

PROCEEDINGS A

rspa.royalsocietypublishing.org

Research



Article submitted to journal

Subject Areas:

fluid mechanics, mathematical modelling, applied mathematics

Keywords:

data assimilation, ensemble Kalman filter, elliptic inverse problems, incompressible fluid mechanics, observation-informed dimension reduction.

Author for correspondence:

Mathieu Le Provost

e-mail:

mathieuleprovost1@gmail.com

A low-rank ensemble Kalman filter for elliptic observations

Mathieu Le Provost¹, Ricardo Baptista²,
Youssef Marzouk² and Jeff D. Eldredge¹¹Mechanical & Aerospace Engineering Department,
University of California, Los Angeles,
Los Angeles, CA 90095, USA²Department of Aeronautics and Astronautics,
Massachusetts Institute of Technology
Cambridge, MA, 02139, USA

We propose a regularization method for ensemble Kalman filtering (EnKF) with elliptic observation operators. Commonly used EnKF regularization methods suppress state correlations at long distances. For observations described by elliptic partial differential equations, such as the pressure Poisson equation (PPE) in incompressible fluid flows, distance localization should be used cautiously, as we cannot disentangle slowly decaying physical interactions from spurious long-range correlations. This is particularly true for the PPE, in which distant vortex elements couple nonlinearly to induce pressure. Instead, these inverse problems have a low effective dimension: low-dimensional projections of the observations strongly inform a low-dimensional subspace of the state space. We derive a low-rank factorization of the Kalman gain based on the spectrum of the Jacobian of the observation operator. The identified eigenvectors generalize the source and target modes of the multipole expansion, independently of the underlying spatial distribution of the problem. Given rapid spectral decay, inference can be performed in the low-dimensional subspace spanned by the dominant eigenvectors. This low-rank EnKF is assessed on dynamical systems with Poisson observation operators, where we seek to estimate the positions and strengths of point singularities over time from potential or pressure observations. We also comment on the broader applicability of this approach to elliptic inverse problems outside the context of filtering.

In data assimilation, observational data from a real dynamical system are used to correct model predictions of the system's state. This paper focuses on the assimilation of elliptic observations—i.e., observations related to the state via an elliptic partial differential equation—in a dynamically evolving process. *Filtering* is a classic problem in data assimilation, where we aim to estimate the state of a dynamical system at time k , $\mathbf{x}_k \in \mathbb{R}^n$, given all observations available up to that time. State-of-the-art results for high-dimensional filtering problems are often obtained with the ensemble Kalman filter (EnKF) [1]. The EnKF recursively updates a set of M realizations of the state, called “particles” or “ensemble members” $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$, to form an empirical approximation of the filtering distribution. Within each assimilation cycle, we perform a *forecast step* followed by an *analysis step*. In the forecast step, the filtering ensemble is propagated through the dynamical model for a finite time to generate samples of the forecast distribution. The analysis step updates the forecast ensemble by conditioning on newly available observations of the true system. The EnKF performs the analysis step by using the forecast ensemble to form a Monte Carlo approximation of the Kalman gain. For typical applications, a *limited* number of ensemble members ($M \sim 100$) are tracked compared to the dimension of the state variable ($n \sim 10^4 - 10^9$) and the dimension of the observation ($d \sim 10^2 - 10^5$) [2]. Hence, the empirical state and observation covariance matrices are rank-deficient, and suffer from both sampling errors and spurious long-range correlations. While in many high-dimensional problems the EnKF can successfully track the state with very limited samples, this success is predicated on an appropriate regularization of the empirical Kalman gain. Predominant regularization techniques assume that the observations are *local*, i.e., that an observation only provides information about a subset of the state variables which are close in physical distance. This assumption is supported by the rapid decay of the correlations between state and observation variables. In this setting, distance localization cuts off long-range correlations [3].

In this paper, we are interested in filtering problems where the observations correspond to *non-local* functions of the state, such as integrals of linear and nonlinear functions of the state [4]. As examples, we can cite the radiance measured by satellites, heat or mass fluxes through surfaces, forces measured on a body immersed in a fluid—or, as we address this paper, solutions of elliptic partial differential equations (PDEs). An elliptic PDE is given by

$$\mathcal{L}u_k(\mathbf{r}) = \mathbf{q}_k(\mathbf{r}; \mathbf{x}_k), \quad \mathbf{r} \in \Omega, \quad (0.1)$$

where Ω is the physical space (a subset of \mathbb{R}^2 or \mathbb{R}^3), $\mathbf{r} \in \Omega$ denotes the point of evaluation of (0.1), \mathcal{L} is an elliptic linear operator (e.g., the Laplacian ∇^2 , in which case (0.1) is a Poisson equation), and \mathbf{q}_k is a spatially distributed forcing term that in general depends nonlinearly on the state \mathbf{x}_k . We will consider filtering problems where, at each assimilation step k , we seek to estimate the state \mathbf{x}_k from a small number of pointwise evaluations of the solution u_k of the elliptic PDE (0.1); in other words, we will observe u_k with noise at only a few points $\mathbf{r} \in \Omega$. Up to a homogeneous part, the solution u_k can be obtained by convolution of the Green's function of the elliptic operator \mathcal{L} with the forcing term \mathbf{q}_k over Ω . The observations thus inherit the non-locality and the nonlinearity of the Green's function, along with further nonlinear state dependence in the forcing term \mathbf{q}_k .

For a Poisson equation, the Green's function of the Laplacian decays logarithmically or algebraically (based on the dimension of the physical space) as a function of distance, so we expect long-range physical interactions between the state and observation variables. Distance localization should be used carefully, as it will remove *all* long-range correlations, whether physical or spurious. In this paper, we develop a novel regularization technique to unambiguously assimilate such *non-local* observations by exploiting the known structure of the solutions of (0.1). We emphasize that Kalman-based ensemble filters, including the present work, treat the dynamical model as a black box, and only need to evaluate it at a set of samples in the forecast step. Thus, our methodology is readily applicable to a wide range of filtering problems with elliptic observations such as [5,6] in electromagnetism or [7–9] in incompressible fluid flows.

In this paper, we focus our discussion on the representative context of incompressible fluid mechanics. In these problems, the flow field is most compactly represented by the vorticity, or curl of the velocity field. We use a low-dimensional Lagrangian representation of the flow field by tracking the positions and strengths of a small collection of point vortices. Inviscid point vortex models have been a long-standing tool to model and explain incompressible fluid phenomena [10–12]. In this paper, we focus on filtering problems where we seek to estimate the characteristics of point singularities over time from limited potential or pressure observations. Despite their relative low dimensionality, these problems can be particularly challenging as the transport equation for point vortices (the Biot–Savart law) and the observation model (the pressure Poisson equation) involve the resolution of Poisson equations like (0.1), whose forcing terms nonlinearly couple all the singularities’ contributions.

In this work, we introduce a new regularization approach motivated by the low-rank approximation of *inverse problems*. Indeed, given the indirect relationship between state and observations, each assimilation step in our current setting can be viewed as the solution of an inverse problem. It is natural to cast this inverse problem in a Bayesian setting, where the parameters of interest comprise the state x_k , the prior distribution is approximated empirically by the forecast ensemble, and the forward model (and hence the likelihood function) follow from (0.1). In a wide range of inverse problems, observations are informative—relative to the prior—about only a low-dimensional subspace of the parameter space [13–16]. There are several reasons for this structure: information about the state is lost through the smoothing action of the forward operator, as well as its nonlinearity; observations may be few in number, relative to the state dimension; and noise in the observations makes it difficult to extract small-amplitude features of the parameters. In the linear–Gaussian setting, Spantini et al. [14] exploit these ideas by identifying the optimal (in a certain precise sense) prior-to-posterior update of any given dimension in the parameter space, based on an eigensystem that combines the forward operator, the prior covariance matrix, and the noise covariance matrix. Similarly, in the linear–Gaussian setting, Giraldo et al. [17] identify low-dimensional projections of the observations that are optimal in the sense of preserving mutual information with the parameters. A duality between parameter and observation reduction in the linear–Gaussian setting is further elucidated in [18, Sec. 3.6.2]. Cui et al. [16] generalize the parameter reduction approach of [14] to the *nonlinear/non-Gaussian setting*, based on integrals of the Jacobian of the nonlinear forward model. Here, the strict duality between parameter and observation space eigensystems is lost, and observation reduction requires additional effort. Inspired by [16], here we propose an analogous method to identify the most informative directions in the observation space, for the nonlinear/non-Gaussian setting. This companion dimension reduction is motivated by the potential for redundancy and low observability of some sensors, as well as the spatial structure of the data. Ideas for such a *joint* dimension reduction have already been exploited in the *fast multipole method (FMM)*, which accelerates the solution of the Poisson equation by clustering interactions between “sources” and “targets” [19,20]. Our algorithm provides a systematic method to identify these clusters of variables from the Jacobian of the observation model, without knowledge of the spatial organization of the singularities and the evaluation points.

We then use this identified structure to *regularize the EnKF*: specifically, we construct a low-rank factorization of the Kalman gain based on the identified modes for the states and observations. We observe that for elliptic observation operators, a few modes can very accurately approximate the row and column spaces of the Kalman gain. The analysis step of the proposed observation-informed low-rank EnKF (LREnKF) can be summarized by the following sequence of steps: identify and project the state and observation variables on the leading directions in the state and observation spaces, compute the Kalman gain and apply the linear update of the EnKF in these reduced coordinates, and lift the result to the original variable space.

The rest of the paper is organized as follows. Section 1 gives a basic outline of the model equations of potential flow theory, and highlights the nonlinear and elliptic properties of the dynamical and observation models. Section 2 reviews the classical treatment of the EnKF and

presents our low-rank ensemble Kalman filter. Example problems are treated in Section 3. Concluding remarks follow in Section . For the sake of conciseness, certain details about the model equations and the proposed low-rank ensemble Kalman filter, including a pseudocode, are provided in the Supplementary Material.

1. Model equations of potential flow

This section gives a minimal outline of the model equations governing an incompressible flow described by advecting point singularities. Further details of the equations are given in Supplementary Material A and can be found in many references, such as [11]. Though we focus in this work on two-dimensional fluid mechanics, the equations are representative of problems in many areas of physics that follow the template of equation (0.1), in two or three dimensions, and the techniques we develop are easily extendable to these other contexts.

In this study, we use both vector notation and complex notation, the former for emphasizing the generality and the latter for compactly representing some of the basic equations. A point in space will either be referred to in vector notation as $\mathbf{r} = (x, y) \in \mathbb{R}^2$, or in complex notation as $z = x + iy$. The conjugate of a complex number $z = \Re(z) + i\Im(z)$ is denoted as $\bar{z} \equiv \Re(z) - i\Im(z)$. We denote the canonical basis of \mathbb{R}^3 as $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$.

In general, a vector field $\mathbf{v}(\mathbf{r}, t)$ such as fluid velocity can be obtained from the sum of the gradient of a scalar potential $\phi(\mathbf{r}, t)$ and the curl of a vector potential (or streamfunction) $\psi(\mathbf{r}, t)$, each of which satisfies a Poisson equation, with right-hand sides given respectively by the divergence of \mathbf{v} and the (negative) curl of \mathbf{v} , i.e., the vorticity. A potential flow refers to a flow field in which those right-hand-side quantities occupy, at most, singular regions of zero volume [11]. We consider a set of N singularities located at $\{z_1, \dots, z_N\}$ with complex strengths $\{S_1, \dots, S_N\}$, subject to a uniform flow with velocity $\mathbf{U}_\infty = (U_\infty, V_\infty)$ or in complex notation $W_\infty = U_\infty - iV_\infty$. We should stress the classical conjugation in the definition of the equivalent complex velocity for consistency with the Cauchy-Riemann equations [11]. The strengths are represented as $S_J = Q_J - i\Gamma_J$ where $Q_J, \Gamma_J \in \mathbb{R}$ denote the volume flux and the circulation of the J th singularity, respectively. By linearity, the solution of the Poisson equation is described by a complex potential $F = \phi + i\psi$, which at a location $z \in \mathbb{C}$ is given by

$$F(z) = \sum_{J=1}^N \frac{S_J}{2\pi} \log(z - z_J) + W_\infty z; \quad (1.1)$$

the velocity field \mathbf{v} is simply the derivative with respect to z . One can show that the dynamics of the singularities are governed by the following set of ordinary differential equations:

$$\frac{dz_J}{dt} = \overline{w_{-J}}(t) = \sum_{K=1, K \neq J}^N \frac{\overline{S_K}}{2\pi} \frac{1}{\bar{z}_J - \bar{z}_K} + \overline{W_\infty}, \quad \frac{dS_J}{dt} = 0, \quad J = 1, \dots, N, \quad (1.2)$$

where w_{-J} in complex notation, or \mathbf{v}_{-J} in vector notation, is the Kirchhoff velocity of the J th singularity, obtained by evaluating the velocity field at singularity J and omitting its own contribution.

For potential flows, the pressure can equivalently be computed by two means: from the inversion of the pressure Poisson equation or from the unsteady Bernoulli equation; see Supplementary Material B. Each formulation provides a different perspective on the mechanisms at play in the pressure response. The pressure Poisson equation is formed by taking the divergence of the Euler equations:

$$\nabla^2 \left(p(\mathbf{r}, t) + \frac{1}{2} \rho \|\mathbf{v}(\mathbf{r}, t)\|^2 \right) = \rho \sum_{J=1}^N [Q_J \mathbf{v}_{-J} \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_J) - \Gamma_J \mathbf{e}_z \cdot (\mathbf{v}_{-J} \times \nabla \delta(\mathbf{r} - \mathbf{r}_J))]. \quad (1.3)$$

Equivalently to (1.3), the pressure induced by point singularities can be obtained in closed form from the unsteady Bernoulli equation. For a fixed evaluation point $z' \in \mathbb{C}$, we get

$$p(z', t) - B(t) = -\frac{1}{2}\rho \left| \sum_{j=1}^N \frac{S_j}{2\pi} \frac{1}{z' - z_j} + W_\infty \right|^2 + \rho \Re \left(\sum_{j=1}^N \frac{S_j}{2\pi} \frac{\overline{w_{-j}(t)}}{z' - z_j} \right), \quad (1.4)$$

where $B(t)$ is function of time only. Eqns. (1.1) and (1.4) are two examples relating the state (i.e., singularity positions and strengths) and observation (i.e., velocity potential or pressure measurements) that serve as prototypes for the larger class of problems in eqn. (0.1) that motivate this work. From their elliptic nature, we expect long-range nonlinear interactions between the state and observation. Furthermore, (1.4) shows that, by the appearance of the Kirchhoff velocity, every vortex affects the manner in which every other vortex induces pressure. For these reasons, estimating the characteristics of the singularities from limited observations poses many challenges.

2. Low-rank ensemble Kalman filter

Before we present the traditional ensemble Kalman filter and our low-rank extension, we introduce our conventions. Serif fonts refer to random variables, e.g., \mathbf{Q} on \mathbb{R}^n or Q on \mathbb{R} . Lowercase roman fonts refer to realizations of random variables, e.g., \mathbf{q} on \mathbb{R}^n or q on \mathbb{R} . $\pi_{\mathbf{Q}}$ denotes the probability density function for the random variable \mathbf{Q} , and $\mathbf{q} \sim \pi_{\mathbf{Q}}$ means that \mathbf{q} is a realization of \mathbf{Q} . $\mu_{\mathbf{Q}}$ and $\Sigma_{\mathbf{Q}}$ denote the mean and the covariance matrix of the random variable \mathbf{Q} , respectively. $\Sigma_{\mathbf{Q}, \mathbf{R}}$ denotes the cross-covariance matrix of the random variables \mathbf{Q} and \mathbf{R} . Empirical quantities are differentiated by carets above each symbol.

A generic nonlinear discrete state-space model can be described by a pair of a Markov processes for a latent state variable $(\mathbf{X}_k)_{k \geq 0} \in \mathbb{R}^n$, and an observed stochastic process $(\mathbf{Y}_k)_{k \geq 1} \in \mathbb{R}^d$ that is conditionally independent given the states at all times. The evolution of the state \mathbf{X}_k is given by an initial distribution $\pi_{\mathbf{X}_0}$ and a *dynamical system* to propagate the state from one time step to the next, i.e.,

$$\mathbf{X}_k = \mathbf{f}_k(\mathbf{X}_{k-1}) + \mathbf{W}_k, \quad (2.1)$$

where $\mathbf{f}_k: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a deterministic function that represents the mean of the discrete-time dynamics and \mathbf{W}_k is an independent *process noise*. In our case, the state \mathbf{X}_k is composed of the positions and strengths of a collection of point singularities, while \mathbf{f}_k corresponds to the time-marching of their positions and strengths from discretizing (1.2). In this work, the observations are given as a nonlinear mapping of the state \mathbf{X}_k with additive noise of the following form:

$$\mathbf{Y}_k = \mathbf{h}_k(\mathbf{X}_k) + \mathbf{\mathcal{E}}_k, \quad (2.2)$$

where $\mathbf{h}_k: \mathbb{R}^n \rightarrow \mathbb{R}^d$ is the *observation operator* and $\mathbf{\mathcal{E}}_k$ is an *observation noise*. Eq. (2.2) is called the *observation model*. In our case, the observation \mathbf{Y}_k is a vector of velocity potential or pressure evaluations at d discrete locations.

2.a Ensemble Kalman filter

For the remainder of this section, we focus on a single analysis step of the EnKF. Our treatment of the analysis step is built on the idea that there is an underlying transformation \mathbf{T}_k , called the *prior-to-posterior transformation* or *analysis map*, that directly maps samples from the forecast (i.e., prior) density $\pi_{\mathbf{X}_k} | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}$ to the filtering (i.e., posterior) density $\pi_{\mathbf{X}_k} | \mathbf{Y}_{1:k} = \mathbf{y}_{1:k}$ [21, 22]. More precisely, $\mathbf{T}_k: \mathbb{R}^{n+d} \rightarrow \mathbb{R}^n$ maps the *joint forecast distribution* of observations and states, $\pi_{\mathbf{Y}_k, \mathbf{X}_k} | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}$, to the filtering distribution $\pi_{\mathbf{X}_k} | \mathbf{Y}_{1:k} = \mathbf{y}_{1:k}$. Further details on the filtering problem and the analysis map can be found in Supplementary Material C.

Linear transform filters [23] apply the following linear map to random variables $(\mathbf{Y}_k, \mathbf{X}_k)$ representing the joint forecast distribution $\pi_{\mathbf{Y}_k, \mathbf{X}_k} | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}$,

$$\mathbf{T}_k(\mathbf{Y}_k, \mathbf{X}_k) = \mathbf{X}_k - \Sigma_{\mathbf{X}_k, \mathbf{Y}_k} \Sigma_{\mathbf{Y}_k}^{-1} (\mathbf{Y}_k - \mathbf{y}_k^*), \quad (2.3)$$

where \mathbf{y}_k^* is the observation to be assimilated, $\Sigma_{\mathbf{X}_k, \mathbf{Y}_k} \in \mathbb{R}^{n \times d}$ is the cross-covariance matrix of the state and observation, while $\Sigma_{\mathbf{Y}_k}^{-1} \in \mathbb{R}^{d \times d}$ is the precision matrix (inverse of the covariance matrix) of the observation's marginal distribution. The linear operator $\mathbf{K}_k = \Sigma_{\mathbf{X}_k, \mathbf{Y}_k} \Sigma_{\mathbf{Y}_k}^{-1} \in \mathbb{R}^{n \times d}$ is called the Kalman gain and maps observation discrepancies $(\mathbf{y}_k - \mathbf{y}_k^*)$ to the state correction. The mean and covariance of \mathbf{T}_k coincide with the standard Kalman update in the case of a Gaussian forecast distribution and linear-Gaussian observation model [23, Corollary 1]. In other words, the transformation yields the exact Bayesian update in the case of jointly Gaussian $(\mathbf{Y}_k, \mathbf{X}_k)$. More generally, however, the mean of \mathbf{T}_k is the linear least-squares estimate of the state at time k . The *stochastic ensemble Kalman filter* (sEnKF) introduced by Evensen [1] estimates the transformation (2.3) by replacing the covariances with empirical covariances that are computed from the joint samples of the observations and states $\{\mathbf{y}_k^i, \mathbf{x}_k^i\} \sim \pi_{\mathbf{Y}_k, \mathbf{X}_k} | \mathbf{y}_{1:k-1}$, i.e.,

$$\hat{\mathbf{T}}_k(\mathbf{y}_k, \mathbf{x}_k) = \mathbf{x}_k - \hat{\Sigma}_{\mathbf{X}_k, \mathbf{Y}_k} \hat{\Sigma}_{\mathbf{Y}_k}^{-1} (\mathbf{y}_k - \mathbf{y}_k^*). \quad (2.4)$$

2.b Assimilation in low-dimensional subspaces

For elliptic inverse problems, the conditional structure of the joint density of the observation and the state $\pi_{\mathbf{Y}_k, \mathbf{X}_k} | \mathbf{y}_{1:k-1}$ is not localized, i.e., there is no rapid decay of the correlations as a function of the distance between variables. Hence, distance based localization in this problem results in biased estimators for the Kalman gain. Instead, we will exploit another kind of low-dimensional structure in the joint distribution. Our treatment draws inspiration from the fast multipole method (FMM) [19], which uses a hierarchical clustering of “source” and “target” elements to accelerate the calculation of the potential field from a large set of point singularities [20]. These clusterings are typically based only on spatial distance, but here we will use information from the prior distribution and observation model (i.e., observation operator and observation noise) to infer these clusters automatically in a way that works independently of the spatial distribution of the singularities and observation locations. To regularize the EnKF, we identify important directions in the state and observation spaces, perform the assimilation in these low-dimensional subspaces, and finally, lift the result to the original state space. First, we present the treatment of low-rank structure in the case of a linear-Gaussian observation model in 2.b(i). The main result of this derivation is a factorization of the Kalman gain that exploits the existence of this low-dimensional subspace. Then, we extend this decomposition to the nonlinear-Gaussian setting in 2.b(ii). An algorithm that summarizes the overall low-rank assimilation procedure is provided in Supplementary Material F.

Remark: For convenience, we drop the time-dependent subscripts from the variables in the rest of this paper, since the analysis step does not involve time propagation.

(i) Low-rank assimilation for the linear-Gaussian case

In this section, we consider the inference problem for the linear-Gaussian observation model:

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathcal{E}, \quad (2.5)$$

where the state is given by $\mathbf{X} \sim \mathcal{N}(\mu_{\mathbf{X}}, \Sigma_{\mathbf{X}})$, the observational error is given by $\mathcal{E} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathcal{E}})$, and \mathcal{E} is independent of \mathbf{X} . The matrix $\mathbf{H} \in \mathbb{R}^{d \times n}$ is called the observation matrix. In order to identify the important assimilation directions, we first define the whitened variables $\tilde{\mathbf{X}} = \Sigma_{\mathbf{X}}^{-1/2} (\mathbf{X} - \mu_{\mathbf{X}}) \in \mathbb{R}^n$, $\tilde{\mathcal{E}} = \Sigma_{\mathcal{E}}^{-1/2} \mathcal{E} \in \mathbb{R}^d$. We use $\mathbf{B}^{1/2}$, $\mathbf{B}^{-1/2}$ to denote a square root of the matrix \mathbf{B} and its inverse, respectively. These whitened variables satisfy $\tilde{\mathcal{E}} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, $\tilde{\mathbf{X}} \sim$

$\mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$. Applying the same whitening to the observation variable, the observation model becomes

$$\tilde{\mathbf{Y}} = \Sigma_{\mathcal{E}}^{-1/2} (\mathbf{Y} - \mu_{\mathbf{Y}}) = \tilde{\mathbf{H}}\tilde{\mathbf{X}} + \tilde{\mathcal{E}}, \quad (2.6)$$

where $\mu_{\mathbf{Y}} = \mathbf{H}\mu_{\mathbf{X}}$ and $\tilde{\mathbf{H}} = \Sigma_{\mathcal{E}}^{-1/2} \mathbf{H} \Sigma_{\mathbf{X}}^{1/2} \in \mathbb{R}^{d \times n}$ is the whitened observation matrix. In this derivation we assume that $d \leq n$. We can now identify the important assimilation directions with a singular value decomposition (SVD) of the whitened observation matrix $\tilde{\mathbf{H}} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^{\top}$, where $\mathbf{U} \in \mathbb{R}^{d \times d}$ and $\mathbf{V} \in \mathbb{R}^{n \times d}$ are the left and right singular vectors, and $\mathbf{\Lambda} \in \mathbb{R}^{d \times d}$ is the diagonal matrix of singular values. From this decomposition, we can rotate and project the whitened state on the subspace spanned by the columns of $\mathbf{V}^{\top} \rightarrow \tilde{\mathbf{X}} = \mathbf{V}^{\top} \tilde{\mathbf{X}} \in \mathbb{R}^d$ — and similarly rotate the whitened observational error and observation on the subspace spanned by the columns of $\mathbf{U}^{\top} \rightarrow \tilde{\mathcal{E}} = \mathbf{U}^{\top} \tilde{\mathcal{E}} \in \mathbb{R}^d$, $\tilde{\mathbf{Y}} = \mathbf{U}^{\top} \tilde{\mathbf{Y}} \in \mathbb{R}^d$. The SVD of the whitened observation matrix simultaneously identifies a pair of orthogonal bases for the state and observation spaces and an ordering for these directions. In Supplementary Material D, we show that the Kalman gain in the rotated space (where $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ live) is given by:

$$\check{\mathbf{K}} = \mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1}. \quad (2.7)$$

Thus, the linear analysis map in the rotated space $\check{\mathbf{T}}$ is given by:

$$\check{\mathbf{T}}(\check{\mathbf{y}}, \check{\mathbf{x}}) = \check{\mathbf{x}} - \check{\mathbf{K}}(\check{\mathbf{y}} - \check{\mathbf{y}}^*) = \check{\mathbf{x}} - \mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1}(\check{\mathbf{y}} - \check{\mathbf{y}}^*), \quad (2.8)$$

where $\check{\mathbf{y}}^* = \mathbf{U}^{\top} \Sigma_{\mathcal{E}}^{-1/2} \mathbf{y}^*$ denotes the realization of the assimilated data in the rotated observation space. In these rotated coordinates, the analysis map is local, i.e., each component of the rotated observation only updates one associated component of the rotated state. Using (2.7) and (2.8), we show in D that the Kalman gain $\mathbf{K} \in \mathbb{R}^{n \times d}$ in the original space factorizes as

$$\mathbf{K} = \Sigma_{\mathbf{X}}^{1/2} \mathbf{V} \mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1} \mathbf{U}^{\top} \Sigma_{\mathcal{E}}^{-1/2}. \quad (2.9)$$

This factorization of \mathbf{K} and its application in the analysis map of the Kalman filter (2.3) provide a concise summary of the inference process in the low-rank informative subspaces: the innovation term $(\mathbf{y} - \mathbf{y}^*)$ is whitened and rotated by applying $\mathbf{U}^{\top} \Sigma_{\mathcal{E}}^{-1/2}$, assimilated in this new set of coordinates using (2.7), and finally lifted to the original state space by applying $\Sigma_{\mathbf{X}}^{1/2} \mathbf{V}$. In the whitened space (where $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ live), it is easy to see that $\mathbf{V} \mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1} \mathbf{U}^{\top}$ constitutes the singular value decomposition of the Kalman gain. From the decay of the singular values $\mathbf{\Lambda}$, we can instead use a truncated singular value decomposition of $\tilde{\mathbf{H}} \approx \mathbf{U}_r \mathbf{\Lambda}_r \mathbf{V}_r^{\top}$, where $r \leq d$, $\mathbf{U}_r \in \mathbb{R}^{r \times r}$, $\mathbf{V}_r \in \mathbb{R}^{n \times r}$ are the first r left and right singular vectors, and $\mathbf{\Lambda}_r \in \mathbb{R}^{r \times r}$. A rank- r approximation of the Kalman gain is then given by $\mathbf{K}_r = \Sigma_{\mathbf{X}}^{1/2} \mathbf{V}_r \mathbf{\Lambda}_r(\mathbf{\Lambda}_r^2 + \mathbf{I}_r)^{-1} \mathbf{U}_r^{\top} \Sigma_{\mathcal{E}}^{-1/2}$. We should emphasize that, without the truncation of the SVD, eq. (2.9) is an exact factorization of the Kalman gain. With the truncated SVD, the state and observation variables are no longer just rotated but also projected to a subspace of dimension $r \leq \min(d, n)$.

(ii) Low-rank assimilation for the nonlinear–Gaussian case

In this section, we consider the nonlinear observation model of eq. (2.2), recalled for reference:

$$\mathbf{Y} = \mathbf{h}(\mathbf{X}) + \mathcal{E}, \quad (2.10)$$

where the state is $\mathbf{X} \sim \mathcal{N}(\mu_{\mathbf{X}}, \Sigma_{\mathbf{X}})$, the observational error is given by $\mathcal{E} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathcal{E}})$, and \mathcal{E} is independent of \mathbf{X} . To handle a nonlinear observation operator with the linear prior-to-posterior transformation (2.3), one simple approximation is to use the Jacobian of the observation operator about the prior mean as the observation matrix [3]. Unfortunately, this treatment is only viable for functions which are well approximated by a linear function over the bulk of the prior distribution. This section generalizes the treatment of the previous section for a nonlinear–Gaussian observation model; see Supplementary Material E for further details.

In the nonlinear and non-Gaussian setting, Cui et al. [16] showed that the most important assimilation directions in the whitened state space can be identified by the eigenvectors of the

state space Gramian, which in the case of Gaussian \mathcal{E} reduces to:

$$\mathbf{C}_{\mathbf{X}} = \int \left(\Sigma_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \Sigma_{\mathbf{X}}^{1/2} \right)^{\top} \left(\Sigma_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \Sigma_{\mathbf{X}}^{1/2} \right) d\pi_{\mathbf{X}}(\mathbf{x}) \in \mathbb{R}^{n \times n}, \quad (2.11)$$

where the expectation is taken over the prior distribution. As expected, eq. (2.11) reverts to $\tilde{\mathbf{H}}^{\top} \tilde{\mathbf{H}}$ in the linear–Gaussian case. $\mathbf{C}_{\mathbf{X}}$ is *positive semi-definite* (PSD) and its eigendecomposition can be written as $\mathbf{C}_{\mathbf{X}} = \mathbf{V} \Lambda_{\mathbf{X}}^2 \mathbf{V}^{\top}$, where $\mathbf{V} \in \mathbb{R}^{n \times n}$ is an orthogonal basis for the whitened state space with associated eigenvalues $\Lambda_{\mathbf{X}}^2 \in \mathbb{R}^{n \times n}$.

Inspired by the treatment in the state space, we propose to use the eigenvectors of the *observation space Gramian* $\mathbf{C}_{\mathbf{Y}}$ to select the important assimilation directions in the whitened observation space:

$$\mathbf{C}_{\mathbf{Y}} = \int \left(\Sigma_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \Sigma_{\mathbf{X}}^{1/2} \right) \left(\Sigma_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}) \Sigma_{\mathbf{X}}^{1/2} \right)^{\top} d\pi_{\mathbf{X}}(\mathbf{x}) \in \mathbb{R}^{d \times d}. \quad (2.12)$$

The matrix $\mathbf{C}_{\mathbf{Y}}$ is also PSD and has the eigendecomposition $\mathbf{C}_{\mathbf{Y}} = \mathbf{U} \Lambda_{\mathbf{Y}}^2 \mathbf{U}^{\top}$, where $\mathbf{U} \in \mathbb{R}^{d \times d}$ is an orthonormal basis for the whitened observation space with associated eigenvalues $\Lambda_{\mathbf{Y}}^2 \in \mathbb{R}^{d \times d}$. For convenience, we assume that the eigenvectors of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$ are ordered by decreasing eigenvalues. For a nonlinear observation model, the eigenvalues of the state and observation Gramians can be different. In practice, we use Monte-Carlo approximations of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$ that are estimated using the prior samples to identify the important subspaces. Depending on the inference problem, the Jacobian $\nabla \mathbf{h} \in \mathbb{R}^{d \times n}$ of the observation operator \mathbf{h} with respect to the state components can be computed either analytically, with automatic differentiation, with complex step differentiation, or, in the worst case, from finite differences. To perform the low-rank assimilation, we only retain the first $r_{\mathbf{X}} \leq n$ eigenmodes for the state space, and $r_{\mathbf{Y}} \leq d$ eigenmodes for the observation space. The ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ can be tuned independently based on the decay of $\Lambda_{\mathbf{X}}^2$ and $\Lambda_{\mathbf{Y}}^2$; typically we recommend setting these ranks to achieve a threshold $\alpha \in [0, 1]$ for the cumulative normalized energy of the eigenvalue spectra, e.g., in the state space, we set $r_{\mathbf{X}} = \min \{r_{\alpha} \in \llbracket 1, n \rrbracket \mid E_{\mathbf{X}, r_{\alpha}} \geq \alpha\}$, where $E_{\mathbf{X}, r_{\alpha}} = \sum_{i=1}^{r_{\alpha}} \lambda_{\mathbf{X}, i}^2 / \sum_{i=1}^n \lambda_{\mathbf{X}, i}^2$. The new projected state and projected observation variables are defined as $\tilde{\mathbf{X}} = \mathbf{V}_{r_{\mathbf{X}}}^{\top} \mathbf{X} = \mathbf{V}_{r_{\mathbf{X}}}^{\top} \Sigma_{\mathbf{X}}^{-1/2} (\mathbf{X} - \mu_{\mathbf{X}}) \in \mathbb{R}^{r_{\mathbf{X}}}$ and $\tilde{\mathbf{Y}} = \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \mathbf{Y} = \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \Sigma_{\mathbf{Y}}^{-1/2} (\mathbf{Y} - \mu_{\mathbf{Y}}) \in \mathbb{R}^{r_{\mathbf{Y}}}$, where $\mathbf{V}_{r_{\mathbf{X}}}, \mathbf{U}_{r_{\mathbf{Y}}}$ denote the first $r_{\mathbf{X}}, r_{\mathbf{Y}}$ columns of \mathbf{V}, \mathbf{U} , respectively.

We now outline how to obtain the factorization of the Kalman gain in the nonlinear setting. In the linear–Gaussian case, the Gramians in (2.11) and (2.12) revert to $\mathbf{C}_{\mathbf{X}} = \tilde{\mathbf{H}}^{\top} \tilde{\mathbf{H}}$ and $\mathbf{C}_{\mathbf{Y}} = \tilde{\mathbf{H}} \tilde{\mathbf{H}}^{\top}$, respectively. From the definition of the SVD of the whitened observation matrix $\tilde{\mathbf{H}}$, we have $\tilde{\mathbf{H}} \mathbf{v}_i = \lambda_i \mathbf{u}_i$ for some left singular vector \mathbf{u}_i , singular value λ_i and right singular vector \mathbf{v}_i . This results in the eigendecompositions $\mathbf{C}_{\mathbf{X}} \mathbf{u}_i = \lambda_i^2 \mathbf{u}_i$ and $\mathbf{C}_{\mathbf{Y}} \mathbf{v}_i = \lambda_i^2 \mathbf{v}_i$. Therefore, if the triplet $\{\mathbf{u}_i, \lambda_i, \mathbf{v}_i\}$ is obtained from the eigendecomposition of $\mathbf{C}_{\mathbf{X}}$ and $\mathbf{C}_{\mathbf{Y}}$, the sign of the singular vectors is lost, i.e., we only have $\tilde{\mathbf{H}} \mathbf{v}_i = \pm \lambda_i \mathbf{u}_i$. For a nonlinear observation model, the factorization of the Kalman gain given in eq. (2.9) is no longer applicable and needs to be generalized. In the rotated and whitened space, the analysis map is given by:

$$\tilde{\mathbf{T}}(\tilde{\mathbf{y}}, \tilde{\mathbf{x}}) = \tilde{\mathbf{x}} - \tilde{\mathbf{K}}(\tilde{\mathbf{y}} - \tilde{\mathbf{y}}^*), \quad (2.13)$$

where $\tilde{\mathbf{K}} = \Sigma_{\tilde{\mathbf{X}}\tilde{\mathbf{Y}}} \Sigma_{\tilde{\mathbf{Y}}}^{-1} \in \mathbb{R}^{r_{\mathbf{X}} \times r_{\mathbf{Y}}}$ is the Kalman gain in the informative space with $\Sigma_{\tilde{\mathbf{X}}\tilde{\mathbf{Y}}} \in \mathbb{R}^{r_{\mathbf{X}} \times r_{\mathbf{Y}}}$ and $\Sigma_{\tilde{\mathbf{Y}}}^{-1} \in \mathbb{R}^{r_{\mathbf{Y}} \times r_{\mathbf{Y}}}$. In the original space, we get

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = \mathbf{x} - \Sigma_{\mathbf{X}}^{1/2} \mathbf{V} \tilde{\mathbf{K}} \mathbf{U}^{\top} \Sigma_{\mathcal{E}}^{-1/2} (\mathbf{y} - \mathbf{y}^*). \quad (2.14)$$

This new factorization of the Kalman gain nicely generalizes eq. (2.9) to an inference problem with a nonlinear observation model. We emphasize that the sign issue presented above is obviated by absorbing it into the definition of $\tilde{\mathbf{K}}$. Let us remark that the definition of $\tilde{\mathbf{K}}$ reduces to the diagonal matrix $\Lambda(\Lambda^2 + \mathbf{I})^{-1}$ in the linear–Gaussian case, where the matrices \mathbf{U}, \mathbf{V} are obtained from the SVD of $\tilde{\mathbf{H}}$. Indeed, eq. (2.14) constitutes a *change of coordinates for the Kalman gain* between the

original space and the informative space. We denote the algorithm applying the linear update in (2.14) in each analysis step as the low-rank EnKF (LREnKF). A pseudo-code for the proposed low-rank EnKF is presented in F.

In comparison, the sEnKF estimates the Kalman gain in the original space, a linear operator of dimensions $n \times d$, where n and d are the dimensions of the space and observation spaces. Leveraging the dimension reduction offered by the state and observation Gramians, the LREnKF only has to estimate the Kalman gain in the informative subspaces of \mathbf{X} and \mathbf{Y} , whose dimension is only $r_{\mathbf{X}} \times r_{\mathbf{Y}}$, where $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ can be significantly smaller than n and d , respectively. To appreciate the benefit of this dimension reduction in the inference problem, it is useful to recall the classical bias-variance trade-off in any statistical learning problem [24]: Given an ensemble size M , the error in the estimation of the Kalman gain \mathbf{K} can be decomposed into a variance term and a bias term. For a limited ensemble size (compared to n and d), the estimated Kalman gain from the vanilla sEnKF will have a large variance, potentially leading to a filter with diverging state estimation error, as is the case for small ensemble size in the example problems. We show later in this paper that with the proposed dimension reduction, the Kalman gain in the informative subspaces $\tilde{\mathbf{K}}$ contains a much smaller number of entries than \mathbf{K} . This results in smaller variance. The bias in computing $\tilde{\mathbf{K}}$ can be controlled by setting the ranks of the state and observation projections. Due to the rapid decay of the spectrum of the state and observation Gramians, increasing the rank of the projected subspace beyond a certain ratio of the cumulative energy will only marginally reduce the bias, but greatly increase the variance. In other words, we choose the ranks $r_{\mathbf{X}}, r_{\mathbf{Y}}$ to capture the column/row space of the Kalman gain and make the bias small for any finite-sample-estimator of \mathbf{K} . Finally, we should stress that without dimension reduction (i.e., $r_{\mathbf{X}} = n$ and $r_{\mathbf{Y}} = d$), the sEnKF and the LREnKF are equivalent, up to the rotations of the variables. Indeed, the sEnKF estimates the Kalman gain $\mathbf{K} \in \mathbb{R}^{n \times d}$ in the original coordinates, while the LREnKF estimates $\tilde{\mathbf{K}} \in \mathbb{R}^{n \times d}$ in the informative coordinates and lifts it to the original coordinates according to (2.14).

2.c Estimation of the leading directions from samples

In practice, we only have access to limited samples $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ from the prior distribution to estimate the state and observation Gramians $\mathbf{C}_{\mathbf{X}}, \mathbf{C}_{\mathbf{Y}}$ and their associated eigenvectors. To understand how well we can approximate these leading eigenvectors from the empirical Gramians $\hat{\mathbf{C}}_{\mathbf{X}}, \hat{\mathbf{C}}_{\mathbf{Y}}$, we recall the following corollary of the Davis–Kahan theorem [25,26]:

Theorem 2.1 (Corollary of the Davis–Kahan theorem [25,26]¹). *Let $\mathbf{\Gamma}$ and $\hat{\mathbf{\Gamma}}$ be $q \times q$ positive semi-definite matrices with eigendecompositions $\mathbf{\Gamma} \mathbf{w}_i = \lambda_i^2 \mathbf{w}_i$ and $\hat{\mathbf{\Gamma}} \hat{\mathbf{w}}_i = \hat{\lambda}_i^2 \hat{\mathbf{w}}_i$, respectively, with eigenvalues sorted in descending order: $\lambda_1^2 \geq \lambda_2^2 \geq \dots \geq \lambda_m^2 \geq 0$, $\hat{\lambda}_1^2 \geq \hat{\lambda}_2^2 \geq \dots \geq \hat{\lambda}_q^2 \geq 0$. For $1 \leq r < q$ with $\lambda_r^2 - \lambda_{r+1}^2 > 0$, we define $\mathbf{W}_r = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r) \in \mathbb{R}^{q \times r}$ and $\hat{\mathbf{W}}_r = (\hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2, \dots, \hat{\mathbf{w}}_r) \in \mathbb{R}^{q \times r}$. Then, the distance between the subspaces spanned by the columns of \mathbf{W}_r and $\hat{\mathbf{W}}_r$, denoted $d(\mathbf{W}_r, \hat{\mathbf{W}}_r)$, satisfies*

$$d(\mathbf{W}_r, \hat{\mathbf{W}}_r) := \left\| \mathbf{W}_r \mathbf{W}_r^\top \left(\mathbf{I} - \hat{\mathbf{W}}_r \hat{\mathbf{W}}_r^\top \right) \right\|_F \leq \frac{\|\mathbf{\Gamma} - \hat{\mathbf{\Gamma}}\|_F}{\lambda_r^2 - \lambda_{r+1}^2}. \quad (2.15)$$

Let us now apply the Davis–Kahan theorem to the pairs $(\mathbf{\Gamma}, \hat{\mathbf{\Gamma}}) = (\mathbf{C}_{\mathbf{X}}, \hat{\mathbf{C}}_{\mathbf{X}})$, and $(\mathbf{C}_{\mathbf{Y}}, \hat{\mathbf{C}}_{\mathbf{Y}})$. Two important conclusions can be drawn from this result. First, it follows from Theorem 2.1 that the error in the subspace spanned by the first r eigenvectors of $\hat{\mathbf{\Gamma}}$ converges to zero at the same rate that $\hat{\mathbf{\Gamma}}$ converges to $\mathbf{\Gamma}$ in the Frobenius norm, as $M \rightarrow \infty$. See Tropp [27] for relevant results on the convergence rates of random matrices. Second, the error in the estimated r -dimensional leading eigenspace depends on the difference between consecutive eigenvalues $\lambda_r^2 - \lambda_{r+1}^2$, which is called the spectral gap. Specifically, the inverse of the spectral gap is a useful indicator of the estimation error in the informative subspaces. Intuitively, two eigenvectors with close eigenvalues are difficult to distinguish. For the fluid dynamics problems considered in

¹A slightly more general version is presented in [26].




Figure 1. Schematic of the setup. Orange dots depict the location of the point sources. The yellow square (located at $(0, 0)$) depicts the location of the centroid of the point sources. Grey diamonds depict the location of evaluation points of the velocity potential ϕ .

this paper (see examples in Section 3), the spectra of the state and observation Gramians decay rapidly, and hence the leading eigenvalues are well separated. This suggests that we can estimate the leading eigendirections of C_X, C_Y with relatively few samples M , and that the row and column spaces of the Kalman gain can be accurately captured using only a small number of such eigendirections.

3. Examples

3.a The leading directions of the state and observation Gramians and the multipole expansion

In this example, we develop intuition for the ideas presented in this paper by connecting the leading eigenvectors of the state and observation space Gramians discussed in eqns. (2.11) and (2.12) to the multipole expansion. We consider a set of N point sources located at $\{z_1, \dots, z_N\}$ with strengths $\{Q_1, \dots, Q_N\}$, and a set of d evaluation points located at $\{z'_1, \dots, z'_d\}$. We assume that the strengths are perfectly known, but we seek to estimate the positions of the singularities from observations of the velocity potential ϕ at the evaluation points corrupted by additive Gaussian noise $\epsilon \sim \mathcal{N}(\mathbf{0}_d, \Sigma_\epsilon)$. To ensure a convergent multipole expansion, we place all the evaluation points outside a circle of radius R where $R = \max_j |z_j|$. We use the following notations: $\mathbf{x} = [z_1, \dots, z_N]^\top \in \mathbb{C}^N$, $\mathbf{S} = [Q_1, \dots, Q_N]^\top \in \mathbb{R}^N$ and $\mathbf{z}' = [z'_1, \dots, z'_d]^\top \in \mathbb{C}^d$. The superscript \top denotes the transpose of a complex vector/matrix without conjugation, while the superscript H denotes the transpose with conjugation of a complex vector/matrix.

The i th component of the observation vector $\mathbf{y} \in \mathbb{R}^d$ is given by $y_i = \phi(z'_i; \mathbf{x}) + \epsilon_i, i = 1, \dots, d$, where the notation $\phi(z'_i; \mathbf{x})$ is used to highlight the dependence on the positions of the point sources. We seek to estimate the position vector \mathbf{x} from the noisy observations \mathbf{y} . The observation model reads in vector form as $\mathbf{y} = \mathbf{h}(\mathbf{x}) + \epsilon$, where $\mathbf{h}: \mathbb{C}^N \rightarrow \mathbb{R}^d$, $\mathbf{x} \mapsto [\phi(z'_1; \mathbf{x}), \dots, \phi(z'_d; \mathbf{x})]^\top$. From section 2.b, the informative subspaces are identified from the Jacobian of the observation operator \mathbf{h} with respect to the state variable \mathbf{x} , denoted $\nabla_{\mathbf{x}} \mathbf{h}(\mathbf{x})$. If $|z_j|/|z'_k| < 1$ for all k, j , then we get from (1.1)

$$\nabla_{\mathbf{x}} \mathbf{h} = \frac{1}{4\pi} \begin{bmatrix} \frac{-Q_1}{z'_1 - z_1} & \cdots & \frac{-Q_N}{z'_1 - z_N} \\ \vdots & & \vdots \\ \frac{-Q_1}{z'_d - z_1} & \cdots & \frac{-Q_N}{z'_d - z_N} \end{bmatrix} \approx \underbrace{\begin{bmatrix} \frac{1}{z'_1} & \cdots & \frac{1}{z'_1{}^p} \\ \vdots & & \vdots \\ \frac{1}{z'_d} & \cdots & \frac{1}{z'_d{}^p} \end{bmatrix}}_{\mathbf{A}_{\mathbf{z}'}} \underbrace{\begin{bmatrix} 1 & \cdots & 1 \\ z_1 & \cdots & z_N \\ \vdots & & \vdots \\ z_1^p & \cdots & z_N^p \end{bmatrix}}_{\mathbf{B}_{\mathbf{x}}^\top} \text{diag} \left(-\frac{\mathbf{S}}{4\pi} \right), \quad (3.1)$$

where the approximation comes from truncating the multipole expansion at the p -th order. This factorization decouples the contribution of the position of the evaluation points ($\mathbf{A}_{\mathbf{z}'}$), the position of the point sources ($\mathbf{B}_{\mathbf{x}}^\top$), and the volume fluxes of the point sources ($\text{diag}(-\mathbf{S})$). Unfortunately, it is not possible to compute analytically the Gramians C_X, C_Y for this inference

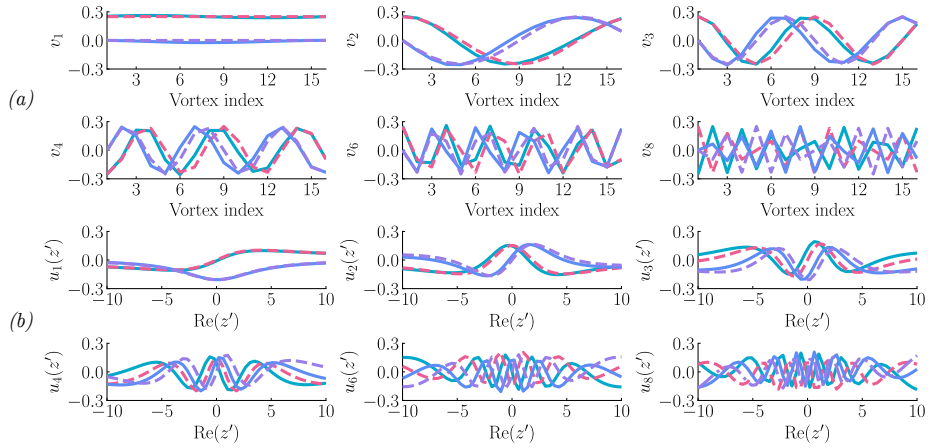


Figure 2. Top panel (a): Comparison of the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th state modes (v_i s) obtained from the SVD of $\nabla_{\mathbf{x}} \mathbf{h}$ (solid turquoise and solid blue lines), and from orthogonalization of the columns of $\mathbf{Q}_{\mathbf{x}}$ (dashed red and dashed purple lines). Lower panel (b): Comparison of the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th observation modes (u_i s) obtained from the SVD of $\nabla_{\mathbf{x}} \mathbf{h}$ (solid turquoise and solid blue lines), and from the columns of $\mathbf{Q}_{\mathbf{z}'}$ (dashed red and dashed purple lines). The vectors have unit norm.

problem, even with a Gaussian distribution for the x and y coordinates of the point source locations. To highlight the connection between the eigenvectors of these matrices and the multipole expansion, we compare the singular vectors obtained from the SVD of $\nabla_{\mathbf{x}} \mathbf{h}$ with the orthonormal bases obtained by orthonormalization of the matrices $\mathbf{A}_{\mathbf{z}'}$, $\mathbf{B}_{\mathbf{x}}$ for a particular position vector \mathbf{x} . We recall that the left singular vectors \mathbf{u}_i of $\nabla_{\mathbf{x}} \mathbf{h}$ are also the eigenvectors of the Gram matrix $\nabla_{\mathbf{x}} \mathbf{h}^H \nabla_{\mathbf{x}} \mathbf{h}$, and similarly, the right singular vectors \mathbf{v}_i of $\nabla_{\mathbf{x}} \mathbf{h}$ are also the eigenvectors of the Gram matrix $\nabla_{\mathbf{x}} \mathbf{h} \nabla_{\mathbf{x}} \mathbf{h}^H$. To compute an orthonormal basis for the columns of $\mathbf{A}_{\mathbf{z}'}$ and $\mathbf{B}_{\mathbf{x}}$, we extract the \mathbf{Q} factor of the QR factorizations of these matrices, denoted by $\mathbf{Q}_{\mathbf{z}'} \in \mathbb{C}^{d \times p}$ and $\mathbf{Q}_{\mathbf{x}} \in \mathbb{C}^{n \times p}$, respectively.

We place a set of $N = 16$ point sources distributed according to $z_J = \rho \exp(2\pi i J/N)$ for $J = 1, \dots, N$, where $\rho = 0.3$. The volume flux of the point sources is set to 4π . The centroid of this set of point sources is at the origin. This is critical for a sensible comparison with the factorization of (3.1). We use $d = 50$ evaluation points distributed along a horizontal line as $z'_k = -10 + (k - 1)\Delta s + 4i$, for $k = 1, \dots, d$ with interspace $\Delta s = 20/(d - 1)$. Fig. 1 depicts the configuration of the different elements. We truncate the multipole expansion at the $p = 30$ th order to get a machine precision approximation of $\nabla_{\mathbf{z}} \mathbf{h}$ in eq. (3.1).

Fig. 2 (a) compares the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th modes obtained from the right singular vectors of $\nabla_{\mathbf{x}} \mathbf{h}$, and from the corresponding columns of the unitary matrix $\mathbf{Q}_{\mathbf{x}}$. Fig. 2 (b) compares the real and imaginary parts of the 1, 2, 3, 4, 6 and 8th modes obtained from the left singular vectors of $\nabla_{\mathbf{x}} \mathbf{h}$, and from the corresponding columns of the unitary matrix $\mathbf{Q}_{\mathbf{z}'}$. The first state mode v_1 is $\pm \mathbf{1}/\sqrt{N}$, where $\mathbf{1}$ denote a vector of ones of length N . Therefore, the leading state correction corresponds to a rigid body translation of the point sources. Three modes capture more than 99.999999% of the cumulative energy of $\nabla_{\mathbf{z}} \mathbf{h}^H \nabla_{\mathbf{z}} \mathbf{h}$. Overall, the state and observation modes obtained from these two procedures show strong similarities, especially for the first three modes. However, the modes are not expected to coincide exactly. The differences in the vectors obtained from these two approaches may be attributed to $\mathbf{Q}_{\mathbf{z}'}$, only relying on the positions of the evaluation points, and $\mathbf{Q}_{\mathbf{x}}$ only relying on the positions of the point sources. In contrast, the SVD simultaneously has access to the locations of the evaluation points and the point sources to identify the left and right singular vectors. As a result, the SVD constructs an *optimized*

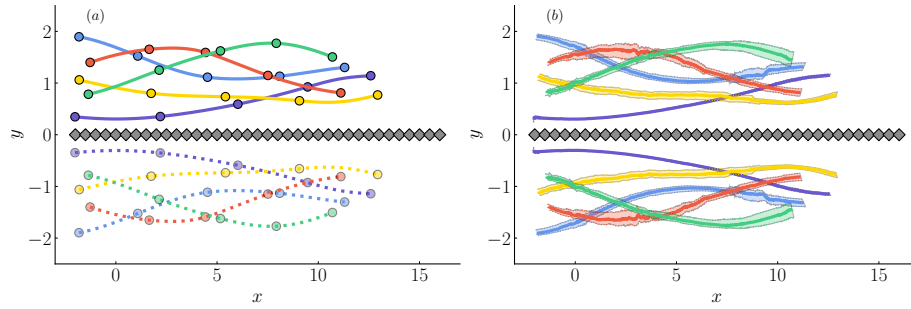


Figure 3. (a): Schematic of the setup. True trajectories of the point vortices (for one realization of the system) are represented by colored lines. The location of the point vortices, sampled every three convective times, is depicted by colored dots. Fainter dashed lines and fainter dots refer to the image point vortices. Location of the pressure sensors are depicted by grey diamonds. (b): Estimation of the trajectories of the vortices with the LREnKF for $M = 40$ with the ranks $r_{\mathbf{x}}$ and $r_{\mathbf{y}}$ set to capture 99% of the cumulative energy spectra. Solid lines depict the time history of the median posterior estimate for the position of the different point vortices. Fainted areas show the 5% and 95% quantiles of the posterior estimate for the position of the point vortices.

“multipole expansion” of the Jacobian (see the Eckart–Young theorem [28]), specifically designed for the particular set of the evaluation points and the point sources.

3.b Inference of the properties of point vortices from pressure observations along a wall

In this second example, we compare the LREnKF with the sEnKF for estimating the positions and strengths of a collection of point vortices advecting along a horizontal wall. We rely on pressure observations collected along the wall to estimate the state. These observations were generated from the same observation model used for inference, thereby making this a *twin experiment* [3]. We consider a set of $N = 5$ point vortices located at $\{z_1, \dots, z_N\}$ with circulations $\{\Gamma_1, \dots, \Gamma_N\}$, and a set of d evaluation points located at $\{z'_1, \dots, z'_d\}$, see Fig. 3 (a). The state variable \mathbf{x} contains the positions and circulations of the N point vortices: $\mathbf{x} = [x_1, y_1, \Gamma_1, \dots, x_N, y_N, \Gamma_N]^\top \in \mathbb{R}^{3N}$. To avoid singular interactions between nearby vortices, we replace the singular Cauchy kernel $k(z) = 1/(2\pi z)$ used to compute the Kirchhoff velocities w_{-j} in (1.2) and (1.4) with the regularized algebraic blob kernel $k_\epsilon(z) = \bar{z}/(2\pi(|z|^2 + \epsilon^2))$, where ϵ is called the blob radius [11], set here to 5×10^{-2} . To enforce the no-flow-through condition along the x axis, we use the method of images and augment our collection of vortices with another set of N vortices at the conjugate positions $\{\bar{z}_1, \dots, \bar{z}_N\}$ with opposite circulation $\{-\Gamma_1, \dots, -\Gamma_N\}$. We emphasize that these mirrored vortices are only an artifice to enforce the no-flow-through in the forecast step, and play no role in the analysis step. We add a freestream flow directed along increasing x values with velocity $\mathbf{U}_\infty = [1, 0]^\top$. The initial position of the j th point vortex is generated randomly with the form $z_j + \rho_r \exp(i\theta)$ where $\{z_1, \dots, z_N\} = \{-2.0 + 0.3i, -1.9 + 1.9i, -1.8 + 1.1i, -1.3 + 1.4i, -1.4 + 0.8i\}$, ρ_r is drawn from $\mathcal{N}(0.0, 0.1)$ and θ is drawn from the uniform distribution on $[0, \pi]$. The initial circulation of the point vortices is drawn from $\mathcal{N}(0.4, 0.1)$. The collection of point vortices is advanced using a forward Euler scheme with time step $\Delta t = 1 \times 10^{-3}$. The observation vector \mathbf{y} consists of pressure observations at $d = 37$ locations linearly distributed with an interspace of 0.5 along the segment $[-2, 16]$, see Fig. 3 (a). We use a Gaussian observation noise with zero mean and covariance 1×10^{-4} . We obtain the Jacobian of the observation operator used in the state and observation Gramians $\mathbf{C}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{y}}$ by analytical differentiation of eq. (1.4).

We assess the performance of the two filters with a *twin experiment*. We draw a random initial condition for the positions and strengths of the different point vortices. Then we simulate the dynamics of the point vortices over the time interval $[0, 12]$ (corresponding to 12000 assimilation




Figure 4. Median spectrum of the state Gramian C_X (left panels (a)–(c)) and observation Gramian C_Y (right panels (d)–(f)) over the time interval $[0, 12]$. Statistics are obtained from a run of the sEnKF with $M = 1000$. Panels [(a), (d)]: Median eigenvalues of C_X and C_Y . Panels [(b), (e)]: Median normalized cumulative energy $E_i = \sum_{j=1}^i \lambda_j^2 / \sum_j \lambda_j^2$ of C_X and C_Y . Panels [(c), (f)]: Median spectral gap $\lambda_i^2 - \lambda_{i+1}^2$ of C_X and C_Y . The abscissa axis is in log scale. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles.




Figure 5. 1, 2, 3 and 8th eigenvectors of the observation Gramian C_Y , i.e., observation modes u_i s, at $t = 0.1$ (left column), $t = 7.0$ (middle column), and $t = 10.0$ (right column). (a)–(c): first mode, (d)–(f): second mode, and (g)–(i): third mode, (j)–(l): eighth mode. The dashed grey vertical line depicts the x component of the position of the centroid of the vortices. The grey horizontal line corresponds to 0 ordinate. Statistics are obtained from a run of the sEnKF with $M = 1000$. The eigenvectors have unit norm.

cycles with the choice of time step Δt). The true state at the time step k is denoted by \mathbf{x}_k^* . At every time step, we generate the noisy pressure observation at the sample locations based on the true state \mathbf{x}_k^* . The realization of the true observation vector at the time step is denoted by \mathbf{y}_k^* . In the twin experiment, we seek to estimate the true state \mathbf{x}_k^* only from the knowledge of the noisy and indirect observations \mathbf{y}_k^* and the prior.

First, we assess the spectra and leading subspaces of the state and observation Gramians. Fig. 4 depicts the median eigenvalues of the state and observation Gramians over the time interval $[0, 12]$. We run the sEnKF with ensemble size $M = 1000$ to generate the samples used to compute empirical Gramians in each analysis step. We choose this larger ensemble size to reduce sampling errors. The median ranks r_X to capture 80%, 90%, 95%, and 99% of the cumulative energy of C_X are 1, 2, 3, and 5, respectively. Similarly, the median ranks r_Y to capture 80%, 90%, and 99% of the cumulative energy of C_Y are 1, 2, and 4, respectively. The sharp decay of the spectra of C_X and C_Y supports our hypothesis of low-rank structure existing in the prior-to-posterior update. A low-dimensional subspace of the pressure observations is only informative along a limited number of directions in the state space. Fig. 5 shows the 1, 2, 3 and 8th eigenvectors of the observation Gramian at the three times: $t = 0.1, 7.0$, and 10.0 . These modes clearly illustrate how information is extracted from the different pressure discrepancies during the analysis step. The




Figure 6. Left column (a): Time-averaged evolution of the median RMSE with the ensemble size M (computed over 50 realizations) with the sEnKF (blue), and with the LREnKF for different ratios of the cumulative energy: 85% (yellow), 95% (orange), and 99% (green). The tracking performance of the sEnKF is unstable for $M < 40$. Dashed lines depict the 25% and 75% quantiles. Right column [(b)-(c)]: Time-history of the median ranks r_X and r_Y of the LREnKF for $M = 30$ (computed over 50 realizations) for different ratios of the cumulative energy of C_X , C_Y : 85% (yellow), 95% (orange), and 99% (green). The dimension of the state and observation spaces, namely n and d , are depicted for comparison in blue.

first observation mode corresponds to a spatially weighted average of the pressure discrepancies about the mean x coordinate of the position of the vortices (depicted by the dashed vertical line on the panels of Fig. 5). The second mode captures differences between the pressure observations collected on the left and right side of the centroid location. The higher modes follow the same pattern and act as refined stencils with growing support to extract higher order features from the pressure observations. The leading eigenvectors of the state Gramian are more difficult to interpret, due to the manner in which vortices are coupled in the determination of pressure.

Next, we compare the performance of the LREnKF with the stochastic EnKF (sEnKF). The comparison of these two filters is natural, as the LREnKF without dimension reduction (i.e., for $r_X = n = 3N$ and $r_Y = d$) reverts to the sEnKF. For a given ensemble size M , one could perform a parametric study to determine the ranks r_X and r_Y which give the best performance of the LREnKF. We found, however, that using fixed ranks for all assimilation cycles is suboptimal as we face one of two possible scenarios: either the ranks are too large, leading to additional variance, or the ranks are too small, leading to additional bias. Furthermore, we expect the dimension of the informative subspaces may vary over time as the flow field evolves. Instead, we select the ranks r_X and r_Y adaptively at each time step based on a predetermined ratio of the cumulative energy of the Gramians C_X and C_Y .

The performance of each filter is assessed using the *root mean squared error (RMSE)* for tracking the true state \mathbf{x}_k^* . The best estimate of the true state at a given assimilation time is given by the mean of the analysis ensemble at this step, that we denote by $\bar{\mathbf{x}}_k^a$. We define the RMSE at one time between the true state \mathbf{x}_k^* and the estimate $\bar{\mathbf{x}}_k^a$ as $\text{RMSE}_k = \|\mathbf{x}_k^* - \bar{\mathbf{x}}_k^a\|_2 / \sqrt{3N}$, where N is the number of point vortices. The mean RMSE of each filter is computed over the time interval $[8, 12]$ (i.e., the last 4000 assimilation steps), to remove any influence of the initial conditions and to ensure empirical stationarity of the filters. The first 8000 assimilation steps are called the spin-up phase and are discarded. The uncertainty in the RMSE is quantified by the 5%, 25%, 75% and 95% quantiles of the time-averaged RMSE over 50 realizations of the same experiment for different initial ensembles. The total computation time for one assimilation step of the sEnKF and the LREnKF with $M = 50$ are 2.3 ms and 8.0 ms, respectively. Given that the LREnKF achieves a lower RMSE for small M , a fairer comparison is to determine the ensemble size needed by these two filters to achieve a tolerated tracking error. For instance, the sEnKF needs $M = 60$ ensemble members for a median RMSE of 0.07, while the LREnKF with a 99% energy ratio only needs $M = 20$ ensemble members. Thus, the performances of the filters are fairly similar given that their

computation times scale linearly with the ensemble size. We stress that the reported computation time of the LREnKF corresponds to the worst case scenario where the entire Jacobian of the observation operator is evaluated at all samples to form the Gramians C_X and C_Y . We provide directions to reduce the computational cost of the LREnKF in the conclusion.

Fig. 6(a) shows the time-averaged RMSE for the LREnKF and the sEnKF. We assess the performance of the LREnKF for 85%, 95%, and 99% of the normalized cumulative energy of C_X and C_Y . For $M < 40$, the sEnKF is unstable, while the LREnKF accurately estimates the true state, even for $M = 10$. For $M \in [40, 60]$, the RMSE of the sEnKF significantly increases as M decreases. For $M = 40$, the median RMSE of the sEnKF is 0.9, while the median RMSEs of the LREnKF with the different energy ratios are smaller than 0.16. For large M , the RMSE of the LREnKF decreases with the energy ratio, but remains larger than the RMSE of the sEnKF. The proposed low-rank approximation of the Kalman gain can lead to a small bias for large ensemble sizes. The regularization technique proposed in this work, however, is designed for the small ensemble size regime; one can always increase the energy ratio (beyond 99%) to recover the performance of the sEnKF for large ensemble sizes. The RMSE of the LREnKF with the different energy ratios decreases for $M \in [10, 40]$, then plateaus. In the small ensemble size regime, adding samples improves the estimate of the dominant directions (see the Davis–Kahan theorem 2.1), leading to a reduction in variance. As we continue to increase the number of samples for a given number of dimensions to estimate, the variance becomes smaller than the bias from using a truncated basis to approximate the Kalman gain. For $M = 10$, the LREnKF with 85% energy ratio performs better than the LREnKF with 99% energy ratio. For $M = 10$, the ranks needed to achieve 99% is large, leading to additional variance. On the other hand, 85% energy requires a much smaller number of entries to estimate, leading to less variance compared to the bias of this approximation. We recommend choosing the target energy ratio based on the available ensemble size.

Fig. 6(b)–(c) show the history of the median ranks r_X and r_Y (computed over 50 realizations) for the different ratios of the cumulative energy for $M = 30$. The ranks plateau over the time window and remain small compared to the dimension of the state and observation spaces. For the LREnKF with 99% of the cumulative energy, r_X, r_Y are smaller than 8, 6, respectively. Fig. 3(b) shows one estimate of the trajectories of the vortices over the time interval $[0, 12]$. The results are obtained from the LREnKF with $M = 50$ with the ranks r_X and r_Y set to capture 99% of the energy spectra. We quantify the uncertainty with the 5% and 95% quantiles of the posterior for the point vortex positions. We observe an excellent agreement with the true trajectories with a time-averaged RMSE of 0.038.

In Section 2.b, we argue that distance localization schemes can be harmful to regularize inference problems with elliptic observation models, as studied herein. We conclude this section by providing an *a posteriori* justification based on the results of the sEnKF run with a large ensemble ($M = 1000$). Fig. 7(a)–(c) shows the magnitude of the cross-covariance between the positions and strengths of the different point vortices and the pressure observations at $t = 1.0$. We notice that the cross-covariance entries decay very quickly on a small support about each point vortex. Past a certain radius, however, the cross-covariance entries quickly increase, and beyond this point, they only decay algebraically. More precisely, we observe a decay as the inverse square of the distance between the point vortex and the pressure observation. This decay rate is expected from the pressure field derived in (1.4). In Fig. 7(c), the covariance between the circulation of three point vortices and pressure observations at a distance of 5 and 10 units away is still about 20% and 10% of its maximal value, respectively. This algebraic decay clearly violates the assumption of rapidly decaying correlations of distance localization schemes. The cross-covariance of the components of the position and the strength for the different point vortices can have different algebraic decay rates as a function of the distance. Even for a particular vortex, the variations of the cross-covariance for different components of the position and the strength are drastically different. In unreported results, we observe that the decay rate of the correlations also varies over time as the vortices evolve. In the best case, this suggests that multiple different localization radii (i.e., three for the state and one for the data) are needed in principle and would have to be tuned

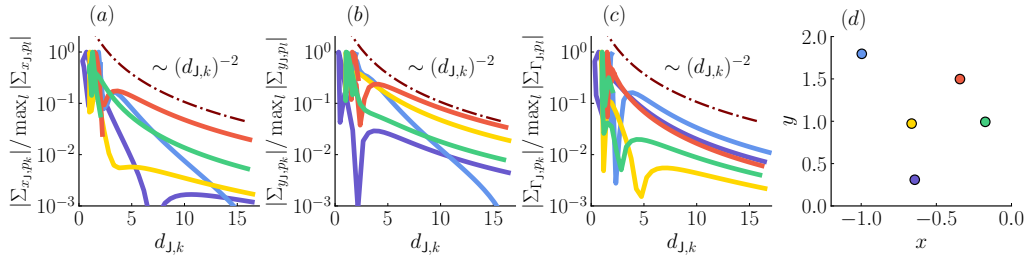


Figure 7. Panel (a), (b), (c) depicts the magnitude of the empirical cross-covariance between the x , y coordinate and strength of the J point vortex with the k th pressure observation at $t = 1.0$, respectively. The curves are plotted against the distance $d_{J,k} = |z_J - z'_k|$, where $z_J = x_J + iy_J$ is the (complex) position of the J th point vortex, and z'_k is the (complex) location of the k th pressure sensor. The magnitude of the cross-covariances are normalized by the maximum cross-covariance (in magnitude) between the x , y coordinate or strength of the J point vortex with the different pressure observation, respectively. The dashed grey vertical line depicts the x component of the position of the centroid of the vortices. Red dashed and dotted curve depicts the algebraic decay $O(1/d_{J,k}^2)$. (d) Mean position of the different point vortices at $t = 1.0$. The same color is used to depict the properties of a point vortex on the panels (a)-(d). The results are obtained from a run of the sEnKF for $M = 1000$.

independently, which is typically impractical. These different facts support our *a priori* hypothesis that distance localization is not uniformly well suited to regularize filtering problems with elliptic observation operators. Moreover, while distance localization can be successful in regularizing the sEnKF for some problems, its underlying assumptions are clearly violated by the algebraic decay in this example. Thus, the application of distance localization should be carefully considered. In unreported results that motivated this work², distance localization was detrimental and worsened the performance of the EnKF in some aerodynamic estimation problems; see [12] for details on the setting.

3.c Inference of a vortex patch advected along a wall from pressure observations.

In this third example, we infer the properties of a circular vortex patch advected along the x axis. A circular vortex patch \mathcal{V} is a disk of uniform vorticity ω_0 . As in 3.b, the vortex patch is constrained to advect along the x axis by the method of images [11]. The total circulation of the vortex patch \mathcal{V} , denoted $\Gamma_{\mathcal{V}}$, is given by $\Gamma_{\mathcal{V}} = \int_{\mathcal{V}} \omega_0 dS(\mathbf{x})$ and the vorticity centroid, denoted $\mathbf{x}_{\mathcal{V}}$, is given by $\mathbf{x}_{\mathcal{V}} = \Gamma_{\mathcal{V}}^{-1} \int_{\mathcal{V}} \mathbf{x} \omega_0 dS(\mathbf{x})$. The problem is parameterized by the ratio of the radius $r_{\mathcal{V}}$ of the vortex patch to the distance $d_{\mathcal{V}}$ between the centroid of the vortex patch and its image centroid.

We discretize each vortex patch with a collection of regularized point vortices; further details are given in [11]. In this study, the vortex patch is discretized by $N_r = 4$ concentric rings of blobs, leading to a subdivision of the vortex patch into $N_{pv} = 1 + 4N_r(N_r - 1) = 49$. The initial position of the J th point vortex is generated randomly with the form $z_J + \rho_r \exp(i\theta)$, where z_J denotes the nominal position of the J th vortex. ρ_r is drawn from $\mathcal{N}(0.0, \sigma_r^2)$, where σ_r corresponds to 10% of the radius between two concentric rings of point vortices, and θ is drawn from the uniform distribution on $[0, \pi]$. The initial circulation of the point vortices is drawn from $\mathcal{N}(\Gamma_{\mathcal{V}}/N_{pv}, 10^{-4})$. For a blob at z_J with circulation Γ_J in the vortex patch \mathcal{V} , there is an image blob located at \bar{z}_J with circulation $-\Gamma_J$. The state is defined, as in the previous example, by the positions and strengths of the point vortices in the upper half-plane. For N_{pv} blobs, the state dimension is $n = 3N_{pv} = 147$. We use a forward Euler model scheme with time step $\Delta t = 5.0 \times 10^{-3}$. The vortex patches are evolved over the time interval $[0, 12]$. The observation vector \mathbf{y} consists of pressure observations collected at $d = 24$ locations with interspace 0.5 along the segment $[-1.5, 10.0]$ of the x axis. The

²See <https://github.com/mleprovost/LocalizedVortex.jl.git> for the numerical results of the localized EnKF.

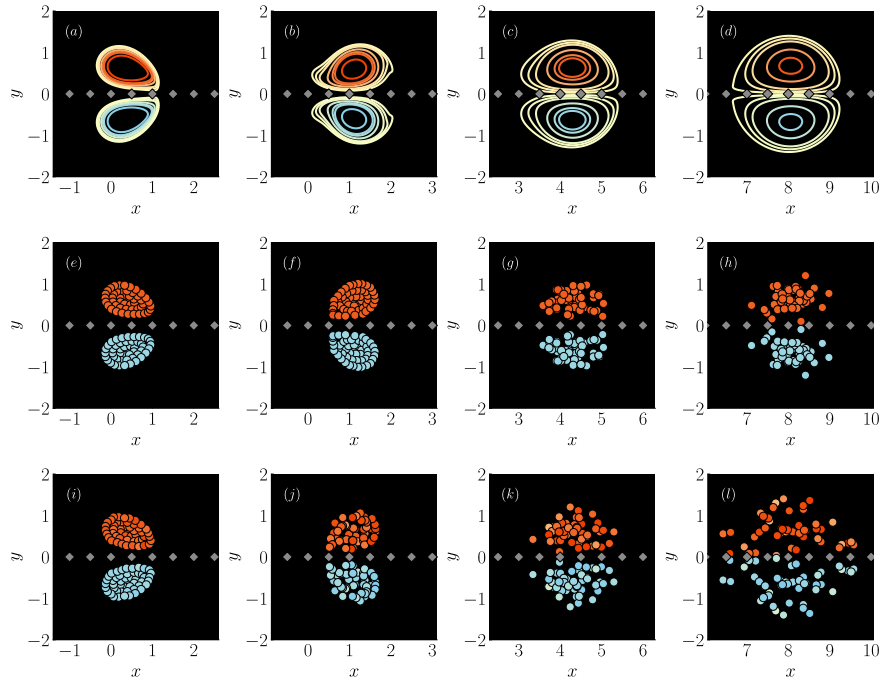


Figure 8. Snapshots of the vorticity distribution at $t = 0.5$ (left column), $t = 1.5$ (second column), $t = 6.0$ (third column), and $t = 12.0$, for a vortex patch and its image, predicted from [(a)-(d)] high-fidelity numerical simulation at Reynolds number 1000, [(e)-(h)] inviscid vortex model with LREnKF with $M = 30$, and [(i)-(l)] inviscid vortex model with sEnKF with $M = 30$. The ranks r_X and r_Y of the LREnKF are set to capture 99% of the normalized cumulative energy of C_X and C_Y , respectively. Pressure sensors are depicted with grey diamonds. Orange denote positive vorticity, while blue vorticity denotes negative vorticity.

observation noise is Gaussian with zero mean and covariance 4×10^{-2} (corresponding to a ratio of peak pressure amplitude to standard deviation of the noise equal to 10).

In this example, the dynamical model that generated the observations is given by solving the incompressible Navier-Stokes equations at a Reynolds number $\Gamma_V/\nu = 1000$, where ν is the kinematic viscosity of the fluid. The true solution contains viscous effects, which are not explicitly included in the forecast model we use for inference. The initial condition is constructed by discretizing the two vortex patches on a uniform Cartesian grid with grid spacing $\Delta x = 0.01$ for the same geometry (i.e., r_V , x_V , and d_V are identical) and the same total circulation Γ_V . We use a high-fidelity Navier Stokes solver with lattice Green function (LGF) [29,30]. The true pressure observations are generated by inverting the pressure Poisson equation (1.3) with LGF. Panels (a)-(d) of Fig. 8 depict the truth vorticity field at $t = 0.5, 1.5, 6.0$ and 12.0 .

To the best of our knowledge, there is no straightforward way to compare the discrete vorticity distribution generated by the ensemble filter with the true continuous vorticity distribution obtained from the Navier-Stokes solver. Instead, we compare the pressure distribution along the x axis from the high-fidelity simulation with the evaluation of the observation model (2.2) at the posterior ensemble generated by the LREnKF, or the sEnKF (i.e., posterior predictive samples for the pressure distribution). The performance of the LREnKF is assessed for 85%, 95% and 99% of the normalized cumulative energy for C_X and C_Y . We assess the performance of the filter on the segment $[-1.5, 11.0]$ with a finer interspace distance of 0.01. We define the median-squared-error (MSE) at time k between the true pressure distribution y_k^* and the median posterior pressure distribution median(y_k^a) (i.e., the median pressure distribution for the posterior ensemble at time k) as $\text{MSE}_k = \|y_k^* - \text{median}(y_k^a)\|_2 / \dim(y_k^*)$, where $\dim(v)$ denotes the dimension of a vector v . Due to the singular nature of the pressure field, we found that the MSE is a more robust estimator

of the performance of each filter than the RMSE due to the presence of outliers in the ensemble estimate of the sEnKF. The spin-up phase of the filters is $[0.0, 8.0]$. The reported MSE is time-averaged over the remaining time interval $[8.0, 12.0]$. The uncertainty in the MSE is quantified by the 5%, 25%, 75% and 95% quantiles of the MSE over 50 realizations of the same experiment with different samples for the initial condition. The computation time for one assimilation step of the sEnKF and the LREnKF with $M = 50$ are 9.8 ms and 42.5 ms, respectively. To achieve a median MSE of at most 0.04, the sEnKF needs $M = 30$ ensemble members, while the LREnKF with the different energy ratios studied only needs $M = 12$ ensemble members.

Fig. 9 shows the median spectra of C_X and C_Y over the entire time interval. The statistics are obtained from a run of the sEnKF for a large ensemble ($M = 1000$). As in the previous example, the inference problem possesses low-rank informative structure. The median ranks r_X needed to capture 80%, 90%, 95%, and 99% of the cumulative energy of C_X are 2, 3, 4, and 16, respectively. Similarly, the median ranks r_Y to capture 80%, 90%, and 99% of the cumulative energy of C_Y are 2, 3, and 5, respectively. These ranks are small compared to the state dimension $n = 147$, and the observation dimension $d = 24$.

Fig. 10 (a) reports the evolution of the median MSE (computed over 50 realizations) for the sEnKF and the LREnKF with the ensemble size M . We assess the performance of the LREnKF for 85%, 95%, and 99% of the normalized cumulative energy of C_X and C_Y . We recall that the LREnKF reverts to the sEnKF when the dimensions are not reduced. For $M \geq 60$, there is no significant difference in the MSE of the sEnKF and the LREnKF (for the different energy ratios). For $M \in [30, 50]$, the MSE of the sEnKF significantly increases as M decreases. Over this interval, the MSE of the LREnKF shows no major variation. Overall, capturing 85% of the cumulative energy of C_X and C_Y is sufficient to yield stable inference results. For $M < 30$, the MSE of the sEnKF diverges, while the LREnKF leads to a reasonable estimate of the pressure distribution for ensemble size as small as 10. This clearly demonstrates the benefit of our regularization for the EnKF. Fig. 10(b)-(c) shows the time history of the median ranks r_X and r_Y required to capture at least 85, 95, 99% of the cumulative energy (computed over 50 realizations) for $M = 20$. The rank r_X initially increases with time, until it reaches a plateau at about $t = 3.0 - 7.0$ depending on the energy ratio. The increase of the rank r_X can be related to the growing role of viscosity in the true pressure response. The rank r_Y remains close to 5 over time for the different energy ratios.

Fig. 8 compares the vorticity distribution at four times, $t = 0.5, 1.5, 6.0$, and 12.0 , from the truth with the distribution of vortex elements of the posterior mean of the LREnKF and the sEnKF. The ensemble estimates are obtained from one realization of each filter for $M = 30$. The ranks of the LREnKF are set to capture at least 99% of the normalized cumulative energy. The large-scale vortex features are well captured by the inviscid vortex model, despite the obvious absence of the viscous effects in our flow representation. The state estimate is informed of the viscous effects through the assimilation of the pressure observations [12]. Overall, the LREnKF provides a more physically consistent estimate of the vorticity distribution than the sEnKF. With a limited ensemble size, the sEnKF cannot distinguish between physical and spurious long-range correlations between the pressure observations and the vortices. As a result, the sEnKF inconsistently displaces the vortices. At $t = 12.0$, the vorticity distribution of the LREnKF contains a dense core of vortices, with a few vortex satellites to capture the viscous diffusion of the vortex patch. However, the sEnKF poorly estimates the truth vorticity distribution for $t > 6.0$. The spatial structure of the vortices is lost, and they occupy a much larger support than the true vorticity distribution.

Fig. 11 compares the history of the pressure from the truth with the posterior mean estimate of the LREnKF and the sEnKF for the same realization of the filters. Over time, the pair of vortex patches is advected along the x axis. Concomitant to the diffusion of the vortex patches, the peak pressure (in magnitude) decreases over time. The true pressure is globally well approximated by the two filters. As we pointed in the previous paragraph, the sEnKF creates spurious displacements of the vortices over time, leading to a growing error in the posterior predictive pressure at later times (the red regions in Fig. 11(e) for $t > 6.0$).




Figure 9. Median spectrum of the state Gramian C_X (left panels (a)–(c)) and observation Gramian C_Y (right panels (d)–(f)) over the time interval $[0, 12]$. Statistics are obtained from a run of the sEnKF with $M = 1000$. Panels [(a), (d)]: Median eigenvalues of C_X and C_Y . Panels [(b), (e)]: Median normalized cumulative energy $E_i = \sum_{j=1}^i \lambda_j^2 / \sum_j \lambda_j^2$ of C_X and C_Y . Panels [(c), (f)]: Median spectral gap $\lambda_i^2 - \lambda_{i+1}^2$ of C_X and C_Y . The abscissa axis is in log scale. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles.




Figure 10. Left column (a): Time-averaged evolution of the median MSE of the posterior pressure with the ensemble size M (computed over 50 realizations) with the sEnKF (blue), and with the LREnKF for different ratios of the cumulative energy: 85% (yellow), 95% (orange), and 99% (green). The sEnKF is found unstable for $M < 30$. Dashed lines depict the 25% and 75% quantiles. Dotted lines depict the 5% and 95% quantiles. Right column [(b)–(c)]: Time-history of the median ranks r_X and r_Y of the LREnKF for $M = 20$ (computed over 50 realizations) for different ratios of the cumulative energy of C_X , C_Y : 85% (yellow), 95% (orange), and 99% (green). The dimension of the state and observation spaces, namely n and d , are depicted for comparison in blue.

4. Conclusion

In this work, we presented a new regularization technique for the EnKF with elliptic observation models. In these problems, the strength of interactions between states (i.e., sources) and the observations (i.e., targets) decays slowly (typically algebraically) with physical distance. As a result, distance localization should be considered with care, as we cannot dissociate true/physical long-range correlations from the spurious long-range correlations induced by limited ensemble sizes. Instead, we observe that it is sufficient to update a limited number of directions of the state space using a low-dimensional projection of the observations. We propose to leverage this structure to regularize the EnKF. This structure yields a low-rank factorization of the Kalman gain in linear–Gaussian observation models, which is computed from the SVD of a whitened observation matrix. We extended this factorization to nonlinear–Gaussian observation models with the help of two tools. First, we proposed a methodology to identify the informative directions of the state and observation spaces for a nonlinear observation model. These directions are obtained as the leading eigenvectors of the state Gramian C_X and observation Gramian C_Y , which are based on the Jacobian of the observation operator h . These eigenvectors revert to the singular vectors of the whitened observation model (up to a sign) in the linear–Gaussian case.

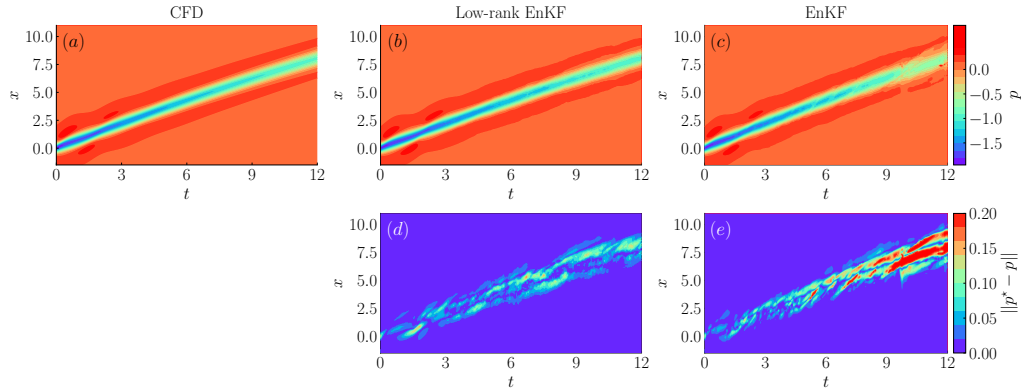


Figure 11. Top line [(a), (b), (c)]: Spatiotemporal map of the pressure induced by a pair of two co-rotating vortex patches from (a) high-fidelity numerical simulation at Reynolds number 1000, and the mean posterior of an inviscid vortex model for one realization of the LREnKF (b), and the sEnKF (c) with $M = 30$. Bottom line [(d), (e)]: Spatiotemporal magnitude of the error between the truth pressure and the mean posterior pressure for one realization of the LREnKF (d), and the sEnKF (e) with $M = 30$. The ranks r_X and r_Y of the LREnKF are set to capture 99% of the normalized cumulative energy of C_X and C_Y , respectively. The same color levels are used for all the top panels, and another set for the bottom panels.

We show how to use these informative directions to define a prior-to-posterior transformation that updates a low-dimensional subspace of the states based on a projection of the observations. This transformation requires estimating a lower-dimensional Kalman gain in the span of the informative directions, thereby resulting in estimators with *lower variance* than the sEnKF. Furthermore, for elliptic observation operators, the spectra of C_X and C_Y are rapidly decaying. This makes the bias from low-rank approximations small for any budget of samples. Setting the ranks r_X and r_Y a priori is not desired, as the dimension of the informative subspace may vary over time (see Examples 3.b, 3.c). Instead, we showed that the ranks r_X and r_Y can be adaptively selected to capture a given ratio of the cumulative energy of C_X and C_Y . We stress that these reduced coordinates are not limited to the estimation of low-dimensional linear transformations. Le Provost et al. [31] used this informative subspace to identify parsimonious and interpretable *nonlinear* prior-to-posterior transformations for aerodynamic applications [22,32].

In Example 3.a, we showed that the leading eigenvectors of C_X and C_Y are reminiscent of the modes from multipole expansions. Nonetheless, the proposed methodology is more general, as it only relies on the Jacobian of the observation operator and works even when observations are taken near the sources. In Example 3.b, the first eigenvectors of C_X form a hierarchical basis of discrete stencils to extract information from the pressure observations. We assessed the LREnKF and the sEnKF on two problems where we sought to estimate the positions and strengths of point vortices from pressure observations. In the second example, the pressure observations were obtained from a high-fidelity simulation of the Navier-Stokes equations. In these two examples, we showed that the LREnKF significantly improved the posterior estimate of the sEnKF by yielding lower RMSE for smaller ensemble sizes.

One limitation of this work is the computational cost of forming the Gramians C_X and C_Y . Fortunately, two key ideas can be used to reduce this cost. First, we do not need the entire set of eigenpairs of C_X and C_Y . The LREnKF only requires a basis for the leading eigenspaces of C_X and C_Y . Given rapid spectral decay of these Gramians, the subspaces spanned by their leading eigenvectors can be estimated with a relatively limited number of samples. Therefore, it may not be necessary to evaluate the Jacobian of the observation operator for every member of the ensemble; instead one could consider a subset of the ensemble, selected by random subsampling or perhaps clustering.

Second, computation of the full Jacobian matrix for each sample is not necessary. By leveraging ideas from randomized numerical linear algebra [33,34], one can construct accurate estimates of the leading eigenspaces of $C_{\mathbf{X}}$ and $C_{\mathbf{Y}}$ using only matrix-vector products.

We also emphasize that the proposed regularization of the EnKF is by no means limited to potential flow problems, nor fluid mechanics more generally. Indeed, elliptic observation models are found across a spectrum of data assimilation problems in engineering and science: geoscience, electromagnetism, and heat transfer, to name a few. For all these problems, regularization of the EnKF has usually been *ad hoc* [6,35]. The principled regularization of the EnKF described in this paper is applicable to any differentiable observation model, without any tuning required from the user—the only free parameters being the ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ which can be set adaptively, as in the examples we presented.

Finally, we comment on the link between our methods and *inverse problems*. While the focus of this paper was on *filtering* problems and their solution using ensemble methods, every assimilation step with an elliptic observation operator is essentially a Bayesian inverse problem, as noted in the Introduction. The inverse problems viewpoint inspires the regularization approaches we develop, and in fact these regularization methods are immediately applicable to ensemble algorithms [36,37] for a wide range of Bayesian inverse problems, as well as to other solution algorithms that could benefit from joint dimension reduction of the data and parameters [38]. In other words, our developments are not restricted to filtering. That said, the specific examples and problem settings that we consider in this paper, as well as the performance metrics and comparisons we use for evaluation (e.g., RMSE, small ensemble sizes), are representative of filtering applications. Evaluation of these methods for the typical settings of Bayesian inverse problems is left to future work.

Data Accessibility. All the computational results are reproducible and code is available at: <https://github.com/mleprovost/LowRankVortex.jl.git>

Authors' Contributions. All authors contributed to the calculations, analysis, and writing and revision of the manuscript. The final version of the manuscript was approved by all authors.

Funding. MLP and JE gratefully acknowledge support of the Air Force Office of Scientific Research Grant No. FA9550-18-1-0440. MLP acknowledges support of the Dissertation Year Fellowship of the Graduate Division of the University of California, Los Angeles. RB and YM acknowledge support from the Department of Energy, Office of Advanced Scientific Computing Research, AEOLUS (Advances in Experimental design, Optimal control, and Learning for Uncertain complex Systems) center.

Appendix

A. The transport equations for singularities

Here, we present the derivation of the transport equations for singularities in potential flow. We recall for reference the complex potential F induced at a location $z \in \mathbb{C}$ by a collection of N singularities located at $\{z_1, \dots, z_N\}$ with complex strengths $\{S_1, \dots, S_N\}$ subject to a freestream velocity given in complex notation $W_\infty = U_\infty - iV_\infty$.

$$F(z) = \sum_{j=1}^N \frac{S_j}{2\pi} \log(z - z_j) + W_\infty z. \quad (\text{A.1})$$

The derivative of the complex potential is the complex velocity w given by:

$$w(z, t) = \frac{dF}{dz} = \sum_{j=1}^N \frac{S_j}{2\pi} \frac{1}{z - z_j} + W_\infty. \quad (\text{A.2})$$

For a velocity field given in vector notation by $\mathbf{v}(\mathbf{x}, t) = (u(\mathbf{x}, t), v(\mathbf{x}, t))$, we stress the classical conjugation (i.e., minus sign) in the definition of the equivalent complex velocity $w(z, t) = u(z, t) - iv(z, t)$. Using the induced velocity (A.2), we can write the transport equation for the positions and strengths of the singularities. By simple inspection of (A.2), the velocity field is singular at the locations of the singularities. Per Kirchhoff's law, a point singularity cannot induce velocity on itself [11]. Instead, a singularity is transported by the total induced velocity on itself minus its own contribution. The regularized velocity of the J th singularity is called the Kirchhoff velocity, denoted \mathbf{v}_{-J} in vector notation, or w_{-J} in complex notation. The transport equations for the singularities follow:

$$\frac{dz_J}{dt} = \overline{w_{-J}}(t) = \sum_{K=1, K \neq J}^N \frac{\overline{S_J}}{2\pi} \frac{1}{\overline{z_J} - \overline{z_K}} + \overline{W_\infty}, \quad \frac{dS_J}{dt} = 0, \quad J = 1, \dots, N. \quad (\text{A.3})$$

B. The pressure Poisson equation and the Bernoulli equation

In this section, we present the derivation of the pressure Poisson equation for potential flows and the analytical expression for the pressure field obtained from the unsteady Bernoulli equation. The Euler equation and the associated divergence and curl conditions on the velocity field are given by:

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho \mathbf{v} \cdot \nabla \mathbf{v} = -\nabla p, \quad (\text{B.1a})$$

$$\nabla \cdot \mathbf{v}(\mathbf{r}, t) = \sum_{J=1}^N Q_J \delta(\mathbf{r} - \mathbf{r}_J), \quad (\text{B.1b})$$

$$\nabla \times \mathbf{v}(\mathbf{r}, t) = \sum_{J=1}^N \Gamma_J \delta(\mathbf{r} - \mathbf{r}_J). \quad (\text{B.1c})$$

The pressure Poisson equation is formed by taking the divergence of the Euler equation (B.1a) and using the identity $\mathbf{v} \cdot \nabla \mathbf{v} = \nabla(\|\mathbf{v}\|^2/2) - \mathbf{v} \times \boldsymbol{\omega}$:

$$\nabla^2 \left(p + \frac{1}{2} \rho \|\mathbf{v}\|^2 \right) = \rho \nabla \cdot (\mathbf{v} \times \boldsymbol{\omega}) - \rho \frac{\partial \nabla \cdot \mathbf{v}}{\partial t}, \quad (\text{B.2})$$

where $\boldsymbol{\omega}(\mathbf{x}, t)$ is the vorticity. Using (B.1b) and (B.1c), we obtain the pressure Poisson equation:

$$\nabla^2 \left(p(\mathbf{r}, t) + \frac{1}{2} \rho \|\mathbf{v}(\mathbf{r}, t)\|^2 \right) = \rho \sum_{J=1}^N [Q_J \mathbf{v}_{-J} \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_J) - \Gamma_J e_z \cdot (\mathbf{v}_{-J} \times \nabla \delta(\mathbf{r} - \mathbf{r}_J))]. \quad (\text{B.3})$$

For a fixed evaluation point denoted \mathbf{r}' in vector notation or z' in complex notation, the unsteady Bernoulli equation reads [11]:

$$p(\mathbf{r}', t) + \rho \frac{1}{2} \|\mathbf{v}(\mathbf{r}', t)\|^2 + \rho \frac{\partial \phi(\mathbf{r}', t)}{\partial t} = B(t), \quad (\text{B.4})$$

where $B(t)$ is the Bernoulli constant. The second term of (B.4) is called the quadratic term, while the third term is called the unsteady term. Using the results derived in A, we get

$$p(z', t) - B(t) = -\frac{1}{2} \rho \left| \sum_{J=1}^N \frac{S_J}{2\pi} \frac{1}{z' - z_J} + W_\infty \right|^2 + \rho \Re \left(\sum_{J=1}^N \frac{S_J}{2\pi} \frac{1}{z' - z_J} \overline{w_{-J}}(t) \right). \quad (\text{B.5})$$

C. The filtering problem

In this section, we give a basic outline of the filtering problem. We recall for reference the dynamical and observation models used in this work. The dynamical model is given by

$$\mathbf{X}_k = \mathbf{f}_k(\mathbf{X}_{k-1}) + \mathbf{W}_k, \quad (\text{C.1})$$

where $\mathbf{f}_k: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a deterministic function and \mathbf{W}_k is an independent *process noise*. The observation model is given by a nonlinear observation operator $\mathbf{h}_k: \mathbb{R}^n \rightarrow \mathbb{R}^d$ corrupted by an additive observation noise \mathcal{E}_k :

$$\mathbf{Y}_k = \mathbf{h}_k(\mathbf{X}_k) + \mathcal{E}_k. \quad (\text{C.2})$$

In the *filtering problem*, we aim to characterize the *filtering distribution* $\pi_{\mathbf{X}_k | \mathbf{Y}_{1:k} = \mathbf{y}_{1:k}}$ (also called the *posterior distribution*), i.e., the probability distribution of the state at time k conditioned on the realizations $\mathbf{y}_1, \dots, \mathbf{y}_k$ of the observation process up to that time. To tackle this problem in high dimensions, we consider a Monte Carlo approximation of the filtering distribution. To do so, we recursively update a set of M *particles* $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ that approximate the posterior distribution. This defines a generic class of methods known as ensemble filtering methods. At every time step, the recursion takes as input the particle approximation $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ of the posterior distribution at time $k-1$ and seeks a particle approximation of the posterior distribution at the next time. These methods recursively apply a two-step procedure: the *forecast step* and the *analysis step*. In the forecast step, each particle \mathbf{x}^i is propagated in the dynamical equation (C.1) to form a Monte Carlo approximation of the forecast distribution $\pi_{\mathbf{X}_k | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}}$. The analysis step treats the forecast as the prior and solves a Bayesian inference problem to condition on a new observation of the true system \mathbf{y}_k^* . Ensemble methods realize this conditioning by moving particles $\{\mathbf{x}^i\}_{i=1}^M$ that represent the prior/forecast distribution into new positions so that they form an empirical approximation of the posterior. We note that different ensemble filtering methods have a common forecast step but differ in the analysis step. Our treatment of the analysis step builds a transformation \mathbf{T}_k that essentially maps samples from the prior distribution $\pi_{\mathbf{X}_k | \mathbf{Y}_{1:k-1} = \mathbf{y}_{1:k-1}}$ to the posterior distribution $\pi_{\mathbf{X}_k | \mathbf{Y}_{1:k} = \mathbf{y}_{1:k}}$ [21,22]. For a linear–Gaussian state-space model, Kalman [39] derived in closed form an exact linear prior-to-posterior transformation for the analysis step. In more general settings, the prior-to-posterior transformation is necessarily nonlinear, but the linear transformation of the Kalman filter is still a foundation for algorithms such as the *ensemble Kalman filter* (EnKF) introduced by Evensen [1]. In particular, the EnKF and other ensemble filtering methods estimate a linear analysis map \mathbf{T}_k from samples. In order to simplify notation, we omit the time dependence subscripts of the variables in the rest of this Supplementary Material, since the analysis step can be treated as a *static* Bayesian inference problem.

D. Low-rank assimilation for the linear–Gaussian case

In this section, we present a detailed derivation of the factorization of the Kalman gain for a linear–Gaussian observation model, recalled for reference:

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathcal{E}, \quad (\text{D.1})$$

where the state is given by $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{X}})$ and the observational error is given by $\mathcal{E} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathcal{E}})$ where \mathcal{E} is independent of \mathbf{X} . The whitened state and noise observation variables are defined as $\tilde{\mathbf{X}} = \Sigma_{\mathbf{X}}^{-1/2}\mathbf{X} \in \mathbb{R}^n$ and $\tilde{\mathcal{E}} = \Sigma_{\mathcal{E}}^{-1/2}\mathcal{E} \in \mathbb{R}^d$. In the whitened space, the observation model becomes:

$$\tilde{\mathbf{Y}} = \Sigma_{\mathcal{E}}^{-1/2}\mathbf{Y} = \tilde{\mathbf{H}}\tilde{\mathbf{X}} + \tilde{\mathcal{E}}, \quad (\text{D.2})$$

where $\tilde{\mathbf{H}} = \Sigma_{\mathcal{E}}^{-1/2}\mathbf{H}\Sigma_{\mathbf{X}}^{1/2} \in \mathbb{R}^{d \times n}$ is the whitened observation matrix, whose singular value decomposition (SVD) reads:

$$\tilde{\mathbf{H}} = \mathbf{U}\mathbf{A}\mathbf{V}^{\top}, \quad (\text{D.3})$$

where $\mathbf{U} \in \mathbb{R}^{d \times d}$ and $\mathbf{V} \in \mathbb{R}^{n \times d}$ are the left and right singular vectors, and $\mathbf{\Lambda} \in \mathbb{R}^{d \times d}$ is the diagonal matrix of singular values. As in the main text, we consider the case $d \leq n$. We recall the definitions of the projected state, observation noise, and observation variables as $\tilde{\mathbf{X}} = \mathbf{V}^\top \tilde{\mathbf{X}} \in \mathbb{R}^d$, $\tilde{\mathbf{\epsilon}} = \mathbf{U}^\top \tilde{\mathbf{\epsilon}} \in \mathbb{R}^d$, and $\tilde{\mathbf{Y}} = \mathbf{U}^\top \tilde{\mathbf{Y}} \in \mathbb{R}^d$. In the rotated spaces, the observation model can be written as:

$$\tilde{\mathbf{Y}} = \mathbf{\Lambda} \tilde{\mathbf{X}} + \tilde{\mathbf{\epsilon}}, \quad (\text{D.4})$$

where the rotated observation operator $\mathbf{\Lambda}$ and the observation error covariance are diagonal matrices. Hence, the inference in this rotated space can be performed in a fully decoupled manner for each state variable. Using the decoupled observation model (D.4), the Kalman gain in the rotated space is given by:

$$\tilde{\mathbf{K}} = \mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1}. \quad (\text{D.5})$$

Thus, the linear analysis map in the rotated space \tilde{T} is given by:

$$\tilde{T}(\tilde{\mathbf{y}}, \tilde{\mathbf{x}}) = \tilde{\mathbf{x}} - \tilde{\mathbf{K}}(\tilde{\mathbf{y}} - \tilde{\mathbf{y}}^*) = \tilde{\mathbf{x}} - \mathbf{\Lambda}(\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1}(\tilde{\mathbf{y}} - \tilde{\mathbf{y}}^*), \quad (\text{D.6})$$

where $\tilde{\mathbf{y}}^* = \mathbf{U}^\top \Sigma_{\tilde{\mathbf{\epsilon}}}^{-1/2} \mathbf{y}^*$ denotes the whitened and rotated realization of the assimilated data. The d columns of \mathbf{V} span the informative subspace of the whitened state space (i.e., where the random variable $\tilde{\mathbf{X}}$ lives). Thus, $\mathbf{P}_d = \mathbf{V}\mathbf{V}^\top \in \mathbb{R}^{n \times n}$ is an orthogonal projector from the whitened state space to this informative subspace. Note that \mathbf{P}_d has rank d as the subscript suggests. Therefore, the whitened state variable can be decomposed as:

$$\tilde{\mathbf{X}} = \mathbf{P}_d \tilde{\mathbf{X}} + (\mathbf{I}_n - \mathbf{P}_d) \tilde{\mathbf{X}} = \tilde{\mathbf{X}}^\parallel + \tilde{\mathbf{X}}^\perp, \quad (\text{D.7})$$

where $\tilde{\mathbf{X}}^\parallel := \mathbf{P}_d \tilde{\mathbf{X}}$ and $\tilde{\mathbf{X}}^\perp := (\mathbf{I}_n - \mathbf{P}_d) \tilde{\mathbf{X}}$. Since \mathbf{P}_d is an orthogonal projector, the decomposition above is unique and $\tilde{\mathbf{X}}^\parallel \in \mathbb{R}^n$ is orthogonal to $\tilde{\mathbf{X}}^\perp \in \mathbb{R}^n$. We can always complete the orthogonal family of columns \mathbf{V} to form a basis for \mathbb{R}^n with $n - d$ orthogonal columns \mathbf{V}^\perp . The projector on this complementary subspace is given by $\mathbf{P}_{n-d}^\perp = \mathbf{V}^\perp \mathbf{V}^{\perp\top} = \mathbf{I}_n - \mathbf{P}_d$. We can now connect the decomposition (D.7) with the projected state variables on the informative and complementary subspaces:

$$\tilde{\mathbf{X}} = \mathbf{V}\mathbf{V}^\top \tilde{\mathbf{X}} + \mathbf{V}^\perp \mathbf{V}^{\perp\top} \tilde{\mathbf{X}} = \mathbf{V}(\mathbf{V}^\top \tilde{\mathbf{X}}) + \mathbf{V}^\perp(\mathbf{V}^{\perp\top} \tilde{\mathbf{X}}) = \mathbf{V} \tilde{\mathbf{X}} + \mathbf{V}^\perp \mathbf{X}^\perp, \quad (\text{D.8})$$

with $\tilde{\mathbf{X}} \in \mathbb{R}^d$ and $\mathbf{X}^\perp \in \mathbb{R}^{n-d}$. We emphasize that the inference performed with the analysis map in (D.6) only acts on the rotated state variable $\tilde{\mathbf{X}} \in \mathbb{R}^d$, while the complementary component $\tilde{\mathbf{X}}^\perp$ of the whitened state is unaffected by the assimilation. Using eq. (D.8), the low-rank analysis map in the original space is given by:

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = \Sigma_{\tilde{\mathbf{X}}}^{1/2} (\mathbf{V} \tilde{T}(\tilde{\mathbf{y}}, \tilde{\mathbf{x}}) + (\mathbf{I}_d - \mathbf{V}\mathbf{V}^\top) \Sigma_{\tilde{\mathbf{X}}}^{-1/2} \mathbf{x}) = \mathbf{x} - \mathbf{K}(\mathbf{y} - \mathbf{y}^*), \quad (\text{D.9})$$

where \mathbf{K} denotes the Kalman gain. In the original space, we obtain the desired factorization of the Kalman gain $\mathbf{K} \in \mathbb{R}^{n \times d}$ given by

$$\mathbf{K} = \Sigma_{\tilde{\mathbf{X}}}^{1/2} \mathbf{V} \mathbf{\Lambda} (\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1} \mathbf{U}^\top \Sigma_{\tilde{\mathbf{\epsilon}}}^{-1/2}. \quad (\text{D.10})$$

In the whitened space (where $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ live), it is easy to see that $\mathbf{V} \mathbf{\Lambda} (\mathbf{\Lambda}^2 + \mathbf{I}_d)^{-1} \mathbf{U}^\top$ constitutes the singular value decomposition of the Kalman gain. Thus, in the whitened space, exploiting the spectrum of the whitened observation matrix gives us the best low-rank approximation of the Kalman gain. Unfortunately, (D.10) is no longer the SVD of the Kalman gain in the original space as the columns of the matrices $\Sigma_{\tilde{\mathbf{X}}}^{1/2} \mathbf{V}$ and $\Sigma_{\tilde{\mathbf{\epsilon}}}^{-1/2} \mathbf{U}$ are not necessary orthogonal. Nonetheless, the proposed factorization gives us a constructive means to form a low-rank approximation of the Kalman gain in the original space (even if it is not optimal in the canonical Frobenius norm) without having to form the entire Kalman gain. This is clearly a desired feature for large inference problems where the Kalman gain cannot be practically formed and stored in memory.

E. Low-rank assimilation for the nonlinear–Gaussian case

This section provides some intuition for the construction of the state and observation Gramians in the nonlinear–Gaussian setting. In linear–Gaussian case, we show that the singular vectors of the whitened observation matrix (defined in Sec. D) are the directions maximizing Rayleigh quotients of interest. In the state space, it is enlightening to examine the directions that maximize the Rayleigh quotient of the posterior to prior precision in the linear–Gaussian case:

$$\mathcal{R}(w) = \frac{\langle w, \Sigma_{\mathbf{x}}^{-1} \mathbf{y} w \rangle}{\langle w, \Sigma_{\mathbf{x}}^{-1} w \rangle}, \text{ for } w \in \mathbb{R}^n. \quad (\text{E.1})$$

Given that $\Sigma_{\mathbf{x}}^{-1} \mathbf{y} = \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} + \Sigma_{\mathbf{x}}^{-1}$, maximizing this Rayleigh quotient is equivalent to maximizing the Rayleigh ratio of the Hessian of the log-likelihood to the prior precision given by:

$$\mathcal{S}(w) = \frac{\langle w, \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} w \rangle}{\langle w, \Sigma_{\mathbf{x}}^{-1} w \rangle}. \quad (\text{E.2})$$

This equation shows the connection between the directions that maximize the Rayleigh quotient of the posterior to prior, and the directions (in the state space) where the observations are most informative with respect to the prior. With the change of variable $v = \Sigma_{\mathbf{x}}^{-1/2} w \in \mathbb{R}^n$, we also obtain

$$\tilde{\mathcal{S}}(v) = \frac{\langle v, \Sigma_{\mathbf{x}}^{1/2} \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} \Sigma_{\mathbf{x}}^{1/2} v \rangle}{\langle v, v \rangle}. \quad (\text{E.3})$$

We denote the columns of \mathbf{V} that span the image space of the projector \mathbf{P}_d introduced in the previous section as $\{v_1, \dots, v_d\}$. Spantini et al. [14] showed that the vector v_j maximizes the Rayleigh quotient $\tilde{\mathcal{S}}$ over the subspace $\mathbb{R}^n \setminus \text{span}\{v_1, \dots, v_{j-1}\}$, which is the null space of the projector generated by the previous columns vectors $\{v_1, \dots, v_{j-1}\}$. In other words, we are successively identifying the directions in the whitened state space where the observations are the most expressive relative to the prior and which are not in the same subspace as the previous directions. The matrix $\mathbf{C}_{\mathbf{x}} = \Sigma_{\mathbf{x}}^{1/2} \mathbf{H}^\top \Sigma_{\boldsymbol{\varepsilon}}^{-1} \mathbf{H} \Sigma_{\mathbf{x}}^{1/2} \in \mathbb{R}^{n \times n}$ can in fact be rewritten as the inner product of the whitened observation matrix introduced in D, i.e., $\mathbf{C}_{\mathbf{x}} = \tilde{\mathbf{H}}^\top \tilde{\mathbf{H}}$. It is straightforward to show that the vectors $\{v_1, \dots, v_d\}$ are the d eigenvectors associated with the d largest eigenvalues of the positive semi-definite matrix $\mathbf{C}_{\mathbf{x}}$. The following $n - d$ eigenvectors form an orthonormal basis for the non-informative subspace \mathbf{V}^\perp . The interpretation of the vectors $\{v_1, \dots, v_n\}$ as the eigenvectors of the matrix $\mathbf{C}_{\mathbf{x}}$ is an important step in generalizing to the setting with nonlinear observational models. In the nonlinear and non-Gaussian setting, Cui et al. [16] showed that the most important assimilation directions (for any realization of the observations) in the whitened state space can be identified by the eigenvectors of the *state space Gramian*:

$$\mathbf{C}_{\mathbf{x}} = \int \left(\Sigma_{\boldsymbol{\varepsilon}}^{-1/2} \nabla h(x) \Sigma_{\mathbf{x}}^{1/2} \right)^\top \left(\Sigma_{\boldsymbol{\varepsilon}}^{-1/2} \nabla h(x) \Sigma_{\mathbf{x}}^{1/2} \right) d\pi_{\mathbf{x}}(x) \in \mathbb{R}^{n \times n}, \quad (\text{E.4})$$

where the expectation is taken over the prior distribution. As expected, eq. (E.4) reverts to $\tilde{\mathbf{H}}^\top \tilde{\mathbf{H}}$ in the linear–Gaussian case.

To the best of our knowledge, there is no proved procedure to identify the most important directions of the observation space for a nonlinear observation model. We propose a heuristic inspired from the construction of eq. (E.4) for the state space that reverts to the columns of the orthogonal matrix \mathbf{U} in the linear–Gaussian case. It is reasonable to look for the directions in the observation space that maximize the relative ratio of the signal to the observational noise. In the

linear–Gaussian case, we can form the Rayleigh quotient \mathcal{T} that conveys this comparison as

$$\mathcal{T}(q) = \frac{\langle q, \Sigma_Y q \rangle}{\langle q, \Sigma_{\mathcal{E}} q \rangle} - 1 = \frac{\langle q, (\Sigma_Y - \Sigma_{\mathcal{E}}) q \rangle}{\langle q, \Sigma_{\mathcal{E}} q \rangle}, \text{ for } q \in \mathbb{R}^d. \quad (\text{E.5})$$

From eq. (D.1), we have $\Sigma_Y = H \Sigma_X H^\top + \Sigma_{\mathcal{E}}$. With the change of variable $u = \Sigma_{\mathcal{E}}^{-1/2} q \in \mathbb{R}^d$, we obtain:

$$\tilde{\mathcal{T}}(u) = \frac{\langle u, \Sigma_{\mathcal{E}}^{-1/2} H \Sigma_X H^\top \Sigma_{\mathcal{E}}^{-1/2} u \rangle}{\langle u, u \rangle} \quad (\text{E.6})$$

It is easy to show that the directions that maximize the ratio in (E.6) are the eigenvectors of the matrix $C_Y = \Sigma_{\mathcal{E}}^{-1/2} H \Sigma_X H^\top \Sigma_{\mathcal{E}}^{-1/2} = \tilde{H} \tilde{H}^\top \in \mathbb{R}^{d \times d}$. The eigenvectors of C_Y are also the column vectors of the matrix U introduced in the previous section. Inspired by the treatment in the state space, we propose to use the eigenvectors of the *observation space Gramian* C_Y to select the important assimilation directions in the whitened observation space:

$$C_Y = \int \left(\Sigma_{\mathcal{E}}^{-1/2} \nabla h(x) \Sigma_X^{1/2} \right) \left(\Sigma_{\mathcal{E}}^{-1/2} \nabla h(x) \Sigma_X^{1/2} \right)^\top d\pi_X(x) \in \mathbb{R}^{d \times d}. \quad (\text{E.7})$$

Similarly to C_X , eq. (E.7) reverts to $\tilde{H} \tilde{H}^\top$ in the linear–Gaussian case.

F. Algorithm for the analysis step of the low-rank ensemble Kalman filter

In this section, we present the complete algorithm for each iteration of the low-rank ensemble Kalman filter. The algorithm transforms a set of forecast samples to analysis samples after assimilating data y^* .

References

1. Evensen G. 1994 Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans* **99**, 10143–10162.
2. Evensen G. 2009 The ensemble Kalman filter for combined state and parameter estimation. *IEEE Control Systems Magazine* **29**, 83–104.
3. Asch M, Bocquet M, Nodet M. 2016 *Data assimilation: methods, algorithms, and applications* vol. 11. SIAM.
4. Van Leeuwen PJ. 2019 Non-local observations and information transfer in data assimilation. *Frontiers in Applied Mathematics and Statistics* **5**, 48.
5. Gwirtz K, Morzfeld M, Kuang W, Tangborn A. 2021 A testbed for geomagnetic data assimilation. *Geophysical Journal International* **227**, 2180–2203.
6. Lang M, Browne P, Van Leeuwen PJ, Owens M. 2017 Data assimilation in the solar wind: Challenges and first results. *Space Weather* **15**, 1490–1510.
7. Colburn C, Cessna J, Bewley T. 2011 State estimation in wall-bounded flow systems. Part 3. The ensemble Kalman filter. *Journal of Fluid Mechanics* **682**, 289–303.
8. da Silva AF, Colonius T. 2020 Flow state estimation in the presence of discretization errors. *Journal of Fluid Mechanics* **890**.
9. Li S, He C, Liu Y. 2022 A data assimilation model for wall pressure-driven mean flow reconstruction. *Physics of Fluids* **34**, 015101.
10. Sarpkaya T. 1975 An inviscid model of two-dimensional vortex shedding for transient and asymptotically steady separated flow over an inclined plate. *Journal of Fluid Mechanics* **68**, 109–128.
11. Eldredge JD. 2019 *Mathematical Modeling of Unsteady Inviscid Flows*. Springer International Publishing.
12. Le Provost M, Eldredge JD. 2021 Ensemble Kalman filter for vortex models of disturbed aerodynamic flows. *Physical Review Fluids* **6**, 050506.

Algorithm 1: `lowrankenkf`($\mathbf{y}^*, \mathbf{h}, \{\mathbf{x}^i\}, r_{\mathbf{X}}, r_{\mathbf{Y}}$) assimilates the data \mathbf{y}^* in the prior samples $\{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ with the low-rank EnKF. The integers $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$ determine the rank of the projected subspace in the state space and in the observation space, respectively.

Input: $\mathbf{y}^* \in \mathbb{R}^d$, $\mathbf{h}: \mathbb{R}^n \rightarrow \mathbb{R}^d$, M samples $\{\mathbf{x}^i\}$ from $\pi_{\mathbf{X}}$, ranks $r_{\mathbf{X}}$ and $r_{\mathbf{Y}}$

Output: M samples $\{\mathbf{x}_a^i\}$ from $\pi_{\mathbf{X}|\mathbf{Y}}(\cdot | \mathbf{y}^*)$

```

/* Evaluate the observation operator for the prior samples, and
   sample from the noise distribution. */
1 for  $i = 1 : M$  do
2    $\mathbf{z}^i \leftarrow \mathbf{h}(\mathbf{x}^i)$ 
3    $\boldsymbol{\epsilon}^i \sim \mathcal{N}(\mathbf{0}_d, \boldsymbol{\Sigma}_{\mathcal{E}})$ 
/* Compute a Monte-Carlo approximation of  $\mathbf{C}_{\mathbf{X}}$  and  $\mathbf{C}_{\mathbf{Y}}$ . */
4  $\hat{\mathbf{C}}_{\mathbf{X}} \leftarrow \frac{1}{M-1} \sum_{i=1}^M \left( \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)^{\top} \left( \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)$ 
5  $\hat{\mathbf{C}}_{\mathbf{Y}} \leftarrow \frac{1}{M-1} \sum_{i=1}^M \left( \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right) \left( \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} \nabla \mathbf{h}(\mathbf{x}^i) \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \right)^{\top}$ 
/* Perform a low-rank eigendecomposition of  $\hat{\mathbf{C}}_{\mathbf{X}}$  and  $\hat{\mathbf{C}}_{\mathbf{Y}}$ . */
6  $\mathbf{V}_{r_{\mathbf{X}}} \leftarrow \text{eigenvector}(\hat{\mathbf{C}}_{\mathbf{X}}, r_{\mathbf{X}})$ 
7  $\mathbf{U}_{r_{\mathbf{Y}}} \leftarrow \text{eigenvector}(\hat{\mathbf{C}}_{\mathbf{Y}}, r_{\mathbf{Y}})$ 
/* Whiten and project the prior samples  $\{\mathbf{x}^i\}$ , observation samples
    $\{\mathbf{z}^i\}$ , and the observation noise samples  $\{\boldsymbol{\epsilon}^i\}$ , by subtracting
   the empirical mean  $\hat{\boldsymbol{\mu}}$ , and rotating the samples by  $\mathbf{V}\boldsymbol{\Sigma}^{-1/2}$ . */
8 for  $i = 1 : M$  do
9    $\check{\mathbf{x}}^i \leftarrow \mathbf{V}_{r_{\mathbf{X}}}^{\top} \boldsymbol{\Sigma}_{\mathbf{X}}^{-1/2} (\mathbf{x}^i - \hat{\boldsymbol{\mu}}_{\mathbf{X}})$ 
10   $\check{\mathbf{z}}^i \leftarrow \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} (\mathbf{z}^i - \hat{\boldsymbol{\mu}}_{\mathbf{Z}})$ 
11   $\check{\boldsymbol{\epsilon}}^i \leftarrow \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} (\boldsymbol{\epsilon}^i - \hat{\boldsymbol{\mu}}_{\mathcal{E}})$ 
/* Form the perturbation matrices for the whitened state
    $\mathbf{A}_{\check{\mathbf{X}}} \in \mathbb{R}^{r_{\mathbf{X}} \times M}$ , the whitened observation  $\mathbf{A}_{\check{\mathbf{Z}}} \in \mathbb{R}^{r_{\mathbf{Y}} \times M}$ , and the
   whitened observation noise  $\mathbf{A}_{\check{\boldsymbol{\epsilon}}} \in \mathbb{R}^{r_{\mathbf{Y}} \times M}$ . */
12 for  $i = 1 : M$  do
13    $\mathbf{A}_{\check{\mathbf{X}}}[:, i] \leftarrow \frac{1}{\sqrt{M-1}} (\check{\mathbf{x}}^i - \hat{\boldsymbol{\mu}}_{\check{\mathbf{X}}})$ 
14    $\mathbf{A}_{\check{\mathbf{Z}}}[:, i] \leftarrow \frac{1}{\sqrt{M-1}} (\check{\mathbf{z}}^i - \hat{\boldsymbol{\mu}}_{\check{\mathbf{Z}}})$ 
15    $\mathbf{A}_{\check{\boldsymbol{\epsilon}}}[:, i] \leftarrow \frac{1}{\sqrt{M-1}} (\check{\boldsymbol{\epsilon}}^i - \hat{\boldsymbol{\mu}}_{\check{\boldsymbol{\epsilon}}})$ 
/* Apply the Kalman gain in the projected space based on the
   representer [40]. Solve the linear system for  $\check{\mathbf{b}} \in \mathbb{R}^{r_{\mathbf{Y}} \times M}$ : */
16  $(\mathbf{A}_{\check{\mathbf{Z}}} \mathbf{A}_{\check{\mathbf{Z}}}^{\top} + \mathbf{A}_{\check{\boldsymbol{\epsilon}}} \mathbf{A}_{\check{\boldsymbol{\epsilon}}}^{\top}) \check{\mathbf{b}} = \mathbf{U}_{r_{\mathbf{Y}}}^{\top} \boldsymbol{\Sigma}_{\mathcal{E}}^{-1/2} (\mathbf{y}^* \mathbf{1}_M^{\top} - (\mathbf{E}_{\mathbf{Z}} + \mathbf{E}_{\mathcal{E}})),$ 
/* where  $\mathbf{E}_{\mathbf{Z}} \in \mathbb{R}^{d \times M}$  and  $\mathbf{E}_{\mathcal{E}} \in \mathbb{R}^{d \times M}$  are the ensemble matrices of the
   samples  $\{\mathbf{z}^i\}$  and  $\{\boldsymbol{\epsilon}^i\}$ .  $\mathbf{1}_M$  denotes a vector of ones of length  $M$ .
   Lift result to the original space */
17 for  $i = 1 : M$  do
18    $\mathbf{x}_a^i \leftarrow \mathbf{x}^i + \boldsymbol{\Sigma}_{\mathbf{X}}^{1/2} \mathbf{V}_{r_{\mathbf{X}}} (\mathbf{A}_{\check{\mathbf{X}}} \mathbf{A}_{\check{\mathbf{Z}}}^{\top}) \check{\mathbf{b}}[:, i]$ 
19 return  $\{\mathbf{x}_a^i\}$ 

```

13. Cui T, Martin J, Marzouk YM, Solonen A, Spantini A. 2014 Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems* **30**, 114015.
14. Spantini A, Solonen A, Cui T, Martin J, Tenorio L, Marzouk Y. 2015 Optimal low-rank approximations of Bayesian linear inverse problems. *SIAM Journal on Scientific Computing* **37**, A2451–A2487.
15. Zahm O, Cui T, Law K, Spantini A, Marzouk Y. 2022 Certified dimension reduction in nonlinear Bayesian inverse problems. *Mathematics of Computation* **in press**. arXiv:1807.03712.
16. Cui T, Zahm O. 2021 Data-free likelihood-informed dimension reduction of Bayesian inverse problems. *Inverse Problems* **37**, 045009.
17. Giraldi L, Le Maître OP, Hoteit I, Knio OM. 2018 Optimal projection of observations in a Bayesian setting. *Computational Statistics & Data Analysis* **124**, 252–276.
18. Jagalur-Mohan J, Marzouk Y. 2021 Batch greedy maximization of non-submodular functions: Guarantees and applications to experimental design. *Journal of Machine Learning Research* **22**, 1–62.
19. Greengard L, Rokhlin V. 1987 A fast algorithm for particle simulations. *Journal of computational physics* **73**, 325–348.
20. Ethridge F, Greengard L. 2001 A new fast-multipole accelerated Poisson solver in two dimensions. *SIAM Journal on Scientific Computing* **23**, 741–760.
21. Marzouk Y, Moselhy T, Parno M, Spantini A. 2016 Sampling via measure transport: An introduction. *Handbook of Uncertainty Quantification* pp. 1–41.
22. Spantini A, Baptista R, Marzouk Y. 2022 Coupling techniques for nonlinear ensemble filtering. *SIAM Review* **to appear**. arXiv:1907.00389.
23. Luenberger DG. 1997 *Optimization by vector space methods*. John Wiley & Sons.
24. Friedman J, Hastie T, Tibshirani R. 2001 *The elements of statistical learning* vol. 1. Springer series in statistics New York.
25. Davis C, Kahan WM. 1970 The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis* **7**, 1–46.
26. Yu Y, Wang T, Samworth RJ. 2015 A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika* **102**, 315–323.
27. Tropp JA. 2015 An introduction to matrix concentration inequalities. *arXiv preprint* arXiv:1501.01571.
28. Eckart C, Young G. 1936 The approximation of one matrix by another of lower rank. *Psychometrika* **1**, 211–218.
29. Taira K, Colonius T. 2007 The immersed boundary method: a projection approach. *Journal of Computational Physics* **225**, 2118–2137.
30. Liska S, Colonius T. 2017 A fast immersed boundary method for external incompressible viscous flows using lattice Green’s functions. *Journal of Computational Physics* **331**, 257–279.
31. Le Provost M, Baptista R, Marzouk Y, Eldredge J. 2021 A low-rank nonlinear ensemble filter for vortex models of aerodynamic flows. In *AIAA Scitech 2021 Forum* p. 1937.
32. Baptista R, Zahm O, Marzouk Y. 2020 An adaptive transport framework for joint and conditional density estimation. *arXiv preprint* arXiv:2009.10303.
33. Halko N, Martinsson PG, Tropp JA. 2011 Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* **53**, 217–288.
34. Eftekhari A, Wakin MB, Li P, Constantine PG. 2016 Randomized learning of the second-moment matrix of a smooth function. *arXiv preprint* arXiv:1612.06339.
35. Bocher M, Fournier A, Coltice N. 2018 Ensemble Kalman filter for the reconstruction of the Earth’s mantle circulation. *Nonlinear Processes in Geophysics* **25**, 99–123.
36. Garbuno-Inigo A, Hoffmann F, Li W, Stuart AM. 2020 Interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler. *SIAM Journal on Applied Dynamical Systems* **19**, 412–441.
37. Reich S, Weissmann S. 2021 Fokker–Planck Particle Systems for Bayesian Inference: Computational Approaches. *SIAM/ASA Journal on Uncertainty Quantification* **9**, 446–482.
38. Kovachki N, Baptista R, Hosseini B, Marzouk Y. 2020 Conditional sampling with monotone GANs. *arXiv preprint* arXiv:2006.06755.
39. Kalman RE. 1960 A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* **82**, 35–45.
40. Burgers G, Jan van Leeuwen P, Evensen G. 1998 Analysis scheme in the ensemble Kalman filter. *Monthly weather review* **126**, 1719–1724.