examples

Four examples are discussed in greater detail, among them one popular benchmark problem of the SP and DA literature each.

7.1 A scalar linear Gaussian model

The first example illustrates the tendency of the EnKF to underestimate the state uncertainty. A related example is studied in [38]. We compare the EnKF variance $\bar{P}_{k|k}$ to the $P_{k|k}$ of the KF via Monte Carlo simulations on the simple scalar state-space model

$$x_{k+1} = x_k + v_k, \tag{33a}$$

$$y_k = x_k + e_k. \tag{33b}$$

The initial state x_0 , the process noise v_k , and the measurement noise e_k are specified by the probability density functions

$$p(x_0) = \mathcal{N}(x_0; 0, 0.1), \tag{33c}$$

$$p(v_k) = \mathcal{N}(v_k; 0, 0.1), \tag{33d}$$

$$p(e_k) = \mathcal{N}(e_k; 0, 0.01). \tag{33e}$$

A trajectory of (33) is simulated and a KF is used to compute the optimal variances $P_{k|k}$. Because the model is time-invariant, the $P_{k|k}$ quickly converge to a constant value. For $k > 3$, $P_{k|k} = 0.0092$ is obtained.

Next, 10,000 Monte Carlo experiments with a sampling-based EnKF with $N = 5$ are performed. The distribution of obtained $\bar{P}_{k|k}$ for $k = 10$ is illustrated in Fig. 1. The




Fig. 1 Distribution of EnKF variances $\bar{P}_{k|k}$ with $k = 10$ and $N = 5$ ensemble members for 10,000 runs on the same trajectory. Also shown is the mean and median of all outcomes and the desired KF variance $P_{k|k}$

vertical lines indicate the $P_{k|k}$ of the KF and the median and mean of the $\bar{P}_{k|k}$ outcomes.

The average $\bar{P}_{k|k}$ over the Monte Carlo realizations is close to the desired $P_{k|k}$. However, there is a large spread among the $\bar{P}_{k|k}$ and the distribution is skewed toward zero with its median below $P_{k|k}$. Although $N > n$, there is a tendency to underestimate $P_{k|k}$.

In order to clarify the reason for this behavior and whether it has to do with the coupling between the EnKF \bar{K}_k and the ensemble members, we repeat the experiment with an EnKF that uses the gain of the stationary KF for all k . The resulting outcomes are illustrated in Fig. 2. Now, the average $\bar{P}_{k|k}$ is correct. However, the median shows that there is still more probability mass below $P_{k|k}$. The tendency to underestimate $P_{k|k}$ and the remaining spread must be due to random sampling errors. For larger N , the effect vanishes, and the median and mean of $\bar{P}_{k|k}$ appear similar for $N \geq 10$.

7.2 The particle filter benchmark

In the second example, we show that the EnKF does not converge to the Bayesian filtering solution in nonlinear systems as $N \rightarrow \infty$ [31]. A well-known benchmark




Fig. 2 Distribution of EnKF variances $\bar{P}_{k|k}$ but computed with the correct Kalman gain. Otherwise, similar to Fig. 1

problem from the PF literature [6] is used. The model is specified by

$$x_{k+1} = \frac{x_k}{2} + 25 \frac{x_k}{1 + x_k^2} + 8 \cos(1.2(k + 1)) + v_k, \tag{34a}$$

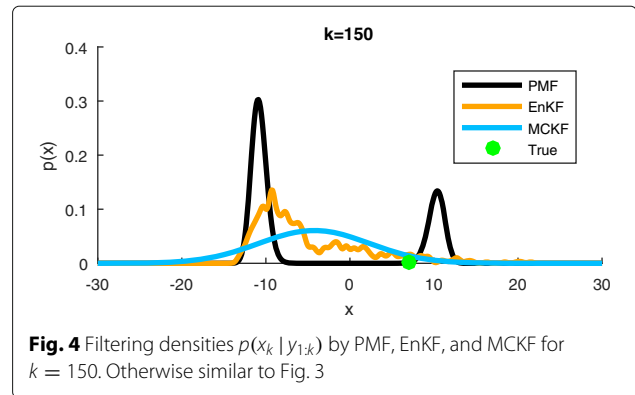
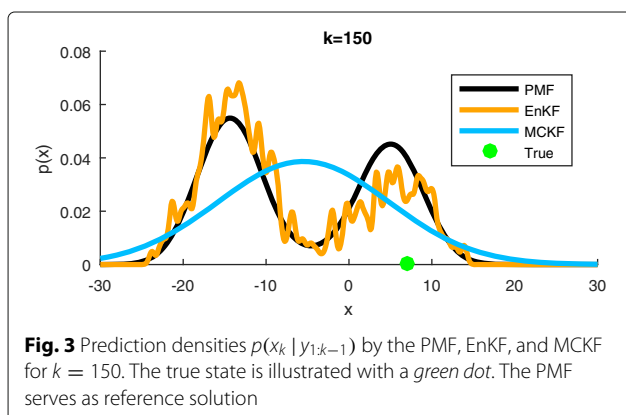
$$y_k = \frac{1}{20} x_k^2 + e_k, \tag{34b}$$

with independent $v_k \sim \mathcal{N}(0, 10)$, $e_k \sim \mathcal{N}(0, 1)$, and $x_0 \sim \mathcal{N}(0, 1)$. Because the model is scalar, the Bayesian filtering densities $p(x_k | y_{1:k})$ can be computed numerically using point mass filters (PMF) [77]. A sampling based EnKF with $N = 500$ is tested and kernel density estimates are used to obtain approximations of $p(x_k | y_{1:k})$ from the ensembles. For comparison, we include a closely related sigma point KF variant that uses Monte Carlo integration with $N = 500$ samples [5]. The only difference to the EnKF is that this Monte Carlo KF (MCKF) carries out the KF measurement update (5) to propagate a mean and a variance. We illustrate the results as Gaussian densities.

Figure 3 shows the prediction results for $k = 150$. The PMF reference solution is bimodal with one mode close to the true state. The reason for this lies in the squared x_k in (34b). The EnKF prediction resembles the PMF well except for the random variations in the kernel density estimate. The MCKF cannot represent the multimodality but the Gaussian bell covers the relevant regions.

The filtering results for $k = 150$ are in Fig. 4. The PMF reference solution has much narrower peaks after including y_k . The EnKF provides a skewed density that does not resemble $p(x_k | y_{1:k})$ even though the EnKF prediction approximated $p(x_k | y_{1:k-1})$ well. This is the main take-away result and confirms [31]. Again, the MCKF exhibits a large variance.

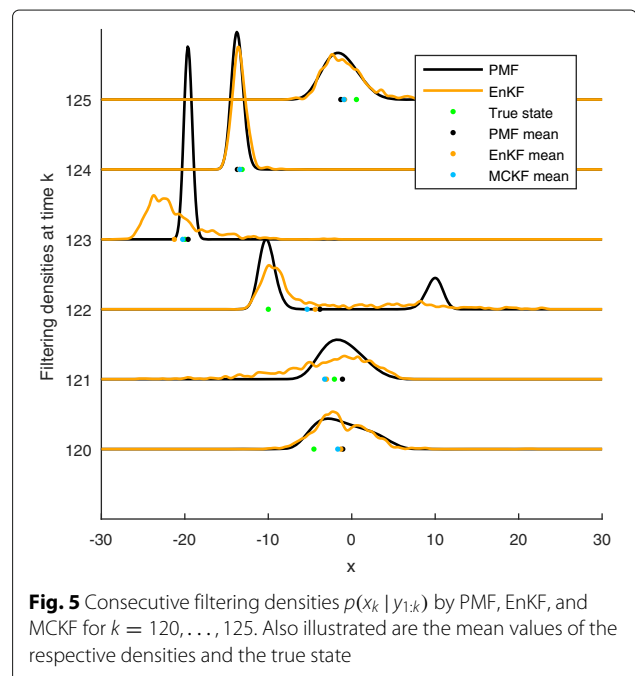
Further filtering results for the PMF and EnKF are shown in Fig. 5. It can be seen that the EnKF solutions sometimes resemble the PMF very well but not always. Similar statements can be made for the prediction results.

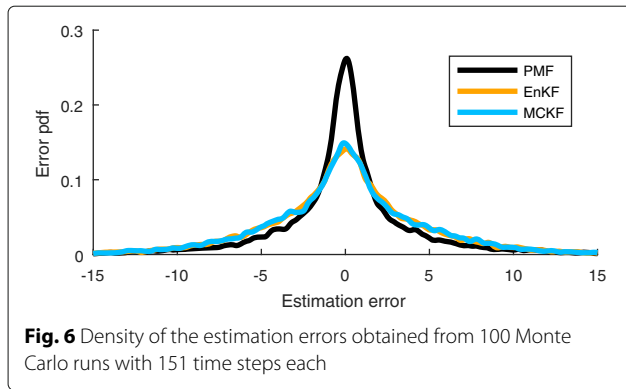


Dots in Fig. 5 illustrate the mean values as state estimates. Especially for the PMF, it can be seen that the mean (though optimal in a minimum variance sense [3]) is debatable for multimodal densities. Often, all estimates are quite close. Figure 6 provides the estimation error densities obtained from 100 Monte Carlo experiments with 151 time steps each. The PMF mean estimates exhibit a larger peak around 0. The estimation errors for the EnKF and MCKF appear similar. This is surprising because the latter employs a Gaussian approximation at each time step. Both error densities have heavier tails than the PMF density. All estimation errors appear unbiased.

7.3 Batch smoothing using the EnKF

We here show how to use the EnKF as smoothing algorithm by estimating batches of states. This allows us to compare its performance for $N < n$ in problems of arbitrary dimension.





First, we formulate an “augmented state” that comprises an entire trajectory of $L + 1$ steps,

$$\xi = [x_0^T \ x_1^T \ \dots \ x_L^T]^T, \tag{35}$$

with dimension $n = (L + 1)n_x$. Second, we note that the measurements y_k , $k = 1, \dots, L$, have uncorrelated measurement noise and known relations to the components of ξ . For linear systems (1), the predicted mean and covariance of ξ can be easily derived, and smoothed estimates of all x_k , $k = 0, \dots, L$, can be obtained by sequentially processing all y_k in KF measurement updates for ξ .

Also, other smoothing variants and the Rauch-Tung-Striebel (RTS) algorithm can be derived from state augmentation approaches [3]. Due to its sequential nature, however, the RTS smoother does not provide joint covariance matrices of x_k and x_{k+i} for $i \neq 0$. Except for this and the higher computational complexity of working with ξ , the batch and RTS smoothers are equivalent for (1).

An EnKF approach to batch smoothing mimics the above. A prediction ensemble for ξ is obtained by simulating N trajectories for random process noise and initial state realizations. This can also be carried out for non-linear models (2). Then, sequential EnKF measurement updates are performed for all y_k .

For our experiments, we use a tracking problem with a constant velocity model [67] and position measurements. The low-dimensional state is given by

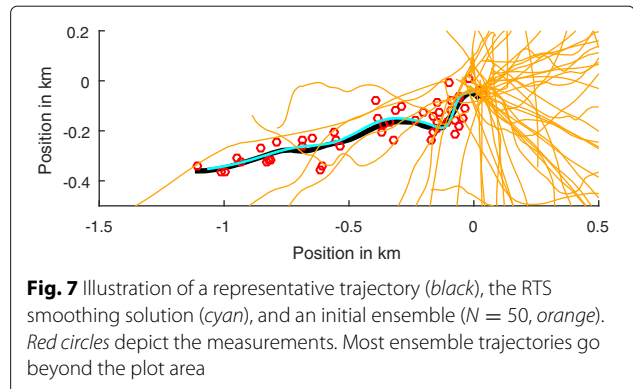
$$x = [x \ y \ \dot{x} \ \dot{y}]^T \tag{36a}$$

and comprises the Cartesian position [m] and velocity [m/s] of an object. The parameters of (1) are given by

$$F = \begin{bmatrix} I_2 & \mathbb{T}I_2 \\ 0 & I_2 \end{bmatrix}, \quad G = \begin{bmatrix} \frac{\mathbb{T}^2}{2}I_2 \\ \mathbb{T}I_2 \end{bmatrix}, \quad H = [I_2 \ 0], \tag{36b}$$

with $\mathbb{T} = 1$ s. The initial state x_0 is Gaussian distributed with

$$\hat{x}_0 = [0 \ 0 \ 15 \ -10]^T, \quad P_0 = \text{diag}(50^2, 50^2, 20^2, 20^2), \tag{36c}$$



and the process and measurement noise covariances are

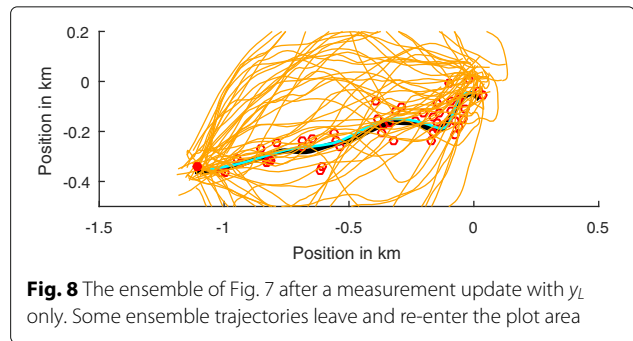
$$Q = \text{diag}(10, 50), \quad R = \begin{bmatrix} 2000 & 1000 \\ 1000 & 1980 \end{bmatrix}. \tag{36d}$$

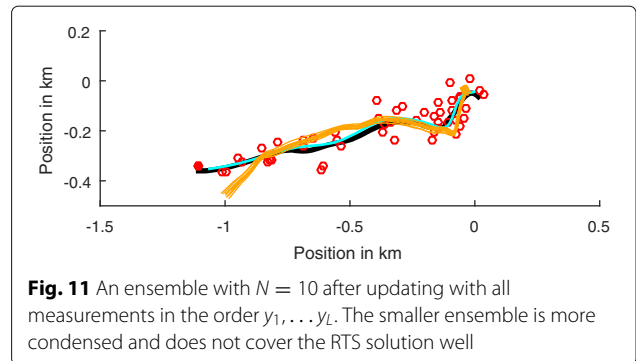
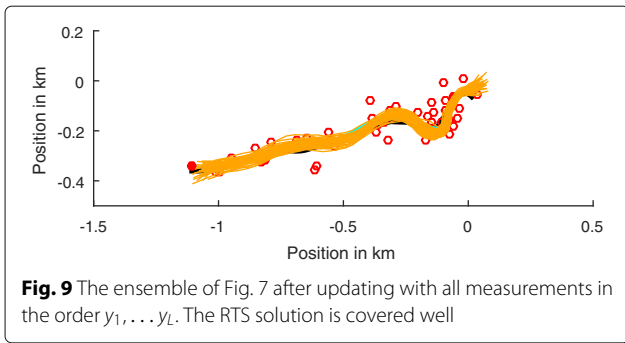
With $n_x = 4$ and $L = 49$ we obtain $n = 200$ as dimension of ξ . The RTS solution is compared to EnKF of ensemble size $N = \{10, 20, 50\}$. Monte Carlo errors are reduced using (21) in the gain computations.

A realization of a true trajectory and its measurements is provided in Fig. 7 together with the RTS estimate and an ensemble of $N = 50$ trajectories. The latter are the initial ensemble of an EnKF. The ensemble is well gathered around the initial position but fans out wildly.

Figure 8 shows the ensemble after an update with y_L only. The measurement at the end of the trajectory provides an anchor point and quickly reduces the spread of the ensemble. Figure 9 shows the result after processing all measurements in sequential order from first to last. The true trajectory and the RTS estimate are mostly covered well by the ensemble. The EnKF with $N = 50$ appears consistent in this respect. Position errors for the RTS and the EnKF are provided in Fig. 10. The EnKF performs slightly worse than the RTS but still gives good results for $N = 50$, without extra inflation or localization.

The next experiment explores the EnKF for $N = 10$. Figure 11 shows the ensemble after processing all measurements. The ensemble is compactly gathered but





does not cover the true trajectory well. The EnKF is overconfident.

A last experiment explores how well an EnKF with $N = 20$ captures the uncertainty of the state estimate. Furthermore, we discuss effects of the order in which the measurements are processed. Specifically, we compare the ensemble covariance of the positions x_k to the exact $\text{cov}(x_k, x_i), i, k = 0, \dots, L$, obtained by KF updates for the augmented state ξ .

The exact covariance after processing all measurements is shown in Fig. 12. Row k in the matrix defines the covariance function between x_k and the remaining x positions. The banded structure indicates that subsequent positions are more related than, say, x_0 and x_L . Figure 13 shows the corresponding EnKF covariance after processing the measurements from y_1 to y_L . The off-diagonal elements do not decay uniformly as in Fig. 12, and spurious positive and negative correlations appear. Furthermore, the correct temporal order of measurements entails an unwanted structure. Later x_k are rated more uncertain according to the lighter areas in the lower right corner of Fig. 13. A covariance after processing the measurements in random order is shown in Fig. 14. The spurious correlations persist but the diagonal elements appear more homogeneous.

From the above experiments, we conclude that the EnKF can provide good estimates for ensembles with $N < n$. However, there is a minimum N required to obtain

consistent results without further measures such as localization or inflation. We have shown adverse effects such as ensembles with too little spread and spurious correlations. As a final note, the alert reader will recognize parallels between the above example and ensemble smoothing methods as presented in [17].

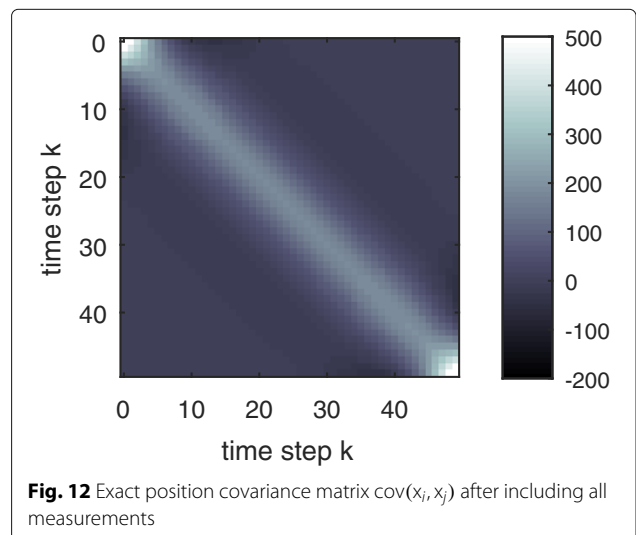
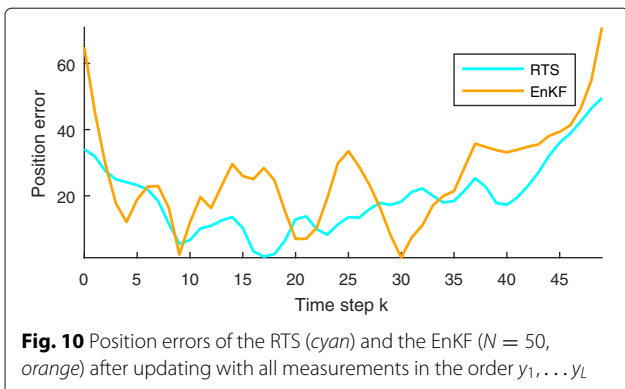
7.4 The 40-dimensional Lorenz model


Our final example is a benchmark problem from the EnKF literature. We investigate the 40-dimensional Lorenz-96 model³ from [53] that is used in, e.g., [36, 38, 42, 50, 52, 63, 69].

The state x mimics an atmospheric quantity at equally spaced locations along a circle. Its evolution is specified by the nonlinear differential equation

$$\dot{x}(j) = (x(j+1) - x(j-2))x(j-1) - x(j) + F(j), \tag{37}$$

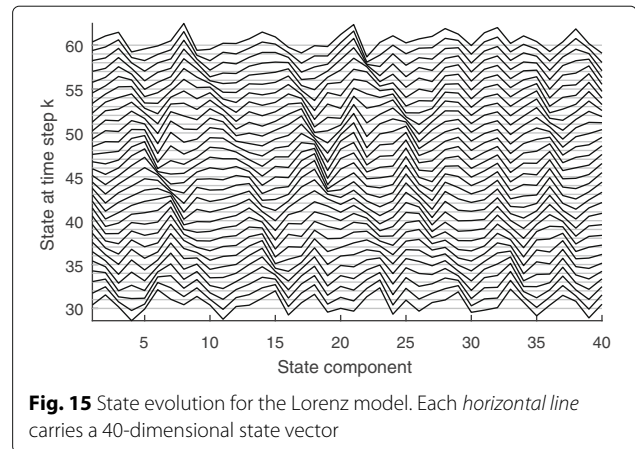
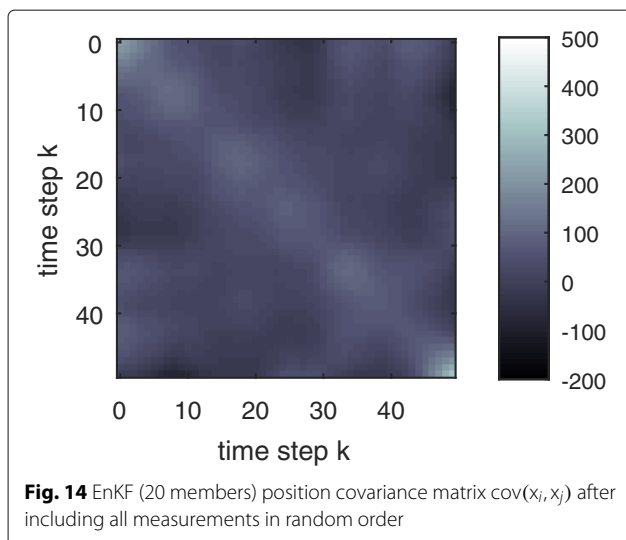
where $j = 1, \dots, 40$ indexes the components of x , with the convention that $x(0) = x(40)$ etc. Instead of the commonly used forcing term $F(j) = 8$, we assume time-dependent $F_k(j) \sim \mathcal{N}(8, 1)$ that are constant for time





intervals $T = 0.05$ only and act as process noise. A Runge-Kutta method (RK4) is used to discretize (37) to obtain the nonlinear state difference Eq. (2a) with $x_k = x_k$ and $v_k = F_k$. The step size T corresponds to about 6 h if x were an atmospheric quantity on a latitude circle of the earth [53]. Although the model (37) is said to be chaotic, the effects are only mild for short integration times T . In our experiments, all $n = 40$ states are measured with additive Gaussian noise $e_k \sim \mathcal{N}(0, I)$. The initial state is Gaussian with $x_0 \sim \mathcal{N}(0, P_0)$, where P_0 is drawn from a Wishart distribution with seed matrix I_n and n degrees of freedom.

Figure 15 illustrates how the state evolves over several time steps. There is a tendency for peaks to move “westwards” as k increases.



We note that there are also alternative approaches for estimating x , for example, by first linearizing and then discretizing (37). However, we adopt the RK4 discretization of the EnKF literature that yields a state transition that is easy to evaluate but difficult to linearize. Because of this, the EKF [3] cannot be applied easily and we obtain a challenging benchmark problem.

We use sampling-based EnKF to estimate long state sequences of $L = 10^4$ time steps. Following [38, 42], the performance is assessed by the error

$$\varepsilon_k = \sqrt{\frac{1}{n} (\hat{x}_{k|k} - x_k)^T (\hat{x}_{k|k} - x_k)}, \tag{38}$$

where $\hat{x}_{k|k}$ is the ensemble mean. We use the average ε_k for $k = 100, \dots, L$, denoted by $\bar{\varepsilon}$, as quantitative performance measure for different EnKF. Useful EnKF must yield $\bar{\varepsilon} < 1$, which is the error when simply taking $\hat{x}_{k|k} = y_k$.

First, we compute a reference solution using an EnKF with $N = 1000$. Without any localization or inflation $\bar{\varepsilon} = 0.29$ is achieved. Figure 16 shows the sample covariance $\bar{P}_{k|k-1}$ of a prediction ensemble $X_{k|k-1}$, our best guess of the true covariance. The banded structure reveals that the problem is suitable for localization. Hence, we construct a matrix ρ for covariance tapering from a compactly supported correlation function [70] that is also used in [14, 26, 38, 43] and appears to be the standard choice. The chosen ρ is a Toeplitz matrix because the components of x_k are at equidistant locations and shown in Fig. 17.

Next, EnKF with different ensemble sizes N , covariance inflation factors c , with or without tapering, are compared. The obtained errors $\bar{\varepsilon}$ are summarized in Table 1. For $N = n = 40$, we obtain a worse $\bar{\varepsilon}$ than for $N = 1000$. While inflation without tapering does reduce the error slightly, the covariance tapering even yields a better result than the EnKF with $N = 1000$. Further improvements are obtained by combining inflation and tapering. Figure 18 shows the estimation error $x_k - \hat{x}_{k|k}$ for $k = 10^4$, $N = 40$, $c = 1.02$, and tapering with ρ . In the background, the

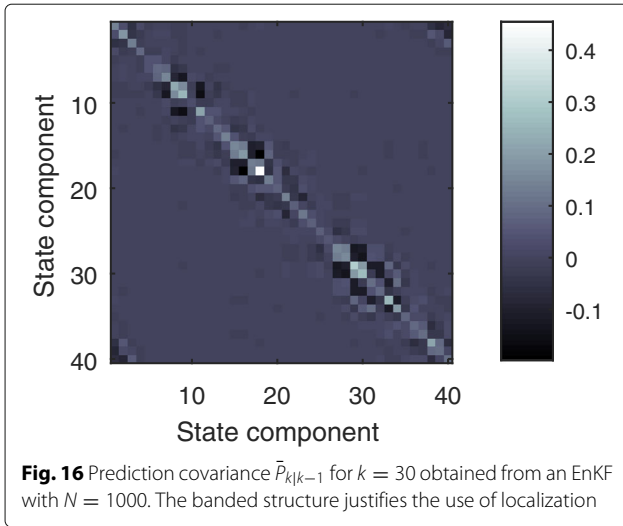


Fig. 16 Prediction covariance $\tilde{P}_{k|k-1}$ for $k = 30$ obtained from an EnKF with $N = 1000$. The banded structure justifies the use of localization

ensemble deviations $\tilde{X}_{k|k}$ are illustrated. The estimation error is mostly contained in the intervals spanned by the ensemble; hence, the EnKF is consistent. Tests on EnKF with $N = 20$ reveal convergence problems, even with inflation the initial estimation error persists. With the help of tapering, however, a competitive error can be achieved. Even further reduction to $N = 10$ is possible with tapering and inflation. The required inflation factor c must be increased to counteract the lack of ensemble spread. Similar to Figs. 18 and 19 illustrates the estimation error and deviation ensemble for $k = 10^4$, $N = 10$, $c = 1.05$, and tapering with ρ . Although the obtained error is larger than for $N = 40$, the ensemble deviations represent the estimation uncertainty well.

A number of lessons have been learned from related experiments. As alternative to the ρ in Fig. 17, a simpler taper that contains only ones and zeros to enforce the

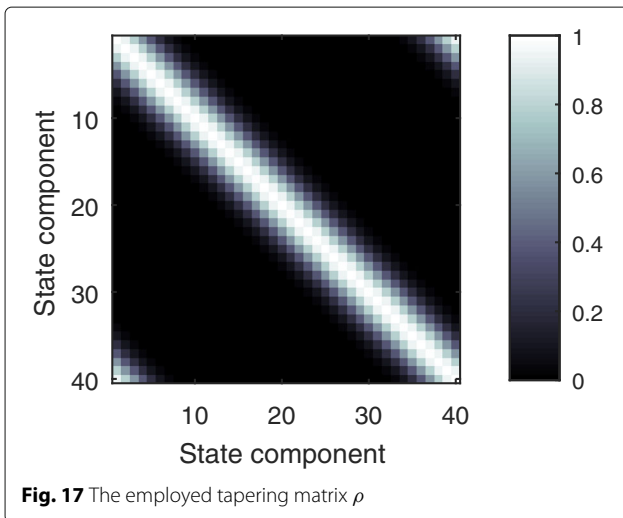


Fig. 17 The employed tapering matrix ρ

Table 1 Averaged errors $\bar{\epsilon}$ for different EnKF

N	c	ρ	$\bar{\epsilon}$
1000	1	No	0.29
40	1	No	0.44
40	1.05	No	0.33
40	1	Yes	0.29
40	1.02	Yes	0.28
20	> 1	No	> 1
20	1.01	Yes	0.3
10	1.05	Yes	0.34

banded structure was used. Although this ρ was indefinite, a reduction in $\bar{\epsilon}$ was achieved without any numerical issues. Hence, the specific structure of ρ appears secondary. The smooth ρ of Fig. 17 remains preferable in terms of $\bar{\epsilon}$, though. Sequential processing of the measurements did not degrade the performance. Experiments without process noise give the lower errors $\bar{\epsilon}$ from, e.g., [38, 42].

8 Conclusions

With this paper, we have given a comprehensive and easy to understand introduction to the EnKF for signal processing researchers. The origin of the EnKF in the KF and its simple implementation have been demonstrated. The unique literature review provides quick access to the most relevant papers in the plethora of geoscientific EnKF publications. Furthermore, we have discussed the challenges related to small ensembles for high-dimensional states, $N < n$, and the available solutions such as localization or inflation. Finally, we have tested the EnKF on signal processing and EnKF benchmark problems.

With its scalability and simple implementation, even for nonlinear and non-Gaussian problems, the EnKF stands out as viable candidate for many state estimation problems. Furthermore, localization ideas and advanced concepts for estimating covariance matrices and the EnKF




Fig. 18 The estimation error $x_k - \bar{x}_{k|k}$ for $k = 10^4$ with the deviation ensemble $\tilde{X}_{k|k}$ in the background for an EnKF with $N = 40$, covariance localization, and inflation factor $c = 1.02$




Fig. 19 The estimation error $x_k - \bar{x}_{k|k}$ for $k = 10^4$ with the deviation ensemble $\tilde{x}_{k|k}$ in the background for an EnKF with $N = 10$, covariance localization, and inflation factor $c = 1.05$

gain from the limited information in the ensembles provide new research directions for the EnKF and high-dimensional filters in general, hopefully with an increased participation from the signal processing community.

Endnotes

¹ With over 3000 citations between 1994 and 2016.

² We assume that the components can be processed sequentially.

³ Also known as the Lorenz-96, L95, L96, or L40 model.

Acknowledgements

This work was supported by the project Scalable Kalman Filters granted by the Swedish Research Council (VR).

Authors' contributions

MR wrote the majority of the text and performed the majority of the simulations. GH and CF contributed text to earlier versions of the manuscript and helped with the simulations. GH, CF, and FG commented on and approved the manuscript. FG initiated the research on ensemble Kalman filters. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Publisher's Note

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

Received: 28 February 2017 Accepted: 19 July 2017

Published online: 03 August 2017

References

1. E Kalnay, *Atmospheric modeling, data assimilation and predictability*. (Cambridge University Press, New York, 2002)
2. RE Kalman, A new approach to linear filtering and prediction problems. *J. Basic Eng.* **82**(1), 35–45 (1960)
3. BD Anderson, JB Moore, *Optimal filtering*. (Prentice Hall, Englewood Cliffs, 1979)
4. S Julier, J Uhlmann, H Durrant-Whyte, in *Proceedings of the American Control Conference 1995* vol.3. A new approach for filtering nonlinear systems (IEEE, Seattle, 1995), pp. 1628–1632
5. M Roth, G Hendeby, F Gustafsson, Nonlinear Kalman filters explained: a tutorial on moment computations and sigma point methods. *J. Adv. Inf. Fusion.* **11**(1), 47–70 (2016)
6. NJ Gordon, DJ Salmond, AF Smith, Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar Signal Process. IEE Proc. F.* **140**(2), 107–113 (1993)
7. F Gustafsson, Particle filter theory and practice with positioning applications. *IEEE Aerosp. Electron. Syst. Mag.* **25**(7), 53–82 (2010)
8. G Evensen, Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res. Oceans.* **99**(C5), 3–10162 (1014)
9. G Burgers, JP van Leeuwen, G Evensen, Analysis scheme in the ensemble Kalman filter. *Mon. Weather Rev.* **126**(6), 1719–1724 (1998)
10. H Durrant-Whyte, T Bailey, Simultaneous localization and mapping: Part I. *IEEE Robot. Autom. Mag.* **13**(2), 99–110 (2006)
11. M Baum, UD Hanebeck, Extended object tracking with random hypersurface models. *IEEE Trans. Aerosp. Electron. Syst.* **50**(1), 149–159 (2014)
12. N Wahlström, E Özkan, Extended target tracking using Gaussian processes. *IEEE Trans. Signal Proc.* **63**(16), 4165–4178 (2015)
13. PL Houtekamer, HL Mitchell, Data assimilation using an ensemble Kalman filter technique. *Mon. Weather Rev.* **126**(3), 796–811 (1998)
14. PL Houtekamer, HL Mitchell, A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Weather Rev.* **129**(1), 123–137 (2001)
15. AH Jazwinski, *Stochastic processes and filtering theory*. (Academic Press, New York, 1970)
16. G Evensen, The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dyn.* **53**(4), 343–367 (2003)
17. G Evensen, *Data assimilation: the ensemble Kalman filter*, 2nd ed. (Springer, Dordrecht, New York, 2009)
18. TM Hamill, in *Predictability of Weather and Climate*. Ensemble-based atmospheric data assimilation (Cambridge University Press, Cambridge, 2006)
19. PL Houtekamer, HL Mitchell, Ensemble Kalman filtering. *Q. J. R. Meteorol. Soc.* **131**(613), 3269–3289 (2005)
20. JS Whitaker, TM Hamill, X Wei, Y Song, Z Toth, Ensemble data assimilation with the NCEP global forecast system. *Mon. Weather Rev.* **136**(2), 463–482 (2008)
21. GP Compo, JS Whitaker, PD Sardeshmukh, N Matsui, RJ Allan, X Yin, BE Gleason, RS Vose, G Rutledge, P Bessemoulin, S Brönnimann, M Brunet, RI Crouthamel, AN Grant, PY Groisman, PD Jones, MC Kruk, AC Kruger, GJ Marshall, M Mauerer, HY Mok, O Nordli, TF Ross, RM Trigo, XL Wang, SD Woodruff, SJ Worley, The twentieth century reanalysis project. *Q. J. R. Meteorol. Soc.* **137**(654), 1–28 (2011)
22. S Lakshminarayanan, D Stensrud, Ensemble Kalman filter. *IEEE Control. Syst.* **29**(3), 34–46 (2009)
23. J Anderson, Ensemble Kalman filters for large geophysical applications. *IEEE Control. Syst.* **29**(3), 66–82 (2009)
24. G Evensen, The ensemble Kalman filter for combined state and parameter estimation. *IEEE Control. Syst.* **29**(3), 83–104 (2009)
25. J Mandel, J Beezley, J Coen, M Kim, Data assimilation for wildland fires. *IEEE Control. Syst.* **29**(3), 47–65 (2009)
26. R Furrer, T Bengtsson, Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *J. Multivar. Anal.* **98**(2), 227–255 (2007)
27. M Butala, J Yun, Y Chen, R Frazin, F Kamalabadi, in *15th IEEE International Conference on Image Processing*. Asymptotic convergence of the ensemble Kalman filter (IEEE, San Diego, 2008), pp. 825–828
28. J Mandel, L Cobb, JD Beezley, On the convergence of the ensemble Kalman filter. *Appl. Math.* **56**(6), 533–541 (2011)
29. M Frei, *Ensemble Kalman Filtering and Generalizations* (Dissertation, ETH, Zürich, 2013). nr. 21266
30. M Katzfuss, JR Stroud, CK Wikle, Understanding the ensemble Kalman filter. *Am. Stat.* **70**(4), 350–357 (2016)
31. F Le Gland, V Monbet, V Tran, in *The Oxford Handbook of Nonlinear Filtering*, ed. by D Crisan, B Rozovskii. Large sample asymptotics for the ensemble Kalman filter (Oxford University Press, Oxford, 2011), pp. 598–634
32. M Butala, R Frazin, Y Chen, F Kamalabadi, Tomographic imaging of dynamic objects with the ensemble Kalman filter. *IEEE Trans. Image Process.* **18**(7), 1573–1587 (2009)
33. J Dunik, O Straka, M Simandl, E Blasch, Random-point-based filters: analysis and comparison in target tracking. *IEEE Trans. Aerosp. Electron. Syst.* **51**(2), 1403–1421 (2015)
34. S Gillijns, O Mendoza, J Chandrasekar, B De Moor, D Bernstein, A Ridley, in *American Control Conference, 2006*. What is the ensemble Kalman filter and how well does it work? (IEEE, Minneapolis, 2006), pp. 4448–4453

35. M Roth, C Fritsche, G Hendeby, F Gustafsson, in *European Signal Processing Conference 2015 (EUSIPCO 2015)*. The ensemble Kalman filter and its relations to other nonlinear filters (IEEE, France, 2015)
36. JL Anderson, An ensemble adjustment Kalman filter for data assimilation. *Mon. Weather Rev.* **129**(12), 2884–2903 (2001)
37. CH Bishop, BJ Etherton, SJ Majumdar, Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Weather Rev.* **129**(3), 420–436 (2001)
38. JS Whitaker, TM Hamill, Ensemble data assimilation without perturbed observations. *Mon. Weather Rev.* **130**(7), 1913–1924 (2002)
39. MK Tippett, JL Anderson, CH Bishop, TM Hamill, JS Whitaker, Ensemble square root filters. *Mon. Weather Rev.* **131**(7), 1485–1490 (2003)
40. JL Anderson, SL Anderson, A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Weather Rev.* **127**(12), 2741–2758 (1999)
41. PJ van Leeuwen, Comment on “data assimilation using an ensemble Kalman filter technique”. *Mon. Weather Rev.* **127**(6), 1374–1377 (1999)
42. E Ott, BR Hunt, I Szunyogh, AV Zimin, EJ Kostelich, M Corazza, E Kalnay, DJ Patil, JA Yorke, A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A.* **56**(5), 415–428 (2004)
43. TM Hamill, JS Whitaker, C Snyder, Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Weather Rev.* **129**(11), 2776–2790 (2001)
44. PN Raanes, On the ensemble Rauch-Tung-Striebel smoother and its equivalence to the ensemble Kalman smoother. *Q. J. R. Meteorol. Soc.* **142**(696), 1259–1264 (2016)
45. M Zupanski, Maximum likelihood ensemble filter: theoretical aspects. *Mon. Weather Rev.* **133**(6), 1710–1726 (2005)
46. TM Hamill, C Snyder, A hybrid ensemble Kalman filter–3d variational analysis scheme. *Mon. Weather Rev.* **128**(8), 2905–2919 (2000)
47. PJ van Leeuwen, A variance-minimizing filter for large-scale applications. *Mon. Weather Rev.* **131**(9), 2071–2084 (2003)
48. C Snyder, T Bengtsson, P Bickel, J Anderson, Obstacles to high-dimensional particle filtering. *Mon. Weather Rev.* **136**(12), 4629–4640 (2008)
49. PJ van Leeuwen, Particle filtering in geophysical systems. *Mon. Weather Rev.* **137**(12), 4089–4114 (2009)
50. PJ van Leeuwen, Nonlinear data assimilation in geosciences: an extremely efficient particle filter. *Q. J. R. Meteorol. Soc.* **136**(653), 1991–1999 (2010)
51. M Frei, HR Künsch, Bridging the ensemble Kalman and particle filters. *Biometrika.* **100**(4), 781–800 (2013)
52. J Poterjoy, A localized particle filter for high-dimensional nonlinear systems. *Mon. Weather Rev.* **144**(1), 59–76 (2015)
53. EN Lorenz, in *Predictability of Weather and Climate*, ed. by T Palmer, R Hagedorn. Predictability—a problem partly solved (Cambridge University Press, Cambridge, 2006), pp. 40–58
54. DT Pham, Stochastic methods for sequential data assimilation in strongly nonlinear systems. *Mon. Weather Rev.* **129**(5), 1194–1207 (2001)
55. X Luo, I Moroz, Ensemble Kalman filter with the unscented transform. *Phys. D Nonlinear Phenom.* **238**(5), 549–562 (2009)
56. SJ Julier, JK Uhlmann, Unscented filtering and nonlinear estimation. *Proc. IEEE.* **92**(3), 401–422 (2004)
57. P Sakov, Comment on “ensemble Kalman filter with the unscented transform”. *Phys. D Nonlinear Phenom.* **238**(22), 2227–2228 (2009)
58. AS Stordal, HA Karlsen, G Nævdal, HJ Skaug, B Vallès, Bridging the ensemble Kalman filter and particle filters: the adaptive Gaussian mixture filter. *Comput. Geosci.* **15**(2), 293–305 (2011)
59. I Hoteit, X Luo, DT Pham, Particle Kalman filtering: a nonlinear Bayesian framework for ensemble Kalman filters. *Mon. Weather Rev.* **140**(2), 528–542 (2011)
60. M Frei, HR Künsch, Mixture ensemble Kalman filters. *Comput. Stat. Data Anal.* **58**, 127–138 (2013)
61. LN Trefethen, D Bau III, *Numerical linear algebra*. (SIAM, Philadelphia, 1997)
62. M Nørgaard, NK Poulsen, O Ravn, New developments in state estimation for nonlinear systems. *Automatica.* **36**(11), 1627–1638 (2000)
63. P Sakov, PR Oke, Implications of the form of the ensemble transformation in the ensemble square root filters. *Mon. Weather Rev.* **136**(3), 1042–1053 (2008)
64. DM Livings, SL Dance, NK Nichols, Unbiased ensemble square root filters. *Phys. D Nonlinear Phenom.* **237**(8), 1021–1028 (2008)
65. WG Lawson, JA Hansen, Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth. *Mon. Weather Rev.* **132**(8), 1966–1981 (2004)
66. O Leeuwenburgh, G Evensen, L Bertino, The impact of ensemble filter definition on the assimilation of temperature profiles in the tropical pacific. *Q. J. R. Meteorol. Soc.* **131**(613), 3291–3300 (2005)
67. Y Bar-Shalom, XR Li, T Kirubarajan, *Estimation with applications to tracking and navigation: Theory Algorithms and Software*. (Wiley-Interscience, New York, 2001)
68. P Sakov, L Bertino, Relation between two common localisation methods for the EnKF. *Comput. Geosci.* **15**(2), 225–237 (2010)
69. BR Hunt, EJ Kostelich, I Szunyogh, Efficient data assimilation for spatiotemporal chaos: a local ensemble transform Kalman filter. *Phys. D Nonlinear Phenom.* **230**(1–2), 112–126 (2007)
70. G Gaspari, SE Cohn, Construction of correlation functions in two and three dimensions. *Q. J. R. Meteorol. Soc.* **125**(554), 723–757 (1999)
71. CE Rasmussen, CKI Williams, *Gaussian processes for machine learning*. (The MIT Press, Cambridge, Mass, 2005)
72. T Hastie, R Tibshirani, J Friedman, *The elements of statistical learning: data mining, inference, and prediction*, 2nd ed. (Springer, New York, NY, 2011)
73. Y Zhang, DS Oliver, Improving the ensemble estimate of the Kalman gain by bootstrap sampling. *Math. Geosci.* **42**(3), 327–345 (2010)
74. I Myrseth, J Sætrum, H Omre, Resampling the ensemble Kalman filter. *Comput. Geosci.* **55**, 44–53 (2013)
75. N Papadakis, E Mémin, A Cuzol, N Gengembre, Data assimilation with the weighted ensemble Kalman filter. *Tellus A.* **62**(5), 673–697 (2010)
76. S Reich, A Nonparametric ensemble transform method for Bayesian inference. *SIAM J. Sci. Comput.* **35**(4), A2013–A2024 (2013)
77. M Roth, F Gustafsson, in *42nd International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. Computation and visualization of posterior densities in scalar nonlinear and non-Gaussian Bayesian filtering and smoothing problems (IEEE, New Orleans, 2017)

Submit your manuscript to a SpringerOpen® journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com