

Data Assimilation of Coherent Structures in Parametric and Gridded Spaces: An Idealized Ocean Eddy Study

Solène Dealbera^{1,2}, Stéphane Raynaud³, Carlos Granero Belinchon^{1,2}, Brahim Boussidi⁴, Clément Le Goff⁴, and Pierre Tandeo^{1,2}

¹IMT Atlantique, Lab-STICC, UMR 6285, 29238, CNRS, Brest, France

²ODYSSEY Team-Project, INRIA Ifremer IMT-Atl., 35042, CNRS, Brest, France

³Shom, 29200 Brest, France

⁴LEGOS Observatoire Midi-Pyrénées, 14 avenue Édouard Belin, 31400 Toulouse, France

Correspondence: Solène Dealbera (solene.dealbera@imt-atlantique.fr)

Abstract. Data assimilation (DA) combines observations and model background to estimate the state of geophysical systems. Coherent structures, such as ocean eddies, are localized geophysical features characterized by a small set of properties: position, amplitude, and shape for instance. When linear and gaussian DA is applied directly to a gridded field without prior structural adjustment, the linear combination of background and observation can distort these properties, particularly when displacement uncertainty dominates. In this work, we explore a DA framework of coherent structures in a reduced parametric space, where the state is defined by a set of parameters describing the structure. Using an idealized one-dimensional ocean eddy, we compare single-step analyses performed both in gridded and parametric spaces. Three tutorial experiments with prescribed parametric uncertainties illustrate this distortion mechanism directly. A fourth experiment, starting from noisy gridded observations, provides a more balanced evaluation using the inverse nonlinear mapping from grid to parameter space. In these idealized experiments, the parametric approach maintains a physically consistent shape after assimilation. The parametric DA requires reliable parameter extraction, and its benefits depend on the number of structures relative to grid resolution.

1 Introduction

Data assimilation (DA) is widely used in geosciences to combine observations with information from a dynamical model (the background) in a statistically optimal way (Carrassi et al., 2018). In its standard formulation, DA aims to estimate the full system state on a spatial grid.

In this study, we do not focus on the full gridded state, but rather on the estimation of specific geophysical structures (GSs), such as ocean eddies or cyclones, viewed as coherent objects characterized by a small set of parameters. When classic, linear-Gaussian DA methods, such as optimal interpolation or the Kalman filter, are applied directly to a gridded field without prior structural adjustment, the presence of a coherent structure can cause the linear combination of background and observation to distort the reconstructed structure: spatially displaced or differently-shaped realizations are averaged rather than reconciled, which smears the structure whenever displacement or shape, rather than amplitude, uncertainty dominates (Chen and Snyder, 2007; Feyeux et al., 2018).

Several strategies have been proposed to introduce structural information explicitly into the assimilation problem, spanning a spectrum from field-based to more object-oriented representations. A first family augments the state vector with structure parameters directly: (Chen and Snyder, 2007) added vortex position to the ensemble Kalman filter (EnKF) state, allowing the filter to separate displacement from amplitude errors. This was extended to tropical cyclones by including position, translation speed, and axisymmetric wind structure (Wu et al., 2010), and further to operationally relevant track, wind-radius, and intensity parameters (Kunii, 2015). These approaches define state variables that are directly interpretable as object properties, though they remain tied to a predefined parametric model, in a spirit close to the parametric representation adopted here. A second family instead treats structural errors through geometric transformations or alternative measures of distance in the field representation. Alignment and morphing approaches modify the field geometry to reconcile displaced or differently shaped structures (Ravela et al., 2007; Beezley and Mandel, 2008; Zhen et al., 2025), while displacement data assimilation explicitly estimates the displacement field required for such alignment while enforcing area preservation (Rosenthal et al., 2017). Alternatively, optimal-transport-based approaches modify the metric used to compare background and observations, using the Wasserstein distance to better account for positional errors of coherent structures (Feyeux et al., 2018). A third family operates on adaptive or moving computational meshes that track coherent structures rather than resolving the full domain at uniform resolution: (Aydøğdu et al., 2019) developed ensemble DA for adaptive, non-conservative moving-mesh models, and mesh-node positions have since been included directly within the assimilation state itself (Sampson et al., 2021), though in this case the assimilated geometry reflects the numerical discretization rather than the structure’s own properties. Feature-informed and Lagrangian-based approaches extract collective descriptors of coherent structures directly from observations, rather than assimilating raw trajectories or fields (Srivastava et al., 2023; Maclean et al., 2017). Non-Gaussian methods such as particle filters (van Leeuwen, 2010; Poterjoy, 2016) can, in principle, handle the nonlinearity associated with displacement and shape uncertainty; part of the difficulty in relating a gridded and a reduced representation, however, lies not solely in this nonlinearity, but in correctly marginalizing, rather than merely interpolating, uncertainty between the two spaces, a distinction examined further in Sect. 2.1.

Correctly linking uncertainty across representation spaces is itself motivated by a broader principle: Rao-Blackwellization (Casella and Robert, 1996) formalizes the idea of analytically marginalizing a conditionally-linear-Gaussian subset of a state while treating the remaining nonlinear component separately, seeking a representation in which a problem becomes, and remains, Gaussian-tractable. This principle underlies Rao-Blackwellized particle filters (Chen and Liu, 2000), including applications to environmental data assimilation that combine Kalman-type updates with particle-based treatment of low-dimensional non-Gaussian parameters (Šmídl and Hofman, 2011). The parametric transformation explored in this study follows this approach: rather than operating directly on the full gridded field, it seeks a lower-dimensional, physically interpretable representation in which relating background and observation uncertainty is more tractable, provided this relationship is established by propagating the full uncertainty through the mapping between spaces, rather than by a local linear approximation of it.

The goal of this study is to investigate, in a controlled and idealized framework, how the choice of representation space influences the behavior of a standard Gaussian linear DA update when applied to coherent structures with uncertain position

and shape. In particular, we compare assimilation performed in the classic gridded space with assimilation performed directly in a reduced parametric space describing the structure.

The idea of working in a low-dimensional parameter space rather than on a full-state vector has a long history in the control and stochastic filtering literature (Jazwinski, 1970). However, in these frameworks, the parameters to be estimated are typically internal coefficients of a dynamical model, rather than geometric descriptors of coherent structures. The present work applies a related principle to an explicitly geometric and physically interpretable representation of a coherent structure, defined here by amplitude, position and radius.

This parametric perspective is also related to reduced-order and subspace approaches, where the system state is represented in a lower-dimensional space, such as empirical orthogonal functions or learned latent variables (Hannachi et al., 2007; Alvera-Azcárate et al., 2025). In contrast to these approaches, which define reduced coordinates in an abstract basis, the present work considers an explicit physically interpretable parameterization of the structure. This allows a direct interpretation of assimilation updates in terms of physically meaningful geometric properties.

We emphasize that this framework does not aim to replace existing Lagrangian, feature-based, or non-Gaussian DA methods, but rather to provide a simplified setting to analyze how standard Gaussian updates behave when applied either in gridded space or in a low-dimensional parameter space representing coherent structures. To the best of our knowledge, this specific comparison of the assimilation behavior between gridded space and an explicitly defined geometric parameter space of a coherent structure has not been analyzed in the literature. In this sense, the study is primarily methodological and diagnostic, focusing on the mechanism of error representation rather than on operational forecasting performance.

In the proposed framework, the state of the system is fully described by the parameters of a single idealized eddy, which allows a direct analysis of how uncertainties in amplitude, position, and radius propagate through the assimilation step in both representations. The computational implications of this lower-dimensional representation are discussed in Sect. 4.4.

We present a set of idealized numerical experiments designed to compare a single analysis step performed in the classic gridded space and in this reduced parametric space. The eddy representation and assimilation methodology are detailed in Sect. 2, the experimental setup in Sect. 3, and the results are presented in Sect. 4. Conclusions and perspectives are discussed in Sect. 5.

2 Methodology

The objective of this paper is to compare two methods of blending two sources of information, a background and an observation, using a basic DA algorithm. We focus exclusively on the analysis step and do not perform repeated assimilation cycles (forecast-analysis-reanalysis). The stability of the system over multiple cycles is beyond the scope of this study and will be investigated in future work.

We consider a simplified one-dimensional cross-shore section of an ocean eddy, described by three parameters: peak amplitude a , peak position p , and maximal velocity radius r , such that $\mathbf{x}^{\text{param}} = (a, p, r)$. In practice, these parameters can be

Table 1. Symbols and associated spaces.

Symbol	Meaning	Space
$\mathbf{x}^{t,\text{param}}$	True state vector	param
$\mathbf{x}^{t,\text{grid}}$	True state vector	grid
$\mathbf{x}^{b,\text{param}}$	Background state vector	param
$\mathbf{x}^{b,\text{grid}}$	Background state vector	grid
$\mathbf{x}^{a,\text{param}}$	Analysis state vector	param
$\mathbf{x}^{a,\text{grid}}$	Analysis state vector	grid
$\mathbf{y}^{\text{param}}$	Observation vector	param
$\mathbf{y}^{\text{grid}}$	Observation vector	grid
$\mathbf{B}^{\text{param}}$	Background error covariance	param
$\mathbf{B}^{\text{grid}}$	Background error covariance	grid
$\mathbf{R}^{\text{param}}$	Observation error covariance	param
$\mathbf{R}^{\text{grid}}$	Observation error covariance	grid
$h(\cdot; \cdot)$	Mapping from param to grid	n/a

estimated using standard eddy detection and fitting algorithms, such as local maxima identification combined with Gaussian profile fitting (e.g., Chelton et al., 2011).

For our idealized experiments, however, the parameters are prescribed manually rather than obtained from a detection algorithm, in order to control uncertainties and systematically study the assimilation behavior. The physical profile along the cross-section is given by the mapping

$$\mathbf{x}^{\text{grid}} = h(\mathbf{lon}; \mathbf{x}^{\text{param}}) = a \exp \left(-\frac{1}{2} \left(\frac{\mathbf{lon} - p}{r} \right)^2 \right), \quad (1)$$

where $\mathbf{lon}$ is the vector of longitudinal positions along the cross-section. Here, h serves as a mapping function from the parametric space $\mathbf{x}^{\text{param}}$ to the physical gridded space $\mathbf{x}^{\text{grid}}$. It is neither an observation operator nor a numerical dynamical model, but simply defines the physical profile corresponding to a given set of parameters. We note that the Gaussian shape assumed here for the eddy profile is a modeling choice, distinct from the Gaussian statistical assumptions made elsewhere in this study (e.g. for the background and observation error distributions); every detected or analyzed eddy is consequently constrained to this Gaussian shape by construction. The framework could in principle be generalized to other parametric profile shapes, which would only change the form of h and its Jacobian in Eq. (12) and Sect. 2.1.

Two spaces are considered for the eddy representation: the *parametric space* (`param`), of dimension $n = 3$ corresponding to (a, p, r) , and the *gridded space* (`grid`), of dimension m equal to the length of the longitude vector. The notation used throughout this manuscript is summarized in Table 1.

The background state in each space is modeled as a perturbation around a reference value:

$$\mathbf{x} = \mathbf{x}^b + \boldsymbol{\eta}^b, \quad (2)$$

where $\mathbf{x}^b$ is the reference background state and $\boldsymbol{\eta}^b \sim \mathcal{N}(\mathbf{0}, \mathbf{B})$ is the background error, assumed Gaussian with covariance matrix $\mathbf{B}$. The observation is related to the true state by:

$$\mathbf{y} = \mathcal{H}(\mathbf{x}) + \varepsilon, \quad (3)$$

110 where $\mathcal{H}$ is the observation operator and $\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ is the observation error with covariance $\mathbf{R}$. Equations (2) and (3) are written generically and apply separately in each space with the appropriate superscript (`param` or `grid`) on all variables (see Table 1).

In our case, since the state vector is fully observed in both spaces, the operator $\mathcal{H}$ is the identity matrix, leading to a simple linear observation equation.

115 The data assimilation is performed in each space with the appropriate superscript (`param` or `grid`) on all variables using the following cost function J :

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \frac{1}{2}(\mathbf{y} - \mathbf{x})^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{x}). \quad (4)$$

Under the assumptions of linearity and Gaussian errors, the minimizer of J is given analytically by the Best Linear Unbiased Estimator (BLUE). No iterative minimization algorithm is therefore required; the solution is computed directly for each configuration. Let $\mathbf{X}^b = [\mathbf{x}_1^b, \dots, \mathbf{x}_N^b]$ and $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_N]$ denote the background and observation ensemble matrices, each of size
120 $n \times N$ (or $m \times N$ in `grid` space), where N is the ensemble size. The covariance matrices are estimated from these ensembles as:

$$\mathbf{B} = \text{Cov}(\mathbf{X}^b, \mathbf{X}^b), \text{ and } \mathbf{R} = \text{Cov}(\mathbf{Y}, \mathbf{Y}), \quad (5)$$

following the Ensemble Optimal Interpolation (EnOI) framework. The BLUE update of each ensemble member i reads:

$$125 \quad \mathbf{x}_i^a = \mathbf{x}_i^b + \mathbf{K}(\mathbf{y}_i - \mathbf{x}_i^b), \quad i = 1, \dots, N, \quad (6)$$

from which the analysis ensemble mean $\bar{\mathbf{x}}^a = \frac{1}{N} \sum_i \mathbf{x}_i^a$ is derived, and the analysis error covariance $\mathbf{P}^a$ is given by:

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{K})\mathbf{B}, \quad (7)$$

where $\mathbf{K} = \mathbf{B}(\mathbf{B} + \mathbf{R})^{-1}$ is the Kalman gain. Equations (5)-(7) apply separately in each space with the appropriate superscript (`param` or `grid`) on all variables.

130 Two experimental frameworks are considered, differing in how the uncertainty is introduced and how the covariance matrices are specified, as described in the following subsections. In the tutorial framework (Sect. 2.1), the uncertainty originates in parameter space and is propagated to grid space via h . In the realistic framework (Sect. 2.2), the uncertainty originates in grid space and must be inverted back to parameter space through a fitting procedure, providing a more balanced end-to-end comparison.

This framework is designed as a controlled tutorial example to isolate the mechanism by which linear Gaussian updates can distort coherent structures in gridded space. It is not intended as a balanced end-to-end comparison of practical DA workflows, since the uncertainty is generated in parameter space and then mapped to grid space, which inherently favors the parametric approach.

140 In this framework, the parameters are assumed to be fully independent and to have equal uncertainties for both the background and the observation. The error covariance matrices in parameter space are therefore diagonal:

$$\mathbf{B}^{\text{param}} = \mathbf{R}^{\text{param}} = \text{diag}(\sigma_a^2, \sigma_p^2, \sigma_r^2) = \sigma^2 \mathbf{I}_n, \quad (8)$$

where $\sigma = (\sigma_a, \sigma_p, \sigma_r)$ contains the prescribed standard deviations for amplitude, position and radius, respectively.

Ensemble members in `param` space are generated by independently sampling each parameter from a Gaussian distribution
 145 around its reference value with the prescribed standard deviation. The corresponding ensemble members in `grid` space are then obtained by applying the mapping h (Eq. 1) to each member, so that the gridded covariance matrices $\mathbf{B}^{\text{grid}}$ and $\mathbf{R}^{\text{grid}}$ are fully determined by the nonlinear push-forward of the parametric uncertainties. As a result, these gridded covariances are in general non-diagonal and reflect the nonlinear geometry of the eddy profile, even though the parametric uncertainties are uncorrelated.

150 An analytic alternative to this ensemble-based propagation would be a first- or second-order Taylor expansion of h around the reference parameter vector. We tested this local linearization and found it inadequate in the nonlinear regime relevant to this study. Because $\partial h / \partial p$ vanishes at the profile peak by symmetry, the Taylor-propagated standard deviation drops to zero with a slope discontinuity at that location for both prior covariance matrices, $\mathbf{B}^{\text{grid}}$ and $\mathbf{R}^{\text{grid}}$, each centered on its own reference position ($p = -5$ and $p = +5$ km, respectively, in Experiment #2, see Sect. 3.1). In contrast, the Monte Carlo
 155 ensemble produces a smooth, non-zero local minimum there, reflecting the non-Gaussian structure induced by a displacement comparable to the eddy radius. More critically, when the background and observation are linearized separately around distinct mean positions, the resulting covariance matrices are rank-deficient and mutually asymmetric; the corresponding Kalman gain produces an unphysical negative undershoot in the analyzed height over a substantial fraction of the domain, and the posterior covariance to collapse to numerically negligible values almost everywhere (Appendix A, Fig. A1). The ensemble (Monte Carlo)
 160 approach described above does not exhibit this pathology and was therefore retained throughout this study.

Figure 1 illustrates the overall workflow comparing assimilation in the parametric and gridded spaces.

2.2 Realistic framework

In realistic applications, observations are available as gridded fields rather than direct eddy parameters. In this second frame-
 work, both the background and the observation start as noisy gridded profiles, and the parametric representation must be
 165 inferred from them through a fitting procedure. This constitutes a more balanced end-to-end test, since the uncertainty now originates in grid space and the inverse nonlinear mapping from grid to parameter space is explicitly accounted for.




Figure 1. Schematic of the data assimilation framework comparing gridded space (left column, red) and parametric space (right column, purple). Top row: background (green) and observation (blue) eddies represented as gridded profiles (top left) and as parameter vectors (top right). The *profile fit* arrow indicates the nonlinear least-squares mapping from gridded observations to parameter space, as used in Experiment #4. Bottom row: analysis state from the analysis step in gridded space $\mathbf{x}^{a,\text{grid}} = [h(\text{lon}_1), \dots, h(\text{lon}_m)]$ (bottom left) and in parametric space $\mathbf{x}^{a,\text{param}} = [a, p, r]$ (bottom right). The *mapping* arrow indicates the projection of the parametric analysis back to gridded space via h (Eq. 1).

The observed and background gridded fields are defined as:

$$\mathbf{y}^{\text{grid}} = h(\mathbf{lon}; \mathbf{y}^{\text{param},*}) + \boldsymbol{\varepsilon}^{\text{grid}}, \quad (9)$$

$$\mathbf{x}^{b,\text{grid}} = h(\mathbf{lon}; \mathbf{x}^{b,\text{param},*}) + \boldsymbol{\eta}^{b,\text{grid}}, \quad (10)$$

170 where $\mathbf{y}^{\text{param},*}$ and $\mathbf{x}^{b,\text{param},*}$ are the reference observation and background parameter vectors, and $\boldsymbol{\varepsilon}^{\text{grid}} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}^{\text{grid}})$, $\boldsymbol{\eta}^{b,\text{grid}} \sim \mathcal{N}(\mathbf{0}, \mathbf{B}^{\text{grid}})$ are the respective gridded noise terms. The gridded covariance matrices are prescribed using an exponentially decaying correlation model:

$$\mathbf{B}_{ij}^{\text{grid}} = \mathbf{R}_{ij}^{\text{grid}} = \sigma_h^2 \exp\left(-\frac{|\text{lon}_i - \text{lon}_j|}{\lambda}\right), \quad (11)$$

175 where σ_h^2 is the surface-height variance and λ is the correlation length scale. In the gridded DA approach, $\mathbf{y}^{\text{grid}}$ is assimilated directly using Eqs. (6)-(7) with these prescribed covariances.

For the parametric DA approach, the parameter vectors are first estimated from the noisy gridded fields through a nonlinear least-squares fit of the Gaussian profile h . Specifically, for a given noisy gridded profile $\mathbf{u}$, the estimated parameter vector $\hat{\boldsymbol{\theta}}$

minimizes:

$$\mathcal{L}(\boldsymbol{\theta}) = \|\mathbf{u} - h(\mathbf{lon}; \boldsymbol{\theta})\|_2^2. \quad (12)$$

180 The minimization is performed using the `scipy.optimize.curve_fit` algorithm with bounds $a \in [0, 10]$ m, $p \in [-50, 50]$ km, and $r \in [0.1, 50]$ km, ensuring physical solutions (positive amplitude and positive radius). The initial values for the fit are set as following: $[\max \mathbf{u}, \mathbf{lon}_{\arg \max \mathbf{u}}, \frac{1}{2} \|\mathbf{lon}_{\arg \max \mathbf{u}} - \mathbf{lon}_{\arg \min \mathbf{u}}\|]$ where the amplitude is initialized from the maximum value of the profile, the center from the longitude at which this maximum occurs, and the width from half the distance between the longitudes of the maximum and minimum values. This simplified initialization is appropriate given the idealized framework of
185 the present experiment. In practice, initial values could be provided by a prior detection algorithm.

The fitting procedure is repeated over an ensemble of N noisy realizations drawn from Eqs. (9)-(10). The parametric covariance matrices are then estimated as the sample covariance of the fitted parameter ensembles, determined by the nonlinear push-forward of the gridded height uncertainties. Unlike the tutorial framework, these covariance matrices are in general non-diagonal, reflecting the nonlinear coupling between parameters introduced by the fitting procedure. The parametric DA is then
190 performed using Eqs. (6)-(7) in `param` space with these estimated covariances.

3 Experiments

Four experiments are presented, organized into two sets corresponding to the two frameworks described in Sect. 2. Experiments #1, #2 and #3 follow the tutorial framework (Sect. 2.1) and are designed to isolate the effect of a single uncertain parameter on the assimilation outcome. Experiment #4 follows the realistic framework (Sect. 2.2) and provides a more balanced end-to-end
195 evaluation by starting from noisy gridded observations. The parameter values for all four experiments are given in Table 2.

3.1 Tutorial experiments (#1, #2, #3)

In each of the three tutorial experiments, a single parameter carries the uncertainty, while the other two are set to their true values and are identical in the background and the observation. The uncertain parameter is, respectively: amplitude a (Experiment #1), position p (Experiment #2), and radius r (Experiment #3). This design isolates the effect of each parameter's nonlinearity on
200 the assimilation update (see Table 2).

This setup inherently favors the parametric approach (Sect. 2.1); these experiments therefore isolate the mechanism by which a nonlinearly entering parameter causes gridded-space updates to smear the structure, rather than providing a balanced workflow comparison.

The background and observation are constructed symmetrically around a true state $\mathbf{x}^{t, \text{param}}$ by applying shifts of equal
205 magnitude and opposite sign:

$$\mathbf{x}^{b, \text{param}} = \mathbf{x}^{t, \text{param}} - \boldsymbol{\delta}, \text{ and } \mathbf{y}^{\text{param}} = \mathbf{x}^{t, \text{param}} + \boldsymbol{\delta}, \quad (13)$$

Table 2. True parameter values $\mathbf{x}^{t,\text{param}}$ and reference parameter values with prescribed standard deviations ($\pm\sigma$) for the background and observation in each experiment. The background and observation are constructed symmetrically around the truth via Eq. (13). For Experiment #4, no standard deviation is prescribed in parameter space; the uncertainty is instead introduced as additive gridded noise with amplitude $\sigma_h = 0.2$ m and correlation length $\lambda = r/2 = 3.5$ km (see Sect. 2.2).

Exp.	Eddy	a [m]	p [km]	r [km]
#1	Truth	1.5	0	7
	Background	1.0 ± 0.7	0	7
	Observation	2.0 ± 0.7	0	7
#2	Truth	1.0	0	7
	Background	1.0	-5 ± 7	7
	Observation	1.0	5 ± 7	7
#3	Truth	1.0	0	11
	Background	1.0	0	8 ± 4
	Observation	1.0	0	14 ± 4
#4	Truth	1.0	0	7
	Background	1.0	-5	7
	Observation	1.0	5	7

where δ is the prescribed shift vector, with only the uncertain parameter having a non-zero component. The associated true gridded profile is $\mathbf{x}^{t,\text{grid}} = h(\mathbf{lon}; \mathbf{x}^{t,\text{param}})$. This symmetric construction ensures that the background and observation errors are equal in magnitude, so that the expected BLUE analysis in parameter space recovers the truth exactly: $\overline{\mathbf{x}}^{a,\text{param}} = \mathbf{x}^{t,\text{param}}$.

210 An ensemble of $N = 10000$ members is used for both the background and the observation. In `param` space, ensemble members are generated by independently sampling each parameter from a Gaussian distribution around its reference value with the prescribed standard deviation σ_a , σ_p , or σ_r (the non-perturbed parameters having zero spread). The corresponding `grid` ensemble members are obtained by applying h (Eq. 1) to each parametric member, as described in Sect. 2.1. The radius is clipped to $r \geq 10^{-6}$ km to prevent non-physical values.

215 3.2 Realistic experiment (#4)

Experiment #4 uses the same reference parameter values as Experiment #2 (position perturbation), but the uncertainty is now introduced directly in grid space following the realistic framework of Sect. 2.2. The background and observation gridded fields are generated as:

$$\mathbf{x}^{b,\text{grid}} = h(\mathbf{lon}; (-5, 1, 7)) + \boldsymbol{\eta}^{b,\text{grid}}, \text{ and } \mathbf{y}^{\text{grid}} = h(\mathbf{lon}; (5, 1, 7)) + \boldsymbol{\epsilon}^{\text{grid}}, \quad (14)$$

220 where the noise terms are drawn from the exponentially correlated distribution of Eq. (11) with $\sigma_h = 0.2$ m and $\lambda = r/2 = 3.5$ km.

For the parametric DA, an ensemble of $N = 10000$ noisy realizations is generated for both the background and the observation. Each realization is fitted independently using the procedure described in Sect. 2.2, yielding an ensemble of parameter vectors from which $\mathbf{B}^{\text{param}}$ and $\mathbf{R}^{\text{param}}$ are estimated. To assess the robustness of the results with respect to the noise realization, the full experiment is repeated $n_{\text{iter}} = 10$ times with independent noise draws; quantitative diagnostics are reported as means and standard deviations over these 10 noise realizations.

3.3 Evaluation

For all four experiments, two comparisons are performed:

- *Surface height reconstruction*: The analysis mean is projected into `grid` space and compared to the true profile $\mathbf{x}^{t,\text{grid}}$. For the parametric analysis, the projection is done by applying h (Eq. 1) to the analyzed parameter vector $\bar{\mathbf{x}}^{a,\text{param}}$. For the gridded analysis, the profile $\bar{\mathbf{x}}^{a,\text{grid}}$ is used directly. Two metrics are reported: the Root Mean Square Error (RMSE) of the analysis mean relative to $\mathbf{x}^{t,\text{grid}}$, and the Continuous Ranked Probability Score (CRPS) of the analysis ensemble relative to $\mathbf{x}^{t,\text{grid}}$, which jointly measures accuracy and sharpness of the probabilistic forecast.
- *Parameter estimation*: The estimated parameters of each analysis are compared to the true values $\mathbf{x}^{t,\text{param}}$. For the parametric analysis, the estimated parameters are directly $\bar{\mathbf{x}}^{a,\text{param}}$. For the gridded analysis, the analyzed profile is fitted with the Gaussian function h using the least-squares procedure of Eq. (12), yielding estimated parameters $\hat{\boldsymbol{\theta}}^{\text{grid}}$. The absolute error of each analysis relative to the truth is reported for each parameter separately, defined as $|\hat{\theta}_k^{\text{grid}} - x_k^{t,\text{param}}|$ and $|\bar{x}_k^{a,\text{param}} - x_k^{t,\text{param}}|$ for $k \in \{a, p, r\}$.

4 Results and discussion

4.1 Tutorial experiments (#1, #2, #3)

The surface height profiles are shown in Figs. 2, 3 and 4 for the amplitude, position and radius perturbation experiments, respectively. The quantitative diagnostics for all experiments are reported in Tables 3 and 4.

4.1.1 Experiment #1 — Amplitude perturbation.

Figure 2a shows that the analyzed profiles are identical in both `param` and `grid` spaces: the analyzed amplitude is 1.5 m, the exact mean of the background (1.0 m) and observation (2.0 m), with position and radius unchanged. Figure 2b shows identical a posteriori standard deviation profiles in both spaces. This agreement reflects the fact that amplitude is a linear parameter of h (Eq. 1): the nonlinear push-forward that distinguishes the two approaches has no effect in this case.

This is confirmed by the metrics in Table 3: the RMSE is 0.00 in both spaces, and the CRPS values are 0.03. The absolute error is 0.00 for all parameters (see Table 4).




Figure 2. Results for Experiment #1 (amplitude perturbation). Left: true surface height (black), mean surface height of the observed (blue), background (green), gridded analysis (red), and parametric analysis (dashed purple) eddies along the longitude. Right: corresponding a posteriori standard deviation $\sigma_h = \sqrt{\text{diag}(\mathbf{P}^a)}$.

250 4.1.2 Experiment #2 — Position perturbation.

By contrast, Fig. 3a shows a clear difference between the two analyses. The `grid` analyzed profile is widened and flattened relative to both the background and the observation. A Gaussian fit of this profile yields $(a, p, r) \approx (0.77, 0.02, 9.05)$ m/km/km: the amplitude is underestimated by 23% and the radius overestimated by 29% relative to the true values, even though the position is correctly recovered. This deformation is a direct consequence of the nonlinear dependence of h on position: the
 255 linear Gaussian update in grid space averages the two displaced profiles, producing a broader, lower peak, rather than a correctly shaped one.

In contrast, the `param` analysis recovers $(a, p, r) = (1.0, 0.0, 7.0)$ m/km/km, the exact midpoint in parameter space, with all three parameters correctly estimated. The a posteriori standard deviation in `grid` space (Fig. 3b) is also spatially underestimated relative to the `param` counterpart, indicating that the gridded analysis is overconfident about the reconstructed profile.

260 These differences are captured by the metrics: the RMSE increases to 0.10 m in the `grid` space, while it remains at 0.00 m in the `param` space. The CRPS is slightly better in the `param` space. The absolute error shows higher values for the estimation of the amplitude (0.24 m) and the radius (2.15 km) in the `grid` space, while the position is well-estimated (Table 4).

4.1.3 Experiment #3 — Radius perturbation.

Figure 4a shows that the `grid` analyzed profile has a narrower peak with a wider tail, compared to the `param` profile. A
 265 Gaussian fit yields $(a, p, r) \approx (0.96, 0.00, 11.09)$ m/km/km: the radius is correctly recovered (11.09 km, close to the expected mean of 11 km), but the amplitude is slightly underestimated (0.96 m vs. 1.0 m). More importantly, the `grid` analyzed profile is no longer strictly Gaussian: the profile shape is sub-Gaussian, with a narrower core and heavier tails than a Gaussian of the




Figure 3. Results for Experiment #2 (position perturbation). Same color and panel convention as Fig. 2.




Figure 4. Results for Experiment #3 (radius perturbation). Same color and panel convention as Fig. 2.

same standard deviation, a consequence of the nonlinear averaging of two Gaussian profiles of different widths in grid space. In the `param` space, the estimated radius is 11 km (the mean of 8 and 14 km) and the profile remains Gaussian. The `grid` RMSE is slightly higher (0.03 m), while the CRPS value is close (0.02 for the `grid` space and 0.03 for the `param` space).

4.2 Realistic experiment (#4)

Figure 5 shows the results for Experiment #4. The qualitative behavior is consistent with Experiment #2: the `grid` analysis produces a widened and underestimated profile (Gaussian fit: $a \approx 0.79$ m, $p \approx 0.04$ km, $r \approx 8.96$ km), while the `param` analysis recovers parameters close to the expected mean. These values are stable across the 10 noise realizations.

Table 3. Gridded field diagnostics for all experiments. RMSE: root-mean squared error of the analysis mean relative to the true profile $\mathbf{x}^{t,\text{grid}}$ (lower is better); CRPS: continuous ranked probability score of the analysis ensemble relative to $\mathbf{x}^{t,\text{grid}}$ (lower is better). For Experiment #4, values are means over 10 noise realizations.

Experiment	Space	RMSE [m]	CRPS
#1	grid	0.00	0.03
	param	0.00	0.03
#2	grid	0.10	0.05
	param	0.00	0.04
#3	grid	0.03	0.02
	param	0.00	0.03
#4	grid	0.09	0.05
	param	0.02	0.02

Table 4. Parameter estimation diagnostics for all experiments. Values are the signed errors of the analyzed parameters relative to the truth $\mathbf{x}^{t,\text{param}}$: $|\hat{\theta}_k^{\text{grid}} - x_k^{t,\text{param}}|$ for the gridded analysis (after Gaussian fitting) and $|\bar{x}_k^{a,\text{param}} - x_k^{t,\text{param}}|$ for the parametric analysis, for $k \in \{a, p, r\}$. For Experiment #4, values are means over 10 noise realizations.

Experiment	Space	Amplitude	Position	Radius
		[m]	[km]	[km]
#1	grid	0.00	0.00	0.00
	param	0.00	0.00	0.00
#2	grid	0.24	0.01	2.15
	param	0.00	0.04	0.00
#3	grid	0.00	0.00	0.01
	param	0.00	0.00	0.09
#4	grid	0.21	0.00	1.94
	param	0.02	0.00	0.24

275 The a posteriori standard deviations differ from Experiment #2 (Fig. 5b). In the `grid` space, the standard deviation is spatially uniform at approximately $0.2/\sqrt{2}$ m, reflecting the prescribed exponentially correlated noise structure. In the `param` space, the standard deviation is spatially varying, with larger values in the region spanned by the background and observation positions: this spatial structure directly reflects the positional uncertainty of the eddy, and is a distinctive feature of the parametric representation.

280 The mean correlation matrix of the fitted parameters, estimated over the 10 noise realizations, is reported in Table 5. The dominant off-diagonal correlation is between amplitude a and radius r (-0.47 ± 0.01): wider fits tend to produce lower peak amplitudes, reflecting a geometric coupling introduced by the nonlinear fitting procedure. The correlation between position




Figure 5. Results for Experiment #4 (realistic framework, position perturbation). Same color and panel convention as Fig. 2. Results are shown for a representative single noise realization; quantitative diagnostics in Tables 3 and 4 are means over 10 noise realizations.

Table 5. Mean parameter correlation matrix $\overline{\mathbf{P}}^{a,\text{param}}$ estimated from the fitting ensemble, averaged over 10 noise realizations (Experiment #4). Values are mean and standard deviations across realizations.

	a	p	r
a	1.00 ± 0.00	0.00 ± 0.01	-0.47 ± 0.01
p	0.00 ± 0.01	1.00 ± 0.00	-0.01 ± 0.02
r	-0.47 ± 0.01	-0.01 ± 0.02	1.00 ± 0.00

p and the other parameters is negligible ($< |0.01|$), consistent with the symmetry of the Gaussian profile. This non-diagonal covariance structure means that $\mathbf{B}^{\text{param}}$ and $\mathbf{R}^{\text{param}}$ are no longer diagonal, unlike in the tutorial experiments.

285 After assimilation, the posterior covariance $\mathbf{P}^{a,\text{param}}$ shows variances approximately halved relative to the prior, consistent with the expected uncertainty reduction for equal background and observation errors. Diagnostic tests on the fitted parameter ensembles further reveal that this fitting-induced structure is not fully Gaussian: the radius parameter, in particular, shows significant skewness and excess kurtosis (Shapiro-Wilk normality test $p < 10^{-30}$), although the ensemble mean used by the EnOI update exhibits only a small bias relative to the generating value.

290 The metrics for Experiment #4 (Tables 3 and 4) confirm the qualitative observations. The RMSE increases to 0.09 m in `grid` (compared to 0.02 m in the `param` space). The CRPS in both spaces remains low, with still a higher difference between the `grid` space (0.05) and the `param` space (0.02). The parameter estimation in the `grid` space shows similar values compared to Experiment #2, with slightly improved performance (0.21 m instead of 0.24 m for the amplitude and 1.94 km instead of 2.15 km for the radius). Compared with the analytically prescribed covariance, estimating the parameter covariance from a
295 noisy fitting procedure results in larger errors for the amplitude (0.02 m) and radius (0.24 km).

Table 6. EKE [$kJ.m^{-2}$]: total integrated eddy kinetic energy computed from the analysis mean profile, reported as a physical illustration. For Experiment #4, values are means over 10 noise realizations.

Experiment	Truth	grid	param
#1	137	137	137
#2	607	275	607
#3	386	356	387
#4	607	304	600

Overall, Experiment #4 shows that the parametric approach recovers the coherent structure properties when starting from noisy gridded observations.

4.3 Physical illustration: Eddy Kinetic Energy

The profile deformations observed in Experiments #2, #3 and #4 have consequences for the physical properties of the reconstructed eddy. As an illustration of this effect, we compute the geostrophic velocity and associated Eddy Kinetic Energy (EKE) from the analyzed height profiles. The geostrophic velocity is given by:

$$v_g(x) = -\frac{g}{f} \frac{\partial h}{\partial x}, \quad (15)$$

where $g = 9.81 \text{ m s}^{-2}$ is gravitational acceleration and $f = 10^{-4} \text{ s}^{-1}$ is the Coriolis parameter. The EKE per unit volume and the total integrated EKE are then:

$$\text{EKE}(x) = \frac{\rho}{2} v_g(x)^2, \quad \text{EKE}_{\text{tot}} = \int_{\text{lon}} \text{EKE}(x) dx, \quad (16)$$

with $\rho = 1000 \text{ kg m}^{-3}$.

This diagnostic is not used here as a validation metric, but as a concrete illustration of how structural distortions in the gridded analysis propagate into physically relevant quantities. Because EKE scales as the square of the height gradient, even moderate profile deformations can have a disproportionate impact on this quantity. Values are given in Table 6.

In Experiment #2, the widened and flattened `grid` analyzed profile produces smaller height gradients and therefore weaker geostrophic currents. The resulting total EKE is reduced by 55% relative to the value computed from the true profile. In Experiment #3, the sub-Gaussian shape similarly reduces the peak gradient, leading to an 8% reduction in total EKE relative to the true value. For both experiments, EKE estimations from the `param` projected profile have the same values as the truth. Similarly to Experiment #2, the flattened `grid` analyzed profile induces an EKE reduction of 50% while the slightly lower peak of the `param` projected profile reduces the EKE by only 1%. These reductions would directly affect estimates of lateral transport, air-sea interactions, and the mesoscale energy budget.

4.4 Computational complexity

We provide a theoretical, order-of-magnitude comparison of the computational costs associated with the Eulerian, Arbitrary Lagrangian-Eulerian (ALE) (Hirt et al., 1974), and parametric representations considered in this study. This is not intended
 320 as a numerical benchmark, since a fair comparison of wall-clock times would require complete implementations of the three approaches.

Consider a coherent structure of characteristic radius r in a domain of size L , discretized with a numerical scheme of order q , the relative discretization error scales as

$$\varepsilon \sim \left(\frac{\Delta \text{lon}}{r} \right)^q, \quad (17)$$

325 so that the grid spacing required to achieve a tolerance ε scales as

$$\Delta \text{lon} \sim r \varepsilon^{1/q}. \quad (18)$$

In d dimensions, the number of grid points per spatial direction is $L/\Delta \text{lon} \sim (L/r) \varepsilon^{-1/q}$, so a fixed Eulerian grid covering the full domain therefore requires

$$m_{\text{Eulerian}}(d) \sim (L/r)^d \varepsilon^{-d/q}, \quad (19)$$

330 degrees of freedom.

An ALE representation (Aydøğdu et al., 2019), which concentrates the resolution around each of K tracked structures, requires instead

$$m_{\text{ALE}}(d) \sim K c^d \varepsilon^{-d/q}, \quad (20)$$

where $c = O(1)$ characterizes the size of the refined region relative to the structure radius.

335 The parametric representation used here reduces the state to $n_{\text{param}} = 3K$ parameters, independently of the grid resolution and physical dimension. Thus, in terms of state dimension, the parametric representation is substantially smaller than both the Eulerian and ALE representations.

The computational cost of the assimilation step nevertheless depends on the ensemble size and on the dimension of the state being updated. For an ensemble-based linear update implemented in the N -dimensional ensemble subspace, rather than
 340 by explicitly forming and inverting the full $m \times m$ covariance matrix (e.g. via the Sherman–Morrison–Woodbury identity), a representative scaling is $O(N^2 m)$ (Evensen, 2003; Nino Ruiz et al., 2015), compared to $O(m^3)$ for a direct, non-ensemble-subspace implementation. The corresponding costs therefore scale as

$$C_{\text{Eulerian}} = O(N^2 m_{\text{Eulerian}}), \quad C_{\text{ALE}} = O(N^2 m_{\text{ALE}}), \quad (21)$$

whereas the parametric analysis itself costs only

$$345 \quad C_{\text{param}} = O(N^2 n_{\text{param}}). \quad (22)$$

The parametric approach, however, requires an additional parameter-extraction step before the analysis. In the present framework, each ensemble member is fitted to the parametric structure using nonlinear least squares (Eq. 12). If m denotes the resolution of the input grid, the cost of this step can be approximated as

$$C_{\text{fit}} = O(Nmn_{\text{iter}}), \quad (23)$$

350 where m is the resolution of the input grid and n_{iter} is the number of nonlinear solver iterations. The total cost of the proposed approach is therefore

$$C_{\text{total,param}} = O(Nmn_{\text{iter}} + N^2n_{\text{param}}). \quad (24)$$

This comparison highlights the main computational trade-off of the proposed representation. The analysis step becomes substantially cheaper because it is performed in a low-dimensional parameter space, but this reduction comes at the expense of
 355 the additional parameter-extraction step applied to each ensemble member. The overall computational benefit therefore depends on the relative cost of these two stages, and hence on the grid resolution, the ensemble size, and the cost of the nonlinear fit.

5 Conclusions

This study investigates data assimilation (DA) performed in a reduced parametric space as well as classic linear and Gaussian DA in gridded space, focusing on a simplified one-dimensional representation of an ocean eddy and on a single analysis step.

360 The results obtained in this idealized framework indicate that DA in gridded space provides satisfactory reconstructions when uncertainty mainly affects linear parameters such as amplitude, as illustrated by Experiment #1, where both approaches yield identical results (RMSE = 0.00 m, zero parameter error). However, when uncertainty involves nonlinear geometric parameters such as position and radius, the gridded-space update leads to a deformation of the reconstructed structure. In Experiment #2, the gridded analysis underestimates the amplitude by 23% (error of 0.24 m) and overestimates the radius by 29% (error of
 365 2.15 km), while the RMSE relative to the true profile reaches 0.10 m, compared with 0.00 m in parametric space. In Experiment #3, the analyzed profile is no longer Gaussian, and the RMSE in gridded space (0.03 m) exceeds that in parametric space (0.00 m). These structural distortions affect derived physical quantities: the total EKE is reduced by 67% in Experiment #2 and by 11% in Experiment #3 relative to the parametric analysis. These EKE figures are reported as an illustration of the physical consequences of the profile deformation, rather than as validation metrics.

370 The tutorial experiments (#1-3) are useful to isolate and illustrate this mechanism, but they do not constitute a balanced end-to-end comparison, since the uncertainty originates in parameter space and is propagated to grid space. Experiment #4, which starts from noisy gridded observations, provides a more symmetric test. The parametric approach retains its qualitative advantage: the RMSE of the gridded analysis (0.09 m) remains higher than that of the parametric analysis (0.02 m), and the gridded analysis underestimates the amplitude by 0.21 m and overestimates the radius by 1.94 km relative to the truth, compared
 375 to errors of 0.02 m and 0.24 km for the parametric analysis. However, the calibration of the parametric error covariances is reduced compared to the tutorial experiments, due to the additional uncertainty introduced by the nonlinear fitting step.

Conclusions drawn from this study should therefore be understood as pertaining to these specific idealized configurations and not as general statements about the relative performance of gridded and parametric DA approaches.

In contrast, the parametric formulation provides a direct update of the structure-defining parameters in this controlled setting. This leads to a more consistent representation of amplitude, position and radius within the assumptions of the experiments. In this idealized setting, the parametric formulation also operates in a lower-dimensional space than the gridded approach ($n=3$ vs. $m=101$).

In the above experiments the assimilation in the parametric space presents several advantages. The structure is analyzed directly through a small set of physically interpretable parameters, and structural consistency is preserved in this space: the analyzed profile remains Gaussian by construction (Experiment #3), with the associated EKE diagnostic (Table 6) deviating from the truth by at most 1%, compared with reductions of up to 55% in gridded space. The parametric formulation also operates in a markedly lower-dimensional space ($n = 3$ vs. $m = 101$), which reduces the cost of the analysis step. This reduction comes at the expense of the additional parameter-extraction step required for each ensemble member. As discussed in Sect. 4.4, the overall computational benefit therefore depends on the relative cost of parameter extraction and analysis, and hence on the grid resolution and ensemble size.

The Rao-Blackwellization principle introduced in Introduction to motivate the parametric transformation used in this study suggests a concrete direction for extending it. As noted in Sect. 4.2, this transformation only partially achieves Gaussian tractability, the radius parameter, in particular, remains significantly non-Gaussian after fitting, unlike the exact conditional Gaussianity that Rao-Blackwellization guarantees by construction; identifying a transformation that more fully achieves this property, or a marginalized/Rao-Blackwellized particle filter (Šmídl and Hofman, 2011) that treats the residual non-Gaussianity explicitly, is a natural target for future work. We emphasize that the comparison in Sect. 2.1 concerns the propagation of uncertainty through an explicitly known analytic mapping h , for which the failure of local linearization could be directly diagnosed against a Monte Carlo reference. In realistic applications, however, the correspondence between gridded and parametric (or feature-based) representations of a coherent structure is generally not available in closed form. In that setting, more general nonlinear, non-Gaussian approaches would be needed to relate uncertainty between representation spaces directly from data, rather than through explicit parametric fitting as done here. Diffusion maps, which construct data-driven geometric embeddings relating different function spaces (Coifman and Lafon, 2006) and have recently been combined with generative particle systems to produce non-Gaussian samples directly from data (Li and Marzouk, 2025), or transport-map-based methods that characterize non-Gaussian joint distributions across representations via measure transport (Parno et al., 2016), remain candidate tools for this purpose. Exploring either avenue lies beyond the scope of the present idealized study and is identified as a direction for future work.

This work is intentionally limited to a controlled and idealized setting and does not include sequential assimilation cycles or prognostic model dynamics. Future work will investigate the extension of this framework to realistic, time-evolving situations, where eddy parameters must be inferred from observations through detection and fitting procedures, and updated sequentially. In such a context, more general dynamical models for parameter evolution will be required, which may include stochastic parameter models, reduced-order models, tracking-based models, or hybrid approaches, rather than being restricted to any

specific method. In particular, the dynamics of eddy parameters in a sequential context are expected to be strongly nonlinear, which may necessitate ensemble-based or nonlinear assimilation methods. The stability of the system over multiple cycles will also be investigated in future work, when this parametric approach is applied in an operational, sequential assimilation context.

415 Finally, the parametric and gridded formulations provide two complementary perspectives: the former emphasizes the evolution of physically interpretable structure parameters, while the latter operates directly on spatial fields. Depending on the application, analyses can be projected back into gridded space for diagnostic purposes or retained in parameter space for object-oriented applications that directly exploit coherent structure properties.

Code availability. The code used for the experiments is available here: <https://gitlab.imt-atlantique.fr/s24dealb/param-grid-1d-eddy>

420 **Appendix A: Taylor-linearized versus ensemble-based uncertainty propagation**

Figure A1 illustrates the Taylor-linearized propagation discussed in Sect. 2.1, for the position-perturbation configuration (Experiment #2); the corresponding ensemble-based (Monte Carlo) propagation for the same configuration is shown in Fig. 3 (Sect. 4.1) and is not repeated here. The pathology extends to the posterior covariance itself: because $B_{\text{Taylor}}^{\text{grid}}$ and $R_{\text{Taylor}}^{\text{grid}}$ are each rank-1 (only p is uncertain), their sum is at most rank 2, and the resulting posterior covariance collapses to numerically negligible values (diagonal entries of order 10^{-16} – 10^{-18} , occasionally slightly negative due to floating-point cancellation in the near-singular pseudo-inverse) almost everywhere on the grid, a spuriously overconfident analysis, rather than merely a biased one.

425

Author contributions. SD developed the code, performed the experiments and prepared the manuscript, with contributions from co-authors; SR proposed the research topic, provided knowledge and background, performed the seminal experiments, and offered mentorship throughout this research; PT and CGB provided knowledge and assistance during code development; CL and BB helped with the writing of the paper.

430

Acknowledgements. This study was carried out as part of the PTD PROTEVS2, project financed by the French Ministry of Defense / DGA and led by Shom.




Figure A1. Taylor-linearized uncertainty propagation for the position-perturbation configuration (Experiment #2); see Fig. 3 for the corresponding ensemble-based propagation. (a) Background mean, observation mean, and the resulting Taylor-linearized BLUE analysis mean height: the Taylor-linearized BLUE analysis (orange) exhibits an unphysical negative undershoot on 28.7% of grid points. (b) Taylor-propagated standard deviation of the background and observation (prior), and of the resulting analysis (posterior); note the vertical scale extends to twice the maximum of Fig. 3b, reflecting the larger peak values of the Taylor-propagated prior standard deviations. The prior curves drop to zero with a slope discontinuity at their respective peaks ($p = -5$ and $p = +5$ km, where $\partial h / \partial p = 0$ by symmetry), while the posterior standard deviation collapses to numerically negligible values almost everywhere, reflecting the rank-deficiency of the underlying covariance matrices.

References

- Alvera-Azcárate, A., Van Der Zande, D., Barth, A., Dille, A., Massant, J., and Beckers, J.-M.: Generation of super-resolution gap-free ocean colour satellite products using data-interpolating empirical orthogonal functions (DINEOF), *Ocean Science*, 21, 787–805, <https://doi.org/10.5194/os-21-787-2025>, 2025.
- Aydınoğlu, A., Carrassi, A., Guider, C. T., Jones, C. K. R. T., and Rampal, P.: Data assimilation using adaptive, non-conservative, moving mesh models, *Nonlinear Processes in Geophysics*, 26, 175–193, <https://doi.org/10.5194/npg-26-175-2019>, 2019.
- Beezley, J. D. and Mandel, J.: Morphing ensemble Kalman filters, *Tellus*, 60, 2008.
- Carrassi, A., Bocquet, M., Bertino, L., and Evensen, G.: Data assimilation in the geosciences: An overview of methods, issues, and perspectives, *WIREs Climate Change*, 9, e535, <https://doi.org/10.1002/wcc.535>, 2018.
- Casella, G. and Robert, C. P.: Rao-Blackwellisation of sampling schemes, *Biometrika*, 83, 81–94, 1996.
- Chelton, D. B., Schlax, M. G., and Samelson, R. M.: Global observations of nonlinear mesoscale eddies, *Progress in Oceanography*, 91, 167–216, <https://doi.org/10.1016/j.pocean.2011.01.002>, 2011.
- Chen, R. and Liu, J. S.: Mixture Kalman filters, *Journal of the Royal Statistical Society: Series B*, 62, 493–508, 2000.
- Chen, Y. and Snyder, C.: Assimilating Vortex Position with an Ensemble Kalman Filter, *Monthly Weather Review*, 135, 1828–1845, <https://doi.org/10.1175/MWR3351.1>, 2007.
- Coifman, R. R. and Lafon, S.: Diffusion Maps, *Applied and Computational Harmonic Analysis*, 21, 5–30, 2006.

- Evensen, G.: The Ensemble Kalman Filter: theoretical formulation and practical implementation, *Ocean Dynamics*, 53, 343–367, <https://doi.org/10.1007/s10236-003-0036-9>, 2003.
- 450 Feyeux, N., Vidard, A., and Nodet, M.: Optimal transport for variational data assimilation, *Nonlinear Processes in Geophysics*, 25, 55–66, <https://doi.org/10.5194/npg-25-55-2018>, 2018.
- Hannachi, A., Jolliffe, I. T., and Stephenson, D. B.: Empirical orthogonal functions and related techniques in atmospheric science: A review, *International Journal of Climatology*, 27, 1119–1152, <https://doi.org/10.1002/joc.1499>, 2007.
- 455 Hirt, C. W., Amsden, A. A., and Cook, J. L.: An arbitrary Lagrangian-Eulerian computing method for all flow speeds, *Journal of Computational Physics*, 14, 227–253, [https://doi.org/10.1016/0021-9991\(74\)90051-5](https://doi.org/10.1016/0021-9991(74)90051-5), 1974.
- Jazwinski, A. H.: *Stochastic Processes and Filtering Theory*, 1970.
- Kunii, M.: Assimilation of Tropical Cyclone Track and Wind Radius Data with an Ensemble Kalman Filter, *Weather and Forecasting*, 30, 1050–1063, <https://doi.org/10.1175/WAF-D-14-00088.1>, 2015.
- 460 Li, F. and Marzouk, Y.: Diffusion map particle systems for generative modeling, *Foundations of Data Science*, 7, 814–837, <https://doi.org/10.3934/fods.2024054>, 2025.
- Maclean, J., Santitissadeekorn, N., and Jones, C. K.: A coherent structure approach for parameter estimation in Lagrangian Data Assimilation, *Physica D: Nonlinear Phenomena*, 360, 36–45, <https://doi.org/10.1016/j.physd.2017.08.007>, 2017.
- Nino Ruiz, E. D., Sandu, A., and Anderson, J.: An efficient implementation of the ensemble Kalman filter based on an iterative Sherman-Morrison formula, *Statistics and Computing*, 25, 561–577, <https://doi.org/10.1007/s11222-014-9454-4>, 2015.
- 465 Parno, M., Moselhy, T., and Marzouk, Y.: A Multiscale Strategy for Bayesian Inference Using Transport Maps, *SIAM/ASA Journal on Uncertainty Quantification*, 4, 1160–1190, <https://doi.org/10.1137/15M1032478>, 2016.
- Poterjoy, J.: A Localized Particle Filter for High-Dimensional Nonlinear Systems, *Monthly Weather Review*, 144, 59–76, <https://doi.org/10.1175/MWR-D-15-0163.1>, 2016.
- 470 Ravela, S., Emanuel, K., and McLaughlin, D.: Data assimilation by field alignment, *Physica D: Nonlinear Phenomena*, 230, 127–145, <https://doi.org/10.1016/j.physd.2006.09.035>, 2007.
- Rosenthal, W. S., Venkataramani, S., Restrepo, J. M., and Mariano, A.: Displacement Data Assimilation, *Journal of Computational Physics*, 330, 594–614, <https://doi.org/10.1016/j.jcp.2016.10.025>, 2017.
- Sampson, C., Carrassi, A., Ayrødødu, A., and Jones, C. K. R. T.: Ensemble Kalman filter for nonconservative moving mesh solvers with a joint physics and mesh location update, *Quarterly Journal of the Royal Meteorological Society*, 147, 1539–1561, <https://doi.org/10.1002/qj.3980>, 2021.
- 475 Šmídl, V. and Hofman, R.: Marginalized Particle Filtering Framework for Tuning of Ensemble Filters, *Monthly Weather Review*, 139, 3589–3599, <https://doi.org/10.1175/2011MWR3586.1>, 2011.
- Srivastava, A., Kang, W., and Tartakovsky, D. M.: Feature-informed data assimilation, *Journal of Computational Physics*, 494, 112 499, <https://doi.org/10.1016/j.jcp.2023.112499>, 2023.
- 480 van Leeuwen, P. J.: Nonlinear data assimilation in geosciences: an extremely efficient particle filter, *Quarterly Journal of the Royal Meteorological Society*, 136, 1991–1999, <https://doi.org/10.1002/qj.699>, 2010.
- Wu, C.-C., Lien, G.-Y., Chen, J.-H., and Zhang, F.: Assimilation of Tropical Cyclone Track and Structure Based on the Ensemble Kalman Filter (EnKF), *Journal of the Atmospheric Sciences*, 67, 3806–3822, <https://doi.org/10.1175/2010JAS3444.1>, 2010.
- 485 Zhen, Y., Resseguier, V., and Chapron, B.: Alignment of geophysical fields: A differential geometry perspective, *Physica D: Nonlinear Phenomena*, 483, 134 997, <https://doi.org/10.1016/j.physd.2025.134997>, 2025.