

RESEARCH ARTICLE

Comparison of ensemble-based data assimilation methods for sparse oceanographic data

Florian Beiser^{1,2}  | Håvard Heitlo Holm¹ | Jo Eidsvik²¹Mathematics and Cybernetics, SINTEF Digital, Oslo, Norway²Department of Mathematical Sciences, NTNU, Trondheim, Norway**Correspondence**Florian Beiser, Mathematics and Cybernetics, SINTEF Digital, Forskningsveien 1, 0373 Oslo, Norway.
Email: florian.beiser@sintef.no**Funding information**

Norges Forskningsråd, Grant/Award Number: 310515

Abstract

Probabilistic forecasts in oceanographic applications, such as drift trajectory forecasts for search-and-rescue operations, face challenges due to high-dimensional complex models and sparse spatial observations. We discuss localisation strategies for assimilating sparse point observations and compare the implicit equal-weights particle filter and a localised version of the ensemble-transform Kalman filter. First, we verify these methods thoroughly against the analytic Kalman filter solution for a linear advection diffusion model. We then use a nonlinear simplified ocean model to do state estimation and drift prediction. The methods are rigorously compared using a wide range of metrics and skill scores. Our findings indicate that both methods succeed in approximating the Kalman filter reference for linear models of moderate dimensions, even for small ensemble sizes. However, in high-dimensional settings with a nonlinear model, we discover that the outcomes are significantly influenced by the dependence of the ensemble Kalman filter on relaxation and the particle filter's sensitivity to the chosen model error covariance structure. Upon proper relaxation and localisation parametrisation, the ensemble Kalman filter version outperforms the particle filter in our experiments.

KEYWORDS

data assimilation, oceanographic applications, sparse observations, spatio-temporal statistics

1 | INTRODUCTION

Data assimilation plays an essential role in enhancing the reliability of operational oceanographic and atmospheric forecasts by providing a framework to update and calibrate numerical models using observed data (Asch et al. 2016; Evensen 2009). There is typically a wide array of available observations, such as satellite imagery, radar measurements, weather stations, and ocean buoys. These

observations represent various physical quantities and can exhibit diverse connections to the dynamical model, necessitating substantial preprocessing efforts in operational prediction systems. Furthermore, assimilating the observations and running the complex numerical simulations often demand substantial computational resources.

In addition to regularly updated operational forecasts, there is a crucial need to provide targeted predictions for localised and time-sensitive scenarios, including

search-and-rescue operations or sea contamination incidents (Breivik et al. 2013; Röhrs et al. 2018). To gain a better understanding of the local conditions in such situations, it is possible to deploy and gather in-situ observations using drifters. However, owing to the fixed schedules for the operational forecasts, such in-situ observations will rarely be processed immediately by the operational data assimilation cycle. Instead, we aim to concurrently assimilate these observations into ensembles of efficient simplified models. This approach complements the traditional operational framework by enabling rapid predictions with quantification of associated uncertainties, possibly without requiring access to supercomputers. This study aims to investigate efficient data assimilation methods tailored for sparse observations for such scenarios.

We limit our scope to point observations, such as information gathered from buoys and drifters. Although these observations provide valuable insight into the ocean state at their respective locations, they are often separated by significant distances, necessitating spatio-temporal modelling to fill the gaps between sparse data points. In this article, we utilise buoy information for two purposes: (i) to constrain an advection diffusion process for particle concentration (Foss et al. 2021), and (ii) to constrain drift trajectories in a simplified ocean model (Holm et al. 2020). Case (i) involves a linear system in space-time, allowing us to explore the properties of the targeted data assimilation methods for sparse observations by comparing their results to the optimal analytic Kalman filter (KF) solution. In contrast, case (ii) deals with a highly nonlinear dynamical model, prompting us to compare the different approaches using various performance metrics through synthetic simulation studies. Our primary motivation for focusing on case (ii) is to improve short-term predictions for search-and-rescue missions by leveraging the utilisation of spatially sparse buoy observations.

To effectively model and estimate the uncertainty associated with discretised high-dimensional state variables governed by (non-)linear partial differential equations, we base our work on ensemble prediction and ensemble-based data assimilation. From a statistical perspective, data assimilation methods share the objective of representing a conditional distribution given the available observations. In our case, as we are motivated by search-and-rescue applications, we focus on filtering distributions rather than smoothers that also give improved estimates of past states. The ensemble KF (EnKF) proposed by Evensen (1994) is widely used in practice, and various numerical adaptations have been developed to address specific practical challenges and problems. Sparse observations pose challenges related to filtering

efficiency and quality, and our focus centres around the ensemble transform KF (ETKF; Bishop et al., 2001) and explores localisation strategies for KFs. In contrast to EnKF and its variants, particle filters (PFs)—for example, see van Leeuwen (2009) or Chopin and Papaspiliopoulos (2020)—are appealing as they do not rely on the assumptions of a linear model and Gaussian probability distributions, at least in their most basic forms. However, PFs are less commonly employed in high-dimensional real-world applications owing to the issue of degeneracy. Holm et al. (2020) demonstrated a modern PF approach based on the implicit equal-weight particle filter (IEWPF) proposed by Zhu et al. (2016), which shows promising forecasting results for drift trajectories.

In this article, we systematically compare the statistical properties and performance of two ensemble-based data assimilation methods for sparse observations in practical oceanographic applications. We compare the ETKF with localisation for sparse observations against the state-of-the-art IEWPF. The comparison involves appraising their ability to reproduce analytic solutions for the linear model, in the first case by assessing a range of statistical performance metrics and evaluating forecast skills in the nonlinear case. Furthermore, we provide comprehensive discussions on the localisation techniques applied to the ETKF and the inherent localisation mechanism within the IEWPF.

The remainder of this article is organised as follows: Section 2 explicates state-of-the-art ensemble-based data assimilation techniques and reviews localisation strategies in the context of sparse observations. In Section 3, we use the dynamical model based on the advection diffusion equation to verify the relevant ensemble-based filtering methods against an analytical solution. Section 4 outlines the nonlinear simplified ocean model for drift trajectory prediction, and we compare the performance of the data assimilation methods in this context. Finally, closing remarks are in Section 5.

2 | THE DATA ASSIMILATION PROBLEM AND ENSEMBLE-BASED FILTERING

Spatio-temporal quantities are denoted by $\mathbf{x}(t, \mathbf{s})$, for time $t > 0$ and location \mathbf{s} , where we restrict ourselves to $\mathbf{s} \in \mathbb{R}^2$. Upon discretisation of the spatio-temporal domain of interest, the locations are represented at grid nodes of spatial locations $(\mathbf{s}_i)_{i=1}^{N_s}$ and time steps t^n , $n = 1, \dots, N_T$, where we will work with equidistant time steps without loss of generality. The state vector at time t^n is denoted $\mathbf{x}^n \in \mathbb{R}^{N_x}$ and can hold more than one physical variable per location if necessary ($N_x \geq N_s$). In oceanographic applications, the

dimension N_X is usually very high owing to large domains and several physical variables.

The numerical model is embraced in the model operator \mathcal{M} . It propagates the state vector from the previous time step t^{n-1} to the current t^n , defining the so-called forecast state. The model usually describes the physics of the ocean. To account for uncertainty coming from external factors, unknown model parameters and non-modelled physics, a Gaussian model error $\mathbf{v}^n \sim \mathcal{N}(0, \mathbf{Q})$ is added every time step. It is assumed that \mathbf{Q} is known and that the error terms are uncorrelated in time. Starting from initial state \mathbf{x}^0 , the model evolves as

$$\mathbf{x}^n = \mathcal{M}\mathbf{x}^{n-1} + \mathbf{v}^n, \quad n = 1, \dots \quad (1)$$

The oceanographic state \mathbf{x}^n is often only partially observed, and sometimes even indirectly. An observation is extracted from the true state vector by the measurement operator \mathbf{H} . Owing to the nature of the problems within this work, we impose the restriction that this operator is linear and that \mathbf{H} only extracts variables directly from a subset of spatial locations, meaning that the matrix consists of at most one 1 entry per row and otherwise 0 entries. The observation locations could change every time step, but for the sake of eased reading we ignore this in our notation. Measurement and representation inaccuracies are represented by the addition of a zero-mean Gaussian error $\epsilon^n \sim \mathcal{N}(0, \mathbf{R})$ with known covariance matrix \mathbf{R} and no correlations in time. We model observations $\mathbf{y}^n \in \mathbb{R}^{N_Y}$ by

$$\mathbf{y}^n = \mathbf{H}\mathbf{x}^n + \epsilon^n \quad (2)$$

and assume that data come from this observation model employed on the unknown true state. The covariance matrix \mathbf{R} is assumed to be diagonal, representing conditional independence between the observations given the state variables. A characteristic of many oceanographic applications, and a key assumption in our setting, is that in-situ observations are spatially sparse and of low dimension compared with the high-dimensional state vector; that is,

$$N_Y \ll N_X. \quad (3)$$

Sequential data assimilation refers to the workflow of sequentially updating the probability density of the state variables as more data become available. Often, this results in reduced uncertainty, especially in the neighbourhood of the observed variables. This kind of data assimilation is formalised in Bayes' rule. Using the state's density conditioned on all previous observations $p(\mathbf{x}^n | \mathbf{y}^{1:n-1})$ as the forecast (prediction or *prior*), this rule is used to assimilate the new observation \mathbf{y}^n and thus provides the analysis (filtering or *posterior*) density $p(\mathbf{x}^n | \mathbf{y}^{1:n})$ of the state. With the

independence assumption for the observation noise terms, this formalism can be applied recursively as

$$p(\mathbf{x}^n | \mathbf{y}^{1:n}) \propto p(\mathbf{y}^n | \mathbf{x}^n) p(\mathbf{x}^n | \mathbf{y}^{1:n-1}), \quad n = 1, 2, \dots, \quad (4)$$

starting with only prior information at the first time step.

Kalman filtering

In the case of a linear model $\mathcal{M} = \mathbf{M}$ and an initial Gaussian distribution for the state $\mathbf{x}^0 \sim \mathcal{N}(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0)$, all forecast and analysis distributions remain Gaussian. The data assimilation problem is then Gauss-linear and Bayes' formula, Equation (4), takes closed-form solutions for the mean vectors and covariance matrices. Assuming that we have an analysis (superscript "a") distribution characterised by the mean $\boldsymbol{\mu}^{n-1,a}$ and covariance matrix $\boldsymbol{\Sigma}^{n-1,a}$, the forecast (superscript "f") distribution $\mathbf{x}^n | \mathbf{y}^{1:n-1} \sim \mathcal{N}(\boldsymbol{\mu}^{n,f}, \boldsymbol{\Sigma}^{n,f})$ is obtained by evolving the first and second moments from the previous time step to the next by

$$\boldsymbol{\mu}^{n,f} = \mathbf{M}\boldsymbol{\mu}^{n-1,a} \quad (5a)$$

$$\boldsymbol{\Sigma}^{n,f} = \mathbf{M}\boldsymbol{\Sigma}^{n-1,a}\mathbf{M}^T + \mathbf{Q}. \quad (5b)$$

The analysis distribution $\mathbf{x}^n | \mathbf{y}^{1:n} \sim \mathcal{N}(\boldsymbol{\mu}^{n,a}, \boldsymbol{\Sigma}^{n,a})$ is achieved by assimilating the latest observation \mathbf{y}^n via Bayes' rule for the given Gaussian model, resulting in

$$\boldsymbol{\mu}^{n,a} = \boldsymbol{\mu}^{n,f} + \mathbf{K}(\mathbf{y}^n - \mathbf{H}\boldsymbol{\mu}^{n,f}) \quad (6a)$$

$$\boldsymbol{\Sigma}^{n,a} = \boldsymbol{\Sigma}^{n,f} - \mathbf{K}\boldsymbol{\Sigma}^{n,f}\mathbf{K}^T. \quad (6b)$$

Here, $\mathbf{K} = \boldsymbol{\Sigma}^{n,f}\mathbf{H}^T(\mathbf{H}\boldsymbol{\Sigma}^{n,f}\mathbf{H}^T + \mathbf{R})^{-1}$ is the Kalman gain, which in Equation (6a) maps the so-called innovation, $\mathbf{y}^n - \mathbf{H}\boldsymbol{\mu}^{n,f}$, to the state space using the state covariance matrix and the observation model. From the numerical perspective, note that the KF requires storage and propagation of the size $N_X \times N_X$ covariance matrix, which can be infeasible for high-dimensional systems.

Ensemble-based data assimilation

In oceanographic applications, the linearity assumptions of the KF are rarely met, and linearised approaches can suffer from divergence challenges. There is hence a need for more flexible methods, and ensemble-based approaches have been employed as a computationally feasible method to represent statistical solutions of nonlinear systems, even for large N_X .

Therein, the continuous distribution of the state variable is approximated by an ensemble of realisations $(\mathbf{x}_e^n)_{e=1}^{N_e}$

and potentially by corresponding weights $(w_e)_{e=1}^{N_e}$. Following the Monte Carlo idea, the marginal distribution of \mathbf{x}^n at time t^n becomes

$$p(\mathbf{x}^n) \approx \sum_{e=1}^{N_e} w_e^n \delta(\mathbf{x}^n - \mathbf{x}_e^n), \quad (7)$$

where δ is the Dirac delta function and $\sum_{e=1}^{N_e} w_e^n = 1$.

In the statistical literature—for example, see Asch et al. (2016) and Vetra-Carvalho et al. (2018)—there are two popular groups of methods for ensemble-based data assimilation, whose foundations and latest variants for the aforementioned problems are outlined in the following subsections.

2.1 | PFs in oceanographic applications

PFs are ensemble-based methods for solving the data assimilation problem using the Monte Carlo approach. In their simplest form, starting from a weighted ensemble approximation for $p(\mathbf{x}^0)$ or $p(\mathbf{x}^{n-1}|\mathbf{y}^{1:n-1})$ in the form of Equation (7), the forecast distribution $p(\mathbf{x}^n|\mathbf{y}^{1:n-1})$ can be approximated by propagating each ensemble member \mathbf{x}_e^{n-1} individually by the model Equation (1) to obtain \mathbf{x}_e^n . Plugging this into Bayes formula, Equation (4), the weights are updated as

$$w_e^n \propto p(\mathbf{y}^n|\mathbf{x}_e^n)p(\mathbf{x}_e^n|\mathbf{x}_e^{n-1})w_e^{n-1}. \quad (8)$$

Since ensemble members that have weights very close to zero do not contribute to the posterior probability distribution, it is common to combine Equation (8) with a discrete resampling of the ensemble members based on their weights; for example, see van Leeuwen (2009); Chopin and Papaspiliopoulos (2020) for reviews of resampling schemes. In practice, this means that we discard ensemble members with low weights and duplicate those with higher weights, thus ensuring that computational resources are used to describe the non-negligible part of the probability distribution. In high-dimensional applications, however, these basic PFs are prone to degenerate; that is, all but one ensemble member get a weight close to zero, leading to loss of statistical properties (Snyder et al. 2008).

Among other concepts, one way to counteract such degeneracy is to sample from a proposal density q_e instead of evolving the ensemble directly according to $p(\mathbf{x}_e^n|\mathbf{x}_e^{n-1})$ (van Leeuwen et al. 2019). The proposal density can be conditioned on the latest observation \mathbf{y}^n and the previous state \mathbf{x}_e^{n-1} for all ensemble members $e = 1, \dots, N_e$. The

weights are then modified to

$$w_e^{n,*} = \frac{w_e^n}{q_e(\mathbf{x}_e^n|\mathbf{x}_e^{n-1}, \mathbf{y}^n)}.$$

The variance in the weights can be reduced in this way, and the minimal variance is achieved by $q_e(\mathbf{x}^n|\mathbf{x}_e^{n-1}, \mathbf{y}^n) = p(\mathbf{x}^n|\mathbf{x}_e^{n-1}, \mathbf{y}^n)$ as described in Doucet et al. (2000) and often referred to, as the optimal proposal. Under the assumption of Gaussian errors and linear observation operator, which applies to our case, this proposal is a Gaussian distribution $\mathcal{N}(\mathbf{x}_e^{n,\text{opt}}, \mathbf{P})$ with

$$\mathbf{x}_e^{n,\text{opt}} = \mathcal{M}\mathbf{x}_e^{n-1} + \mathbf{Q}\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1} \times (\mathbf{y}^n - \mathbf{H}\mathcal{M}\mathbf{x}_e^{n-1}) \quad (9a)$$

$$\mathbf{P} = \mathbf{Q} - \mathbf{Q}\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{H}\mathbf{Q}. \quad (9b)$$

Note that the expression for $\mathbf{x}_e^{n,\text{opt}}$ is similar to the KF update in Equation (6a), but it uses the covariance structure from the model error \mathbf{Q} instead of the forecast $\Sigma^{n,f}$. Even this formulation of the PF will degenerate for high-dimensional systems; for example, see Morzfeld et al. (2017).

The optimal proposal density filter can be modified such that all posterior ensemble members obtain a certain target weight w_{target}^n , thus avoiding ensemble collapse and the need of resampling. Instead of drawing realisations from the proposal distribution directly, the IEWPF first samples ξ_e and ζ_e from $\mathcal{N}(0, \mathbf{I}_{N_x})$ and next implicitly transforms the samples to a target distribution. This filter, introduced by Zhu et al. (2016) and modified by Skauvold et al. (2019), utilises a version of the optimal proposal density where ξ_e and ζ_e are constructed to be perpendicular and scaled according to factors $\alpha_e^{1/2}$ for $e = 1, \dots, N_e$ and $\beta^{1/2}$ respectively, before being transformed by \mathbf{P} according to

$$\mathbf{x}_e^n = \mathbf{x}_e^{n,\text{opt}} + \mathbf{P}^{1/2}(\alpha_e^{1/2}\xi_e + \beta^{1/2}\zeta_e). \quad (10)$$

Here, the constant β is a tuning parameter that influences the statistical quality of the results, whereas the α_e values are calculated implicitly for each ensemble member to ensure equal weights. To get an intuition of the idea behind IEWPF, we observe that Equation (9a) develops the state using \mathcal{M} but with no model error. These unperturbed states are then used to get $\mathbf{x}_e^{n,\text{opt}}$ by assimilating the observations. We then perturb $\mathbf{x}_e^{n,\text{opt}}$ in Equation (10), but instead of using the model error, which is sampled from $\mathcal{N}(0, \mathbf{Q})$, we sample perturbations from $\mathcal{N}(0, \alpha_e\mathbf{P})$ and $\mathcal{N}(0, \beta\mathbf{P})$ in such a way that we counteract the change of weights in the ensemble. Note that \mathbf{P} is constructed from

\mathbf{Q} as seen in Equation (9b), so that these perturbations can be thought of as targeted sampling of the model error.

The tuning parameter β needs to be selected with care. A small value gives small spread of the ensemble that likely underestimates the variability, whereas a bigger β increases the spread. Holm et al. (2020, appendix A) derived lower and upper bounds for this tuning parameter. In the subsequent experiments, we tune β manually, mainly by calibration of coverage probabilities as suggested in Skauvold et al. (2019). From experience, it seems that β values around 0.5 are a good start. Albeit that the choice of β is independent of the ensemble size, it is influenced by the dynamics of the problem. Hence, one can find a suitable choice of β for a specific kind of scenario and then keep it fixed in future experiments with similar characteristics.

It should be noted that, in contrast to general PFs, the IEWPF requires that the model error is additive and from a Gaussian distribution. Furthermore, there are no guarantees on how the IEWPF performs, even when the ensemble size goes to infinity. Still, the performance tends to be very good in large-size systems. With $\beta = 0$, the implicit transform has a gap that leads to asymptotic bias (Skauvold et al., 2019), but this seems to be adjusted reasonably well by the second part having $\beta > 0$. The IEWPF has recently been shown to be applicable and efficient for assimilating point-based observations into a simplified ocean model based on the shallow-water equations (Holm et al. 2020). Herein, this method represents a state-of-the-art PF and is investigated more thoroughly.

2.2 | EnKFs in oceanographic applications

The EnKF by Evensen (1994, 2009) is an ensemble-based version of the KF, given in Equations (5) and (6). Originally presented as a data assimilation method for nonlinear systems, it also solves the problem of having to store and propagate the $N_X \times N_X$ state covariance matrix Σ .

In the ensemble representation from Equation (7), all weights are kept equal, and the state of each ensemble member is propagated by the model in Equation (1). The forecast state covariance is estimated from the ensemble as

$$\hat{\Sigma}^{n,f} = \frac{1}{N_e - 1} \sum_{e=1}^{N_e} (\mathbf{x}_e^{n,f} - \bar{\mathbf{x}}^{n,f}) (\mathbf{x}_e^{n,f} - \bar{\mathbf{x}}^{n,f})^T, \quad (11)$$

where $\bar{\mathbf{x}}^{n,f}$ denotes the ensemble mean. The ensemble members are then updated along the same linear projection

$$\mathbf{x}_e^{n,a} = \mathbf{x}_e^{n,f} + \hat{\mathbf{K}}(\mathbf{y}^n - \mathbf{H}\mathbf{x}_e^{n,f} - \epsilon_e^n), \quad (12)$$

where the Kalman gain becomes $\hat{\mathbf{K}} = \hat{\Sigma}^{n,f} \mathbf{H}^T (\mathbf{H} \hat{\Sigma}^{n,f} \mathbf{H}^T + \mathbf{R})^{-1}$. In Equation (12), the perturbation $\epsilon_e^n \sim \mathcal{N}(0, \mathbf{R}^n)$ is added to adjust the variance in the solution ensemble, motivated by exact sampling in the linear Gaussian situation. This solution is therefore termed the stochastic EnKF (Burgers et al., 1998; Houtekamer & Mitchell, 1998; van Leeuwen, 2020). The classical stochastic EnKF in Equation (12) requires that we obtain and store the relevant covariances from the ensemble and then factorise matrices to solve a linear system. For high-dimensional applications, this quickly becomes expensive, and it is therefore common to circumvent the covariance assembling (Evensen 2003) or to use so-called deterministic square-root formulations instead (Whitaker and Hamill 2002).

To avoid working in the state space, the ETKF (Bishop et al., 2001) reformulates Equation (12) via linear algebraic identities into a particular example of a deterministic square-root filter, which works in ensemble dimensions instead. Mathematically, let $\mathbf{X}^{n,f} = [\mathbf{x}_1^{n,f}, \dots, \mathbf{x}_{N_e}^{n,f}]$ be the matrix of prior ensemble states, and let $\bar{\mathbf{X}}^{n,f}$ be an $N_X \times N_e$ matrix where all columns are $\bar{\mathbf{x}}^{n,f}$. The ETKF then works on the state perturbation matrix $\mathbf{X}_{\text{pert}}^{n,f} = \mathbf{X}^{n,f} - \bar{\mathbf{X}}^{n,f}$, and calculates the mean of the analysis ensemble

$$\bar{\mathbf{X}}^{n,a} = \bar{\mathbf{X}}^{n,f} + \mathbf{X}_{\text{pert}}^{n,f} \mathbf{A} (\mathbf{H} \mathbf{X}_{\text{pert}}^{n,f})^T \mathbf{R}^{-1} (\mathbf{y}^n - \mathbf{H} \bar{\mathbf{x}}^{n,f}), \quad (13)$$

where

$$\mathbf{A} = ((N_e - 1) \mathbf{I}_{N_e} + (\mathbf{H} \mathbf{X}_{\text{pert}}^{n,f})^T \mathbf{R}^{-1} \mathbf{H} \mathbf{X}_{\text{pert}}^{n,f})^{-1} \quad (14)$$

plays the role of the analysis covariance matrix. The ensemble members are then spread around $\bar{\mathbf{x}}^{n,a}$ according to

$$\mathbf{X}^{n,a} = \bar{\mathbf{X}}^{n,a} + \mathbf{X}_{\text{pert}}^{n,f} ((N_e - 1) \mathbf{A})^{1/2}, \quad (15)$$

where we use a singular-value decomposition to find the square-root of \mathbf{A} . The properties of the ETKF remain the same as for the EnKF, and we refer to Li (2007) for further details on the transform.

The derivation of these methods assumes a linear model, and asymptotic convergence results for increased ensemble size cannot be proved for nonlinear cases. Still, the EnKF and its variants have been prevalent and successfully used in oceanographic applications; for example, see Carrassi et al. (2018).

The error covariance matrix in these kinds of filters is estimated from the ensemble and can lead to systematic underestimation. Typical approaches to counteract this are inflation or localisation. Anderson and Anderson (1999) introduced ‘‘covariance inflation’’ by a multiplicative factor to keep more variability in the ensemble, where

TABLE 1 Number of non-zero elements in the matrix operations used in the implicit equal-weight particle filter.

Equation	Matrix operation	Dimensions	No. non-zeros for any $(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}$	No. non-zeros for diagonal $(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}$
	$\mathbf{Q}^{1/2}$	$N_X \times N_X$	$\leq (2r_Q + 1)^2 N_X$	\leftarrow
	\mathbf{Q}	$N_X \times N_X$	$\leq (4r_Q + 1)^2 N_X$	\leftarrow
	\mathbf{H}	$N_Y \times N_X$	N_Y	\leftarrow
	$(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}$	$N_Y \times N_Y$	$\leq N_Y^2$	N_Y
	$\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}$	$N_X \times N_Y$	$\leq N_Y^2$	N_Y
(9a)	$\mathbf{Q}\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}$	$N_X \times N_Y$	$\leq (4r_Q + 1)^2 N_Y^2$	$\leq (4r_Q + 1)^2 N_Y$
	$\mathbf{H}\mathbf{Q}^{T/2}$	$N_Y \times N_X$	$\leq (2r_Q + 1)^2 N_Y$	\leftarrow
	$(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{H}\mathbf{Q}^{T/2}$	$N_Y \times N_X$	$\leq (2r_Q + 1)^2 N_Y^2$	$\leq (2r_Q + 1)^2 N_Y$
	$\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{H}\mathbf{Q}^{T/2}$	$N_X \times N_X$	$\leq (2r_Q + 1)^2 N_Y^2$	$\leq (2r_Q + 1)^2 N_Y$
S in (16)	$\mathbf{Q}^{1/2}\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{H}\mathbf{Q}^{T/2}$	$N_X \times N_X$	$\leq (2r_Q + 1)^4 N_Y^2$	$\leq (2r_Q + 1)^4 N_Y$

several suggestions for the determination of an adaptive factor exist in literature; for example, see Desroziers et al. (2006); Anderson (2009); Sætrom and Omre (2013); Raanes et al. (2019). Similarly, additive inflation was presented by Ott et al. (2004). However, Li et al. (2009) also point out that covariance inflation may not work appropriately in large, complex models. Hence, we will mainly concentrate on localisation.

2.3 | Sparse observations

The focus of this article is on assimilating spatially sparse point observations. This naturally suggests to look closer on “localisation” in the filters. Although localisation is important for general applications, the sparse observations scenario considered here motivates one to study specific methods with good assimilation quality and algorithmic efficiency.

2.3.1 | Localisation and sparse observation handling in the IEWPF

The need for localisation in EnKF-based schemes arises from the spurious correlations introduced by the term $\hat{\Sigma}^{n,f}\mathbf{H}^T$, which represents the estimated covariance terms between all state variables and all observations in Equation (12). As pointed out in Section 2.1, the optimal proposal distribution in Equation (9a) updates the state vector with a similar expression, but it uses the correlations in the model error, $\mathbf{Q}\mathbf{H}^T$, rather than the empirical $\hat{\Sigma}^{n,f}\mathbf{H}^T$. This means that the optimal proposal filter by design does not lead to spurious correlations. Still, the structure of \mathbf{Q}

of course influences the distribution, and in particular a local structure in \mathbf{Q} that does not overlap between observation sites entails updates in Equation (9a) that are local as well.

To show that the IEWPF updates in Equations (9a) and (10) are local if \mathbf{Q} is local, we consider the pattern of non-zero values in the matrix operations in the two equations. First, let us use r_Q to denote the radius in terms of number of grid cells that information is spread through the model error covariance matrix. This means that $\mathbf{Q}^{1/2}$ contains at most $(2r_Q + 1)^2$ non-zero elements, whereas $\mathbf{Q} = \mathbf{Q}^{T/2}\mathbf{Q}^{1/2}$ spreads the information twice as far, thus having at most $(4r_Q + 1)^2$ non-zero elements. Second, we observe that Equation (9b) can be rewritten as

$$\mathbf{P} = \mathbf{Q}^{T/2}(\mathbf{I}_{N_X} - \underbrace{\mathbf{Q}^{1/2}\mathbf{H}^T(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{H}\mathbf{Q}^{T/2}}_{=: \mathbf{S}})\mathbf{Q}^{1/2}, \quad (16)$$

where we for convenience use \mathbf{S} to represent the longest matrix expression. This means that $\mathbf{P}^{1/2} = (\mathbf{I}_{N_X} - \mathbf{S})^{1/2}\mathbf{Q}^{1/2}$, which is what we need in Equation (10).

In Table 1, we give upper bounds on the number of non-zero elements when stepping through the matrix operations in Equations (9a) and (16). In the rightmost column, we assume that observations are sufficiently sparse so that both $\mathbf{H}\mathbf{Q}\mathbf{H}^T$ and \mathbf{R} are diagonal, whereas we make no such assumption in the second column from the right. We do, however, assume that \mathbf{H} maps a subset of the state variables directly to observational space. If we now consider a single observation, meaning $N_Y = 1$ and scalar $(\mathbf{H}\mathbf{Q}\mathbf{H}^T + \mathbf{R})^{-1}$, we see from Table 1 that the innovation in Equation (9a) is spread in the neighbourhood of the observation location within a radius of $2r_Q$ grid cells.

Thus, $\mathbf{x}_e^{n,\text{opt}}$ differs from $\mathcal{M}\mathbf{x}_e^{n-1}$ only locally around the observation.

To see the resulting non-zero pattern for \mathbf{S} in Equation (16), we realise that $\mathbf{H}\mathbf{Q}^{\text{T}/2}$ consists of a single row from $\mathbf{Q}^{\text{T}/2}$. Furthermore, we have that the leftmost term $\mathbf{Q}^{1/2}\mathbf{H}^{\text{T}}$ is the same as $(\mathbf{H}\mathbf{Q}^{\text{T}/2})^{\text{T}}$. Consequently, \mathbf{S} is simply the (scaled) outer product of the column from $\mathbf{Q}^{1/2}$ corresponding to the observation, and therefore has the local correlation pattern in both its rows and columns, but zero in all rows and columns for which the observed state variable is not correlated with through $\mathbf{Q}^{1/2}$. This then means that $(\mathbf{I}_{N_x} - \mathbf{S})$ differs from the identity only locally to the observed state variable as well, which finally means that $\mathbf{P}^{1/2} = (\mathbf{I}_{N_x} - \mathbf{S})^{1/2}\mathbf{Q}^{1/2}$ differs from $\mathbf{Q}^{1/2}$ only locally to the observations as well.

It should be noted, though, that the values of α_e and β depend on the innovation obtained from all observations in the domain. These parameters are therefore global parameters, but since they are scalars they do not contribute to any change in the local correlation structures. Related discussions on how to utilise local covariance structures and sparse observations for efficient implementations of the IEWPF can be found in Holm (2020).

In addition to demonstrating the built-in localisation in the IEWPF, Table 1 shows how sensitive the IEWPF is to the structure of the model error covariance matrix \mathbf{Q} for spreading observed information in the state space. In the extreme case of a diagonal \mathbf{Q} (meaning $r_Q = 0$) all matrix operations will contain exactly N_Y non-zeros, and only the observed state variables will be affected by the data assimilation.

2.3.2 | Localisation in the EnKF

In a statistical sense, the spurious correlations in the EnKF are due to a poor Monte Carlo approximation of the true covariance matrix; for example, see Houtekamer and Zhang (2016). In the spatio-temporal physical model, information propagates at finite speed and long-distance correlations are unlikely to be significant. Prevailing techniques to counteract these artefacts are covariance or observation localisation as they are outlined in Sakov and Bertino (2011). Both of these exploit the physical distance between two points in space to reduce information propagation effects, and this has been demonstrated to work well in practice; for example, see Soares et al. (2021). For many oceanographic applications, it is important that the geostrophic imbalance introduced by the localisation in the EnKF does not outweigh the natural imbalance—Greybush et al. (2011) provide a discussion and representative experiments on this issue.

For the ETKF, Ott et al. (2004) introduce an efficient localisation scheme, and this is referred to as LETKF. Further developed implementations using parallelisation and observation batching exist; for example, see Hunt et al. (2007). The LETKF variants are also popular in numerical weather prediction; for example, see Szunyogh et al. (2007), where for example global satellite data are common. In the LETKF, one loops over the state locations or sets of state locations in a batch area. Doing so, one updates state variables by means of the ETKF using only a specified set of local observations per batch.

2.3.3 | Sparse observation handling in the ETKF

In an oceanographic scenario with observations at only a few locations, we prefer to use covariance localisation (Houtekamer and Mitchell 2001) to achieve computational control, as we explain later herein. One then defines local domains around each observation site only, where the size N_X^{loc} of a local area is significantly smaller than the full state space. For the choice of the radius of the resulting local domains, several approaches exist—for example, see Kirchgessner et al. (2014)—but we advocate using model-informed radii such as the model error range. We next assume that observations with non-overlapping local areas have negligible correlation, and they can be updated separately. Owing to the motivating case with sparse observations, we expect to have few overlapping areas. The reduced dimension of the local area compared with the full state vector will make computations more efficient.

Still, with nonlinear dynamical models, it is sometimes difficult to predict the possibly undesired effects of local approximations. Using sequential data integration, one can run through the data in multiple assimilation steps, and in doing so one properly accounts for the correlations. In cases of overlapping local observation areas, we therefore recommend splitting the observations into observational batches \mathcal{B}_b , $b = 1, \dots, B$, of assumed uncorrelated observations for serial processing, as originally introduced in Houtekamer and Mitchell (1998). In our context, the batches at each step are constructed from far-apart observations. The local areas and the sequential processing give good control for handling correlations from sparse spatial observations. Nerger (2015) discusses how interactions of localisation and serial observation processing could destabilise the filter, but this effect seems to be minuscule with reasonably set local areas and minimal correlation within a batch.

Figure 1 (left) illustrates the definition of local areas around depicted observation sites. The figure further indicates the splitting into batches (middle and right displays),




FIGURE 1 Schematic decomposition of the physical space (left panel) into local areas $N_X^{\text{loc}} \ll N_X$ around observations, indicated by a filled grid cell, and the separation into two batches of uncorrelated observations (middle and right panels).




FIGURE 2 Schematic construction of the weighting \mathbf{w}_b using the $\mathbf{w}_{\text{loc},j}$ within the groups. For batches 1 and 2 from Figure 1, the weighting vectors \mathbf{w}_1 and \mathbf{w}_2 are built from the Gaspari–Cohn kernel around each observation within the two groups. At the corresponding centre, each $\mathbf{w}_{\text{loc},j}$ equals one and decays towards the boundary of the local domains. By definition, the supports are non-overlapping and the contributions are added up.

with observation sites within a batch being sufficiently far apart from each other.

For the covariance localisation, we consider weighting vectors \mathbf{w}_b and assume a tapering that assures $\mathbf{w}_b = 1$ at an observation site and $\mathbf{w}_b = 0$ outside the local areas. While veering away from an observation site, \mathbf{w} should transit decreasingly monotone and smoothly from one to zero. An example of a kernel fulfilling those requirements locally in continuous space is the Gaspari–Cohn (GC) function introduced in Gaspari and Cohn (1999), which enjoys popularity in EnKF localisation. We let $\mathbf{w}_{\text{loc},j} \in \mathbb{R}^{N_X^{\text{loc}}}$ be a properly scaled discretisation of the GC kernel around observation j , such that its support matches the radius of the local area. With the same notation, the weighting is composed as

$$\mathbf{w}_b = \sum_{j \in \mathcal{B}_b} \mathbf{w}_{\text{loc},j}. \quad (17)$$

For the observation batch from Figure 1, the weighting vectors \mathbf{w}_1 and \mathbf{w}_2 are illustrated in Figure 2 together with the contributions from $\mathbf{w}_{\text{loc},1}, \dots, \mathbf{w}_{\text{loc},5}$.

Within the recursion of batches \mathcal{B}_b , the local analysis states $\mathbf{x}^{n,a}(j) \in \mathbb{R}^{N_X^{\text{loc}}}$ are calculated independently in each local area around the observation sites $j \in \mathcal{B}_b$, whereby the computational overhead in presence of sparse observations is implicitly controlled as the assimilation scheme operates in much smaller dimensions than N_X . Note that the processing of a batch can hence influence the ensemble that is used as the next forecast. The analyses only take values in their respective small local regions, but to avoid cumbersome notation for transformations, we abuse the same notation for their extension to the full state size $\mathbf{x}^{n,a}(j) \in \mathbb{R}^{N_X}$.

Because the ETKF constructs neither the covariance matrix nor the Kalman gain, we do not incorporate the

weights directly into the Kalman update, but rather calculate the ETKF analysis state by Equation (15) and weight it with the forecast afterwards—recall that $\mathbf{x}^{n,a}(j)$ are the columns of $\mathbf{X}^{n,a}$ in Equation (15). For the stochastic EnKF, we see that this corresponds to tapering the Kalman gain, and both approaches boil down to equivalent implementations. The weighting reads

$$\mathbf{x}_e^{n,a,b} = (1 - \mathbf{w}_b)\mathbf{x}_e^{n,a,b-1} + \mathbf{w}_b \sum_{j \in B_b} \mathbf{x}_e^{n,a}(j), \quad (18)$$

where $\mathbf{x}^{n,a,0} = \mathbf{x}^{n,f}$ and after the last batch, we set $\mathbf{x}^{n,a} = \mathbf{x}^{n,a,B}$ to obtain a final analysis state. For the global analysis state, this means that most analysis information is used near observation sites, whereas the forecast with its full spread is retained far away from data. The formulation of covariance location in the form of Equation (18) is furthermore convenient since the ETKF analysis scheme can be used as a black box without interference of the localisation weights.

From a principled statistical perspective, one could process each individual observation in a serial manner, but the collection in prescribed batches reduces iterations. In doing so, one assimilates the spatial data recursively, similar to the assimilation over time, albeit without the dynamical state evolution because all updates happen at the time when the data becomes available. For computational efficiency one again imposes some kind of local routine, and in practice this may rely on GC tapering of the matrices involved. This is commonly done in implementations of kriging, or in applications with sequential uncertainty reduction, where the analysis can depend on the choice of conditioning order; for example, see Nussbaumer et al. (2018).

Complementary to inflation as mentioned in Section 2.2, Zhang et al. (2004) present relaxation-to-prior to counteract overfitting. Introducing a scaling parameter $\phi \in [0, 1]$, the weighting vector can be constructed with neatly integrated relaxation as

$$\mathbf{w}_b^{\text{relax}} = \sum_{j \in B_b} \phi \mathbf{w}_{\text{loc},j}. \quad (19)$$

Here, $\phi = 0$ represents a pure Monte Carlo simulation, and $\phi = 1$ is the previously presented scheme without relaxation.

Algorithm 1 summarises the ETKF with covariance localisation as we will use it later on. We will refer to it as “SparseObsETKF” in order to avoid mixing up with the LETKF. This implementation helps us to keep good control of correlations as well as computational overhead in the presence of sparse observations, and it provides a convenient integration of tapering and relaxation.

Algorithm 1. Analysis scheme with localisation for sparse observations (SparseObsETKF)

Given $\mathbf{X}^{n,f}$. Parameters: localisation radius
and relaxation ϕ
Set $\mathbf{X}^{n,a,0} = \mathbf{X}^{n,f}$
for $b = 1, \dots, B$ **do**
 Allocate $\mathbf{w}_b^{\text{relax}}$ $\triangleright N_X$
 for $j \in B_b$ **do**
 Calculate local $\mathbf{X}^{n,a}(j)$ using the
 ETKF where $\mathbf{X}^{n,a,b-1}$ is the forecast $\triangleright N_X^{\text{loc}}$
 $\mathbf{w}_b^{\text{relax}} += \phi \mathbf{w}_{\text{loc},j}$
 end for
 $\mathbf{X}^{n,a,b} = (1 - \mathbf{w}_b^{\text{relax}})\mathbf{X}^{n,a,b-1} + \mathbf{w}_b^{\text{relax}}$
 $\sum_{j \in B_b} \mathbf{X}^{n,a}(j)$ $\triangleright N_X$
end for
 $\mathbf{X}^{n,a} = \mathbf{X}^{n,a,B}$

3 | COMPARISON AGAINST THE ANALYTICAL KALMAN FILTER IN A LINEAR GAUSSIAN ADVECTION DIFFUSION MODEL

In this section we examine a linear Gaussian spatio-temporal model. As mentioned in Section 2, this means that the analytic KF in Equations (5) and (6) defines the optimal solution. Ensemble-based approximations and localisation effects of the different filtering techniques from Section 2 can be verified against the KF.

3.1 | Advection diffusion model

Inspired by Sigrist et al. (2015), we consider a stochastic advection diffusion equation for state c given by

$$\frac{\partial c(t, \mathbf{s})}{\partial t} = \nabla \cdot d \nabla c(t, \mathbf{s}) - \mathbf{v}_t \cdot \nabla c(t, \mathbf{s}) + \zeta c(t, \mathbf{s}) + W(t, \mathbf{s}). \quad (20)$$

The model’s parameters are $d = 0.25$ for the diffusion, $\mathbf{v} = (1.0, 0.1)^T$ for the advection, and $\zeta = -0.0001$ for the damping. We assume the stochastic error process W has uncorrelated elements over time but smooth dependent spatial components at each time, and that Equation (20) holds for one sampled path (realisation) of W . We consider a rectangular spatial domain $[0, 5] \times [0, 3]$ with periodic boundary conditions, and c will be initialised at time $t = 0$ as a Gaussian random field.

Equation (20) can, for instance, be used to represent marine pollution dynamics (Foss et al. 2021), where the goal is to predict the concentration $c = c(t, \mathbf{s})$ of a contaminant over time and space in the ocean. In that case, the




FIGURE 3 The “truth” realisation of the stochastic process at the initialisation and selected observation times, with black dots marking the observation sites and black crosses signifying two selected locations of interest.




FIGURE 4 For the same times as in Figure 3, the resulting analysis mean (upper row) and standard deviation (lower row) of the Kalman filter.

advection parameter \mathbf{v} would typically come from a full ocean model if vertical currents are ignored.

In the discretised setting, the spatial domain is covered by a uniform Cartesian grid with centre points $(\mathbf{s}_i)_{i=1}^{N_s}$ in quadratic cells of size 0.1×0.1 . The state vector \mathbf{x}^n collects all concentrations $c(t^n, \mathbf{s}_i)$ at regular time steps t^n . The initial state is represented by $\mathbf{x}^0 \sim \mathcal{N}(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0)$ with mean vector $\boldsymbol{\mu}^0$ and covariance matrix $\boldsymbol{\Sigma}^0$ having Matérn-type:

$$\Sigma_{k,l}^0 = \sigma^2(1 + \psi D_{k,l}) \exp(-\psi D_{k,l}),$$

where $\sigma = 0.5$ is the standard deviation (assumed constant at all locations) and $\psi = 3.5$ is the Matérn correlation decay parameter, and $D_{k,l}$ is the distance between \mathbf{s}_k and \mathbf{s}_l . The mean $\boldsymbol{\mu}^0$ equals 10 in the northeast with higher bell-shaped concentration values in the southwest; see Figure 4(left).

For the numerical solution of the stochastic partial differential equation in Equation (20), a temporal forward and spatial central finite-difference scheme is employed such that the model resembles Equation (1) with the linear operator $\mathcal{M} = \mathbf{M}$. With periodic boundary conditions, the low-concentration area leaves the domain on the east boundary and enters from the west. The model error \mathbf{v} is again represented by a Gaussian random field with a covariance matrix \mathbf{Q} of a similar Matérn type. A smaller standard deviation $\sigma = 0.125$ is used, and a larger correlation decay parameter $\psi = 7.0$ leads to model noise with smaller correlation.

3.2 | Experiment design and analytic solution

A single realisation of the advection diffusion generated by the forward model is used to retrieve observations for the filtering; see Figure 3. It is simulated for 250 time steps with $\Delta t = 0.01$ on a grid of size 50×30 . The simulated process is observed at $t^n = 25n$, $n = 1, \dots, 10$ at 15 grid cells marked red in Figure 3. These direct state observations are made with a small observation error $\epsilon^n \sim \mathcal{N}(0, r^2 \mathbf{I})$, $r = 0.1$.

The KF reference solution is depicted for a selection of time steps in Figure 4. As for the truth in Figure 3, the filtering mean (Figure 4, top) shows an east-north-east movement of the concentrations, as expected from the advection term. The standard deviations (Figure 4, bottom) are clearly reduced by the data assimilation, especially around the observation sites and in the advection direction. With time, however, the accuracy of the solution converges, as the corrections from doing data assimilation are balanced out by the dynamic model errors.

3.3 | Numerical results and evaluation metrics

We now solve this concentration advection diffusion problem using the ensemble-based methods IEWPF, ETKF, and SparseObsETKF from Section 2. The parameter




FIGURE 5 Mean error $\mathbf{err}_{\text{mean}}^{\text{KF}}$ at $t = 250$ for assimilation experiments with $N_e = 50$ ensemble members. The root-mean-squared error is given in brackets for each specified method. IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

β emerging in the IEWPF is tuned manually and set to 0.55, and this will be discussed further in relation to some of the results. We set the localisation radius of the SparseObsETKF equal to the correlation range of the model error, which leads to four observational batches. First, we do not use any relaxation in the perfect linear model, as suggested by Raanes et al. (2019). The performance of ensemble-based solutions is compared with the KF reference solution. We use a set of metrics to evaluate different statistical aspects of the data assimilation methods.

3.3.1 | Root-mean-squared error

The ensemble mean $\bar{\mathbf{x}}^a$ is compared with the KF mean $\boldsymbol{\mu}^a$. Here, we consider the state at $t = 250$ after assimilating all available observations. The error in the mean at each position is then the vector $\mathbf{err}_{\text{mean}}^{\text{KF}} = (\boldsymbol{\mu}^a - \bar{\mathbf{x}}^a)$. As a scalar metric to compute the behaviour over all grid cells, we use the root-mean-squared error (RMSE):

$$\text{RMSE} = \|\mathbf{err}_{\text{mean}}^{\text{KF}}\|_2. \quad (21)$$

Figure 5 shows $\mathbf{err}_{\text{mean}}^{\text{KF}}$ at each grid cell for a single data assimilation experiment with $N_e = 50$ ensemble members for each method, with the RMSEs shown in brackets. All three ensemble-based data assimilation methods lead to means that closely resemble the KF reference solution. The mean error is in general low and smoothly distributed for both IEWPF and SparseObsETKF, whereas the errors of ETKF are somewhat bigger. Based on RMSE, IEWPF performs slightly better than SparseObsETKF, whereas the RMSE for ETKF is about twice that of IEWPF.

To deduce reliable conclusions beyond one dataset and single ensembles, we repeat the data assimilation experiment multiple times for several independently generated true states. In Table 2, we report averaged results for 20 replicate synthetic truths and five ensemble-based data

assimilation experiments each. For this relatively small ensemble size of $N_e = 50$, the localisation in the SparseObsETKF halves the RMSE compared with the standard ETKF, and the RMSE of the IEWPF lies in the middle of the ETKF with and without localisation.

3.3.2 | Frobenius covariance difference

We contrast the empirical covariance estimates $\hat{\boldsymbol{\Sigma}}^a$ in Equation (11) with the KF reference $\boldsymbol{\Sigma}^a$. We compute the Frobenius covariance difference (FCD) to compare these covariance matrices:

$$\text{FCD} = \|\boldsymbol{\Sigma}^a - \hat{\boldsymbol{\Sigma}}^a\|_F,$$

where $\|\cdot\|_F$ denotes the Frobenius norm (elementwise sum). Averaged results for the FCD over replicate experiments are presented in Table 2 using $N_e = 50$. Here, we see that the FCD for IEWPF and SparseObsETKF are very similar for all cases. The covariance approximations show smaller errors for the ETKF solution than for the other methods. At a single step, the ETKF approximation to the covariance is unbiased, whereas the other methods have no such guarantee. Still, it is not obvious that the ETKF performs better after many data assimilation steps. Also, when we test the entries close and far from the diagonal of the covariance matrix, we cannot see any other behaviour in the results.

3.3.3 | Integrated quadratic distance

We next study a metric for the marginal distribution mismatch of discrete ensemble-based distribution approximations to the Gaussian KF reference solution. The reference cumulative distribution function (CDF) of the KF is denoted F^a . The empirical CDF (ECDF) of the ensemble-based solutions are denoted \hat{F}^a .

TABLE 2 Metrics for marginal distribution averaged over 20 synthetic truths and five ensemble realisations, each with $N_e = 50$. Standard deviations are given in parentheses.

Method	RMSE	FCD	$10^{-2} \times d_{IQ}$ at s_1	$10^{-2} \times d_{IQ}$ at s_2
Monte Carlo	8.27 (2.88)	47.0 (8.35)	12.8	13.5
IEWPF	1.67 (0.43)	2.77 (0.14)	2.51	2.58
ETKF	2.14 (0.40)	2.14 (0.04)	2.57	2.86
SparseObsETKF	1.15 (0.24)	2.79 (0.15)	1.29	1.68

Abbreviations: ETKF, ensemble transform Kalman filter; FCD, Frobenius covariance difference; IEWPF, implicit equal-weight particle filter; RMSE, root-mean-squared error; SparseObsETKF, ETKF with covariance localisation.




FIGURE 6 The cumulative distribution function (CDF) of the Kalman filter is compared with the empirical CDFs (ECDFs) of IEWPF [0.0242, 0.0238], ETKF [0.0164, 0.0254], and SparseObsETKF [0.0093, 0.0117], with $N_e = 50$, for two distinct positions: (a) observation site s_1 and (b) far away from observation sites s_2 . The integrated quadratic difference d_{IQ} for s_1 and s_2 are shown within the brackets. IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

In the analysis, two specific locations shown in the far right panel of Figure 3 are studied based on their different characteristics: $s_1 = (0, 0)$ is an observation site and $s_2 = (2.5, 1.5)$ is as far away from observation data as possible.

In Figure 6, the CDF of the KF is depicted in comparison with the ECDFs of IEWPF, SparseObsETKF, and ETKF at the two different locations for a small ensemble size. First, since the scales of the x -axis in both displays are the same, it becomes obvious that the standard deviation at an observation site is much smaller than at an unobserved location. Next, we see that the different filtering methods differ in quality when compared with the analytic solution. For the observation site s_1 there is no clear qualitative difference, but at s_2 one may already identify a slight divergence in ETKF's ECDF, whereas SparseObsETKF and IEWPF still approximate the reference CDF quite well.

The tuning parameter β steers the spread in the analysis ensemble of the IEWPF, and among the evaluation metrics presented the ECDF reveals the scale the best. We used it to optimise the filtering distribution manually and found 0.55 as best choice. For smaller values, the ECDF gets too sharp, and for higher values the spread gets too large. Similarly, the variance in the SparseObsETKF ensemble usually increases as the relaxation parameter ϕ decreases. When using $\phi < 1$ we observed that the spread in the ensemble becomes too big compared with the CDF, and the best match is achieved for $\phi = 1$.

Thorarinsdottir et al. (2013) suggest a proper divergence function to compare marginal CDFs, condensing

the error into a scalar number. The integrated quadratic difference is defined by

$$d_{IQ} = \int (F^a - \hat{F}^a)^2 dx,$$

where the quadratic error is integrated over the sample space of the variable. Errors captured in d_{IQ} can originate from either a lack of Gaussianity, a wrong scaling, a bias, a combination thereof.

Table 2 shows averaged results for d_{IQ} at s_1 and s_2 for the three ensemble methods. IEWPF and ETKF produce similar results, whereas SparseObsETKF clearly gives the best results. The reason is twofold: First, the IEWPF and ETKF update the entire field at each data assimilation time, even at locations like s_2 that are far away. With the limited ensemble size, this is likely to induce some undesired bias and variability far from data. Second, the SparseObsETKF is rather accurate near the observations sites, like the other filters; and because the advection and diffusion are known, the local updating propagates reasonably over time to the far location s_2 .

3.3.4 | Probability coverage level

Based on the mean and variance of the ensemble-based solutions, we can check how often a prediction interval covers the true realisation. For the KF reference, we will have near-nominal coverage because the truth is simulated




FIGURE 7 Estimated $CP_{1.64}^1$ using 500 replication experiments for the Kalman filter (KF) and the ensemble-based methods with $N_e = 50$. The brackets show the averaged estimated coverage probabilities. The colour scale is centred at the target probability of 90%. IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

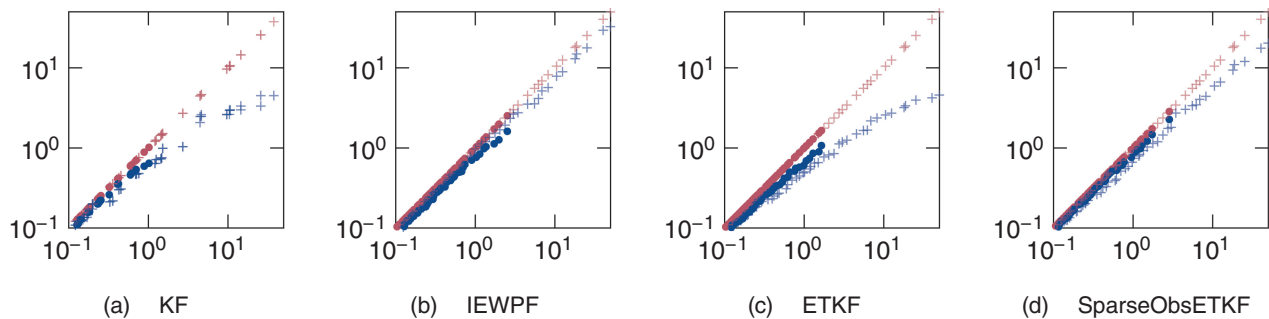


FIGURE 8 Eigenvalue update in the very first (crosses) and very last (dots) data assimilation step. Prior-versus-prior on the diagonal in red, prior-versus-posterior in blue. KF: Kalman filter; IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

from the same model. Coverage probabilities (CPs) of the analysis after the first observation time are

$$CP_{1.64}^1 := P(\mathbf{x}_{\text{true}}^1 \in [\boldsymbol{\mu}^{1,a} \pm 1.64\boldsymbol{\sigma}^{1,a}]) \approx 0.90.$$

This means that the probability that the truth is covered in the interval of 1.64 standard deviations from the mean is 90%. For all methods, we use replicated synthetic truths and estimate the CPs.

Figure 7 shows the estimates of the CPs averaged over 500 runs. The KF results cover the nominal 90% very well, with a variability as expected from 500 replicates. Even though the KF is analytic, the truths are generated stochastically and we see the Monte Carlo error here. The IEWPF and the SparseObsETKF also give very good estimates even though the ensemble size of 50 members is low, and there are no visible structures around observation sites. In contrast, the ETKF without localisation suffers from strong undercoverage in this experiment. The CPs are around 0.9 near observation sites but fall to lower levels (which are outside the truncation interval of the plot) away from these.

3.3.5 | Eigenvalues analysis

An eigendecomposition of the covariance matrix yields eigenvalues representing the variability among orthogonal

axes of linear combinations of state variables. The largest (first) eigenvalue is the variance in the direction of the first eigenvector, and for a Gaussian distribution this represents the largest half-axis in the ellipsoid defined via the quadratic form. Further, for a Gaussian distribution, the entropy (disorder) is defined via the log-determinant of the covariance matrix, which is the sum of the log-eigenvalues. By studying the eigenvalues of the covariance matrices of the different data assimilation methods, we hence gain insight into the variability reduction and the disorder of the distributions. Since it does not make sense to average eigenvalues over multiple realisations, we leave the spectral analysis to a qualitative view with cross-plots of eigenvalues for the different methods.

Figure 8 shows cross-plots of the covariance matrix eigenvalues of forecast (first axis) and assimilated (second axis). In all displays, the crosses represent the first data assimilation step and the dots are at the last data assimilation step. For the KF, eigenvalues are computed directly from the covariance matrix, and for the ETKF variants these are the eigenvalues of the estimated covariance matrix in Equation (11). In the case of the IEWPF, the prior ensemble is without model error, and hence we perturb the prior before plotting.

In Figure 8, the forecast-versus-forecast points lie on a straight line, but we still notice that all the corresponding dots are closer to the origin than some of the

crosses. Hence, the eigenvalues are clearly smaller at the last step, indicating that the data assimilation over time gives smaller entropy. Going from forecast (red) to assimilated (blue), the eigenvalues of the covariance matrix are reduced. Taking the KF as a benchmark, this reduction is particularly large for the biggest eigenvalues, indicating the updating is not only local but also shrinks the variability of dominating linear combinations and the entropy of the distribution. At the first step (crosses), the ETKF updating appears very similar to that of the KF. At the last step, the eigenvalues of the ETKF are larger and not reduced quite like for the KF. This indicates that, even though the ETKF undercovers the distribution (Figure 7), there is not quite sufficient reduction in the largest eigenvalues. Both IEWPF and SparseObsETKF get smaller reductions in eigenvalues than the KF. At the first data assimilation step the reduction is larger for SparseObsETKF than for the IEWPF, whereas the IEWPF has more reduction at the last data assimilation step.

3.3.6 | Spatial connectivity

Although the previous metrics have considered the marginal solution at one time step only, the correlation between different time steps and between different spatial positions gives further insight into the statistical quality of the filtering methods.

Given data up to time t^{n-1} , the correlation between the concentration at \mathbf{s}_k at $t^{n-1} = 225$ and \mathbf{s}_l at t^n can then be calculated from the KF results via

$$\begin{aligned} \text{Corr}(\mathbf{x}_k^{n-1,a}, \mathbf{x}_l^{n,f}) &= \frac{\text{Cov}(\mathbf{x}_k^{n-1,a}, \mathbf{x}_l^{n,f})}{\sigma_k^{a,n-1} \sigma_l^{f,n}} \\ &= \frac{\mathbf{M}\Sigma_{k,l}^{a,n-1}}{\sqrt{\Sigma_{k,k}^{a,n-1}} \sqrt{\Sigma_{l,l}^{f,n}}}. \end{aligned} \quad (22)$$

Similarly, we can estimate these correlations from the ensemble-based methods by

$$\begin{aligned} \widehat{\text{Corr}}(\mathbf{x}_k^{n-1,a}, \mathbf{x}_l^{n,f}) &= \frac{1}{N_e - 1} \frac{1}{\hat{\sigma}_k^{n-1,a} \hat{\sigma}_l^{n,f}} \\ &\times \sum_{e=1}^{N_e} (\mathbf{x}_{e,k}^{n-1,a} - \bar{\mathbf{x}}_k^{n-1,a})(\mathbf{x}_{e,l}^{n,f} - \bar{\mathbf{x}}_l^{n,f}). \end{aligned}$$

In Figure 9, we show the correlation fields of the state between a reference grid cell at $t^9 = 225$ and all other grid cells at $t^{10} = 250$. As reference locations \mathbf{s}_k , we consider \mathbf{s}_2 positioned far away from any observations (in the top row) and \mathbf{s}_3 at an observation site in the middle of the domain (in the bottom row). First, from the KF solution (left), we

recognise the advection field in the model that transports information towards eastnortheast from both locations, as well as the diffusion causing the correlation to have longer range than the model error correlation radius. Second, the maximal correlation to the reference point is higher when \mathbf{s}_k is not an observation site. In spatial statistics, conditioning on data breaks up some of the prior correlations. Since most of the update from the data assimilation occurs near the observation locations, the conditional correlation tends to be smaller in the proximity of data. At locations that are far from observations, more of the prior correlation remains.

In the three rightmost columns of Figure 9 we see the correlations estimated with the three ensemble-based methods using an ensemble with $N_e = 250$ members. We see that all methods capture a similar correlation structure with respect to the advection and diffusion, and the relative balance between prior model and information from the observation. With $N_e = 250$, the results are less smooth than the KF solution. In the upper scenario, the area of high analytical correlations becomes less apparent among background noise, whereas in the lower scenario the respective regions are easier to identify in all methods.

The spatial error in the approximation of correlation between two consecutive model steps for the reference location \mathbf{s}_k is evaluated collectively across all grid cells as

$$\begin{aligned} \text{CE}(\mathbf{s}_k)^2 &= \sum_{l=1}^{N_s} |\text{Corr}(\mathbf{x}_k^{n-1,a}, \mathbf{x}_l^{n,f}) \\ &\quad - \widehat{\text{Corr}}(\mathbf{x}_k^{n-1,a}, \mathbf{x}_l^{n,f})|^2. \end{aligned} \quad (23)$$

The CE values for the specific data assimilation run shown in Figure 9 are given in brackets in the figure. These results for a single run with $N_e = 250$ already reveal that the contribution to CE can come from multiple sources, such as over- or underestimation of the actual correlations and from spurious correlations. The final CE does not qualitatively expose which of the error sources are present to what extent, but it quantifies how well the analytical structure is approximated. The chosen Corr and $\widehat{\text{Corr}}$ suggest that spurious correlations are a present error source in all methods, but IEWPF and SparseObsETKF tend to overestimate the high correlations. Meanwhile, ETKF underestimates the correlations, thus leading to a smaller CE.

3.4 | Sensitivity to localisation

The localisation properties in the IEWPF, are steered by the structure of the model error covariance \mathbf{Q} and in the SparseObsETKF by the definition of the localisation radius, which we again defined according to the




FIGURE 9 Correlations $\mathbb{C}\text{orr}(\mathbf{x}_k^9, \mathbf{x}_l^{10})$ and $\widehat{\mathbb{C}\text{orr}}(\mathbf{x}_k^9, \mathbf{x}_l^{10})$ between fixed locations s_k and all other grid points in the domain s_l , for $k = 3$ in the upper row and $k = 2$ in the lower row. The ensemble-based methods use $N_e = 250$ ensemble members, and the respective CE estimates are given in brackets. KF: Kalman filter; IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

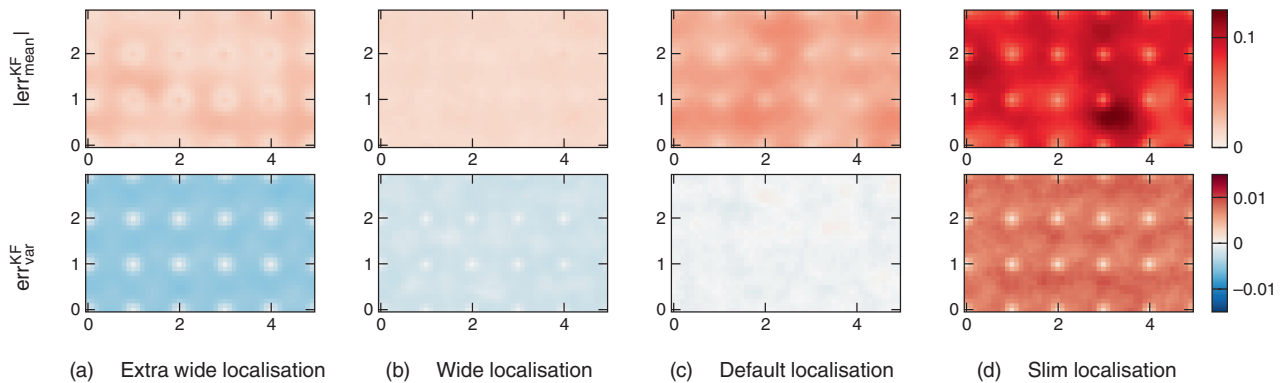


FIGURE 10 Different localisation parametrizations for the IEWPF (implicit equal-weight particle filter). Averaged results over 20 truth realisations and five ensemble initialisations each.

correlation range of \mathbf{Q} ; see Section 2.3. In the rest of the section those choices are held fixed, but here we showcase their influence. To do so, we consider four different ψ for the model error covariance matrix, while keeping everything else unaltered; that is, the standard derivation and the IEWPF tuning parameter β . Analogously, we vary the localisation radius in the SparseObsETKF. The parameter $\psi = 3.0$ corresponds to no localisation and $\psi = 5.0$ and 11.0 represent roughly a doubling and halving of the localisation radius respectively. Importantly, it should be noted that the GC and the Matérn-type covariance kernels decay with different rates such that the results are not one-to-one comparable, but we can still record trends within each method.

Figures 10 and 11 allow one to get an impression of the RMSE together with spatial effects. Since the localisation is in-built into the IEWPF, one has to be careful with the

interpretation of Figure 10a. Nevertheless, there are clear unintended artefacts in mean and variance errors. Even though the mean for $\psi = 5.0$ in Figure 10b is very well calibrated in the entire field, the variance is slightly underestimated and the error shows structures around observation locations. Note that an increase of β may counteract the underestimation. For the standard IEWPF, this is the opposite way around; here, the errors in the variances are minuscule, but the mean is not equally well calibrated. The IEWPF with reduced localisation radius performs poorly when further away from observation locations. The ETKF without localisation underestimates the variance of the KF. The SparseObsETKF with a localisation radius that spans the full y -extent of the domain (Figure 11b) performs slightly better for mean and variance estimation than the parametrisation as we have chosen in the rest of the section. We recall that the localisation usually used




FIGURE 11 Different localisation parametrizations for the SparseObsETKF (ensemble transform Kalman filter with covariance localisation). Averaged results over 20 truth realisations and five ensemble initialisations each.

yields a computational advantage owing to reduced analysis dimensions. Again, the smallest localisation parametrization fails to assimilate the entire domain.

We note that the localisation parametrization as used in the rest of the section work reasonably well. Nevertheless, we notice that there is potential to fine-tune both methods further, but in the interest of a limited number of tuning parameters we continue with the native parameters.

3.5 | Discussion of evaluation metrics

The set of comparative metrics from the previous subsection has given us a collection of metrics that quantify some statistical qualities of the ensemble-based data assimilation methods in reference to the analytical KF solution.

Table 2 shows the statistically averaged results for these performance scores at $t = 250$. These results are obtained across five data assimilation runs for 20 different synthetic truths and are therefore more reliable than the single realisations demonstrated in Figures 5, 6, and 9. We have used $N_e = 50$ ensemble members here for each run. In addition to comparing the data assimilation methods against each other, we have also included the results using pure Monte Carlo simulations without observations (top row). These serve to demonstrate the worst-case scenario for each metric, and we see how all three data assimilation methods clearly outperform this, as expected. In the experiments, we have observed that the IEWPF takes several assimilation steps until it is sufficiently calibrated, which is respected by the choice of t here such that the comparison stays fair; see Section 4 for details.

Based on the results in Table 2, there is no method that clearly dominates on all individual criteria. For RMSE, it seems that SparseObsETKF is much better than ETKF, but this is not as clear when considering FCD, where ETKF scores best. Maybe more surprising, SparseObsETKF is

significantly better than ETKF when measuring the error in the ECDF at \mathbf{s}_2 far from the observation, but not at \mathbf{s}_1 at an observation site. A plausible explanation is that, when updating the state far from an observation, all covariances are relatively weak, which means that spurious correlations more easily dominate data assimilation. With localisation, we ensure that only the most relevant small correlations are considered, thus improving the result. This effect will then be less at an observation site, as the most important correlations are stronger. IEWPF gets all metrics between ETKF and SparseObsETKF. We observe that a worse FCD has no influence on the d_{IQ} at the positions considered.

3.5.1 | Effects of ensemble size

In Figure 12, we study how different ensemble sizes influence RMSE, d_{IQ} , and correlations for the three ensemble-based methods. We use ensemble sizes $N_e \in \{25, 50, 100, 250, 1000, 5000\}$, and the results are averaged across multiple experiments for each of these sizes. In general, we expect that increased ensemble size leads to more accurate statistical estimates, and thereby better results. This is clearly the case for ETKF for all metrics; and SparseObsETKF shows the same trend, but not as strongly. We see that SparseObsETKF outperforms ETKF with respect to RMSE and d_{IQ} for small ensemble sizes; but ETKF is better with large ensembles, as the performance of SparseObsETKF stagnates for $N_e > 250$. SparseObsETKF improves less than ETKF with larger ensembles because it ignores correlations, and this gives bias in the analysis. The IEWPF yields results between ETKF and SparseObsETKF for small sample sizes, but there is slower convergence as the ensemble size increases. Unlike ETKF, which converges to the true Gaussian distribution in this case, there is no such guarantee for the IEWPF.

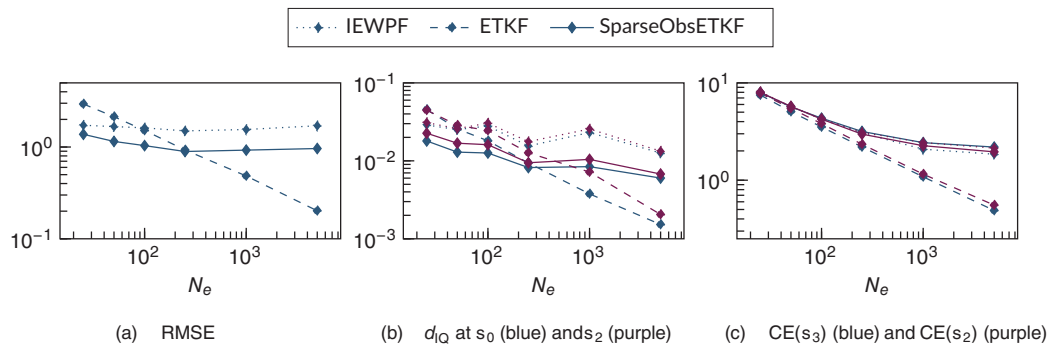


FIGURE 12 Evolution of the comparison measures as a function of the ensemble size N_e : (a) root-mean-squared error (RMSE); (b) d_{IQ} at s_0 (blue) and s_2 (purple); (c) $CE(s_2)$ (blue) and $CE(s_3)$ (purple). IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

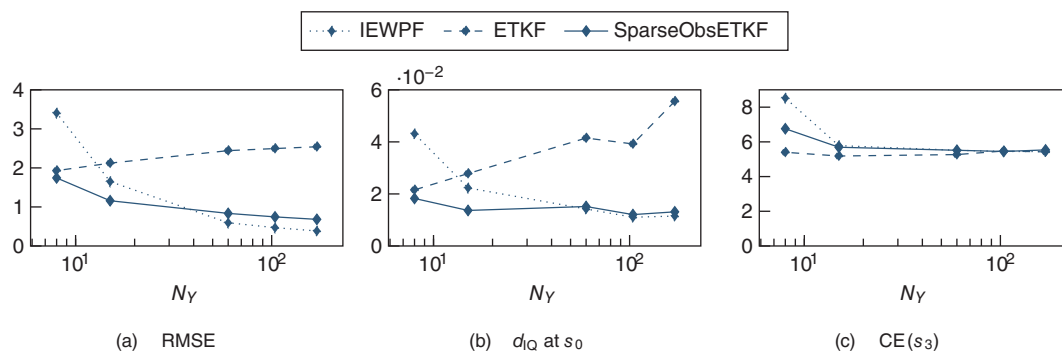


FIGURE 13 Evolution of the comparison measures as we change the sparsity/density of observation data through the number of observations N_Y . The ensemble size is fixed at $N_e = 50$. (a) Root-mean-squared error (RMSE); (b) d_{IQ} at s_0 ; (c) $CE(s_3)$. IEWPF: implicit equal-weight particle filter; ETKF: ensemble transform Kalman filter; SparseObsETKF: ETKF with covariance localisation.

Since the second-stage perturbation step of the IEWPF is designed to reduce a systematic bias and help performance, fine-tuning the choice of the β scaling parameter could improve convergence for some properties, but maybe not similarly so for all the desired scores. The correlations mismatch compared with the KF in Figure 9 is slightly different depending on the fixed reference point, but they converge with increasing sample size, especially so for the ETKF, which has curves going faster to zero. For both SparseObsETKF and IEWPF there seem to be a remaining mismatch in this CE score even for thousands of ensemble members.

3.5.2 | Effects of sparsity of observational data

In a regime dominated by the sparsity of observations, we also want to stress test all methods with respect to the amount of observational data. For this purpose, we repeat the case study using $N_Y \in \{8, 15, 60, 104, 170\}$ regularly

placed observation sites. These numbers are chosen such that the observation locations have distance of 15, 10, 5, 4, and 3 grid cells apart from each other, respectively. We use $N_e = 50$ ensemble members. Of course, the localisation scheme for the SparseObsETKF is not designed for dense data and will get computationally very inefficient due to a high number of batches that are processed serially. The localisation radius is not modified.

Figure 13 shows the same averaged metrics as before, with respect to a growing number of observations. Note that it no longer makes sense to distinguish between locations near and far from observations, since the observation sites get denser over the domain. The increase in observation data leads to a sharpening in the reference distribution calculated from the KF. For the ETKF, we observe that RMSE does not improve and its ECDF approximation gets worse, compared with the KF. This is because of the underestimation in variance and a slight bias that strongly penalises the d_{IQ} . Both SparseObsETKF and IEWPF improve their quality for increasing observation data size. This is surprising and noteworthy for IEWPF, as

PFs tend to collapse for high-dimensional observations. In contrast to RMSE and d_{IQ} , the CE does not depend on the observation sparsity and is practically constant on the level that we saw in Figure 12c.

3.5.3 | Summary

In this case study, we verify the ensemble-based methods from Section 2 for a linear Gaussian data assimilation problem with the analytical KF reference. The SparseObsETKF and IEWPF include localisation, either explicitly or implicitly, that is connected to the correlation radius of the model error. Beyond the verification, we can in particular record that SparseObsETKF outperforms ETKF and IEWPF for smaller ensemble sizes (about $N_e \leq 250$). Though the ETKF reduces the spectral radius in every data assimilation step more than the other methods, it requires a large ensemble size to obtain a reasonable approximation of the full covariance matrix. The localised version that ignores large-distance correlations is performing well for small ensemble sizes, but it does not improve much more for larger sizes. Similar tendencies are seen with the IEWPF. The approximation of the correlations between different time steps depend mostly on the ensemble size—the model error plays a major role in the evaluation, and this criterion requires a higher ensemble size for a sufficient representation. For reasonable ensemble sizes, say 100, both IEWPF and SparseObsETKF operate well for any density of observation data. For most criteria we tested in this example with sparse point data, the SparseObsETKF tends to give slightly better performance than the IEWPF. Based on this extensive statistical evaluation, we hence recommend considering the SparseObsETKF for similar kinds of applications with sparse data and limited ensemble sizes.

4 | COMPARISON FOR DRIFT TRAJECTORY FORECASTING IN A SIMPLIFIED OCEAN MODEL

We now increase both dimensionality and complexity as we turn to a nonlinear simplified ocean model. This gives insight into the behaviour of the ensemble-based data assimilation methods on a challenging case with applied relevance. The practical purpose of this configuration is to use ensembles of computationally efficient simplified ocean models instead of, or complementary to, single realisations of complex operational ocean models in time-critical situations. The simplified models allow for larger ensembles, and hence facilitate uncertainty quantification. Such an approach can be useful

in search-and-rescue operations, where drifters released by the vessel or relevant anchored buoys (also called moorings) can give sparse in-situ observations during the operation. These point observations can then be assimilated into the ensemble-based representation to improve the drift trajectory forecasts that specify a search area.

Owing to the nonlinearity of such a model, there is no analytical reference solution for the ensemble distributions available. We can nevertheless compare SparseObsETKF and IEWPF by studying their predictive properties with the ground truth in a simulation study. We base our numerical experiments on those presented in Holm et al. (2020), where the IEWPF was successfully tailored for efficient graphics-processing-unit-accelerated assimilation of point observations of a chaotic shallow-water model. We expand on the numerical result from that work by evaluating more skill scores, and by providing an in-depth comparison between IEWPF and SparseObsETKF for state estimation and drift trajectory forecasts. In particular, this will also serve as a more thorough evaluation of the applicability of IEWPF in this context.

4.1 | Simplified ocean model

The simplified ocean model is represented by the rotational shallow-water equations given by

$$\begin{aligned} \eta_t + (hu)_x + (hv)_y &= 0, \\ (hu)_t + \left(hu^2 + \frac{1}{2}gh^2\right)_x + (huv)_y &= fhv, \\ (hu)_t + (huv)_x + \left(hv^2 + \frac{1}{2}gh^2\right)_y &= -fhu. \end{aligned} \quad (24)$$

This is a nonlinear two-dimensional hyperbolic conservation law, which models conservation of mass through the deviation η from equilibrium sea level, and conservation of momentum through hu and hv , which are vertically integrated ocean currents in x - and y -direction respectively. By denoting the equilibrium depth of the ocean by H , we get the total depth as $h = H + \eta$. Furthermore, g is the acceleration due to gravity, and f is the Coriolis parameter that accounts for the rotating frame of reference.

We solve Equation (24) using the high-resolution central-upwind finite-volume scheme proposed by Chertock et al. (2018). In our notation from the model equation, Equation (1), the state vector \mathbf{x}^n consists of the cell averaged values $(\eta_i^n, (hu)_i^n, (hv)_i^n)$ at time t^n for all cells i in the discretised domain. The \mathcal{M}^n operator then applies the finite-volume scheme to evolve the state from \mathbf{x}^{n-1} to \mathbf{x}^n . Note that the time step used by the numerical method can be chosen independently from the model time step,




FIGURE 14 State of the synthetic truth after (a) 3, (b) 6, and (c) 10 simulation days. The arrows indicate the direction and strength of the ocean currents derived from u and v respectively. The background visualises the magnitude of velocity in metres per second. The turquoise dots mark the fixed-point buoy positions.

meaning that \mathcal{M} might consist of multiple iterations of the numerical scheme.

We apply a small-scale Gaussian model error $\mathbf{v} \sim \mathcal{N}(0, \mathbf{Q})$. It is constructed from a coarse-scale perturbation of η , which is smoothed by a second-order autoregressive function and projected onto the numerical grid. The model error for hu and hv is then inferred according to geostrophic balance to ensure physical feasibility. Further details about this model are available in Brodtkorb and Holm (2021) and Holm et al. (2020).

4.2 | Experiment design

In the following, we use the same experimental design as in Holm et al. (2020). We consider a rectangular domain covering $1,100 \times 666 \text{ km}^2$ that is discretised as a uniform Cartesian 500×300 grid. The domain has periodic boundary conditions and constant equilibrium depth $H = 230.0 \text{ m}$. The initial conditions, for the ground truth as well as for all ensemble members, consist of a westward jet in the north part of the domain and an eastward jet in the south, with $hv = 0$. Both jets are balanced according to geostrophy by η so that the initial conditions are in steady state. This steady state is unstable, however, and slight perturbations, such as those from the model error \mathbf{v} , cause chaotic behaviour.

As an example of the turbulent behaviour, Figure 14 shows the water velocities for one realisation that is labelled as the synthetic truth \mathbf{x}_{true} . Here, the model error is added every 60 s, and the model error correlation radius is approximately 40 km. From Figure 14, we see that the jets in \mathbf{x}_{true} are still quite regular after 3 days, but grow more irregular after 6 and 10 days. It should be noted that the mean state from a pure Monte Carlo experiment without data assimilation will result in $hv \approx 0$ even after 10 days. This indicates that it is challenging to correctly capture where and how the turbulent behaviour will develop.

From \mathbf{x}_{true} , we extract direct observations of only (hu, hv) at 60 locations in the domain every 5 min between day 3 and day 10, with observational noise sampled from $N(0, \mathbf{I})$. The turquoise dots in Figure 14 show the observation sites. In total, the experiment is characterised by 450,000 state variables versus only 120 very sparse noisy observations. After day 10, three drifters are released in the domain, and advected according to the simulated currents at every time step of the numerical scheme using a simple Euler scheme. Part of the challenge for the data assimilation methods is to forecast the trajectories of these drifters.

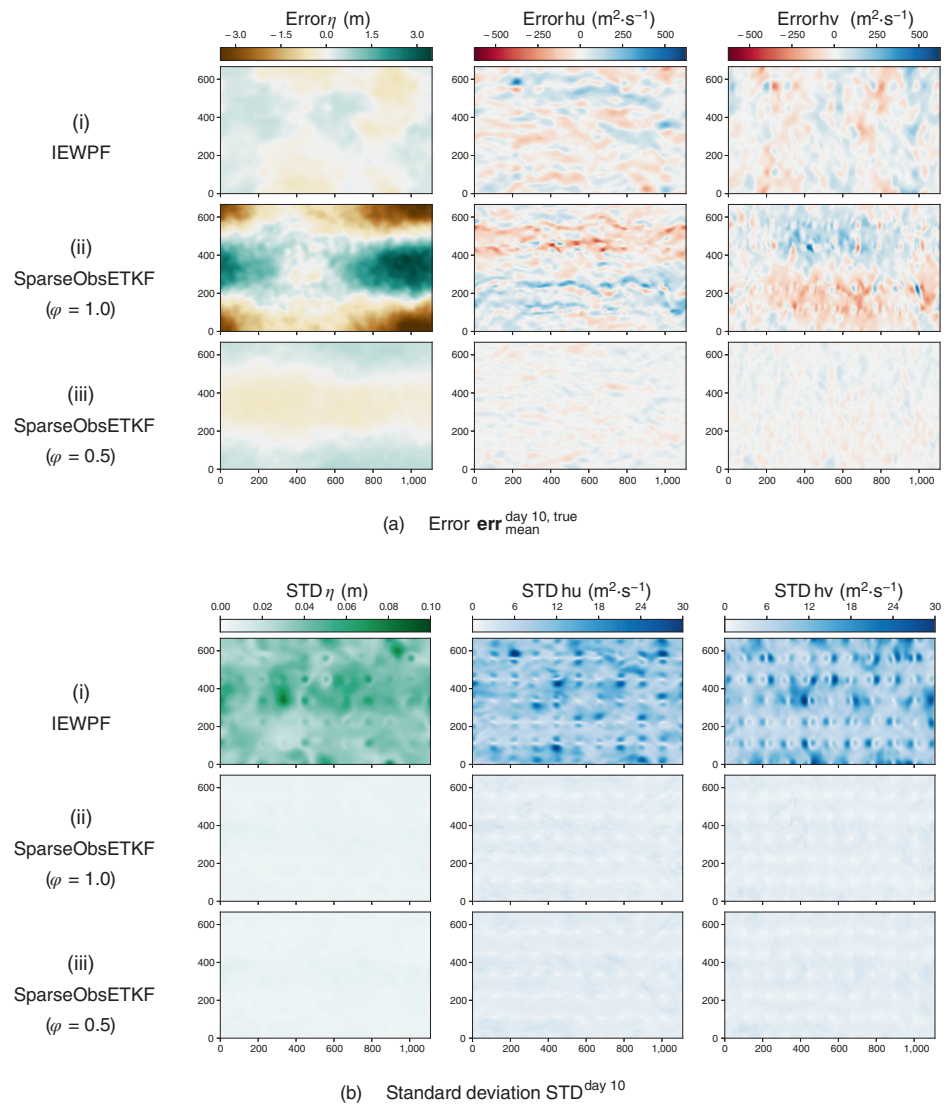
The data assimilation starts at simulation day 3 after each ensemble member has been spun up from the steady state through independently sampled model errors. Even though all ensemble members are visually very similar at this stage, they have started to develop internal instabilities that will grow over time unless the observations are successfully assimilated.

This case is much more challenging than the advection diffusion model in Section 3: The shallow-water model is nonlinear, there are unobserved variables, and it has significantly higher dimensionality. Critically, the nonlinear dynamics of the shallow-water model are challenging to capture. In the advection diffusion model, the state converges towards an equilibrium due to the diffusion, whereas our shallow-water equation case gets chaotic dynamics that make the ensemble naturally diverge in time.

4.3 | Numerical results

Classical EnKF approaches like the ETKF lead to useless results for this difficult case, and only results of the IEWPF and SparseObsETKF are shown in the comparison. We use $N_e = 100$ as a compromise between computational effort and statistical quality. Based on our experiments on this high-dimensional nonlinear model, the IEWPF

FIGURE 15 Properties of the state estimation for the different physical variables in this simplified ocean model (sea-surface elevation η and momenta hu and hv) after day 10 measured in (a) the error of the mean versus the truth and (b) the standard deviation presented for the IEWPF (implicit equal-weight particle filter) and the SparseObsETKF (ensemble transform Kalman filter with covariance localisation) without and with inflation.



performance is not very sensitive to the explicit choice of β and we use the maximal allowed value. The localisation radius for the SparseObsETKF is chosen slightly larger than the model error correlation radius. In contrast to Section 3, we now also investigate the influence of relaxation in the SparseObsETKF. We present results for the SparseObsETKF without relaxation ($\phi = 1.0$) and for the SparseObsETKF whose weights in the localisation are scaled by $\phi = 0.5$. We compare data assimilation methods with the simulated truth using a number of skill scores that refer to this ground truth.

4.3.1 | State estimation

We first look at deviations of the ensemble mean from the truth by

$$\mathbf{err}_{\text{mean}}^{\text{day 10, true}} = \bar{\mathbf{x}}^{\text{day 10}} - \mathbf{x}_{\text{true}}^{\text{day 10}}, \quad (25)$$

which represents the error in the correct physical unit. We also investigate the standard deviation in the ensemble

$$\text{STD}^{\text{day 10}} = \frac{1}{N_e - 1} \sqrt{\sum_{e=1}^{N_e} (\mathbf{x}_e^{\text{day 10}} - \bar{\mathbf{x}}^{\text{day 10}})^2}, \quad (26)$$

which gives insight about the ensemble spread around its mean.

Figure 15a shows the mismatch between the truth and the ensemble means of the conserved variables after assimilating the final observations on day 10. Significant differences become clear in the error of the sea-surface elevation η (left): Whereas the IEWPF has some moderate, relatively smooth error over the entire domain, the mean of the SparseObsETKF is far off in half of the domain. In particular, the rims in the error field are very sharp, also recognisable in the error spots of the currents at the edges of the jets. This indicates that the ensemble

produces very fast changing ocean fields with the tendency to non-physical members. However, relaxation with $\phi = 0.5$ (bottom) impressively fixes some of those issues and the error fields become much smoother and closely calibrated, even though there is still a recognisable, but weak, inherited pattern in the error for elevation.

There are structured artefacts identifiable around the observation sites for IEWPF in Figure 15b. Even though Figure 15a shows that the ensemble mean gives a very precise description of the ground truth, the ensemble variance is large. As discussed in Holm et al. (2020), the IEWPF updates the momentum locally by inducing a corrective current formed by the structure in the model error covariance matrix \mathbf{Q} . In this case, \mathbf{Q} induces geostrophically balanced dipole structures, which means that while improving the state at the observation site we risk deteriorating the solution in its vicinity. This illustrates a weakness of IEWPF, showing that its quality is only as good as the structure of \mathbf{Q} .

In the standard deviations for the SparseObsETKF in Figure 15b (middle row), there is an expected pattern of low values around observation sites. Since the localisation only corrects around the buoys and leaves the forecast otherwise unchanged, the variance in one data assimilation step is mainly reduced in local areas. With the dynamical model over time, the variance reduction is disseminated over the entire domain. Furthermore, the standard deviation in the SparseObsETKF is on a very low level. Having areas of low error together with the sudden changes towards big errors raises suspicions of overfitting. Also, the relaxed SparseObsETKF (bottom row) achieves a similarly low standard deviation after day 10.

4.3.2 | Drift trajectory forecasting

To further compare the practical applicability of IEWPF and SparseObsETKF, we look at forecasts of drift trajectories starting at day 10.

Figure 16 demonstrates the forecasted trajectories of drifters that are realised after ten simulation days in the simplified ocean model. The three drop locations are selected to capture different characteristics in the currents: Drifter 1 (Figure 16a) starts in the middle of a rather weak and big east stream, drifter 2 (Figure 16b) starts in a rather strong west stream, and drifter 3 (Figure 16c) starts in a turbulent area in between the dominating streams. For the first 2 days of forecast, we show the true trajectory along with the trajectories for all ensemble members and the ensemble mean, whereas for the third day (right) we show the estimated kernel density (Scott 1992) of the final drifter locations along with the true trajectory.

For drifter 1, all trajectories have an eastwards drift, but the IEWPF members fan out from the beginning whereas the SparseObsETKF trajectories stay close together. Without relaxation, the truth becomes an outlier in the SparseObsETKF forecast. With relaxation, the truth stays within the forecast. The trajectories from the IEWPF catch the truth in a high-probability area, but their spread covers almost the entire extent of the domain in the y -direction.

Even though drifter 2 starts within a jet, it drifts only shortly westwards before it takes a sudden turn towards the north. Here, we can again see the turbulent behaviour of this nonlinear model. The trajectories of the IEWPF again spread out widely and, therefore, do not reveal any consistent dynamical pattern in the underlying currents. SparseObsETKF not only misses the true trajectory completely, it also shows some wriggling trajectories, which indicates that there are unbalanced gravitational waves in the ensemble. Relaxation increases the spread in the trajectories up to day 2, and most of the ensemble members capture the sudden turn in the truth, even though this happens a day after assimilating the final observations.

Drifter 3, which is released in an unstable area, follows what is almost a rotation-like pattern. Here, the IEWPF is unable to estimate a clear direction even for the first 24 hr, and after day 10 the drifter distribution stretches out across almost half the simulation domain. In contrast, for SparseObsETKF with and without relaxation the ensemble gives a precise forecast for the first day, only showing a spread for the two last days. The truth is well represented by the ensemble for both experiments, but we see that the spread is remarkably reduced when using relaxation.

In general, we see that even though IEWPF is able to give a good state estimation through the mean, the spread in the underlying ocean state is too large to facilitate precise drift trajectory forecasts. Furthermore, LETKF without relaxation shows clear signs of overfitting, as the forecasts have low spread and do not match the ground truth. Introducing relaxation into the SparseObsETKF reduces this overfitting such that the true trajectories are correctly forecasted and uncertainty is better represented. Even more important, giving more weight to the forecast that comes from the physical model prevents the ensemble from unintended anomalies. The drift trajectories estimation draws attention away from the ocean states towards dynamic visual characteristics in the ensemble.

4.4 | Discussion of skill scores

Complementary to drift trajectory forecasts, we look into characteristics of both methods during the data assimilation phase between day 3 and day 10. We compare

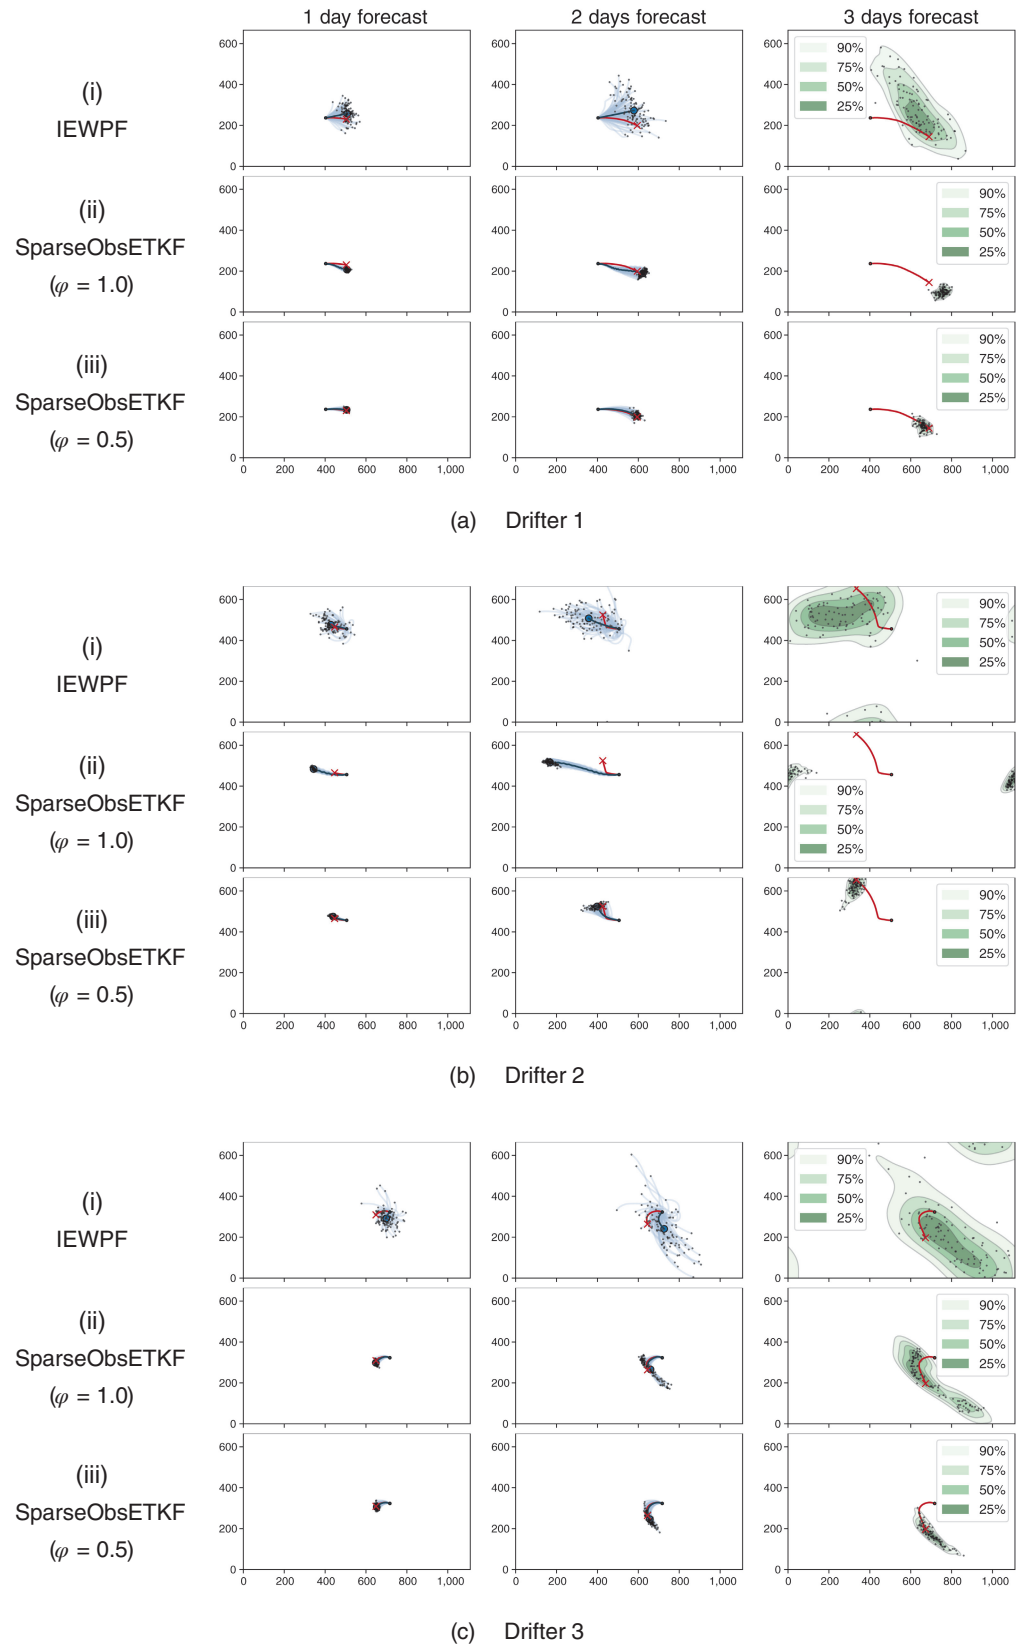


FIGURE 16 Drift trajectory forecasts for three different starting positions: (a) drifter 1; (b) drifter 2; (c) drifter 3. True trajectories represented in red, and for the first 2 days the trajectories of ensemble members are in light blue and the ensemble mean in dark blue. For the third day, the forecasted drifter positions of all ensemble members are shown with black dots and selected levels of the estimated kernel density are visualised. IEWPF: implicit equal-weight particle filter; SparseObsETKF: ensemble transform Kalman filter with covariance localisation.

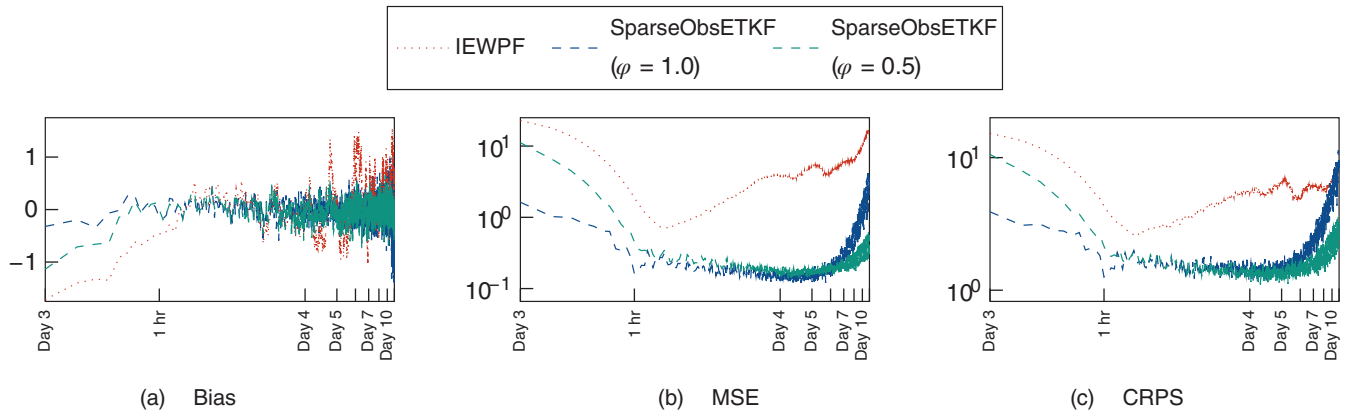


FIGURE 17 Evolution of skill scores for the IEWPF (red), SparseObsETKF without relaxation (blue), and SparseObsETKF with relaxation (turquoise) for the data assimilation phase in the experiment: (a) bias; (b) mean square error (MSE); (c) continuous ranked probability score (CRPS). IEWPF: implicit equal-weight particle filter; SparseObsETKF: ensemble transform Kalman filter with covariance localisation.

statistical properties of the ensemble against the observation data. In this setting, the key idea of skill scores is to evaluate how reliably the ensemble can forecast the next observation. An illustrative introduction with a lot of examples from atmospheric weather forecasting can be found in Wilks (2005, Chapter 7). Mathematically speaking, in this setting with data comparison, a score is

$$s(\hat{F}^{n,f}, \mathbf{y}^n) \in \mathbb{R}, \quad (27)$$

which in our case quantifies some property of the empirical distribution from an ensemble forecast $(\mathbf{H}\mathbf{x}_e^{n,f})_{e=1}^{N_e}$ against the true observation \mathbf{y}^n , meaning hu_j and hv_j for all $j = 1, \dots, N_Y$. We consider three different skill scores to judge the performance, the bias, the mean square error (MSE), and the continuous ranked probability score (CRPS):

Bias After asserting the calibration of the full analysis mean in Figure 15a, we investigate this further by evaluating the bias of the forecast as

$$s_1^n = \frac{1}{N_Y} \sum_{j=1}^{N_Y} \left[\overline{hu}_j^{n,f} - \mathbf{y}_{j,1}^n + \overline{hv}_j^{n,f} - \mathbf{y}_{j,2}^n \right], \quad (28)$$

where \overline{hu} and \overline{hv} are the ensemble means. The bias discovers systematic trends off in the estimator.

MSE We further investigate the distance of each ensemble member individually from the data by measuring the mean-squared error (MSE) as

$$s_2^n = \frac{1}{N_e} \sum_{e=1}^{N_e} \left[\frac{1}{N_Y} \sum_{j=1}^{N_Y} \left| hu_{e,j}^{n,f} - \mathbf{y}_{j,1}^n \right|^2 + \left| hv_{e,j}^{n,f} - \mathbf{y}_{j,2}^n \right|^2 \right]. \quad (29)$$

The MSE equals zero only when all ensemble members predict the observation exactly. However, this is, of course,

not desired from a probabilistic forecast representing associated uncertainty. Nevertheless, a small MSE is desired and yields an accurate fit to the data with regards to the standard deviation in the observation error.

CRPS Similar to the integrated quadratic differences, which compared distribution forecasts, we use a scoring rule that analyses the distribution of the ensemble members with the observation; see Gneiting and Raftery (2007). The continuous ranked probability score (CRPS) is here defined by

$$s_3^n = \frac{1}{N_Y} \sum_{j=1}^{N_Y} \left[\frac{1}{N_e} \sum_{e=1}^{N_e} |hu_{e,j}^{n,f} - \mathbf{y}_{j,1}^n| + |hv_{e,j}^{n,f} - \mathbf{y}_{j,2}^n| - \frac{1}{2N_e^2} \sum_{e=1}^{N_e} \sum_{k=1}^{N_e} |hu_{e,j}^{n,f} - hu_{k,j}^{n,f}| + |hv_{e,j}^{n,f} - hv_{k,j}^{n,f}| \right]. \quad (30)$$

Large CRPS values can originate from bias (first terms) or the spread in the ensemble (last terms). Together with the scores for the bias and MSE, this allows one to identify the source of ensemble errors and to infer the properties of the ensemble.

Figure 17 presents the evolution of these skill scores for each data assimilation time step. These results are obtained from the same run as in Section 4.3. When we assimilate the first observation after spin-up on day 3, note that the spread in the ensemble is relatively large by construction for all methods. It should be noted that hu and hv take values up to $500 \text{ m}^2 \cdot \text{s}^{-1}$, which means that all methods have a relatively small bias. In the starting phase, the SparseObsETKF immediately calibrates to the observations, whereas the IEWPF and relaxed SparseObsETKF require several data assimilation steps to correct the bias. We see, however, that the bias for all methods grows over time, but with the relaxed SparseObsETKF keeping the

smallest values. The systematic bias from the relaxation becomes negligible, as it is sufficiently often reduced by repeated weighting with the unbiased analysis.

Looking at the MSE and CRPS, we see that both SparseObsETKF versions improve during the first few assimilation steps and stabilise at a certain level. As expected, the initial improvement with relaxation is slower than without relaxation, but this gap is closed already after one simulation hour, which corresponds to 12 data assimilation cycles. The quality of both SparseObsETKF versions is then stable until approximately simulation day 5, when the model dynamics become more turbulent. At this point, the solution without relaxation starts to deteriorate due to the overfitting. Note that we see a similar trend for the relaxed SparseObsETKF later in the experiment. By relaxing even more ($\phi = 0.25$, not shown), we confirm the trend with even slower convergence in the beginning and later divergence at the end.

Similar to the relaxed SparseObsETKF, IEWPF also converges during the initial data assimilation cycles, but the skill scores do not stabilise and instead diverge slowly. The slow initial convergence was also pointed out in Section 3, where we had to run the data assimilation sufficiently long to reach a stable level before being able to provide a fair comparison.

Note that the ensemble variance can be derived from MSE and bias. Then we see that the forecast variance behaves qualitatively similar as the MSE (not shown here). For a full assessment of the skills of a data assimilation method, a single skill score gives only limited information. But, for instance, the combination of bias and CRPS broadens the insights, since the bias helps to explain the contributions in the CRPS. However, the differences in the skill score results, especially between SparseObsETKF without relaxation and IEWPF, do not seem substantial, whereas we have seen contrary properties in the drift trajectories that stay concealed in the monitoring of the skill scores. In general, this discussion tells us that the SparseObsETKF assimilates the ensemble much stronger towards data than the IEWPF does and exemplifies the effects of relaxation.

4.4.1 | Rank histograms

We next look at rank histograms to analyse the adequacy of the ensemble spread. A short time span in the simulation is repeated multiple times and the rank of the simulation truth in the ensemble ordering is monitored at six dynamically independent locations. Rank histograms then present the frequency of which a certain rank is reported among the N_e realisations of the ensemble, and the shape of the histograms is used as a diagnostic tool to identify shortcomings of methods Saetra et al. (2004). Flat rank

histograms are commonly understood as an indication for ensemble consistency or reliability of the ensemble, as it means that every ensemble member is sampled from the same distribution as the truth.

In Figure 18, we show rank histograms from repeating our experiment 1,000 times, using $N_e = 40$ and simulating the first hour of data assimilation after the spin-up only. The most striking result is the clear U-shape in hu for SparseObsETKF without relaxation, which indicates that the truth often is an outlier in the ensemble and that the ensemble is underdispersive. Furthermore, we observe that IEWPF produces a slight hill-shaped rank histogram for hu , corresponding to an overdispersive ensemble. Both these observations match well with what we saw in Section 4.3. In comparison, the hu rank histogram for the relaxed SparseObsETKF closely resembles a uniform distribution. Note also that the rank histograms for hv are flatter for all three methods, but with a slight tendency towards overdispersion for the relaxed SparseObsETKF. This might be from the nature of the problem, as almost all dynamics in the case are along the x -axis.

Though the rank histograms give insights on how able the ensemble is to represent the uncertainty and we are able to draw similar assertions from them as we already suspected before, the advice of Hamill (2001) and Wilks (2011) is to be careful with their interpretation, since, for example, spatial effects between the different locations become hidden.

4.4.2 | Summary

Based on these results for the nonlinear model, we see that the CRPS together with the bias are a good start for an analysis of the ensembles during the data assimilation phase. These scalar scores are simple to include in any data assimilation sequence. By also analysing statistics over all state variables, we are able to identify additional spatial artefacts and a fundamentally different variance in the ensemble. Even though the rank histogram for the IEWPF looks reasonably flat and we get a well-calibrated mean for the state estimation, we see through the standard deviation that there are artefacts in the ensemble, leading to a higher spread than what we see for SparseObsETKF. This also made us realise how sensitive IEWPF is to the covariance structure in the model error.

Stress testing the SparseObsETKF in these high-dimensional nonlinear experiments with very sparse data discloses that the SparseObsETKF has a tendency of overfitting to the observations, resulting in an underestimated variance in the ensemble. SparseObsETKF is also not able to correctly estimate the unobserved variable η . In practice,

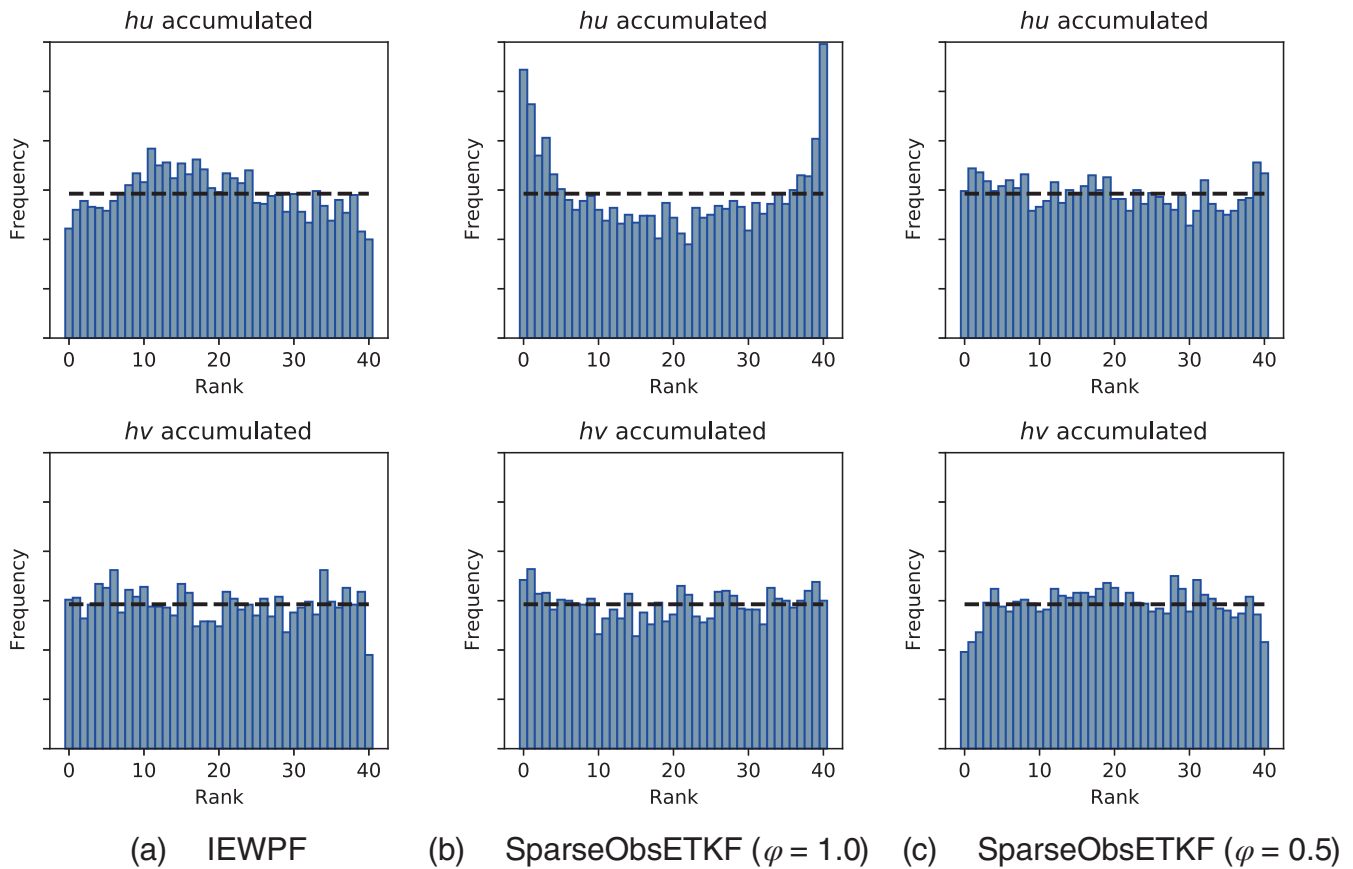


FIGURE 18 Rank histograms recording the rank of the true observation within the ensemble for the observed variables: (a) IEWPF (implicit equal-weight particle filter); (b) SparseObsETKF (ensemble transform Kalman filter with covariance localisation) ($\phi = 1.0$); (c) SparseObsETKF ($\phi = 0.5$). The dashed line indicates the hypothetical uniform distribution.

we see that relaxing the ensemble to prior perturbations is a good remedy for these flaws. It significantly weakens the defects but still inherits the structures from the full SparseObsETKF. We experienced that it is not necessary to fine-tune the relaxation parameter, as results were similar for $\phi = 0.25$ and $\phi = 0.75$. We further point out that we tested classical covariance inflation, but this led to non-physical states for η , whereas the variance of hu and hv was barely affected. Hence, it is fair to use SparseObsETKF and the in-built relaxation with caution.

5 | CONCLUSION

We have compared two conceptually different state-of-the-art ensemble-based data assimilation methods, namely IEWPF and ETKF with localisation, with emphasis on the handling of sparse observation data, and studied how their performances compare with each other. We have shown how localised updates are implemented by design in the IEWPF, provided that the model error

covariance matrix has a local pattern and that observations are only spread in a certain radius around the observation locations and that this is especially enhanced for sparse observations. For the ETKF, we employ an explicit localisation scheme that gives good control of computations in reduced dimensions and of the correlations between observations. We have considered two distinct cases, both motivated by simplified models applicable to oceanography. The first case studied state estimation of a linear Gaussian advection diffusion model, for which we also computed the analytical filtering distribution. This facilitated an in-depth statistical verification of the two methods in terms of estimation of the mean, covariances, distribution coverage, spectral radius, and spatial-temporal connectivity. In the assessment, which also included the standard ETKF, we recorded the performance of the ensemble-based methods in relation to the number of ensemble members and observation size. The second case was a nonlinear shallow-water model used for forecasting of drift trajectories. Here, we compared the performance of SparseObsETKF and IEWPF in terms

of skill scores and forecast abilities. We also discussed relaxation for the SparseObsETKF localisation scheme for this case. The extensive collections of comparison metrics allowed us to analyse several properties in the ensemble representations.

Our results for the first case verified that both the IEWPF and the SparseObsETKF give very good estimates of the analytical reference solution. Additionally, we exhibit effects of different localisation parametrisations for the SparseObsETKF and IEWPF. For moderate ensemble sizes, both methods delivered on par with the KF and clearly outperformed the ETKF in terms of RMSE and coverage probabilities. ETKF was best at estimating the covariance matrices, but it suffers from spurious correlations in the updates. The SparseObsETKF yields small divergences independent of the ensemble size. In the estimation of spatio-temporal model correlations, our results revealed that all three methods performed quite evenly. ETKF converges fastest with respect to the ensemble size. Interestingly, we found that SparseObsETKF and IEWPF only showed minor improvements when increasing the ensemble size, meaning that it is most beneficial to choose one of these methods when computational resources are limited. IEWPF was the scheme benefiting the most from increased number of observations.

In the nonlinear case, we learned that both IEWPF and SparseObsETKF gave estimations of the observed momentum variables with bounded errors, but SparseObsETKF without relaxation did so at the expense of non-physical fields and small spread, causing the drift trajectory forecast in some cases to diverge from the truth. These issues were also seen in the skill scores. IEWPF, on the other hand, showed artefacts around observation sites, indicating that the model error correlation matrix might not always represent the optimal mapping for assimilating the observations. This also caused a large spread in the forecasted drift trajectories. In the case of SparseObsETKF, we showed that applying relaxation clearly improved calibration, resulting in very good general performance. These results were backed up with high-quality results in the skill scores throughout the data assimilation period and precise predictions of the drift trajectories.

To summarise, the most important findings in this article where we consider spatially very sparse observations can be listed as follows:

- Evaluating a broad range of statistical metrics and skill scores proves to be a huge advantage, as it gives a deep insight into the data assimilation methods that are not obtainable through only looking at single metrics. We recommend to start with testing bias and CRPS as one-dimensional quantities, but urge to continue with qualitative analysis for the entire spatial field.

- We strengthen the argument that IEWPF, in contrast to most other PFs, is applicable to high-dimensional applications but that its results are highly dependent on the structure of model error covariance matrix.
- We have seen that the SparseObsETKF works well even for relatively small ensemble sizes, but good calibration of the relaxation is required to retain good results also for complex models. Upon proper relaxation, the results from the SparseObsETKF outperformed those of the IEWPF.

These results, moreover, open up new directions for future research. For instance, it would be interesting to investigate in more detail how sensitive IEWPF is to the structure of the model error covariance matrix. In our work, we tested our implementation of the covariance localisation only with respect to the ETKF. It would be interesting to check whether other variants of the EnKF work equally well, or if they have advantages or disadvantages over ETKF. Beyond this, the sensitivity of the localisation concept to the relaxation parametrisation could be tuned adaptively. Finally, it would also be interesting to test the IEWPF and the SparseObsETKF in a real-world setting by assimilating real observations into an ensemble of simplified ocean models with the aim to predict true drift trajectories.

AUTHOR CONTRIBUTIONS

Florian Beiser: conceptualization; methodology; software; visualization; writing – original draft; writing – review and editing. **Håvard Heitlo Holm:** conceptualization; funding acquisition; methodology; project administration; software; supervision; validation; writing – original draft; writing – review and editing. **Jo Eidsvik:** conceptualization; methodology; supervision; writing – original draft; writing – review and editing.

ACKNOWLEDGEMENTS

We would like to thank Kjetil Olsen Lye for valuable feedback on the manuscript. Furthermore, we thank the two anonymous reviewers for their comments that improved the manuscript further.

DATA AVAILABILITY STATEMENT

The source code and data used to produce the results presented in this article are openly available under a GPL3 licence at <http://doi.org/10.5281/zenodo.10039791> for Section 3 and at <http://doi.org/10.5281/zenodo.10037281> and <http://doi.org/10.5281/zenodo.10039872> for Section 4.

ORCID

Florian Beiser  <https://orcid.org/0000-0001-5579-749X>

REFERENCES

- Anderson, J.L. (2009) Spatially and temporally varying adaptive covariance inflation for ensemble filters. *Tellus, Series A: Dynamic Meteorology and Oceanography*, 61 A, 72–83.
- Anderson, J.L. & Anderson, S.L. (1999) A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127, 2741–2758.
- Asch, M., Bocquet, M. & Nodet, M. (2016) *Data assimilation: methods, algorithms, and applications*. Philadelphia: SIAM.
- Bishop, C.H., Etherton, B.J. & Majumdar, S.J. (2001) Adaptive sampling with the ensemble transform kalman filter. Part i: theoretical aspects. *Monthly Weather Review*, 129, 420–436.
- Breivik, Ø., Allen, A.A., Maisondieu, C. & Olagnon, M. (2013) Advances in search and rescue at sea. *Ocean Dynamics*, 63, 83–88.
- Brodtkorb, A.R. & Holm, H.H. (2021) Coastal Ocean forecasting on the GPU using a two-dimensional finite-volume scheme. *Tellus, Series A: Dynamic Meteorology and Oceanography*, 73, 1–22.
- Burgers, G., Van Leeuwen, P.J. & Evensen, G. (1998) Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review*, 126, 1719–1724.
- Carrassi, A., Bocquet, M., Bertino, L. & Evensen, G. (2018) Data assimilation in the geosciences: an overview of methods, issues, and perspectives. *Wiley Interdisciplinary Reviews: Climate Change*, 9, 1–50.
- Chertock, A., Dudzinski, M., Kurganov, A. & Lukáčová-Medvid'ová, M. (2018) Well-balanced schemes for the shallow water equations with Coriolis forces. *Numerische Mathematik*, 138, 939–973.
- Chopin, N. & Papaspiliopoulos, O. (2020) *An introduction to sequential Monte Carlo methods*. Cham: Springer.
- Desroziers, G., Berre, L., Chapnik, B. & Poli, P. (2006) Diagnosis of observation, background and analysis-error statistics in observation space. *Quarterly Journal of the Royal Meteorological Society*, 131, 3385–3396.
- Doucet, A., Godsill, S. & Andrieu, C. (2000) Methods for Bayesian filtering. *Statistics and Computing*, 10, 197–208.
- Evensen, G. (1994) Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99, 10143–10162.
- Evensen, G. (2003) The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53, 343–367.
- Evensen, G. (2009) *Data assimilation. The ensemble Kalman filter*. Berlin Heidelberg: Springer.
- Foss, K.H., Berget, G.E. & Eidsvik, J. (2021) Using an autonomous underwater vehicle with onboard stochastic advection-diffusion models to map excursion sets of environmental variables. *Environmetrics*, 33, 1–18.
- Gaspari, G. & Cohn, S.E. (1999) Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125, 723–757.
- Gneiting, T. & Raftery, A.E. (2007) Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 359–378.
- Greybush, S.J., Kalnay, E., Miyoshi, T., Ide, K. & Hunt, B.R. (2011) Balance and ensemble Kalman filter localization techniques. *Monthly Weather Review*, 139, 511–522.
- Hamill, T.M. (2001) Interpretation of rank histograms for verifying ensemble forecasts. *Monthly Weather Review*, 129, 550–560.
- Holm, H.H. (2020) *Efficient forecasting of drift trajectories using Simplified Ocean models and nonlinear data assimilation on GPUs*. PhD thesis. Trondheim: Norwegian University of Science and Technology.
- Holm, H.H., Sættra, M.L. & van Leeuwen, P.J. (2020) Massively parallel implicit equal-weights particle filter for ocean drift trajectory forecasting. *Journal of Computational Physics: X*, 6, 100053.
- Houtekamer, P.L. & Mitchell, H.L. (1998) Data assimilation using an ensemble kalman filter technique. *Monthly Weather Review*, 126, 796–811.
- Houtekamer, P.L. & Mitchell, H.L. (2001) A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129, 123–137.
- Houtekamer, P.L. & Zhang, F. (2016) Review of the ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 144, 4489–4532.
- Hunt, B.R., Kostelich, E.J. & Szunyogh, I. (2007) Efficient data assimilation for spatiotemporal chaos: a local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, 230, 112–126.
- Kirchgeßner, P., Nerger, L. & Bunse-Gerstner, A. (2014) On the choice of an optimal localization radius in ensemble Kalman filter methods. *Monthly Weather Review*, 142, 2165–2175.
- Li, H. (2007) *Local ensemble transform Kalman filter*. PhD thesis. College Park: University of Maryland.
- Li, H., Kalnay, E. & Miyoshi, T. (2009) Simultaneous estimation of covariance inflation and observation errors within an ensemble Kalman filter. *Quarterly Journal of the Royal Meteorological Society*, 135, 523–533.
- Morzfeld, M., Hodyss, D. & Snyder, C. (2017) What the collapse of the ensemble Kalman filter tells us about particle filters. *Tellus, Series A: Dynamic Meteorology and Oceanography*, 69, 1–15.
- Nerger, L. (2015) On serial observation processing in localized ensemble Kalman filters. *Monthly Weather Review*, 143, 1554–1567.
- Nussbaumer, R., Mariethoz, G., Gloaguen, E. & Holliger, K. (2018) Which path to choose in sequential Gaussian simulation. *Mathematical Geosciences*, 50, 97–120.
- Ott, E., Hunt, B.R., Szunyogh, I., Zimin, A.V., Kostelich, E.J., Corazza, M. et al. (2004) A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A: Dynamic Meteorology and Oceanography*, 56, 415–428.
- Raanes, P.N., Bocquet, M. & Carrassi, A. (2019) Adaptive covariance inflation in the ensemble Kalman filter by Gaussian scale mixtures. *Quarterly Journal of the Royal Meteorological Society*, 145, 53–75.
- Röhrs, J., Dagestad, K.F., Asbjørnsen, H., Nordam, T., Skancke, J., Jones, C.E. et al. (2018) The effect of vertical mixing on the horizontal drift of oil spills. *Ocean Science*, 14, 1581–1601.
- Sættra, O., Hersbach, H., Bidlot, J.R. & Richardson, D.S. (2004) Effects of observation errors on the statistics for ensemble spread and reliability. *Monthly Weather Review*, 132, 1487–1501.
- Sætrum, J. & Omre, H. (2013) Uncertainty quantification in the ensemble Kalman filter. *Scandinavian Journal of Statistics*, 40, 868–885.
- Sakov, P. & Bertino, L. (2011) Relation between two common localisation methods for the EnKF. *Computational Geosciences*, 15, 225–237.
- Scott, D.W. (1992) *Multivariate density estimation: theory, practice, and visualization*, Second edition. New York: Wiley.
- Sigrist, F., Künsch, H.R. & Stahel, W.A. (2015) Stochastic partial differential equation based modelling of large space-time data

- sets. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 77, 3–33.
- Skauvold, J., Eidsvik, J., van Leeuwen, P.J. & Amezcuca, J. (2019) A revised implicit equal-weights particle filter. *Quarterly Journal of the Royal Meteorological Society*, 145, 1490–1502.
- Snyder, C., Bengtsson, T., Bickel, P. & Anderson, J. (2008) Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136, 4629–4640.
- Soares, R.V., Luo, X., Evensen, G. & Bhakta, T. (2021) Handling big models and big data sets in history-matching problems through an adaptive local analysis scheme. *SPE Journal*, 26, 973–992.
- Szunyogh, I., Satterfield, E.A., Elana, J., Gyarmati, G., Kalnay, E., Hunt, B.R. et al. (2007) The local ensemble transform Kalman filter and its implementation on the NCEP global model at the University of Maryland the LETKF and its implementation on the NCEP GFS model. Proc. ECMWF Workshop on Flow-Dependent Aspects of Data Assimilation, 1–18.
- Thorarindottir, T.L., Gneiting, T. & Gissibl, N. (2013) Using proper divergence functions to evaluate climate models. *SIAM-ASA Journal on Uncertainty Quantification*, 1, 522–534.
- van Leeuwen, P.J. (2009) Particle filtering in geophysical systems. *Monthly Weather Review*, 137, 4089–4114.
- van Leeuwen, P.J. (2020) A consistent interpretation of the stochastic version of the ensemble Kalman filter. *Quarterly Journal of the Royal Meteorological Society*, 146, 2815–2825.
- van Leeuwen, P.J., Künsch, H.R., Nerger, L., Potthast, R. & Reich, S. (2019) Particle filters for high-dimensional geoscience applications: a review. *Quarterly Journal of the Royal Meteorological Society*, 145, 2335–2365.
- Vetra-Carvalho, S., van Leeuwen, P.J., Nerger, L., Barth, A., Altaf, M.U., Brasseur, P. et al. (2018) State-of-the-art stochastic data assimilation methods for high-dimensional non-Gaussian problems. *Tellus, Series A: Dynamic Meteorology and Oceanography*, 70, 1–38.
- Whitaker, J.S. & Hamill, T.M. (2002) Ensemble data assimilation without perturbed observations. *Monthly Weather Review*, 130, 1913–1924.
- Wilks, D.S. (2005) *Statistical methods in the atmospheric sciences*. London: Elsevier Inc.
- Wilks, D.S. (2011) On the reliability of the rank histogram. *Monthly Weather Review*, 139, 311–316.
- Zhang, F., Snyder, C. & Sun, J. (2004) Impacts of initial estimate and observation availability on convective-scale data assimilation with an ensemble kalman filter. *Monthly Weather Review*, 132, 1238–1253.
- Zhu, M., van Leeuwen, P.J. & Amezcuca, J. (2016) Implicit equal-weights particle filter. *Quarterly Journal of the Royal Meteorological Society*, 142, 1904–1919.

How to cite this article: Beiser, F., Holm, H.H. & Eidsvik, J. (2024) Comparison of ensemble-based data assimilation methods for sparse oceanographic data. *Quarterly Journal of the Royal Meteorological Society*, 150(759), 1068–1095. Available from: <https://doi.org/10.1002/qj.4637>