

Iterative Ensemble Smoothers for Data Assimilation in Coupled Nonlinear Multiscale Models

Evensen, Geir; Vossepoel, Femke C.; van Leeuwen, Peter Jan

DOI

[10.1175/MWR-D-23-0239.1](https://doi.org/10.1175/MWR-D-23-0239.1)

Publication date

2024

Document Version

Final published version

Published in

Monthly Weather Review

Citation (APA)

Evensen, G., Vossepoel, F. C., & van Leeuwen, P. J. (2024). Iterative Ensemble Smoothers for Data Assimilation in Coupled Nonlinear Multiscale Models. *Monthly Weather Review*, 152(6), 1277-1301. <https://doi.org/10.1175/MWR-D-23-0239.1>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Iterative Ensemble Smoothers for Data Assimilation in Coupled Nonlinear Multiscale Models

GEIR EVENSEN,^{a,b} FEMKE C. VOSSEPOEL,^c AND PETER JAN VAN LEEUWEN^d

^a Norwegian Research Center (NORCE), Bergen, Norway

^b Nansen Environmental and Remote Sensing Center, Bergen, Norway

^c Department of Geoscience and Engineering, Delft University of Technology, Delft, Netherlands

^d Department of Atmospheric Sciences, Colorado State University, Fort Collins, Colorado

(Manuscript received 26 October 2023, in final form 2 February 2024, accepted 29 February 2024)

ABSTRACT: This paper identifies and explains particular differences and properties of adjoint-free iterative ensemble methods initially developed for parameter estimation in petroleum models. The aim is to demonstrate the methods' potential for sequential data assimilation in coupled and multiscale unstable dynamical systems. For this study, we have introduced a new nonlinear and coupled multiscale model based on two Kuramoto–Sivashinsky equations operating on different scales where a coupling term relaxes the two model variables toward each other. This model provides a convenient testbed for studying data assimilation in highly nonlinear and coupled multiscale systems. We show that the model coupling leads to cross covariance between the two models' variables, allowing for a combined update of both models. The measurements of one model's variable will also influence the other and contribute to a more consistent estimate. Second, the new model allows us to examine the properties of iterative ensemble smoothers and assimilation updates over finite-length assimilation windows. We discuss the impact of varying the assimilation windows' length relative to the model's predictability time scale. Furthermore, we show that iterative ensemble smoothers significantly improve the solution's accuracy compared to the standard ensemble Kalman filter update. Results and discussion provide an enhanced understanding of the ensemble methods' potential implementation and use in operational weather- and climate-prediction systems.

KEYWORDS: Bayesian methods; Inverse methods; Kalman filters; Uncertainty; Coupled models; Data assimilation

1. Introduction

Numerical weather prediction at national and international weather centers uses different data assimilation approaches to initialize ocean and atmosphere models. [de Rosnay et al. \(2022\)](#) and [Laloyaux et al. \(2016\)](#) describe the data assimilation system at ECMWF, which combines the assimilation of ocean observations using a 3DVAR scheme with the assimilation of atmospheric observations using a 4DVAR method. The updates are uncoupled during inner loops but weakly coupled during the short forecast connecting successive assimilation windows. While this approach is less optimal than providing a fully coupled update, it is a practical intermediate step when using variational data assimilation methods with different time windows. The reason for uncoupled inner loops is a fundamental and unsolved problem related to the different time scales of the ocean and atmosphere. An optimal time window for an ocean model might be a week, while it could be a day for the atmosphere at a similar resolution. Hence, one must clarify the optimal window length in a 4DVAR system containing both model components. Furthermore, data latency is a problem because ocean observations tend to have much larger latency than atmospheric observations, mainly for historical reasons. The data latency is one of the reasons

that GFDL uses a two-step ensemble-based filtering algorithm applied to a fully coupled climate model ([Chang et al. 2013](#)), and NCEP assimilates data into partially coupled Earth system components using a 3DVAR-based scheme to incorporate land surface, atmosphere, ocean, and sea ice observations ([Saha et al. 2010](#)).

[Penny et al. \(2017\)](#) provide a comprehensive overview of data assimilation in coupled ocean–atmosphere models. While most operational approaches have a weak coupling between the two subsystems and so-called outer-loop coupling, as mentioned above, more conceptual studies investigate the effect of a strong coupling. For example, [Luo and Hoteit \(2014\)](#) use a multiscale Lorenz-96 model to evaluate the performance of a state-space estimation strategy that assimilates data into each subsystem and correlated quantities from other coupled subsystems. In a slightly more realistic setting, [Han et al. \(2013\)](#) couple a Lorenz-63 model to a pycnocline ocean model.

[Penny et al. \(2019\)](#) compare several data assimilation algorithms on a single quasigeostrophic model for a coupled ocean and atmosphere. [Tondeur et al. \(2020\)](#) used the same model to study coupled data assimilation with the ensemble Kalman filter (EnKF) for a system consisting of a slow ocean coupled to a fast atmosphere. This study illustrates the mechanisms behind information propagation between the system's two components and concludes that cross-component effects are strong from the slow to the fast scale.

We can view the current study as an extension of [Tondeur et al. \(2020\)](#), in which, on the one hand, we will study the information flow between two coupled systems more closely

 Denotes content that is immediately available upon publication as open access.

Corresponding author: Geir Evensen, geev@norce-research.no

and, on the other hand, use more sophisticated data assimilation techniques. To this end, we will use a simplified representation of two coupled Earth system components to explore the information propagation between model components with predominantly large and small spatial scales. The coupled model uses the Kuramoto–Sivashinsky (KS) equations to describe interactions between two systems: one with a longer spatial scale (typically the atmosphere) and one with a shorter spatial scale (typically the ocean). In addition, we evaluate the efficiency of different adjoint-free data assimilation algorithms in coupled data assimilation.

We describe the model in the following section. After that, we introduce and discuss the data assimilation methods before we run multiple experiments to examine the impact and value of coupled data assimilation and various sensitivities to window lengths, number of iterations, etc.

2. Coupled multiscale Kuramoto–Sivashinsky model

We will now describe the nonlinear multiscale model used to study the performance of our data assimilation methods for the coupled data assimilation problem. We need a model with specific properties, including unstable and near-chaotic dynamics with nonlinear saturation of the linear instabilities, similar to oceanic and atmospheric behavior. For computational reasons and ease of interpretation, we search for a one-dimensional and univariate equation for each component of the coupled system. Finally, we require a model that can include different spatial and temporal scales, as we wish to study the assimilation of observations in coupled models with different scales. The model system should resemble the behavior of coupled climate models, where the ocean and atmospheric components have vastly different spatial and temporal scales.

Our choice landed on a variant of the KS equation, from which we derived two model configurations operating on different spatial scales and with a coupling term connecting the two models. In the current study, we only include different spatial scales between the model components, but it is possible to alter the temporal scale of the models in future studies.

The KS equation is a fourth-order partial differential equation initially used to describe diffusive thermal instabilities in laminar flame fronts (Kuramoto 1978; Sivashinsky 1977, 1980). The model can also describe the dynamics of fluid films on inclines (Shlang and Sivashinsky 1982; Tilley et al. 1994), flow in pipes (Chang 1986), and dynamics of chemical reactions. Kuramoto (1978) described how the coupling of an oscillation and a spatial inhomogeneity could produce spatiotemporal chaos and how one can obtain a balance between phase instability and amplitude instability. The KS equation is

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = -\frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4}, \quad (1)$$

which reduces to the inviscid Burgers equation if we ignore the diffusion terms on the right-hand side. The nonlinear advection term uu_x transfers energy between large and small scales and creates shocks. The harmonic diffusion term has a negative sign and acts to enhance any small-scale feature in

the model solution. In contrast, the biharmonic diffusion term with a negative sign is a small-scale selective positive diffusion, thereby controlling any growing instabilities in the model.

Protas et al. (2004) introduced the KS equation for an evaluation of 4DVAR data assimilation methods. Jardak et al. (2010) and Chorin and Krause (2004) evaluated the performance of Bayesian filters using the KS equation, Azouani and Titi (2014) and Lunasin and Titi (2017) used the equations as a testbed for the so-called continuous data assimilation technique, and Waller et al. (2014) applied it for correlated observation error estimation.

We refer to Lak (2023) for a discussion of the numerical implementation and example codes. In short, we have used a Crank–Nicolson/Adams–Bashforth (CNAB) scheme for time stepping the model. The time-stepping method is second order and implicit in time. We discretized space using a Fourier decomposition as this approach renders the inverse matrices in the scheme diagonal, allowing for a highly efficient implementation. Our FORTRAN-90 subroutine for the model integration is available from Evensen (2023b).

In our implementation, we have used two KS models, and we refer to the two model solutions as Atmos and Ocean with the symbols A and O referring to their respective variables. The coupled model equations read

$$\frac{\partial A}{\partial t} = -\frac{1}{2} \frac{\partial A^2}{\partial x} - \frac{\partial^2 A}{\partial x^2} - \frac{1}{2} \frac{\partial^4 A}{\partial x^4} + \alpha_{oa}(O - A), \quad (2)$$

$$\frac{\partial O}{\partial t} = -\frac{1}{2} \frac{\partial O^2}{\partial x} - \frac{\partial^2 O}{\partial x^2} - \frac{\partial^4 O}{\partial x^4} + \omega_{ao}(A - O). \quad (3)$$

We couple the two equations through the coupling terms $\alpha_{oa}(O - A)$ and $\omega_{ao}(A - O)$ where we used coupling coefficients of 0.003 in both equations in all the following experiments, except for one uncoupled simulation where we set the coefficients to zero. Additionally, we halved the biharmonic damping of the Atmos variable to have more structures in the solutions.

We defined two different physical lengths for the Ocean and Atmos model domains to introduce different spatial scales in the two models. We used a periodic model domain with 1024 grid points but assumed a physical domain length of 32 for the Atmos domain and 256 for the Ocean domain. Multiplying the time derivative in one of the equations by a number not equal to one would introduce different time scales in the two model components.

Our multiscale coupled model is suitable for conceptualizing data assimilation in coupled systems. This paper will still refer to the model as KS, although we have modified it from the original KS model in Eq. (1). We can consider it a 1D analog to the 2D Navier–Stokes equations. The model has a unique solution given a set of initial conditions (Tadmor 1986), and it has chaotic behavior and a finite-dimensional global attractor (Papageorgiou and Smyrlis 1991). The 1D KS model simplifies the analysis and interpretation of results, and we avoid using computationally expensive 2D or 3D models.

In nonlinear chaotic models, such as weather prediction models, we can define the model's *predictability time* relative to




FIG. 1. How we split time into discrete time windows is shown. Within each time window l , we integrate a model K time steps and assimilate the data $\mathbf{d}_l$ into the state variable $\mathbf{z}_l$.

the initial uncertainty as how long we can integrate the model before the model solution's uncertainty reaches the climatological variability level. Typically, nonlinear unstable ocean and atmosphere models initially experience exponential error growth before the errors saturate at the climatological variability level due to nonlinear effects.

3. Data assimilation methods

We have adopted data assimilation methods, notation, and formulations from our recent open-access textbook (Evensen et al. 2022), where we split the time dimension into a sequence of assimilation time windows as illustrated in Fig. 1. The definition of finite-length assimilation windows facilitates assimilation of all data available within the window in one update. It makes it possible to use iterative ensemble smoothers to reduce the model nonlinearity's impact and obtain superior results compared to the standard EnKF solution (Bocquet and Sakov 2013a,b). The three methods considered in this paper are versions of the ensemble smoother (ES) that trace back to the ES of van Leeuwen and Evensen (1996), the ES with multiple data assimilation (ESMDA) proposed by Emerick and Reynolds (2012), and the ensemble randomized maximum likelihood (EnRML) method by Chen and Oliver (2012, 2013). We are using the ensemble subspace implementation of Evensen et al. (2019) for the EnRML, which we in the following refer to as the iterative ES (IES). The term “randomized maximum likelihood” is misleading; as we will see below, we solve for a randomized ensemble of maximum a posteriori solutions.

a. A theoretical basis for ensemble methods

The ensemble data assimilation methods attempt to sample the posterior probability density function (pdf) $f(\mathbf{z}|\mathbf{d})$ for the state $\mathbf{z}$ conditional on the measurements $\mathbf{d}$, as defined by Bayes's theorem:

$$f(\mathbf{z}|\mathbf{d}) = \frac{f[\mathbf{d}|\mathbf{g}(\mathbf{z})]f(\mathbf{z})}{f(\mathbf{d})}. \quad (4)$$

Here, $f[\mathbf{d}|\mathbf{g}(\mathbf{z})]$ is the measurement's likelihood function and $f(\mathbf{z})$ is the pdf for the prior estimate of the state $\mathbf{z}$. We have introduced the composite model and measurement operator $\mathbf{g}$ in

$$\mathbf{y} = \mathbf{g}(\mathbf{z}), \quad (5)$$

which maps the state vector $\mathbf{z} \in \mathbb{R}^n$ to the predicted measurements $\mathbf{y} \in \mathbb{R}^m$.

Kitanidis (1995) and Oliver et al. (1996) showed that for a Gaussian prior and likelihood, minimization of an ensemble of cost functions, one for each realization j ,

$$\begin{aligned} \mathcal{J}(\mathbf{z}_j) = & \frac{1}{2}(\mathbf{z}_j - \mathbf{z}_j^f)^T \mathbf{C}_{zz}^{-1}(\mathbf{z}_j - \mathbf{z}_j^f) \\ & + \frac{1}{2}[\mathbf{g}(\mathbf{z}_j) - \mathbf{d}_j]^T \mathbf{C}_{dd}^{-1}[\mathbf{g}(\mathbf{z}_j) - \mathbf{d}_j], \end{aligned} \quad (6)$$

results in an *approximate randomized sampling* of the posterior Bayesian pdf. This “RML” sampling is exact for a linear model and measurement operator when using a Gaussian prior and likelihood. The significance of the approximation depends on the level of nonlinearity of the model in Eq. (5). The cost functions are mutually independent for each realization j and use the random samples $\mathbf{z}_j^f \sim \mathcal{N}(\mathbf{z}^f, \mathbf{C}_{zz})$ for the prior state vector and $\mathbf{d}_j \sim \mathcal{N}(\mathbf{d}, \mathbf{C}_{dd})$ for the perturbed measurements to represent the uncertainties. The superscript f denotes first guess or forecast. The covariances $\mathbf{C}_{zz}$ and $\mathbf{C}_{dd}$ are the error covariances for the prior state vector and the measurements. As shown in van Leeuwen (2020), a more consistent interpretation is to perturb $\mathbf{g}(\mathbf{z}_j)$ instead of the observations, as observations are already perturbed values from the true state, but because of the symmetry in the Gaussian assumption, the results are independent of the perturbation choice.

Note that to solve the assimilation problem for a particular assimilation window, we must assume that measurements are independent between the different windows and that the model is a first-order Markov process. Additionally, we apply a filtering assumption over time windows by only updating the solution in the current assimilation window and ignoring any updates of the past windows. We can easily relax this filtering assumption, where we do not update previous windows, by applying an ensemble Kalman smoother (EnKS) approach to update the solution in previous windows (Evensen and van Leeuwen 2000).

We derive the ensemble methods by setting the gradient of the cost functions in Eq. (6) to zero:

$$\mathbf{C}_{zz}^{-1}(\mathbf{z}_j - \mathbf{z}_j^f) + \nabla_{\mathbf{z}} \mathbf{g}(\mathbf{z}_j) \mathbf{C}_{dd}^{-1}[\mathbf{g}(\mathbf{z}_j) - \mathbf{d}_j] = 0. \quad (7)$$

This equation has no explicit solution in the case when the model $\mathbf{g}(\mathbf{z}_j)$ is nonlinear. However, by introducing a linearization of $\mathbf{g}(\mathbf{z}_j)$ around the prior estimate $\mathbf{z}_j^f$, we obtain an ensemble of Kalman filter update equations.

$$\mathbf{z}_j^a = \mathbf{z}_j^f + \mathbf{C}_{zz} \mathbf{G}_j^T (\mathbf{G}_j \mathbf{C}_{zz} \mathbf{G}_j^T + \mathbf{C}_{dd})^{-1} [\mathbf{d}_j - \mathbf{g}(\mathbf{z}_j^f)]. \quad (8)$$

In the nonlinear case, these equations are only accurate for modest updates when the linear approximation is valid (Bonavita et al. 2018).

Alternatively, we can use the gradient expression from Eq. (7) and the Hessian

$$\nabla_{\mathbf{z}} \nabla_{\mathbf{z}} \mathcal{J}(\mathbf{z}_j) \approx \mathbf{C}_{zz}^{-1} + \nabla_{\mathbf{z}} \mathbf{g}(\mathbf{z}_j) \mathbf{C}_{dd}^{-1} [\nabla_{\mathbf{z}} \mathbf{g}(\mathbf{z}_j)]^T, \quad (9)$$

in a Gauss–Newton iteration to minimize the ensemble of individual and uncoupled cost functions exactly. This approach corresponds to an ensemble-4DVAR type method where we




FIG. 2. Exp. PRED-uncoupled: Uncoupled ensemble-prediction experiment without DA. The upper plots show the Atmos variable with (left) the reference solution, (center) the prior ensemble mean, and (right) the ensemble-predicted standard deviation. The lower plots are the corresponding results for the Ocean variable.

could use, e.g., incremental 4DVAR methods to minimize each independent cost-function realization. We refer to Evensen (2019) for a complete derivation of the resulting iteration.

However, in this paper, we will introduce another approximation that allows us to eliminate the model's tangent linear operator. We will represent the model sensitivity by an ensemble-averaged sensitivity through the linear regression or equivalently the least squares approximation:

$$\nabla_z g(\mathbf{z}_i) \triangleq \mathbf{C}_{zz}^{-1} \mathbf{C}_{zy}, \quad (10)$$

where we replace all individual model sensitivities for the different realizations with a common averaged one. This approximation implies that we no longer solve strictly for the minima of the cost functions in Eq. (6).

The final approximation is to represent all covariances using an ensemble of realizations and we define the ensemble matrix $\mathbf{Z} \in \mathbb{R}^{n \times N}$ as

$$\mathbf{Z} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N). \quad (11)$$

Furthermore, we define the matrix $\mathbf{\Pi} \in \mathbb{R}^{n \times N}$,

$$\mathbf{\Pi} = \left(\mathbf{I} - \frac{1}{N} \mathbf{1} \mathbf{1}^T \right) \sqrt{N-1}, \quad (12)$$

where $\mathbf{1} \in \mathbb{R}^N$ is a vector with all elements equal to one and $\mathbf{I}_N$ is the N -dimensional identity matrix. If we multiply an ensemble matrix with the matrix $\mathbf{\Pi}$, this subtracts the mean from the ensemble and divides the result by $\sqrt{N-1}$.

We can then define the zero-mean and scaled ensemble-anomaly matrix as

$$\mathbf{A} = \mathbf{Z} \mathbf{\Pi}, \quad (13)$$

and the ensemble covariance is

$$\overline{\mathbf{C}}_{zz} = \mathbf{A} \mathbf{A}^T, \quad (14)$$

where the “overbar” denotes that we have an ensemble-covariance matrix.

Correspondingly, we can define an ensemble of perturbed measurements, $\mathbf{D} \in \mathbb{R}^{m \times N}$, when given the real measurement vector, $\mathbf{d} \in \mathbb{R}^m$, as

$$\mathbf{D} = \mathbf{d} \mathbf{1}^T + \sqrt{N-1} \mathbf{E}, \quad (15)$$

where $\mathbf{E} \in \mathbb{R}^{m \times N}$ is the centered measurement-perturbation matrix whose columns are sampled from $\mathcal{N}(0, \mathbf{C}_{dd})$ and




FIG. 3. Exp. PRED-coupled: Coupled ensemble-prediction experiment without DA. See description in Fig. 2.

divided by $\sqrt{N-1}$. We define the ensemble covariance matrix for the measurement perturbations as

$$\bar{\mathbf{C}}_{dd} = \mathbf{E}\mathbf{E}^T. \quad (16)$$

The ensemble algorithms work both with a full-rank $\mathbf{C}_{dd}$ or the ensemble version of the measurement covariance represented by the perturbations in $\mathbf{E}$.

Finally, we define the ensemble of model-predicted measurement anomalies

$$\mathbf{Y} = \mathbf{g}(\mathbf{Z}^f)\mathbf{\Pi}, \quad (17)$$

where we have multiplied the model prediction by the matrix $\mathbf{\Pi}$ to subtract the ensemble mean and divide the resulting anomalies by $\sqrt{N-1}$.




FIG. 4. The plots show the time evolution of the spatially averaged residuals (RMSE) of the prior ensemble-mean predictions from the PRED-uncoupled and PRED-coupled experiments to the reference solution as the dark lines and the ensemble-predicted RMSE as the light lines. We define the ensemble-predicted RMSE as the RMS of the ensemble-predicted standard deviation where the mean is over the spatial coordinate.




FIG. 5. Exp. PRED-uncoupled: Correlations from an uncoupled ensemble prediction without assimilation of data. The upper left plot shows the space–time correlation between a predicted Ocean measurement at the center of the domain (open circle) and the Ocean variable. The upper-right plot shows the corresponding correlations for an Atmos measurement with the Atmos variable. The lower-left plot shows the cross correlation between an Ocean observation and the Atmos variable, and the lower-right plot shows the corresponding cross correlation between an Atmos measurement and the Ocean variable. Note the difference in contour levels for the upper and lower plots.

From Eq. (8), we then obtain the ensemble Kalman filter or smoother update equation formulated entirely in terms of ensemble matrices (Evensen 2003) as

$$\begin{aligned}
 \mathbf{Z}^a &= \mathbf{Z}^f + \mathbf{A}^f \mathbf{Y}^T (\mathbf{Y} \mathbf{Y}^T + \mathbf{E} \mathbf{E}^T)^{-1} [\mathbf{D} - \mathbf{g}(\mathbf{Z}^f)] \\
 &= \mathbf{Z}^f + \mathbf{A}^f \mathbf{W} \\
 &= \mathbf{Z}^f (\mathbf{I} + \mathbf{\Pi} \mathbf{W}) \\
 &= \mathbf{Z}^f (\mathbf{I} + \mathbf{W} / \sqrt{N-1}).
 \end{aligned} \tag{18}$$

The update becomes a “weakly nonlinear” combination of the prior ensemble anomalies, as $\mathbf{W}$ depends on the ensemble of state vectors through the predicted measurements. Still, we will use the notation “linear combination” in the remainder of this paper as this fits well with the form of the update in Eq. (18).

Similarly to the EnKF update, we can write the IES equation using the ensemble matrices. Still, a better alternative is to use the EnRML’s *ensemble subspace* variant developed by Evensen et al. (2019) and Raanes et al. (2019), and we have used the particular ensemble subspace implementation from Evensen et al. (2019) in this work. We refer to these papers and Evensen et al. (2022) for an in-depth discussion of the subspace algorithm. We also remark that the first step in the IES algorithm becomes identical to the ES update if we choose an iteration steplength of one. Hence, we can use the numerical IES implementation to compute the update in ES and IES (and in the ESM DA approach discussed below).

b. Ensemble smoother

When formulated for a single data assimilation window, the update computed by the ensemble smoother is a simple




FIG. 6. Exp. PRED-coupled: Correlations from a coupled ensemble prediction without assimilation of data. See Fig. 5 for a figure description.

extension of the EnKF update by Evensen (1994) and Burgers et al. (1998). In the standard form of EnKF, one updates the model solution instantly when measurements are available. *The ES provides a framework where we can update the solution at any or all time steps in an assimilation window, using simultaneously all measurements available within the window.* When measurements are only available at the end of the window, the ES solution at the end of this window is identical to the EnKF solution. Furthermore, the ES provides a means for using the EnKF formalism with measurements distributed over the assimilation window and updates computed at the end of the window. Thus, we can view the ES as an extension of the EnKF that adds an update in the space–time domain with measurements distributed in space and time over this domain.

c. Iterative ensemble smoother

While the ES computes one linear step in the gradient direction to approximate the cost functions' minima, the IES

uses Gauss–Newton iterations to search for the ensemble of cost functions' minima, which define the posterior ensemble. If the cost functions do not contain local minima, the iterative smoothers should converge to the global minimum of each cost-function realization. But note that we are using the ensemble-averaged model sensitivity from Eq. (10) that slightly changes the gradient Eq. (7) for each realization. Additionally, we use ensemble covariances, which constrain the posterior ensemble to the ensemble subspace spanned by the prior ensemble of realizations.

In its standard formulation, IES attempts to estimate the initial conditions of the assimilation window. After updating the initial conditions, we must reintegrate the ensemble of model realizations over the window to compute the updated gradient in each iteration. Realizing that the final converged transition matrix defines the posterior ensemble solution as a linear combination of the prior ensemble realizations and that we do not need the gradient after the last update, we could

TABLE 1. The coupled vs uncoupled data assimilation experiments.

Experiment	Window length	Num Ocean obs	Num Atmos obs	Δt Ocean obs	Δt Atmos obs
MDA5-w05-d2-O	5	40	0	2	—
MDA5-w05-d2-Osep	5	40	0	2	—
MDA5-w05-d4-A	5	0	15	—	4
MDA5-w05-d4-Asep	5	0	15	—	4
MDA5-w05-d5	5	40	10	5	5
MDA5-w05-d5sep	5	40	10	5	5

use this final transition matrix to update the ensemble directly over the whole assimilation window. These two strategies would give identical results in the linear case without model errors.

d. Ensemble smoother with multiple data assimilation

The ESM DA is an interesting alternative to IES, and it also attempts to approximately sample the posterior from Bayes's theorem. When requiring that a finite number μ of coefficients α_i satisfy the condition

$$\sum_{i=1}^{\mu} \frac{1}{\alpha^i} = 1, \quad (19)$$

we can write Bayes's theorem as

$$\begin{aligned} f(\mathbf{z}|\mathbf{d}) &\propto f[\mathbf{d}|\mathbf{g}(\mathbf{z})]f(\mathbf{z}) \\ &= f[\mathbf{d}|\mathbf{g}(\mathbf{z})] \left(\sum_{i=1}^{\mu} \frac{1}{\alpha^i} \right) f(\mathbf{z}). \end{aligned} \quad (20)$$

This rewriting of the likelihood allows for a gradual introduction of the measurements over a predefined number of update steps. For example, we could use two steps and set $\alpha_1 = \alpha_2 = 2$, which satisfies the condition in Eq. (19). In this case, we would have to solve recursively the two updates

$$f_1(\mathbf{z}|\mathbf{d}) \propto f[\mathbf{d}|\mathbf{g}(\mathbf{z})]^{1/2} f(\mathbf{z}), \quad (21)$$




FIG. 7. Exp. MDA5-w05-d2-O: Assimilation of Ocean data and updating the Ocean and Atmos variables. The upper plots show the Atmos variable with (left) the reference solution, (center) the posterior ensemble mean, and (right) the ensemble-predicted posterior standard deviation. The lower plots are the corresponding results for the Ocean variable. The grid points indicate the location of measurements used in the assimilation.




FIG. 8. Exp. MDA5-w05-d2-Osep: As in Fig. 7, but with separate assimilation where the Ocean data only update the Ocean variable.

$$f(\mathbf{z}|\mathbf{d}) \propto f[\mathbf{d}|\mathbf{g}(\mathbf{z})]^{1/2} f_1(\mathbf{z}|\mathbf{d}). \quad (22)$$

The impact of raising a Gaussian likelihood to a power $1/\alpha$ is an inflation of the measurement error variance where the error covariance matrix becomes multiplied by α . Consequently, we can solve each update step using ES and Eq. (18) but with an inflated measurement error covariance $\alpha \mathbf{C}_{dd} \approx \sqrt{\alpha} \mathbf{E} \sqrt{\alpha} \mathbf{E}^T$. A word of caution is that it is necessary to resample measurement perturbations $\sqrt{\alpha} \mathbf{E}$ in each step to avoid introducing a bias from using dependent samples. This step is easier to understand if we realize we are perturbing the model-predicted observations and the model changes with each update. During the stepwise updating, we must update the initial state of the assimilation window and then reintegrate the model ensemble over the window to obtain the “prior” ensemble of realizations for the next step.

The advantage of ESM DA over ES is that while ES computes one large linear update, ESM DA computes a sequence of small linear updates, reducing the impact of the linearization in the ES scheme. We can view ESM DA as an Euler pseudo-time-stepping in state space with a short step size while ES computes the update over one large time step of length one. ESM DA has become one of the most popular data assimilation methods in the petroleum community, and several companies use the method operationally for parameter estimation in large reservoir models (e.g., Zhao

et al. 2017; Emerick 2018). For linear dynamics and measurements, ESM DA and ES will result in the same solution with increasing ensemble size, independent of the number of ESM DA steps.

When used in sequential data assimilation, ESM DA updates the initial conditions of the assimilation window through a finite number of ES steps. In each ESM DA step, we must rerun the model ensemble over the assimilation window to obtain the prior for the next update step. However, as for IES, we can choose two strategies for the final step in ESM DA. We can update the initial conditions of the assimilation window and rerun the model ensemble to obtain the solution over the window or update the ensemble directly over the whole assimilation window using the ES algorithm without rerunning the ensemble. Both these approaches are valid and consistent as the ESM DA steps are independent ES steps, and as for ES, we can freely choose the update strategy. A critical remark is that the prior ensemble for the last ESM DA step is the “nearly converged” ensemble from the previous update step. In the experiments below, we will show that, as for ES, it is an advantage to compute directly the final ES update over the whole assimilation window, particularly when the assimilation window becomes long compared to the predictability time of the model.

There are fundamental differences between IES and ESM DA. While IES is an iterative method, ESM DA computes a




FIG. 9. Exp. MDA5-w05-d4-A: Assimilation of Atmos data and updating both the Ocean and Atmos. See Fig. 7 for further explanation.

predefined number of update steps. IES minimizes the ensemble of initially defined RML cost functions, and ESM DA solves in each step for the minima of a new resampled ensemble of cost functions using a linear ES update. However, although the final ensemble solutions will be different using IES and ESM DA, both methods consistently attempt to sample the posterior pdf, and in the linear case with increasing ensemble size, they converge to the same pdf.

4. Importance of coupled data assimilation

We will now demonstrate the importance of computing fully coupled data assimilation updates. By the notation “fully coupled,” we understand that measurements of one model component will update all model components through their cross covariances estimated from the ensemble of coupled model realizations. We start in the following section by examining the impact of the coupling term on the KS model and the cross-covariance functions. After that, in section 4b, we will compare coupled and uncoupled data assimilation experiments.

a. Ensemble predictions and covariances

We now present two ensemble-prediction experiments to illustrate some properties of the coupled KS models with our

choice of model parameters. Figure 2 shows the results from the uncoupled prediction experiment, where we run an ensemble of 1000 realizations of the two models in Eqs. (2) and (3) in an uncoupled mode with $\alpha_{oa} = \omega_{ao} = 0$. Hence, the two models evolve independently of each other. We generate the initial samples of smooth Gaussian pseudorandom one-dimensional fields using an inverse fast Fourier transform from a Gaussian spectrum in wave space with random phases for each wavenumber (Evensen 2009, see chapter 11). Note that the realization’s smoothness is not critical as the model quickly develops variability on its inherent scale.

In Fig. 3, we present the coupled reference simulations where $\alpha_{oa} = \omega_{ao} = 0.003$. In both the PRED-uncoupled and PRED-coupled experiments, we notice the apparent chaotic behavior of the two model variables and the differences in spatial scales between the Ocean and Atmos variables, as shown in the left panels of the Figs. 2 and 3. While the two models evolve independently in PRED-uncoupled experiment, we observe a significant impact of the coupling in PRED-coupled experiment, where the fine-scale Ocean features follow the large-scale structures in the Atmos solution, and the latter also changes compared to the uncoupled simulation. The ensemble mean converges quickly toward zero for both compartments in both experiments, as shown in the center panels. Very importantly, the standard deviations




FIG. 10. Exp. MDA5-w05-d4-Bsep: As in Fig. 9, but with separate assimilation where the Atmos data only update the Atmos variable.

shown in the right panels remain at the climatological level of around 1.25 and 1.75 for the two model components but appear to saturate at a slightly lower level in the coupled simulation than in the uncoupled one, suggesting that the coupling slightly stabilizes the model. The higher standard deviation in the Atmos variable, as seen in both simulations, is likely a result of the lower biharmonic diffusion in this model and the difference in scales where the biharmonic diffusion has less effect on the Atmos variable's larger scales.

Figure 4 shows the time evolution of the root-mean-square residuals between the ensemble mean and the reference solution as the dark lines. The light lines are the root-mean-square of the ensemble standard deviations. We have averaged the residuals over the spatial coordinate. With the current ensemble initialization, the error growth saturates after an integration time of about 10–20 units of time, approximately indicating the model's predictability time.

In coupled data assimilation, we are interested in the multivariate ensemble correlations between a predicted observation of a model variable with all model variables in all grid points. In Figs. 5 and 6, we show the space–time correlation functions between an observation of either the Ocean or the Atmos variables located in the center of the domain and the respective model components. The correlation function of an

observation with the measured variable reflects the spatial scales of the variable. However, we note that there is also a correlation in time, which is essential when computing smoother solutions over an assimilation window. We also notice that the correlations from the uncoupled and coupled models differ slightly due to the change in dynamics introduced by the coupling term.

In the uncoupled case, the ensemble correlations between the Atmos and Ocean variables are zero (within the sampling errors). Hence, for an uncoupled model system, an Ocean observation will only influence the Ocean variable and similarly for an Atmos observation. However, for the coupled model, we obtain significant cross correlations between the variables, as seen in Fig. 6. These cross correlations will cause the assimilation of a measurement of one model variable to update the other model variable. We notice that the Ocean observation cross correlation with other Ocean variables has changed compared to the uncoupled case, even at the observation location, because the ocean takes up larger scales from the Atmos variables due to the coupling.

The cross correlations indicate that an Ocean observation will influence the Atmos variable predominantly in the past, while an Atmos observation predominantly influences the Ocean variable in the future. Also, with these cross correlations, an Ocean observation will update the Atmos variable




FIG. 11. Exp. MDA5-w05-d5: Assimilation of Ocean and Atmos data and updating both the Ocean and Atmos. See Fig. 7 for further explanation.

on the Atmos variable's spatial and time scale. In contrast, an Atmos observation will update the Ocean variable on multiple scales. That is, the update introduces a prominent Atmos-scale feature and a finer Ocean-scale variability to the Ocean variable. We conclude from studying the space–time correlations that the Atmos model partly drives the Ocean model, presumably due to the Atmos' larger variability amplitude.

b. Coupled data assimilation experiments

We adopt a general naming convention for the paper's experiments, starting with ES, IES, or MDA, indicating the assimilation method used. For the IES and MDA, an integer follows immediately after the method name and defines the number of iterations in IES or steps in ESMDA. IES without a number means we iterate to convergence. Then, the “w05” tells us that the assimilation window length is five units of time, and “d2” means that we have observations available every $\Delta t = 2$. Note that the first measurements become available at time $t = 50$, and we assume uncorrelated measurement errors with a standard deviation of 0.3 for all the experiments.

We have ignored the impact of model errors in this study as the model is highly nonlinear and unstable, and we will have

strong error variance growth between assimilation steps that likely will dominate the impact of modest model errors. However, if model errors are substantial, they should be included in the data-assimilation problem as unknown forcing fields. In that case, the methods would update the initial conditions and these forcing fields. Evensen (2019, 2021) discussed how to include model errors in iterative ensemble smoothers. The state vector size used in the data assimilation would increase, but depending on how we parameterize the model error, including the model errors might still be feasible. In general, we would argue for including model errors in the data assimilation system, as the weak-constraint data assimilation becomes more stable since observation information can influence the solution at and around the observation time such that the sensitivity to the initial conditions is strongly reduced (see, e.g., Amezcua et al. 2017). Including model errors will again result in a more Gaussian assimilation problem. Hence, we expect all methods to perform better in this case, and while the state-vector size increases, the expected number of iterations decreases.

We have used a 1000-member ensemble in all the following experiments to minimize the impact of sampling errors and to avoid using localization and inflation. In an experiment with localization, it was not possible to visually observe any difference between the global and local analysis. However, in a




FIG. 12. Exp. MDA5-w05-d5sep: As in Fig. 11, but with separate assimilation for the Ocean and Atmos variables.

forthcoming manuscript, we will address localization in experiments with smaller ensemble sizes.

In the experiments discussed in the current section, we use ESM DA with five steps and assimilation windows having a length of five units of time. We varied the spatial measurement coverage to emphasize the results (see Table 1). We start by testing the impact of coupled assimilation of only Ocean data in Exp. MDA5-w05-d2-O and only Atmos data in Exp. MDA5-w05-d2-A, and we present the results in Figs. 7 and 9. Here, we jointly update the Ocean and Atmos variables when we assimilate the Ocean or Atmos data. Complementary to these experiments, we have run Exps. MDA5-w05-d2-Osep and MDA5-w05-d2-Asep, using “uncoupled” data assimilation, where we only update the Ocean variable using the Ocean data and only the Atmos variable when conditioning on the Atmos data (see Figs. 8 and 10).

Assimilation of only Ocean data in experiment MDA5-w05-d2-O constrains the Ocean variable well, as seen in Fig. 7. Interestingly, the information from the Ocean data also improves the Atmos variable significantly. After a few data assimilation updates, we have reduced the residuals for the Atmos variable to a low value compared to what we started with (see Fig. 13). Note that we have a relatively high density of Ocean measurements in space to resolve the spatial scales in the Ocean model variable. In this case, the Atmos solution’s recovery results from the coupled data assimilation. In contrast,

the results of Exp. MDA5-w05-d2-Osep illustrate that when we only update the Ocean variables when assimilating Ocean data, we see hardly any impact on the Atmos variable, and we also observe worsening Ocean results compared to the coupled case. *This result is significant, as it shows that if we improve the atmosphere via ocean observations, the atmosphere couples back to the ocean to improve the ocean simulation further.*

The situation differs for the alternative case, MDA5-w05-d4-A, where we only assimilate Atmos data. We obtained good convergence for the large-scale Atmos model, but with the current spatial resolution in the data, we would need more information to control the Ocean model. We see some impact of the assimilation in the Ocean variable in Fig. 9 where we recover the Ocean’s “large-scale” structures and somewhat reduce the domain-averaged estimated standard deviation. However, compared with Exp. MDA5-w05-d4-Asep in Fig. 10, this weak improvement in the Ocean variable is primarily due to the model coupling and dynamic interaction between the Atmos and Ocean variables. Hence, in our problem setting, the atmospheric observations do not directly update the ocean. Still, the ocean variables benefit indirectly from better atmospheric fields in the forecast stage. Interestingly, the estimated standard deviation is large near sharp gradients in the Ocean solution, for which the Atmos observations are too course.




FIG. 13. The plots show the time evolution of the residuals for the MDA5-w05-d5 experiments listed in Table 1. The dark lines indicate RMSE relative to the reference solution while the light lines are the ensemble-predicted RMSE. The upper plots show results from joint updates of the (left) Ocean and Atmos variables using Ocean observations, (center) Atmos observations, and (right) Ocean and Atmos observations. The lower plots show the corresponding results when we compute separate and independent updates of the Ocean and Atmos variables.

The difference between the coupled assimilation of Ocean and Atmos data in experiment MDA5-w05-d5 and the separate assimilation of Ocean and Atmos data in experiment MDA5-w05-d5sep emphasizes the value of the coupled assimilation. Exp. MDA5-w05-d5 presented in Fig. 11 converges after a few assimilation windows, and we can control the further evolution of the model ensemble. Exp. MDA5-w05-d5sep shown in Fig. 12 converges much slower, and the resulting estimate has larger errors for the Ocean and Atmos variables. The residual plots in Fig. 13 also support the conclusion that combined assimilation of the Ocean and Atmos data captures the interaction between Ocean and Atmos more accurately and can lead to an improvement in the state estimate compared to the results of separate assimilation for the Ocean and Atmos variables.

5. Sensitivity study for ensemble DA methods

We will now further examine the properties of the iterative ensemble smoothers with our coupled model. The focus is not on the coupling but, instead, on how to best configure the assimilation setup for the different assimilation methods in terms of the update strategy for the final iteration or step, the length of the data assimilation window relative to the model's predictability time, the optimal number of steps in ESMDA, and the required number of IES iterations. We will now examine ES, IES, and ESMDA in a recursive setting, where the assimilation updates in future assimilation windows will

benefit from the updates from the previous assimilation steps and, thereby, a reduced nonlinearity and a more Gaussian pdf of the prior ensemble. We note that the current versions of IES and ESMDA were previously only studied for parameter estimation on a single window. Still, from the results reported by Bocquet and Sakov (2013a,b), with their alternative iterative smoother formulation, we expect a significant improvement when using the iterative methods compared to what we can get with EnKF or ES.

Before we discuss the experiments in more detail, we note that ES allows for state updates over the whole assimilation window, even when there are no model errors. The reason is that the ensemble is present over the whole assimilation window, and Bayes's theorem allows us to update the solution at every time step within the window. Of course, there is a direct advantage to updating the state at the end of the window in a perfect model setting, as this will typically lead to a better forecast than updating only at the beginning of the assimilation window. This approach is impossible in strong-constraint variational methods, which define the data assimilation problem as finding the initial conditions that lead to the model solution that minimizes a cost function. The violation of the strong constraint model assumption in ES might seem inconsistent. Still, it is similar to the inconsistency in a strong-constraint 4DVAR over a sequence of assimilation windows: we violate the perfect model assumption at the start of each assimilation window. Van Leeuwen et al. (2015, chapter 1) discussed the possible degradation of the

TABLE 2. The experiments used for testing whether to update initial conditions or the whole ensemble over the data assimilation window. The experiment names ending with ini update the initial conditions of the data assimilation window and then rerun the model ensemble to obtain the final solution over the window.

Experiment	Window length	Num Ocean obs	Num Atmos obs	Δt Ocean obs	Δt Atmos obs
ES-w06-d2	6	40	10	2	2
MDA5-w12-d5	12	40	10	5	5
IES-w05-d2	5	40	10	2	2
ES-w06-d2ini	6	40	10	2	2
MDA5-w12-d5ini	12	40	10	5	5
IES-w05-d2ini	5	40	10	2	2

solution at the end of the assimilation window in a strong-constraint 4DVAR.

a. Window update or model rerun

From section 3, we noticed that we can update the solution at any instant over the assimilation window when using ES. Furthermore, for an assimilation window, the update uses the same linear combination of the prior ensemble realizations independently of the time step we are updating. Hence, for the ES update, we can choose between two strategies: update the solution over the whole window or update only the assimilation window's initial conditions and rerun the model to obtain the solution over the window. We will see that these two strategies lead to substantially different results for nonlinear

dynamical models. At the same time, in the limit of infinite ensemble size, there would be no difference for linear models without model errors.

For ES, it is always better to update the whole window and, particularly, the end time of the window, as this estimate becomes the initial condition for the continued integration over the following window. The importance of updating the whole window depends on the length of the assimilation window relative to the model's predictability time. For short time windows, the error growth caused by the nonlinear and unstable model dynamics will not have time to impact the prediction significantly. In this case, whether we update the window's initial conditions or compute the linear ES update over the whole window yields similar results. However, for




FIG. 14. Exp. ES-w06-d2: ES with an assimilation window with a length of six time units and a data-assimilation update of the final solution over the assimilation window.




FIG. 15. Exp. ES-w06-d2ini: As in Fig. 14, but with an update of the window's initial condition.

longer time windows, the uncertainty at the end of the window obtained by integrating the model ensemble from updated initial conditions typically tends toward climatology, and the ensemble prediction in the next assimilation window starts all over again from a climatological level where the ensemble has forgotten all information from previously assimilated data.

We will now illustrate the impact of updating either the data assimilation window's initial conditions followed by an ensemble integration or updating the ensemble directly over the whole window in the final iteration or step. We have run the experiments listed in Table 2, where we chose the window lengths in the different experiments to emphasize the impact of using the two strategies with the different ensemble smoothers.

In Figs. 14 and 15, we illustrate the difference between the two update strategies when using ES for a case with window lengths of six units of time and measurements every second unit of time. We obtain a faster convergence and lower standard deviations when we update the ensemble over the whole window. This case corresponds to computing an EnKS solution, where we have used measurements within an assimilation window to update the ensemble at the end of the assimilation window before continuing the integration.

ESMDA updates the assimilation window's initial conditions recursively, and following each update, we must rerun

the ensemble over the window to obtain the solution for the current ESM DA step. This procedure increases the ensemble spread and uncertainty over the window due to using a nonlinear and unstable model. However, we can control some uncertainty growth at the final update step by calculating the ESM DA update over the whole window. This approach avoids the final ensemble integration. In the case of long assimilation windows, we obtain a better estimate of the solution at the end of the assimilation window and, thereby, a better estimate of the initial conditions of the following window. Figures 16 and 17 illustrate how computing the solution over the whole window in the final ESM DA step improves the estimate. In particular, we notice a faster convergence for the Atmos variable, and we are better at controlling the dynamic instabilities that develop while integrating the model over the window. In an operational weather or climate prediction setting, we would typically update the solution at the end of the assimilation window at the final ESM DA step before predicting the solution over the following window.

The situation changes entirely in the IES as shown in Figs. 18 and 19. Updating the solution over the whole window in the final iteration significantly degrades the estimate. When using IES, we should update the initial condition of the assimilation window in the final integration.




FIG. 16. Exp. MDA5-w12-d5: ESMDA with an assimilation window with a length of five time units and a data-assimilation update of the final solution over the assimilation window.

The reason ESMDA experiences an improvement while the IES solution degrades, in this case, is the following. *When computing the last update in ESMDA, the solution is a linear combination of the ensemble simulations of the previous MDA step. On the other hand, in IES, the estimate is a linear combination of the prior ensemble simulations.* Each step in ESMDA improves the initial conditions for the assimilation window, leading to an improved ensemble prediction over the assimilation window with lower uncertainty. Thus, the prior ensemble is already close to the posterior solution in the final update step.

The summary in Fig. 20 clearly illustrates the above findings. The most surprising observation from this figure is that the window length used for ESMDA significantly exceeds the one used with IES. We will see in the following section that the possibility of using a standard ES update for the final ESMDA step makes ESMDA less sensitive to the window length when used in sequential data assimilation. ESMDA also tolerates more significant errors in the initial conditions of the ensemble since we can partly correct these errors in the final ES-type update step. Furthermore, while ESMDA computes recursive linear regression updates, IES is a gradient-descent method and becomes sensitive to substantial nonlinearities. We will elaborate more on these topics in the following sections.

1) ENSEMBLE SMOOTHER SENSITIVITY EXPERIMENTS

We will now examine how the ensemble smoothers perform when we vary the length of the assimilation windows. In the case of ES, we repeat the experiments for window lengths increasing from one to nine units of time, and the time interval between the measurements is two for all the experiments. We have 10 and 40 equally spaced measurements of the Atmos and Ocean variables at each measurement time. We denote the experiments ES-w0[1–9]-d2[ini], where ES denotes the ensemble smoother, and the numbers 0[1–9] define the length of the assimilation windows, and d2 tells us that we have measurements available every second unit of time. Finally, an “ini” at the end tells us we reran the model ensemble to obtain the solution. In contrast, the experiments without the ini computed the ES update over the whole window.

Note that the ES and the EnKF will always have identical solutions at the end of the assimilation window; therefore, the prior for the following window will also be the same. We have already shown that we obtain better results from ES when we update the solution at the end of the assimilation window than when we update the window’s initial conditions and rerun the ensemble over the window. From the upper plots in Fig. 20, we notice that ES generally leads to more accurate solutions when updating the solution over the window and not rerunning the ensemble. Updating the window’s initial




FIG. 17. Exp. MDA5-w12-d5ini: As in Fig. 16, but with an update of the window's initial condition.

conditions results in poorer performance for all window lengths, and we can conclude that the standard EnKF with updates at the end of the assimilation window is the way to go when the emphasis is on forecasting.

Regarding the length of the assimilation window, the ES experiments perform well for a short window length of two units of time. Still, the performance deteriorates rapidly for longer assimilation windows. This worsening of the results with increasing window length comes from the model's non-linearity, which results in a prior ensemble prediction over the window that approaches climatology for long windows. Then, the linear update computed by ES will break down, as previously discussed by Evensen and van Leeuwen (2000) in an example using the Lorenz-63 equations. From the upper left plot in Fig. 20, we notice how the standard deviation over the window grows until the simulation approaches climatology with long windows. The net effect is that at the end of the assimilation window, no predictive skill propagates forward into the following window. So, when using ES and EnKF, we should use short data assimilation windows and update the solution at the end of, or the whole, window.

2) IES SENSITIVITY EXPERIMENTS

We now move on to examining IES in more detail. Similarly to the ES experiments in the previous section, we use the

notation IES-w[01–15]-d2[ini] for the IES experiments. With IES, we could extend the length of the assimilation window, and we have run experiments with assimilation window lengths up to 15 units of time. In all the experiments, we used a maximum of 12 iterations, but in most cases, we observed no improvement after about five to eight iterations.

The plots in the second row of Fig. 20 summarize the residuals from using IES with different assimilation window lengths. IES performs equally well for window lengths of two to seven units of time, as seen from the right plot in the second row of Fig. 20, and the results are significantly better than those obtained using ES (see upper left panel in the figure). Thus, an iterative method benefits from a more accurate estimate but comes at a higher computational cost.

As for ES, we repeated all 15 IES-w[01–15]-d2ini experiments by a set of experiments IES-w[01–15]-d2 where we used the IES transition matrix to update the whole window in the final iteration. However, from the residuals in Fig. 20, it is clear that for IES, in contrast to ES, we should rerun the model ensemble in the final iteration. We explained in section 3 that IES computes the updated ensemble of initial conditions as a linear combination of the prior ensemble initial conditions. Furthermore, we recall that the linear combination




FIG. 18. Exp. IES-w05-d2: IES with an assimilation window with a length of 12 time units and a data-assimilation update of the final solution over the assimilation window.

defined by the transition matrix leads to a posterior ensemble prediction that minimizes the ensemble of cost functions in Eq. (6). Thus, we lose the effect of the model nonlinearity by directly updating the posterior ensemble over the data assimilation window. Note that the estimates would become identical when using a linear model.

3) ESM DA SENSITIVITY EXPERIMENTS

We now apply the ESM DA method in experiments MDA-w[01–15]-d2[ini], similar to those from the previous section. The ESM DA formulation for the recursive data assimilation problem assumes that we update the initial state and rerun the model ensemble in each step to recompute the gradient. However, since ESM DA computes a sequence of independent ES updates, we can choose whether to update the initial conditions for the assimilation window or the whole solution over the window in the final ESM DA step.

From the plots in the bottom row of Fig. 20, we notice that ESM DA performs significantly better than ES in both the case with updating the whole window and when we rerun the model ensemble over the window. We can extend the window length significantly compared to when using ES, and there are hardly any differences between the

ESM DA or IES solutions for window sizes up to six units of time. We note, however, that the ESM DA residuals diverge from the ensemble standard deviations for window lengths greater than 14 units of time in the MDAX experiments.

We also note that computing an ES update of the whole window in the final step stabilizes the solution for longer windows and improves the estimate, particularly for long window lengths. This strategy also saves one ensemble integration.

To conclude, we are less impacted by the model's nonlinearity when using ESM DA than when using ES, and ESM DA has the added advantage of significantly extending the data assimilation window length. We also obtained a substantially better result with ESM DA than ES, although at a higher computational cost for the same window length.

b. Sensitivity to number of MDA steps

The number of ESM DA steps can significantly impact the solution (Evensen and Eikrem 2018; Evensen et al. 2021). Ideally, we would like to use as few steps as possible to reduce computational cost since every step implies a simulation of the whole ensemble. Using around four to eight steps is common (Zhao et al. 2017; Evensen and Eikrem




FIG. 19. Exp. IES-w05-d2ini: As in Fig. 18, but with an update of the window's initial condition.

2018), but we currently lack a general rule for how many steps we need to obtain the best estimate. We can only try a different number of steps, and when the solution does not change within the expected sampling errors, we can assume we have converged.

An issue is that every ESM DA step is an independent ES update and introduces a new set of perturbed observations with increasing perturbation magnitudes when using multiple steps. Another problem is that using a limited ensemble size introduces sampling errors. In the current experiments, we have used 1000 realizations to minimize the impact of sampling errors and data assimilation windows of 5, 10, and 15 units of time. We ran nine experiments using one to nine steps, and we plot the residuals in Fig. 21, averaged over the converged solution during the last 100 units of time. The experiments using a single step are equivalent to running ES, and it is clear that using more than one step improves the solution. In the current example, using two MDA steps leads to a significant improvement, and for a window length of five units of time, we obtain nearly identical accuracy when using ESM DA with two to nine update steps. In this case, using more than two to three steps are useless. For the longer assimilation windows of 10 and 15 units of time, we obtain the most consistent solution using only three to seven MDA steps, all with similar performance. We also note that additional steps lead to

divergence between the actual and estimated residuals. In the case with an assimilation window of 15 units of time and nine MDA steps, we experience substantial filter divergence.

Note that running experiments with a relatively small ensemble size makes it useless to run additional ESM DA steps if the method has already converged to a level lower than the sampling errors resulting from the limited ensemble size. A positive result from these experiments is that we can obtain estimates with excellent accuracy using relatively few MDA steps.

6. Comparative performance of ensemble smoothers

Which ensemble method should we use? We can apply the ES with an update of the whole assimilation window or in a pure EnKF setting with an update of the final time step of the window. However, ESM DA and IES will improve upon the ES results in cases with significant nonlinearity at an additional computational cost.

We like the consistency of the IES performance seen in the middle right plot of Fig. 20, where the converged solution shows consistency between actual and estimated residuals up to a window length of eight units of time. On the other hand, ESM DA converged in three to five steps for moderate assimilation window lengths and may be more computationally




FIG. 20. The plots show the time-averaged residuals for the ES, ESMDA, and IES experiments listed in Table 2 as a function of window length for the period 101–200. The dark lines indicate RMSE relative to the reference solution, while the light lines are the ensemble-predicted RMSE.

efficient than IES. ESMDA also works well with longer window lengths due to the final MDA step, where we compute an ES update over the whole window, which results in a higher accuracy and lower uncertainty of the initial conditions for the subsequent window.

In all the previous IES experiments, we ran a maximum of 12 iterations, but in most cases, IES converged in fewer iterations. In a sequential data assimilation system, the prior for each successive assimilation window is typically rather good, so we could reduce the number of iterations

in IES without sacrificing the quality of the results. We also tested IES on the cases IES4-w[02–06]-d2ini using a maximum of four iterations for each update, which in computational cost corresponds to running ESMDA with five steps since ESMDA avoids the final ensemble integration by updating the ensemble over the whole assimilation window in the last step. We note that the final ES update used in ESMDA helps control nonlinear instabilities and leads to an improved initial condition for the following window.




FIG. 21. The plots show time-averaged residuals for the ESMDA as a function of the number of MDA steps for three different window lengths of 5, 10, and 15 units of time for the period 101–200 when we expect the assimilation cycle to have converged. The dark lines indicate RMSE relative to the reference solution, while the light lines are the ensemble-predicted RMSE.

In Fig. 22, we compare the time series of residuals for the IES (IES-w[02–06]-d2ini), IES with four iterations (IES4-w[02–06]-d2ini), and the default case of ESMDA with five steps (MDA5-w[02–06]-d2). We cannot conclude from these plots that one configuration is significantly better or worse. Figure 23 compares the residuals from the different cases over the time interval 101–200, where the experiments have reached a quasi-steady accuracy level following the initial assimilation updates. Note that using a finite ensemble size introduces a small uncertainty. We could have repeated the experiments with different random seeds to obtain more robust estimates and plotted all residuals with an uncertainty estimate, but this becomes more relevant when using smaller ensemble sizes than the 1000 realizations used in all the experiments in this paper.

In addition to the actual performance issues, there could also be preferences regarding the theoretical formulations of the different techniques. Could we assume that ESMDA solves for minuscule ES updates that introduce Gaussianity into the next window's prior? The recursive-in-time updating process in sequential data assimilation may also reduce the RML sampling approximation.

The main cost of running the different ensemble smoothers is associated with integrating the ensemble of model realizations in each update step. While the ES is the most cost efficient as it only computes a single update, ESMDA and IES improve the results but require a few, e.g., three to four iterations to converge. We also note that the shorter the assimilation window, the less nonlinear the problem becomes and the faster the convergence. It is good to realize that ESMDA with, say, 50 localized ensemble members and four iterations have costs similar to, or even less than, an ensemble of 4DVARs. A rough estimate of the computational costs is the number of model runs needed, assuming the actual assimilation update requires significantly less computation than the model runs. In that case, the cost is the number of ensemble members times the number of iterations used.

Additional sensitivity studies could involve varying ensemble sizes, but we have decided to exclude such a study in this

paper since it would also include using localization methods. In future studies, we also want to examine the combined impact of different temporal and spatial scales in the two model components.

7. Summary

In this study, we have demonstrated the possibility of using adjoint-free iterative ensemble methods for sequential data assimilation. We defined a coupled multiscale model system based on the nonlinear Kuramoto–Sivashinsky model. The model system allows for examining data assimilation in multiscale dynamical systems and models with unstable dynamics. An advantageous property of this model is its saturation of prediction uncertainty at a climatologic level, similar to what we observe in atmospheric and ocean models. In addition, the model is one dimensional in space, allowing us to represent spatial variability in the model. In addition to demonstrating the importance and value of coupled data assimilation updates, we have studied the properties of iterative ensemble smoothers. The joint assimilation of data for a model operating at different physical scales and with collective updates of both model components yields superior results compared to treating the model components and data independently. In particular, we demonstrated the importance of assimilating measurements from the component with the small spatial scales into the coupled system for reconstructing the component with the large spatial scales. We also explained the similarity between ensemble 4DVAR and our iterative smoothers. The main differences include the replacement of the tangent-linear operator with an ensemble-averaged model sensitivity, which eliminates the construction of tangent-linear and adjoint models. Additionally, we represent all covariance matrices with an ensemble of model realizations. We demonstrated the effectiveness and efficiency of ESMDA and IES in dealing with nonlinearities in a coupled model. Furthermore, we illustrated the effect of updating over the assimilation window versus updating the window's initial conditions in ensemble smoothers. We believe the methods and formulations, initially developed for petroleum applications, will be valuable for




FIG. 22. The plots show the time evolution of residuals for the ensemble prediction using different window lengths for a converged IES, an IES using only four iterations (IES4), and an ESM DA with five steps (MDA5). The dark lines indicate RMSE relative to the reference solution, while the light lines are the ensemble-predicted RMSE.




FIG. 23. The plots show time-averaged residuals for the IESX, IES4X, and ESMXA experiments as a function of window length for the period 101–200. The dark lines indicate RMSE relative to the reference solution, while the light lines are the ensemble-predicted RMSE.

future data assimilation systems for atmospheric, oceanic, and other coupled Earth system models.

Acknowledgments. Geir Evensen acknowledges financial support from the NORCE research project REMEDY, funded by industry partners Aker BP ASA, Equinor Energy AS, Wintershall Dea Norge AS, and the Research Council of Norway, Grant 336240. Peter Jan van Leeuwen was supported by the National Science Foundation, Grant 2220201.

Data availability statement. The code and input data with instructions to compile and run all experiments and plot the results presented in this paper are available from the GitHub repository (Evensen 2023a).

REFERENCES

- Amezcu, J., M. Goodliff, and P. J. Van Leeuwen, 2017: A weak-constraint 4D EnsembleVar. Part I: Formulation and simple model experiments. *Tellus*, **69**, 1271–1564, <https://doi.org/10.1080/16000870.2016.1271564>.
- Azouani, A., and E. S. Titi, 2014: Feedback control of nonlinear dissipative systems by finite determining parameters - A reaction-diffusion paradigm. *Evol. Equ. Control Theory*, **3**, 579–594, <https://doi.org/10.3934/eect.2014.3.579>.
- Bocquet, M., and P. Sakov, 2013a: An iterative ensemble Kalman smoother. *Quart. J. Roy. Meteor. Soc.*, **140**, 1521–1535, <https://doi.org/10.1002/qj.2236>.
- , and —, 2013b: Joint state and parameter estimation with an iterative ensemble Kalman smoother. *Nonlinear Processes Geophys.*, **20**, 803–818, <https://doi.org/10.5194/npg-20-803-2013>.
- Bonavita, M., P. Lean, and E. Holm, 2018: Nonlinear effects in 4D-Var. *Nonlinear Processes Geophys.*, **25**, 713–729, <https://doi.org/10.5194/npg-25-713-2018>.
- Burgers, G., P. J. van Leeuwen, and G. Evensen, 1998: Analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.*, **126**, 1719–1724, [https://doi.org/10.1175/1520-0493\(1998\)126<1719:ASITEK>2.0.CO;2](https://doi.org/10.1175/1520-0493(1998)126<1719:ASITEK>2.0.CO;2).
- Chang, H.-C., 1986: Nonlinear waves on liquid film surfaces—I. Flooding in a vertical tube. *Chem. Eng. Sci.*, **41**, 2463–2476, [https://doi.org/10.1016/0009-2509\(86\)80032-X](https://doi.org/10.1016/0009-2509(86)80032-X).
- Chen, Y., and D. S. Oliver, 2012: Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Math. Geosci.*, **44** (1), 1–26, <https://doi.org/10.1007/s11004-011-9376-z>.
- , and —, 2013: Levenberg–Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. *Comput. Geosci.*, **17**, 689–703, <https://doi.org/10.1007/s10596-013-9351-5>.
- Chang, Y.-S., S. Zhang, A. Rosati, T. L. Delworth, and W. F. Stern, 2013: An assessment of oceanic variability for 1960–2010 from the GFDL ensemble coupled data assimilation. *Climate Dyn.*, **40**, 775–803, <https://doi.org/10.1007/s00382-012-1412-2>.
- Chorin, A. J., and P. Krause, 2004: Dimensional reduction for a Bayesian filter. *Proc. Natl. Acad. Sci. USA*, **101**, 15 013–15 017, <https://doi.org/10.1073/pnas.0406222101>.
- de Rosnay, P., and Coauthors, 2022: Coupled data assimilation at ECMWF: Current status, challenges and future developments. *Quart. J. Roy. Meteor. Soc.*, **148**, 2672–2702, <https://doi.org/10.1002/qj.4330>.
- Emerick, A. A., 2018: Deterministic ensemble smoother with multiple data assimilation as an alternative for history-matching seismic data. *Comput. Geosci.*, **22**, 1175–1186, <https://doi.org/10.1007/s10596-018-9745-5>.
- , and A. C. Reynolds, 2012: History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations. *Comput. Geosci.*, **16**, 639–659, <https://doi.org/10.1007/s10596-012-9275-5>.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, 10 143–10 162, <https://doi.org/10.1029/94JC00572>.
- , 2003: The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.*, **53**, 343–367, <https://doi.org/10.1007/s10236-003-0036-9>.
- , 2009: *Data Assimilation: The Ensemble Kalman Filter*. 2nd ed. Springer, 307 pp., <https://doi.org/10.1007/978-3-642-03711-5>.
- , 2019: Accounting for model errors in iterative ensemble smoothers. *Comput. Geosci.*, **23**, 761–775, <https://doi.org/10.1007/s10596-019-9819-z>.
- , 2021: Formulating the history matching problem with consistent error statistics. *Comput. Geosci.*, **25**, 945–970, <https://doi.org/10.1007/s10596-021-10032-7>.
- , 2023a: EnKF_MS. GitHub, accessed 15 January 2024, https://github.com/geirev/EnKF_MS.
- , 2023b: EnKF_MS Kuramoto-Sivashinsky model. GitHub, accessed 15 January 2024, https://github.com/geirev/EnKF_MS/blob/main/src/m_model.F90.
- , and P. J. van Leeuwen, 2000: An ensemble Kalman smoother for nonlinear dynamics. *Mon. Wea. Rev.*, **128**,

- 1852–1867, [https://doi.org/10.1175/1520-0493\(2000\)128<1852:AEKSFN>2.0.CO;2](https://doi.org/10.1175/1520-0493(2000)128<1852:AEKSFN>2.0.CO;2).
- , and K. S. Eikrem, 2018: Conditioning reservoir models on rate data using ensemble smoothers. *Comput. Geosci.*, **22**, 1251–1270, <https://doi.org/10.1007/s10596-018-9750-8>.
- , P. N. Raanes, A. S. Stordal, and J. Hove, 2019: Efficient implementation of an iterative ensemble smoother for data assimilation and reservoir history matching. *Front. Appl. Math. Stat.*, **5**, 47, <https://doi.org/10.3389/fams.2019.00047>.
- , and Coauthors, 2021: An international initiative of predicting the SARS-CoV-2 pandemic using ensemble data assimilation. *Found. Data Sci.*, **3**, 413–477, <https://doi.org/10.3934/fods.2021001>.
- , F. C. Vossepoel, and P. J. van Leeuwen, 2022: *Data Assimilation Fundamentals: A Unified formulation for State and Parameter Estimation*. Springer, 254 pp., <https://doi.org/10.1007/978-3-030-96709-3>.
- Han, G., X. Wu, S. Zhang, Z. Liu, and W. Li, 2013: Error covariance estimation for coupled data assimilation using a Lorenz atmosphere and a simple pycnocline ocean model. *J. Climate*, **26**, 10 218–10 231, <https://doi.org/10.1175/JCLI-D-13-00236.1>.
- Jardak, M., I. M. Navon, and M. Zupanski, 2010: Comparison of sequential data assimilation methods for the Kuramoto–Sivashinsky equation. *Int. J. Numer. Methods Fluids*, **62**, 374–402, <https://doi.org/10.1002/ld.2020>.
- Kitanidis, P. K., 1995: Quasi-linear geostatistical theory for inverting. *Water Resour. Res.*, **31**, 2411–2419, <https://doi.org/10.1029/95WR01945>.
- Kuramoto, Y., 1978: Diffusion-induced chaos in reaction systems. *Prog. Theor. Phys. Suppl.*, **64**, 346–367, <https://doi.org/10.1143/PTPS.64.346>.
- Lak, M., 2023: Test case for PDEs: Kuramoto–Sivashinsky (KS). University of California, accessed 15 January 2024, <https://online.kitp.ucsb.edu/online/transturb17/gibson/html/5-kuramoto-sivashinsky.html>.
- Laloyaux, P., M. Balmaseda, D. Dee, K. Mogensen, and P. Janssen, 2016: A coupled data assimilation system for climate reanalysis. *Quart. J. Roy. Meteor. Soc.*, **142**, 65–78, <https://doi.org/10.1002/qj.2629>.
- Lunasin, E., and E. S. Titi, 2017: Finite determining parameters feedback control for distributed nonlinear dissipative systems -A computational study. *Evol. Equ. Control Theory*, **6**, 535–557, <https://doi.org/10.3934/eect.2017027>.
- Luo, X., and I. Hoteit, 2014: Ensemble Kalman filtering with a divided state-space strategy for coupled data assimilation problems. *Mon. Wea. Rev.*, **142**, 4542–4558, <https://doi.org/10.1175/MWR-D-13-00402.1>.
- Oliver, D. S., N. He, and A. C. Reynolds, 1996: Conditioning permeability fields to pressure data. *ECMOR V-Fifth European Conf. on the Mathematics of Oil Recovery*, Leoben, Austria, European Association of Geoscientists and Engineers, 259–270, <https://doi.org/10.3997/2214-4609.201406884>.
- Papageorgiou, D. T., and Y. S. Smyrlis, 1991: The route to chaos for the Kuramoto–Sivashinsky equation. *Theor. Comput. Fluid Dyn.*, **3**, 15–42, <https://doi.org/10.1007/BF00271514>.
- Penny, S. G., and Coauthors, 2017: Coupled data assimilation for integrated Earth system analysis and prediction: Goals, challenges and recommendations. Special Rep. WWRP 2017-3, 45 pp., <https://repository.library.noaa.gov/view/noaa/28431>.
- , E. Bach, K. Bhargava, C.-C. Chang, C. Da, L. Sun, and T. Yoshida, 2019: Strongly coupled data assimilation in multi-scale media: Experiments using a quasi-geostrophic coupled model. *J. Adv. Model. Earth Syst.*, **11**, 1803–1829, <https://doi.org/10.1029/2019MS001652>.
- Protas, B., T. R. Bewley, and G. Hagen, 2004: A computational framework for the regularization of adjoint analysis in multi-scale PDE systems. *J. Comput. Phys.*, **195**, 49–89, <https://doi.org/10.1016/j.jcp.2003.08.031>.
- Raanes, P. N., A. S. Stordal, and G. Evensen, 2019: Revising the stochastic iterative ensemble smoother. *Nonlinear Processes Geophys.*, **26**, 325–338, <https://doi.org/10.5194/npg-26-325-2019>.
- Saha, S., and Coauthors, 2010: The NCEP climate forecast system reanalysis. *Bull. Amer. Meteor. Soc.*, **91**, 1015–1058, <https://doi.org/10.1175/2010BAMS3001.1>.
- Shlang, T., and G. I. Sivashinsky, 1982: Irregular flow of a liquid film down a vertical column. *J. Phys. France*, **43**, 459–466, <https://doi.org/10.1051/jphys:01982004303045900>.
- Sivashinsky, G. I., 1977: Nonlinear analysis of hydrodynamic instability in laminar flames—I. Derivation of basic equations. *Acta Astronaut.*, **4**, 1177–1206, [https://doi.org/10.1016/0094-5765\(77\)90096-0](https://doi.org/10.1016/0094-5765(77)90096-0).
- , 1980: On flame propagation under conditions of stoichiometry. *SIAM J. Appl. Math.*, **39**, 67–82, <https://doi.org/10.1137/0139007>.
- Tadmor, E., 1986: The well-posedness of the Kuramoto–Sivashinsky equation. *SIAM J. Math. Anal.*, **17**, 884–893, <https://doi.org/10.1137/0517063>.
- Tilley, B. S., S. H. Davis, and S. G. Bankoff, 1994: Nonlinear long-wave stability of superposed fluids in an inclined channel. *J. Fluid Mech.*, **277**, 55–83, <https://doi.org/10.1017/S0022112094002685>.
- Tondeur, M., A. Carrassi, S. Vannitsem, and M. Bocquet, 2020: On temporal scale separation in coupled data assimilation with the ensemble Kalman filter. *J. Stat. Phys.*, **179**, 1161–1185, <https://doi.org/10.1007/s10955-020-02525-z>.
- van Leeuwen, P. J., 2020: A consistent interpretation of the stochastic version of the ensemble Kalman filter. *Quart. J. Roy. Meteor. Soc.*, **146**, 2815–2825, <https://doi.org/10.1002/qj.3819>.
- , and G. Evensen, 1996: Data assimilation and inverse methods in terms of a probabilistic formulation. *Mon. Wea. Rev.*, **124**, 2898–2913, [https://doi.org/10.1175/1520-0493\(1996\)124<2898:DAAIMI>2.0.CO;2](https://doi.org/10.1175/1520-0493(1996)124<2898:DAAIMI>2.0.CO;2).
- , Y. Cheng, and S. Reich, 2015: *Nonlinear Data Assimilation*. Frontiers in Applied Dynamical Systems: Reviews and Tutorials, Vol. 2, Springer, 118 pp., <https://doi.org/10.1007/978-3-319-18347-3>.
- Waller, J. A., S. L. Dance, A. S. Lawless, and N. K. Nichols, 2014: Estimating correlated observation error statistics using an ensemble transform Kalman filter. *Tellus*, **66**, 23 294, <https://doi.org/10.3402/tellusa.v66.23294>.
- Zhao, Y., F. Forouzanfar, and A. C. Reynolds, 2017: History matching of multi-facies channelized reservoirs using ES-MDA with common basis DCT. *Comput. Geosci.*, **21**, 1343–1364, <https://doi.org/10.1007/s10596-016-9604-1>.