



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Ensemble Kalman filter for vortex models of disturbed aerodynamic flows

Mathieu Le Provost and Jeff D. Eldredge

Phys. Rev. Fluids **6**, 050506 — Published 12 May 2021

DOI: [10.1103/PhysRevFluids.6.050506](https://doi.org/10.1103/PhysRevFluids.6.050506)

An ensemble Kalman filter for vortex models of disturbed aerodynamic flows

Mathieu Le Provost and Jeff D. Eldredge*

*Mechanical & Aerospace Engineering Department,
University of California, Los Angeles, Los Angeles, CA, 90095, USA*

The task of dynamic flow estimation is to construct an instantaneous approximation of an evolving flow—and particularly, its response to disturbances—using measurements from available sensors. Building from previous work by Darakananda et al. (Phys Rev Fluids 2018), we further develop an ensemble Kalman filter (EnKF) framework for aerodynamic flows based on an ensemble of randomly-perturbed inviscid vortex models of flow about an infinitely-thin plate. In the forecast step, vortex elements in each ensemble member are advected by the flow and new elements are released from each edge of the plate; the elements are modestly aggregated to maintain an efficient representation. The vortex elements and leading edge constraint are corrected in the analysis step by assimilating the surface pressure differences across the plate measured from the truth system. We show that the overall framework can be physically interpreted as a series of adjustments to the position and shape of an elliptical region of uncertainty associated with each vortex element. In this work, we compare the previously-used stochastic EnKF with the ensemble transform Kalman filter (ETKF), which uses a deterministic analysis step. We examine the response of the flat plate at 20° in two perturbed flows, with truth data obtained from high-fidelity Navier–Stokes simulation at Reynolds number 500. In the first case, we apply a sequence of large-amplitude pulses near the leading edge of the plate to mimic flow actuation. In the second, we place the plate in a vortex street wake behind a cylinder. In both cases, we show that the vortex-based framework accurately estimates the pressure distribution and normal force, with no *a priori* knowledge of the perturbations or their structure. We show that, in each case, the ETKF is consistently more robust than the stochastic EnKF and is qualitatively better at representing the coherent structures of the true flow. Finally, we examine the mapping from measurements to state update in the analysis step through singular value decomposition of the Kalman gain.

Keywords: inviscid vortex model, ensemble Kalman filter, disturbed separated flow, flow estimation

* jdeldre@g.ucla.edu

I. INTRODUCTION

The objective of aerodynamic flow estimation is to construct an approximation of the time-varying flow about a lifting surface from the available sensors. This task has a variety of potential uses, principally within the context of an overall closed-loop control strategy, in which the estimation of the evolving flow field enriches the information available to the controller. For example, we may seek to exploit transient flow mechanisms, such as the high lift achievable in dynamic stall [1–5], or to manage the flow in the presence of unknown atmospheric disturbances (gusts) [6, 7]. In such situations, the flow is subjected to transient disturbances—from flow actuators, wing maneuvers, or gusts—and in many instances, the disturbance induces the flow to separate or significantly affects an already separated flow.

The flow estimation problem can be posed in a variety of settings, loosely categorized by the degree to which the estimator is built on a physics-based model. At one extreme, the estimator incorporates little to no physics and a supervised machine learning architecture is trained to map sensor data to some description of the flow field [8, 9]. In contrast, if the estimator relies on a predictive physics-based model, then its job is to use the sensor data to correct that prediction for unrepresented effects. In particular, though the physics model may be outfitted with the geometry and the basic flow state, the sensor data is the only source of information on disturbances. This prediction and measurement update process is the generic task of data assimilation, the focus of this paper.

All estimation strategies seek the state vector \mathbf{x}_k , which contains sufficient information to specify the system’s state at any discrete time step t_k . In flow estimation, \mathbf{x}_k not only contains a finite-dimensional representation of the flow field, but also any boundary conditions or model parameters that are uncertain. This state vector’s evolution is predicted (*forecast*) with a dynamical operator \mathbf{f}_k , based on a time discretization of some underlying physics model,

$$\mathbf{x}_k = \mathbf{f}_k(\mathbf{x}_{k-1}). \quad (1)$$

The dynamical model is paired with an observation operator \mathbf{h}_k ,

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k), \quad (2)$$

which returns a vector of predicted observations, \mathbf{z}_k . The data assimilation framework provides a means of folding new observations from the true system into the model through adjustments of the state vector. Generically, this adjustment (*analysis*) is carried out through an optimization: determine the adjustment that minimizes the difference between new sensor measurements \mathbf{z}^* from the truth system and the model-predicted measurements from (2). Though this optimal adjustment can be derived deterministically, it is important to remember that we anticipate uncertainty in the estimation task due to random noise in both the dynamics and the measurements. Thus, in stochastic estimation, such as in the various forms of the Kalman filter, the adjusted state estimate is the one that minimizes its uncertainty after new measurements have been assimilated [10, 11]. The degree to which we favor the dynamical prediction or the measurements is determined by the relative levels of noise we assume for each.

Under the assumption that the noises are drawn from Gaussian distributions, the system’s uncertainty is completely characterized by the covariance matrix. Both the Kalman filter and its variant for non-linear systems, the extended Kalman filter, require that this covariance matrix be stored and propagated along with the mean state itself. For high-dimensional systems such as in fluid dynamics, this storage and propagation of the covariance matrix is computationally intractable. This challenge was addressed by Evensen with the ensemble Kalman filter (EnKF) [12], a Monte Carlo approach to the two-step filter algorithm. In the EnKF, the system’s mean and covariance are approximated by an ensemble of states randomly sampled from the probability distribution. In the forecast, each state is advanced by the dynamical model; in the analysis step, each ensemble member is updated with the new observation to minimize the posterior covariance. Thus, for state dimension n and ensemble size $M < n$, one need only store Mn state components rather than a $n \times n$ matrix.

The dynamical operator in a flow estimation framework can take various forms, including high-fidelity Navier–Stokes simulation. Indeed, this has been the basis for recent work by da Silva and Colonius [13, 14]. In [13], they advanced the state vector for flow past an airfoil with an operator derived from an immersed boundary projection method [15]. The state vector contained grid velocity data as well as the value of the freestream velocity and its time derivative. The truth system in this case was generated by the same simulations, but with the time-varying freestream velocity prescribed. However, the EnKF-based estimator only used sparse surface pressure measurements from this truth system, from which it successfully estimated the model’s instantaneous freestream velocity.

Since an ostensible goal for flow estimation is real-time control, dynamical models obtained from CFD will generally not be fast enough to update the state. In their subsequent work in [14], da Silva and Colonius utilized a dynamical model for the airfoil flow consisting of cheaper coarse-grid simulations. Their state vector was then augmented with a bias error, which they successfully estimated along with the coarsened grid state from surface observations of the higher-fidelity truth data. Though this coarse-grid approach is certainly faster, it is unclear whether the approach can be extended to very small-dimensional state vectors before aliasing errors would overwhelm the solution of the

governing flow (Navier–Stokes or Euler) equations. Data-based modal decompositions, such as proper orthogonal decomposition (POD), are commonly used to reduce the dimension of grid-based solutions [16], and a Galerkin model based on such a decomposition might provide a suitable approach for computationally-efficient estimation. However, a modal decomposition cannot efficiently capture the influence of large-amplitude flow disturbances if they are missing from the basis modes.

We seek an alternative low-dimensional representation of the flow, with a dynamical model that is computationally efficient but remains sufficiently rich to describe important physical phenomena. Aerodynamics has long relied on inviscid vortex models for representing and interpreting unsteady flow phenomena [17–21]. Vortex models have a particular advantage in a flow estimation framework, because their physical fidelity and numerical accuracy are distinct. For example, even a model consisting of a small number of point vortices is still an exact solution (to within time marching error) of the two-dimensional inviscid equations of motion, including the non-linear convection. Thus, a vortex model of any size might serve as a low-dimensional representation of a more complex flow furthermore, it can readily account for the influence of disturbances, e.g., through Biot–Savart interactions.

Viscosity is obviously explicitly absent from an inviscid vortex model. However, among the roles of viscosity in moderate to large Reynolds number external flows, only the generation of vorticity at the surface of a body is truly essential to the flow’s large-scale dynamics. Classically, this vorticity generation was restricted to vortex shedding from the trailing edge, by regularizing to flow at the edge with the Kutta condition [17, 21]. However, in recent years, a number of inviscid vortex models have been developed for predicting aerodynamic flows with leading-edge separation, at a wide range of angles of attack [22–25]. Ramesh et al. [24] introduced the concept of the leading-edge suction parameter (LESP), a nondimensional measure of the amount of suction (integrated pressure) at the nose of an airfoil. In their criterion [24], which can be interpreted as a generalization of the Kutta condition [21], vortex elements are only shed from the leading edge if LESP exceeds a critical value, $LESP_c$; if $LESP_c$ is set to zero, then the Kutta condition is enforced at that edge. The circulation of the new vortex element is proportional to the amount by which LESP exceeds $LESP_c$. Thus, to complete a basic inviscid vortex model of a separated flow past an airfoil, one need only specify the value of $LESP_c$. Ramesh et al. [24] hypothesized that this critical value is constant and a function only of the airfoil shape; they tuned its value in a thin airfoil model to match the observed separation in a high-fidelity Navier–Stokes simulation of the airfoil. This hypothesis has been confirmed to hold for a wide variety of unsteady flows generated by an airfoil undergoing unsteady maneuvers [26]. However, for flows subjected to disturbances such as gusts, there is no *a priori* reason to believe that the hypothesis still holds, and, without further knowledge, its value should be assumed time-varying and uncertain. Furthermore, because $LESP_c$ is a measure of a critical behavior in the pressure, its value is likely to be observable from surface pressure measurements.

Thus, an inviscid vortex model outfitted with a leading-edge shedding condition based on $LESP_c$ is a strong candidate for the forecast step in aerodynamic flow estimation. A caveat of the model is that its computational complexity is $O(q^2)$, where q is the number of elements. With elements shed from the edges at every time step, the model can easily lose its computational advantage over CFD-based approaches. To restrain the growth of this dimensionality, vortex elements can be aggregated with various conservation constraints, e.g., based on moments of the vortex system [27], the rate of change of linear impulse [25], or the velocity induced at the leading edge [28]. Aggregation addresses another issue that arises when a vortex model is applied to a flow past an airfoil at moderate angle of attack. Such flows exhibit extensive interactions between shed vorticity and the plate. Singular vortex elements thus remain close to the plate and, without the influence of viscosity, create strong and non-physical pressure disturbances that interfere with the model’s predictive utility. Aggregation tends to move vortex elements away from the plate, thereby mitigating this behavior [25].

To exploit these potential benefits of a vortex-based assimilation procedure, Darakananda et al. [29] developed a data-assimilated vortex model in which the state vector consisted of the positions and strengths of vortex elements and the value of $LESP_c$ for a flat-plate airfoil. The vortices were aggregated by the same approach as in [25], which ensured that the number of elements remained lower than 60 in all cases considered. The vortex model was expressed in an EnKF framework, with the observation vector composed from the jump in surface pressures across the flat plate. By assimilating measurements from a high-fidelity simulation at Reynolds number 500, the estimation framework was able to successfully reproduce the force history on the plate.

The current work extends the estimation framework of Darakananda et al. [29] and is motivated by two observations from that work. In one of the examples in [29], the flat plate was subjected to vortical gusts generated upstream. Remarkably, the estimation framework performed well even with no representation for the gust itself in the state vector. In other words, the dynamics of interaction between the gust and the separated flow was entirely accounted for through adjustments made in the analysis steps to vorticity shed from the plate. The estimator needed no knowledge of the gust’s structure, a great advantage since gusts are, by definition, unknown. In this work, we probe this more deeply by focusing exclusively on flows subjected to incident gusts. We purposely apply gusts that are stronger and in faster sequences than in [29], so that their individual influences on the flow cannot easily be decorrelated. Even if we knew each gust’s parametric form, it would be very difficult to unambiguously disentangle their parameters. In this paper,




FIG. 1. Schematic of the two-dimensional infinitely-thin plate.

we explore whether this parameter estimation task can be avoided.

The second observation from [29] was that the estimator’s performance on any given flow varied substantially from one realization to the next, often exhibiting non-physical spikes in the pressure and force at random instants. This lack of consistent behavior is clearly undesirable for practical estimation applications. We attempt to address this by utilizing a variant of the EnKF called the ensemble transform Kalman filter (ETKF). As we review later, the ensemble in the original EnKF (called the stochastic EnKF, or sEnKF) does not reproduce the exact posterior covariance relation of the Kalman filter, causing spurious noise in the resulting estimate. The ETKF, developed by Bishop et al. [30], fixes this with an analysis step that exactly reproduces the expected covariance.

The rest of this paper is organized as follows. Section II reviews the aggregated vortex model and presents a state-space formulation of the model. Section III introduces the ensemble Kalman filter theory, including both the original sEnKF and the ETKF. The results of the data-assimilated aggregated vortex model, using both forms of EnKF, are presented on example problems in IV. Concluding remarks follow in V. For the sake of completeness, pseudo-codes for the two EnKF methods are provided in the Appendix.

II. THE AGGREGATED VORTEX MODEL

This section provides a summary of the aggregated vortex model, which constitutes the dynamical part of our flow estimator. Let $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ be the canonical basis of \mathbb{R}^3 . We focus in this paper on the response of a two-dimensional infinitely-thin plate of chord length c at an angle of attack α , translating impulsively from rest at a velocity $U\mathbf{e}_1$ in a fluid of density ρ , see Fig. 1. Throughout this study, positions will be reported in chord lengths and time will be measured in convective units $t^* = tU/c$.

The flow is modeled by a collection of q regularized point vortices, called blobs, identical to the model used by Darakananda et al. [29]. The vorticity field is given, as usual for a regularized vortex particle method, by

$$\omega(\mathbf{r}, t) = \sum_{j=1}^q \Gamma_j \zeta_\varepsilon(\mathbf{r} - \mathbf{r}_j), \quad (3)$$

where ζ_ε is a blob kernel with radius ε . The purpose of a blob is to de-singularize the Biot–Savart interactions with other nearby elements. Its effect on elements farther than ε is indistinguishable from that of a point vortex. The blob kernel in this work takes the commonly-used algebraic form $\zeta_\varepsilon(\mathbf{r}) = (\varepsilon^2/\pi)(|\mathbf{r}|^2 + \varepsilon^2)^{-2}$. We refer readers to [21] for a review of modeling strategies of two dimensional inviscid flows.

As we have mentioned earlier, the main role of viscosity in moderate to high Reynolds number external flows, is to generate vorticity at the surface of a body. We model this vorticity flux by shedding new vortex elements from the edges of the plate. In particular, we apply the classical Kutta condition at the trailing edge. However, enforcing the Kutta condition at the leading edge at small to moderate angles of attack leads to a non-physical flow. The regularized flow will leave tangentially the leading-edge in the opposite direction to the freestream. Instead, Ramesh et al. [31] found experimentally that the flow at the leading edge can support a finite amount of suction (integrated pressure) before to trigger a flow separation. They defined the *leading edge suction parameter (LESP)* as a nondimensional measure of the amount of suction at the nose of an airfoil. An elegant vortex shedding criterion can be formulated from this definition: vortex elements are shed only if the LESP exceeds a critical value denoted $LESP_c$ [24, 31]. Eldredge [21] showed that the Kutta condition is recovered if $LESP_c$ is set to zero. The circulation of the new vortex element is proportional to the amount by which LESP exceeds $LESP_c$. We apply this shedding criterion at the leading

edge. The LESPC is included with the vortex elements' positions and strengths as part of the system state vector, \mathbf{x} . Following Ramesh et al. [24], the LESPC is forecast to remain constant. Ramesh et al. [26] and Darakananda et al. [29] have shown that LESPC has a strong authority over the vortex dynamics near the leading edge. The true vortex dynamics are encoded in the pressure measurements of the plate, from which we can distill an estimate of the LESPC [9, 29]. In other words, though LESPC is predicted to remain constant, the assimilated pressure measurements will tend to cause it to vary, thereby triggering the release of weaker or stronger vorticity.

We use the same vortex element aggregation scheme developed by Darakananda et al. [29]. By aggregating vortex elements at every time step, the overall population remains modest (smaller than 60 in all cases explored) and the aggregated elements stay relatively farther from the plate than without such treatment, dramatically reducing the occurrence of spurious pressure disturbances.

In our flow estimator, the state variable \mathbf{x}_k contains the positions and circulations of q blobs and the critical leading-edge suction parameter LESPC,

$$\mathbf{x}_k = [x_k^1 \ y_k^1 \ \Gamma_k^1 \ \dots \ x_k^q \ y_k^q \ \Gamma_k^q \ \text{LESPC}_k]^\top. \quad (4)$$

Thus, the state vector dimension is $n = 3q + 1$. The dynamical model (1) applies the following operations:

- (a) Enforce the no-flow-through condition on the plate,
- (b) Introduce new vortices according to the Kutta condition at the trailing edge and the current estimate of LESPC at the leading edge,
- (c) Advect vortices and plate,
- (d) Aggregate vortex elements, and nullify the strength of the source blobs.

More details on \mathbf{f}_k can be found in [29].

In this work, we propose to rely on the assimilation of pressure jump coefficients— $\Delta C_p = 2(p^+ - p^-)/\rho U^2$ where $+$ and $-$ denote the upper and lower side of the plate—obtained from the truth system to correct our aggregated vortex model [29]. The role of the assimilation is twofold: to better account for the viscous effects, and to account for the presence of flow disturbances in the true system that are not modeled into the vortex model. The observation operator \mathbf{h}_k in (2) uses the unsteady Bernoulli equation to predict the pressure jump coefficients at d locations on the plate [32]. In this work, pressure measurements are calculated (and also provided by the truth system) at the following $d = 50$ sampled Chebyshev points:

$$\frac{c}{2} \cos\left(\frac{10i\pi}{512+1}\right) \quad \text{for } i = 1, \dots, d. \quad (5)$$

III. ENSEMBLE KALMAN FILTERING METHODOLOGY

In this section we describe the filtering methodology that underpins our data-assimilated vortex model. We start by reviewing the basic filtering problem and the purpose of addressing this problem with an ensemble approach. We then present two flavors of ensemble Kalman filter: the sEnKF as well as a deterministic variant called the ensemble transform Kalman filter (ETKF) (Bishop et al. [30]). Both will be evaluated in the vortex model applications that follow in Section IV. Then, we will review the important technique of covariance inflation to prevent filter divergence. Finally, we will interpret the data-assimilated vortex model on physical grounds.

A. A review of the filtering problem

In this section, we present a basic outline of the general filtering problem and its connection with Bayesian inference; details on this can be found in several references, including Asch et al. [11], Bishop et al. [30], Carrassi et al. [33], Evensen [34] and Vetra-Carvalho et al. [35]. Though the variants of the ensemble Kalman filter can be explained without this background, it is useful to discuss it in order to justify the use of the ensemble Kalman filter in a filtering problem and to identify some of its limitations. In this section and those that follow, we use the following font conventions. Serif fonts refer to random variables, e.g. \mathbf{A} on \mathbb{R}^n , or A on \mathbb{R} . Lowercase roman fonts refer to realization of random variables, e.g. \mathbf{a} on \mathbb{R}^n , or a on \mathbb{R} . $\pi_{\mathbf{A}}$ denotes the probability distribution for the random variable \mathbf{A} , and $\mathbf{a} \sim \pi_{\mathbf{A}}$ means that \mathbf{a} is a realization of \mathbf{A} . $\pi_{\mathbf{A}|\mathbf{B}=\mathbf{b}}$, or $\pi_{\mathbf{A}|\mathbf{B}}(\cdot|\mathbf{b})$, is the probability distribution of the random variable \mathbf{A} knowing that random variable \mathbf{B} takes the value \mathbf{b} . It is called the *conditional probability* of \mathbf{A} given \mathbf{b} .

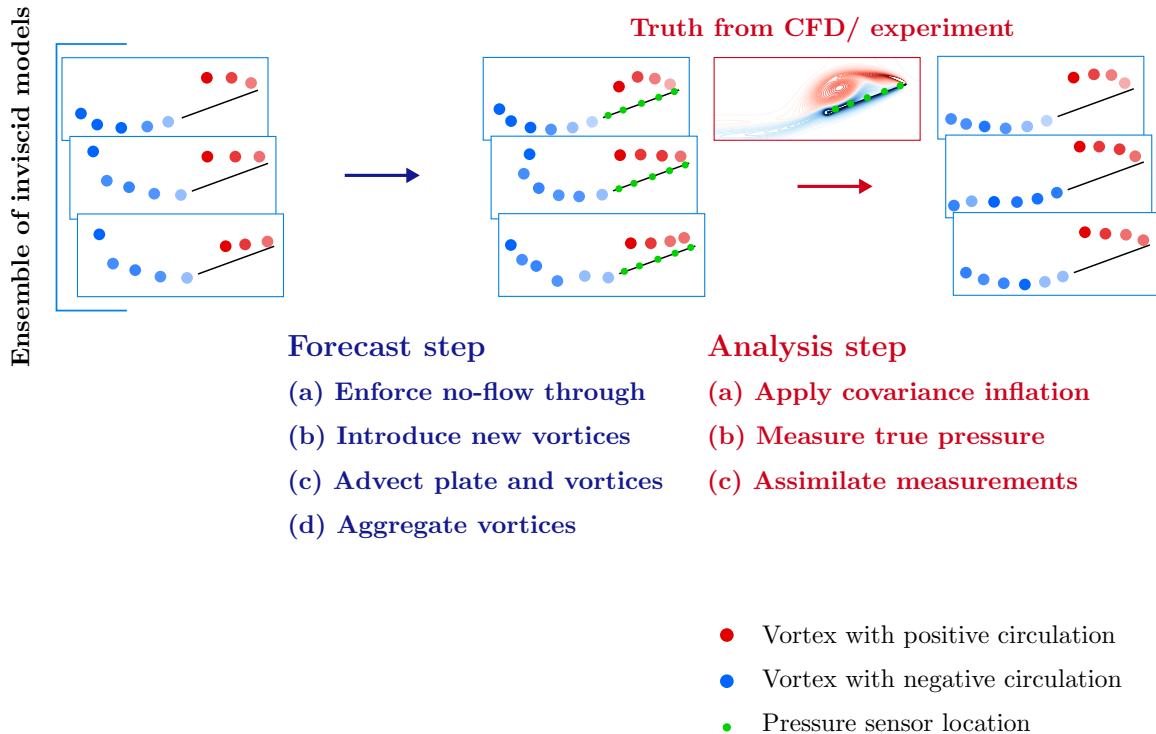


FIG. 2. One time step of the data-assimilated aggregated vortex model

For the sake of generality, we frame our discussion on a generic discrete nonlinear state-space model (1), but we will connect the reader throughout to the present case of a low-order vortex model. Formally, the model is described by the pair of random variables $(\mathbf{X}_k, \mathbf{Z}_k)$ for $k \geq 1$, where \mathbf{X}_k is the state variable of a Markov process in \mathbb{R}^n and $\mathbf{Z}_k \in \mathbb{R}^d$ is an observation of the state \mathbf{X}_k , assumed to be conditionally independent of the state (detailed later in this section). In our case, the state variable is composed of the LESPC and the positions and strengths of the entire set of point vortices, while the observation vector is given by the vector of pressure jump coefficients along the plate. The dynamics of the state \mathbf{X}_k is described by the probability distribution for the initial condition $\pi_{\mathbf{x}_0}$ and a propagation equation, also called the dynamical equation:

$$\mathbf{X}_k = \mathbf{f}_k(\mathbf{X}_{k-1}) + \mathbf{W}_k, \quad (6)$$

where $\mathbf{f}_k : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the forward operator and \mathbf{W}_k is an additive process noise. In the present case, \mathbf{f}_k is the time-discretized vortex model. Propagating the dynamical equation for the state variable \mathbf{X}_{k-1} is equivalent to sampling from the transition distribution $\pi_{\mathbf{x}_k | \mathbf{x}_{k-1}}(\cdot | \mathbf{x}_{k-1})$ where \mathbf{x}_{k-1} is one realization of the random variable \mathbf{X}_{k-1} .

In general, the state \mathbf{X}_k is only indirectly observed in a noisy and nonlinear fashion through \mathbf{Z}_k :

$$\mathbf{Z}_k = \mathbf{h}_k(\mathbf{X}_k) + \mathbf{V}_k, \quad (7)$$

where $\mathbf{h}_k : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is the *observation operator* and \mathbf{V}_k is an additive *measurement noise*. In our case, \mathbf{h}_k applies the unsteady Bernoulli equation to the state vector to obtain a vector of pressure differences, at discrete locations, between the upper and lower surfaces of the plate. The gradient of the observation operator \mathbf{h}_k evaluated at the current state \mathbf{x}_k is called the *tangent linear* of the observation operator, denoted \mathbf{H}_k , i.e., $\mathbf{H}_k = \nabla \mathbf{h}_k$ [11]. It is important to note that we do not restrict ourselves in this section to Gaussian distributions for the initial condition, process noise, or measurement noise. Furthermore, the forward and observation operators can be time-dependent. In general, the dynamical and observation models could include forcing terms. In our case, the system is not subject to known forcing terms, and we have omitted these terms for brevity.

In the filtering problem, we leverage the realizations of the observation process $(\mathbf{z})_{1:k}$ to infer the realization of the state variable \mathbf{x} at time step k . Our main objective is to estimate the posterior distribution of this state—providing us with complete information about the state and its uncertainty—given all of the observations made thus far, $\pi_{\mathbf{x}_k | \mathbf{z}_{1:k}}(\mathbf{x} | \mathbf{z}_{1:k}) = \pi_{\mathbf{x}_k | \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k}(\mathbf{x} | \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k)$. In the present case, we seek to estimate the position and strength of the vortex elements and the LESPC based on all the pressure observations that are available up to this time. Let us consider the estimation of this posterior distribution from the most recent observation, \mathbf{z}_k , which can be computed from Bayes' theorem [11, 33, 35]:

$$\pi_{\mathbf{x}_k | \mathbf{z}_{1:k}}(\mathbf{x} | \mathbf{z}_{1:k}) = \frac{\pi_{\mathbf{z}_k | \mathbf{x}_k}(\mathbf{z}_k | \mathbf{x})\pi_{\mathbf{x}_k | \mathbf{z}_{1:k-1}}(\mathbf{x} | \mathbf{z}_{1:k-1})}{\pi_{\mathbf{z}_k}(\mathbf{z}_k)}, \quad (8)$$

where $\pi_{\mathbf{z}_k | \mathbf{x}_k}(\mathbf{z}_k | \mathbf{x})$ is the likelihood distribution, i.e., the probability of the observation \mathbf{z}_k if we knew the state \mathbf{x} ; $\pi_{\mathbf{x}_k | \mathbf{z}_{1:k-1}}(\mathbf{x} | \mathbf{z}_{1:k-1})$ is the prior distribution, our estimate of the state before knowledge of the new observation \mathbf{z}_k ; and $\pi_{\mathbf{z}_k}(\mathbf{z}_k)$ is the distribution of the observation. In our case, Bayes' theorem tells us how the probability densities of the positions and strengths of the vortex elements and the LESPC value get updated by the assimilation of a new pressure observation. Note that, in the likelihood, we have assumed that the observation errors are independent in time, so it is only conditioned on the current state and not on past observations. Indeed, the Bernoulli equation only depends on the current collection of vortex elements and LESPC value, and not on its past values. Since $\pi_{\mathbf{z}_k}(\mathbf{z}_k)$ can be thought of as a normalizing constant, we can rewrite (8) as:

$$\pi_{\mathbf{x}_k | \mathbf{z}_{1:k}}(\mathbf{x} | \mathbf{z}_{1:k}) \propto \pi_{\mathbf{z}_k | \mathbf{x}_k}(\mathbf{z}_k | \mathbf{x})\pi_{\mathbf{x}_k | \mathbf{z}_{1:k-1}}(\mathbf{x} | \mathbf{z}_{1:k-1}). \quad (9)$$

This equation provides us with a means of assimilating a new observation, \mathbf{z}_k , into our forecast probability distribution for the state [36]. The forecast itself comes from our dynamical model. Under the Markov assumption for the state dynamics, the distribution of the state at time step k depends only on the state at the previous time step $k-1$, i.e., $\pi_{\mathbf{x}_k | \mathbf{x}_{0:k-1}} = \pi_{\mathbf{x}_k | \mathbf{x}_{k-1}}$. The Markov assumption is appropriate in this problem as the time marching of the discrete vortex model only depends on the distribution of vortex elements and the LESPC value at the previous time step. The joint distribution $\pi_{\mathbf{x}_{k-1:k}}$ can thus be factorized into this state transition and the posterior distribution at the end of the previous step:

$$\pi_{\mathbf{x}_{k-1:k} | \mathbf{z}_{1:k-1}} = \pi_{\mathbf{x}_k | \mathbf{x}_{k-1}}\pi_{\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1}}. \quad (10)$$

By recursively applying this equation with (9) substituted for the posterior distribution, it is easy to show that

$$\pi_{\mathbf{x}_{0:k} | \mathbf{z}_{1:k}} \propto \pi_{\mathbf{x}_0} \prod_{i=1}^k \pi_{\mathbf{z}_i | \mathbf{x}_i} \pi_{\mathbf{x}_i | \mathbf{x}_{i-1}}. \quad (11)$$

This equation is central in Bayesian inference since it justifies the use of sequential methods to estimate the posterior distribution [33, 36]. We can update our previous estimate with new observations sequentially without having to restart the calculation at every time step.

The prior distribution $\pi_{\mathbf{x}_k | \mathbf{z}_{1:k-1}}$ can be obtained by marginalizing (10) over all the possible realizations of the state at time $k-1$:

$$\pi_{\mathbf{x}_k | \mathbf{z}_{1:k-1}} = \int_{\mathbf{x}_{k-1}} \pi_{\mathbf{x}_k | \mathbf{x}_{k-1}} \pi_{\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1}} d\mathbf{x}_{k-1}. \quad (12)$$

This equation is called the Chapman–Kolmogorov equation and corresponds to a direct integration of the state transition kernel [33, 36].

Thus, Bayes' theorem (9) coupled with equation (12) provides us with an elegant update rule for incorporating new measurements into the probability distribution, and thus, improving our estimate of the mean and uncertainty of the state. It requires a time-discretized dynamical model (6) for the state transition, $\pi_{\mathbf{x}_k | \mathbf{x}_{k-1}}$, and an observational model (7) for the likelihood $\pi_{\mathbf{z}_k | \mathbf{x}_k}$.

However, there are two primary challenges with applying this in the contexts of interest here. First, we do not generally know the forms of distributions involved in any of these formulas. In what follows in this paper, we will make the typical assumption that all errors are drawn from Gaussian distributions with zero mean, but there are good reasons to doubt that this is reasonable in the present nonlinear context of a vortex model. Second, for high-dimensional problems, the integration of the Chapman–Kolmogorov equation (12) is intractable. Indeed, here, as in many physics contexts, the forecast step corresponds to the time advancement of a partial differential equation (the Euler equations), so it is infeasible that we could advance such an equation over all possible values of the state. In

our case, this corresponds to one time step of the discrete vortex model for all the possible values of the strength and position of the vortex elements and the LESPC. Even if our the discrete vortex model is lower dimensional than the discretized Navier–Stokes system, to be sure, it still contains tens to hundreds of degrees of freedom.

Both of these challenges motivate our use of ensemble filtering methods in this work to sequentially estimate the posterior distribution. We use a set of M particles $\{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M\}$ sampled from the state distribution $\pi_{\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1}}$ and we aim to construct a particle approximation of the posterior distribution $\pi_{\mathbf{x}_k | \mathbf{z}_{1:k}}$. It should be stressed that the term “particle” is used here in its usual sense in stochastic estimation as a member of the ensemble; it does not denote a vortex particle, a term that we avoid in this paper in favor of “vortex element” or “vortex blob”. In the present case, a particle \mathbf{x}^i is a sample of the distribution of interest for the positions and strengths of the vortex elements and the LESPC.

To construct this particle approximation of the posterior distribution, we perform the following two steps: a forecast step and an analysis step. In the forecast step, each particle \mathbf{x}^i is propagated through the dynamical equation (1) (the vortex model) and randomly perturbed with noise (called additive covariance inflation, discussed below) to form samples from the prior distribution $\pi_{\mathbf{x}_k | \mathbf{z}_{1:k-1} = \mathbf{z}_{1:k-1}}$. This ensemble forecast step constitutes a Monte Carlo approximation of the Chapman–Kolmogorov equation (12). Given this sampling $\{\mathbf{x}^i\}$ from the prior distribution, one can easily create a sampling $\{\mathbf{z}^i\}$ from the likelihood distribution $\pi_{\mathbf{z}_k | \mathbf{x}_k}$ by evaluating the observation equation (2) (the Bernoulli equation) for each ensemble member $\mathbf{x}^i, i = 1, \dots, M$.

The analysis step then updates the set of ensemble members by assimilating the new realization \mathbf{z}_k^* of the observation variable \mathbf{Z}_k ; in this paper, \mathbf{z}_k^* represents a new pressure difference measured from the truth system, a high-fidelity Navier–Stokes simulation. With the probability distributions unknown, this would require that we estimate the posterior distribution via Bayes’ theorem (9) given the finite set of samples $\{\mathbf{x}^i, \mathbf{z}^i\}$ from the joint distribution formed by the prior and likelihood. However, the task is made considerably simpler if we assume that the distributions are Gaussian. Under that assumption, and with knowledge of the prior mean and covariance matrix, then Bayes’ theorem leads naturally to the analysis step of the classical Kalman filter [10]. In the particle approximation, in which we estimate this mean and covariance from the finite ensemble statistics that emerge from the forecast, we obtain the ensemble Kalman filter [11]. It should be noted, however, that since the underlying distributions are likely non-Gaussian, *our EnKF framework is not expected to converge in the limit of large ensemble size to the true Bayesian solution* (Mandel et al. [37]). Indeed, the EnKF can only produce a Gaussian approximation of the posterior density. In the following sections we will present two forms of this method, but first we present some important notation.

B. Additional notation

We will use the superscript f to denote prior, forecast quantities and the superscript a to denote analysis, posterior quantities. In particular, the exact prior mean is $\mathbf{x}_k^f = \mathbb{E}[\mathbf{X}_k | \mathbf{z}_{1:k-1} = \mathbf{z}_{1:k-1}] = \mathbb{E}[\mathbf{X}_k | \mathbf{z}_{1:k-1}]$ and exact prior covariance is $\mathbf{P}_k^f = \text{Cov}[\mathbf{X}_k | \mathbf{z}_{1:k-1}]$. The posterior mean and covariance are given by $\mathbf{x}_k^a = \mathbb{E}[\mathbf{X}_k | \mathbf{z}_{1:k}]$ and $\mathbf{P}_k^a = \text{Cov}[\mathbf{X}_k | \mathbf{z}_{1:k}]$, respectively. The notations used in this section are similar to those of Asch et al. [11] and Raanes [38].

Given an ensemble (\mathbf{x}^i) of size M , we define the ensemble matrix $\mathbf{X} \in \mathbb{R}^{n \times M}$ as [11, 38]:

$$\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M]. \quad (13)$$

We use an overbar to denote statistics obtained from the ensemble, such as the sample mean and covariance:

$$\bar{\mathbf{x}} = \frac{1}{M} \sum_{i=1}^M \mathbf{x}^i, \quad \bar{\mathbf{P}} = \frac{1}{M-1} \sum_{i=1}^M (\mathbf{x}^i - \bar{\mathbf{x}}) (\mathbf{x}^i - \bar{\mathbf{x}})^\top, \quad (14)$$

where \top denotes the transpose operator. We define the anomaly matrix $\mathbf{X}' \in \mathbb{R}^{n \times M}$ of an ensemble as

$$\mathbf{X}' = \frac{1}{\sqrt{M-1}} [\mathbf{x}^1 - \bar{\mathbf{x}}, \mathbf{x}^2 - \bar{\mathbf{x}}, \dots, \mathbf{x}^M - \bar{\mathbf{x}}]. \quad (15)$$

The anomaly matrix obviously has zero mean. Using $\mathbf{1}$ to denote the vector of ones of length M , we can conveniently define the sample mean, anomaly matrix, and sample covariance from the ensemble matrix:

$$\bar{\mathbf{x}} = \frac{1}{M} \mathbf{X} \mathbf{1}, \quad \mathbf{X}' = \frac{1}{\sqrt{M-1}} \mathbf{X} (\mathbf{I} - \mathbf{1} \mathbf{1}^\top / M), \quad \bar{\mathbf{P}} = \mathbf{X}' \mathbf{X}'^\top. \quad (16)$$

Finally, we have the relation between the ensemble matrix and the anomaly matrix:

$$\mathbf{X} = \bar{\mathbf{X}} + \sqrt{M-1} \mathbf{X}' \quad (17)$$

where $\bar{\mathbf{X}} = \bar{\mathbf{x}} \mathbf{1}^\top = [\bar{\mathbf{x}}, \bar{\mathbf{x}}, \dots, \bar{\mathbf{x}}] \in \mathbb{R}^{n \times M}$. It is easy to verify that $\mathbf{1}^\top \mathbf{1} = M$, $\mathbf{X}' \mathbf{1} = 0$, and $\bar{\mathbf{X}} \mathbf{1} / M = \bar{\mathbf{x}}$.

C. The stochastic ensemble Kalman filter

The application of the stochastic ensemble Kalman filter for aerodynamic flow estimation was already presented in [13, 29], but we review it here in order to identify an important deficiency that motivates our current treatment. The forecast step of the sEnKF—and, indeed, of all forms of the EnKF—is given by applying (6) (the vortex model) to each ensemble member [34]:

$$\mathbf{x}_k^{f,i} = \mathbf{f}(\mathbf{x}_{k-1}^{a,i}) + \mathbf{w}_k^i \quad \text{for } i = 1, \dots, M, \quad (18)$$

where the random process noise \mathbf{w}_k^i is drawn from the distribution $\mathbf{w}_k^i \sim \pi_{\mathbf{Q}_k}$. It should be noted that this process noise is not inherently part of the dynamical model and must be explicitly introduced; this is discussed in Section III E.

In order to simplify notation, we drop the time dependence subscript of the variables in the rest of this section, since the analysis step does not involve time propagation. The analysis step of the sEnKF seeks a linear transformation of the form:

$$\mathbf{x}^{a,i} = \mathbf{x}^{f,i} + \mathbf{K}(\mathbf{z}^* + \mathbf{v}^i - \mathbf{h}(\mathbf{x}^{f,i})), \quad \text{for } i = 1, \dots, M, \quad (19)$$

where $\mathbf{x}^{f,i}$ and $\mathbf{x}^{a,i}$ are the i th prior and posterior ensemble member, respectively; \mathbf{v}^i is drawn from the measurement noise distribution $\pi_{\mathbf{V}}$; and \mathbf{z}^* is the realization of the observation variable \mathbf{Z} at the current assimilation time. The expression $\mathbf{z}^* + \mathbf{v}^i - \mathbf{h}(\mathbf{x}^{f,i})$ is called the innovation for the i th ensemble member. In our case, (19) shows how discrepancies in the pressure observations are mapped via the Kalman gain \mathbf{K} into linear updates of the position and strength of the vortex elements and the LESPC for each ensemble member.

The algorithm of the sEnKF is provided in Algorithm 1. The original analysis step derived by Evensen [12] did not include the measurement noise term, but this was corrected by Burgers et al. [39] to reflect that \mathbf{z}^* is drawn from a random distribution and to avoid some spurious correlations within the ensemble. The matrix $\mathbf{K} \in \mathbb{R}^{n \times d}$ is called the *Kalman gain* and is identical to its form in the standard Kalman filter, derived to minimize the trace of the posterior covariance matrix \mathbf{P}^a (as we will discuss further below):

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^\top (\mathbf{H} \mathbf{P}^f \mathbf{H}^\top + \mathbf{V})^{-1} \quad (20)$$

with \mathbf{V} the covariance matrix of the measurement noise and \mathbf{H} the tangent linear of the observation operator. The exact prior covariance \mathbf{P}^f is approximated by the sample prior covariance $\bar{\mathbf{P}}^f$ formed from the prior ensemble $\{\mathbf{x}^{f,i}\}$ with (16), $\bar{\mathbf{P}}^f = \mathbf{X}'^f \mathbf{X}'^f{}^\top$; similarly, the posterior covariance is approximated by $\bar{\mathbf{P}}^a = \mathbf{X}'^a \mathbf{X}'^a{}^\top$.

The Kalman gain suggests that we must calculate the tangent linear of the observation operator. Evensen [12] proposed a technique called implicit linearization that approximates $\mathbf{P}^f \mathbf{H}$ and $\mathbf{H} \mathbf{P}^f \mathbf{H}^\top$ given the prior ensemble. First, we construct $\bar{\mathbf{z}}^f$ the mean of the observations $\{\mathbf{h}(\mathbf{x}^{f,i})\}$ for $i = 1, \dots, M$:

$$\bar{\mathbf{z}}^f = \frac{1}{M} \sum_{i=1}^M \mathbf{h}(\mathbf{x}^{f,i}). \quad (21)$$

Let us then define the innovation anomaly matrix \mathbf{Z}'^f , with i th column given by

$$\mathbf{Z}'^{f,i} = \frac{\mathbf{h}(\mathbf{x}^{f,i}) - \bar{\mathbf{z}}^f - \mathbf{v}^i + \bar{\mathbf{v}}}{\sqrt{M-1}}, \quad \text{for } i = 1, \dots, M, \quad (22)$$

and with $\bar{\mathbf{v}}$ the sample mean of $\{\mathbf{v}^i\}$. It should be noted that this mean is itself a random number whose expected value is 0.

The key ingredient of the technique is to make the following approximation [11]:

$$\mathbf{H}(\mathbf{x}^{f,i} - \bar{\mathbf{x}}^f) \simeq \mathbf{h}(\mathbf{x}^{f,i}) - \bar{\mathbf{z}}^f. \quad (23)$$

Using this approximation, it is easy to show that $\mathbf{P}^f \mathbf{H}^\top$ and $\mathbf{H} \mathbf{P}^f \mathbf{H}^\top$ are approximated by $\mathbf{X}'^f \mathbf{Z}'^f{}^\top$ and $\mathbf{Z}'^f \mathbf{Z}'^f{}^\top - \mathbf{V}$, respectively [11], and thus, the Kalman gain (20) takes the simple form

$$\mathbf{K} = \mathbf{X}'^f \mathbf{Z}'^f{}^\top \left(\mathbf{Z}'^f \mathbf{Z}'^f{}^\top \right)^{-1}. \quad (24)$$

To gain more insights on this formulation of the filter, we can rewrite the analysis update (19) with ensemble matrix notations. From (17), the update of each posterior ensemble member is equivalent to a separate update of the

posterior mean $\bar{\mathbf{x}}^a$ and the posterior anomaly matrix \mathbf{X}'^a . From the linear analysis step (19), the posterior anomaly matrix \mathbf{X}'^a is updated according to ([11]):

$$\mathbf{X}'^a = \mathbf{X}'^f - \mathbf{K}\mathbf{Z}'^f; \quad (25)$$

the update equation for the posterior mean state $\bar{\mathbf{x}}^a$ is obtained by taking the expectation of the analysis step (19):

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{K}(\mathbf{z}^* - \bar{\mathbf{z}}^f), \quad (26)$$

where $\bar{\mathbf{z}}^f$ is defined in (21). Equation (26) is the exact Kalman update equation.

Our derivations thus far have omitted the fact that our actual ensemble is finite sized. Because of the assumed linear form of the analysis step, equation (26) is reproduced almost exactly for finite-sized ensembles, except for the addition of a lingering sample mean of the measurement noise, $\bar{\mathbf{v}}$, in parentheses. However, a more problematic discrepancy lies in the relationship between the posterior and prior covariance matrices. For obtaining this relationship, let us first denote by \mathbf{W}' the anomaly matrix of the measurement noise $\mathbf{W}'^i = (\mathbf{v}^i - \bar{\mathbf{v}})/\sqrt{M-1}$; the product $\mathbf{W}'\mathbf{W}'^\top$ approximates the measurement covariance \mathbf{V} . Then, from the update equation (25) and the definitions of the covariances and innovation anomaly matrix, we can construct the equation for the posterior covariance matrix [11]:

$$\bar{\mathbf{P}}^a = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\bar{\mathbf{P}}^f(\mathbf{I}_n - \mathbf{K}\mathbf{H})^\top + \mathbf{K}\mathbf{W}'\mathbf{W}'^\top\mathbf{K}^\top + (\mathbf{I}_n - \mathbf{K}\mathbf{H})\mathbf{X}'^f\mathbf{W}'^\top\mathbf{K}^\top + \mathbf{K}\mathbf{W}'\mathbf{X}'^f{}^\top(\mathbf{I}_n - \mathbf{K}\mathbf{H})^\top. \quad (27)$$

By taking the expectation of this matrix over the measurement noise process and then minimizing its trace over \mathbf{K} , we recover the classical Kalman gain (20) and a simple form of the expected value of the posterior covariance:

$$\mathbb{E}[\bar{\mathbf{P}}^a] = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\bar{\mathbf{P}}^f. \quad (28)$$

Indeed, the classical Kalman filter adopts this expected value in order to propagate the covariance matrix, $\mathbf{P}^a = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\mathbf{P}^f$. However, the finite sample covariance $\bar{\mathbf{P}}^a = \mathbf{X}'^a\mathbf{X}'^a{}^\top$ does not reproduce this ideal relation except in the limit of infinite ensemble sizes. Rather, the covariance of the finite ensemble only strictly satisfies (27), due to spurious correlations between the forecast anomalies and observation noise, $\mathbf{X}'^f\mathbf{W}'^\top$. Thus, the stochastic analysis step of the sEnKF introduces error that degrades the performance of the ensemble Kalman filter [11, 30]. An alternative form of the EnKF will be presented in the next section that addresses this issue.

D. The ensemble transform Kalman filter

In the previous section, it was shown that finite ensembles do not reproduce the expected value of the posterior covariance matrix $\bar{\mathbf{P}}^a$, and that this can cause the performance of the sEnKF to degrade. To circumvent the issue, Bishop et al. [30] developed the *ensemble transform Kalman filter (ETKF)* that exactly reproduces the ideal propagation equation for the covariance. This has been shown to give better performance in other applications [11, 38]. The ETKF belongs to a more general class of ensemble Kalman filters—called *deterministic ensemble Kalman filters*—that verify exactly the correct propagation equation through various analysis schemes.

To generate a posterior anomaly matrix \mathbf{X}'^a , one can start from the desired propagation equation for the covariance and use the factorization of the sample covariance matrix (16):

$$\bar{\mathbf{P}}^a = \mathbf{X}'^a\mathbf{X}'^a{}^\top = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\bar{\mathbf{P}}^f \quad (29)$$

The right-hand side of this equation can also be factorized when $\bar{\mathbf{P}}^f$ is replaced with its own factorization and the ensemble expression (24) of the Kalman gain is introduced. In the following, it should be noted the innovation anomaly matrix \mathbf{Z}'^f is defined as in (22), but now without the observation noise, so that $\mathbf{Z}'^f = \mathbf{H}\mathbf{X}'^f$:

$$\begin{aligned} \mathbf{X}'^a\mathbf{X}'^a{}^\top &= (\mathbf{I}_n - \mathbf{K}\mathbf{H})\mathbf{X}'^f\mathbf{X}'^f{}^\top \\ &= (\mathbf{I}_n - \mathbf{X}'^f\mathbf{Z}'^f{}^\top(\mathbf{Z}'^f\mathbf{Z}'^f{}^\top + \mathbf{V})^{-1}\mathbf{H})\mathbf{X}'^f\mathbf{X}'^f{}^\top \\ &= \mathbf{X}'^f(\mathbf{I}_M - \mathbf{Z}'^f{}^\top(\mathbf{Z}'^f\mathbf{Z}'^f{}^\top + \mathbf{V})^{-1}\mathbf{Z}'^f)\mathbf{X}'^f{}^\top \\ &= \mathbf{X}'^f\mathbf{G}\mathbf{X}'^f{}^\top. \end{aligned} \quad (30)$$

Matrix $\mathbf{G} \in \mathbb{R}^{M \times M}$ is symmetric positive definite. In order to produce an analysis equation for the posterior anomaly matrix, we seek a square-root factorization of \mathbf{G} [11]. In other words, we look for a matrix $\mathbf{G}^{1/2}$ such that $\mathbf{G} = \mathbf{G}^{1/2} \mathbf{G}^{1/2\top} = \mathbf{G}^{1/2\top} \mathbf{G}^{1/2}$. For a positive-definite matrix, the square-root decomposition exists but is not unique. Indeed for an arbitrary orthogonal matrix $\mathbf{U} \in \mathbf{O}(M)$ —i.e., any matrix such that $\mathbf{U}\mathbf{U}^\top = \mathbf{U}^\top\mathbf{U} = \mathbf{I}_M$ —then $\mathbf{G}^{1/2}\mathbf{U}$ is also a square-root of \mathbf{G} [11, 38].

Therefore, the analysis update of the anomaly can be written as a right linear transformation [11]:

$$\mathbf{X}'^a = \mathbf{X}'^f \mathbf{G}^{1/2} \mathbf{U} \quad (31)$$

with some choice of $\mathbf{U} \in \mathbf{O}(M)$, discussed below. Furthermore, the analysis update of the sample mean (25) can be written with the help of (24) as

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{X}'^f \mathbf{Z}'^f \mathbf{Z}'^f \top (\mathbf{Z}'^f \mathbf{Z}'^f \top + \mathbf{V})^{-1} \boldsymbol{\delta}, \quad (32)$$

where $\boldsymbol{\delta} = \mathbf{z}^* - \mathbf{h}(\bar{\mathbf{x}}^f)$ is the mean innovation. Equations (31) and (32) are used together, through (17), to update the ensemble.

It is notable that this analysis step assembles both the posterior anomaly and the update to the mean from the columns of \mathbf{X}'^f . In other words, the analysis is said to be performed in the *ensemble space*: the update to the state vector is a linear combination of the deviations among the forecast ensemble members from the mean. In our case, for example, the position of a certain vortex element will be updated with a linear combination of the ensemble deviations of that element's position from the ensemble mean, after the ensemble has been advanced by the vortex model. This emphasizes the importance of such variance among ensemble members: without it, any discrepancy $\boldsymbol{\delta}$ between the new observation and the predicted observation is simply ignored.

The form of this update is an attractive property of the ETKF since the ensemble size M —which is on the order of 100—is typically similar to the size of the state n encountered in the vortex model, and certainly much smaller than the state vector associated with a CFD simulation [11, 30, 35]. Using Sherman–Morrison–Woodbury identities, the formula for $\mathbf{G}^{1/2}$ can be simplified [35]:

$$\mathbf{G}^{1/2} = (\mathbf{I}_M - \mathbf{Z}'^f \top (\mathbf{Z}'^f \mathbf{Z}'^f \top + \mathbf{V})^{-1} \mathbf{Z}'^f)^{1/2} = (\mathbf{I}_M + \mathbf{Z}'^f \top \mathbf{V}^{-1} \mathbf{Z}'^f)^{-1/2}. \quad (33)$$

In this work, we choose the symmetric square-root factorization, $\mathbf{G} = \mathbf{G}^{1/2} \mathbf{G}^{1/2}$, among the different possibilities. Given an eigendecomposition $\mathbf{G} = \mathbf{R} \boldsymbol{\Sigma} \mathbf{R}^\top$, the symmetric square root $\mathbf{G}^{1/2}$ is given by $\mathbf{R} \boldsymbol{\Sigma}^{1/2} \mathbf{R}^\top$, with $\boldsymbol{\Sigma}^{1/2}$ the entry-wise positive square root of the diagonal matrix $\boldsymbol{\Sigma}$. Several studies [34, 38, 40] have shown that the symmetric root has useful properties, particularly that it does not introduce a spurious mean in the posterior anomaly matrix: $\mathbf{G} \mathbf{1} = \mathbf{1}$, so $\mathbf{1}$ is an eigenvector of \mathbf{G} with unit eigenvalue; thus, this is also true of $\mathbf{G}^{1/2}$, so $\mathbf{X}'^f \mathbf{G}^{1/2} \mathbf{1} = \mathbf{X}'^f \mathbf{1} = \mathbf{0}$. The overall construction of the size M square-root $\mathbf{G}^{1/2}$ is computationally inexpensive.

Sakov and Oke [41] have found that the choice of the symmetric square root leads consistently to better performance. However, for large ensemble, the symmetric square root can lead to the creation of outliers in the ensemble. One can prevent the appearance of these spurious members by right multiplying $\mathbf{G}^{1/2}$ by a random orthogonal matrix \mathbf{U} from time to time. To avoid bias in the posterior anomaly matrix, we must also ensure that \mathbf{X}'^a has zero mean. To do so, \mathbf{U} must preserve the mean, equivalent to requiring that $\mathbf{U} \mathbf{1} = \mathbf{1}$. Some authors [42, 43] have proposed a scheme based on sampling of the standard normal distribution and use of Householder reflections to construct such mean-preserving random rotations; this algorithm is presented in Algorithm 2. For the other assimilation steps, \mathbf{U} is simply the identity \mathbf{I}_M . The overall ETKF is summarized in Algorithm 3. For further reading on the differences between the sEnKF and the ETKF, we refer readers to Katzfuss et al. [44].

E. Covariance inflation

In the different flavors of the EnKF, one seeks to estimate the posterior covariance matrix of a state variable of large dimension n with only a few ensemble members, $M \sim 100$. Therefore, the posterior covariance is usually underestimated with large sampling errors. This rank deficiency typically leads to long-range spurious correlations within the covariance matrix. Over multiple assimilation cycles, the Kalman gain will decrease to zero and the measurements will no longer be assimilated into the ensemble; this is called *filter divergence*. However, efficient regularization techniques such as the *covariance inflation* and *localization* have been developed to mitigate these sampling errors. These techniques underpin the success of the ensemble Kalman filters in high-dimensional filtering problems. Localization is not implemented in this work but will be discussed in section IV to circumvent certain flaws of our flow estimator and will be investigated in future work.

Covariance inflation is typically applied between the forecast and analysis steps. A physical interpretation of the covariance inflation in the context of our vortex modeling is presented in the next section. The *multiplicative inflation* increases the spread of the ensemble about the sample mean, by rescaling the deviation $\mathbf{x}^{f,i} - \bar{\mathbf{x}}^f$ by the multiplicative factor $\beta > 1$ for each ensemble member:

$$\mathbf{x}^{f,i} \leftarrow \bar{\mathbf{x}}^f + \beta(\mathbf{x}^{f,i} - \bar{\mathbf{x}}^f), \text{ for } i = 1, \dots, M. \quad (34)$$

This form of inflation is equivalent to multiplying the sample prior covariance $\bar{\mathbf{P}}^f$ by β^2 ; the best results are typically found when β is set to 1.01 and 1.03 for the sEnKF and ETKF, respectively.

Additive inflation, in contrast, accounts for the process noise that is inherently missing from the dynamical model; such noise is critical for ensuring that the ensemble maintains a variance among its members. This technique adds to each ensemble member a sample from a Gaussian distribution with zero mean and covariance \mathbf{Q} . The additive inflation corresponds to a Tikhonov regularization of the prior covariance matrix:

$$\bar{\mathbf{P}}^f \leftarrow \bar{\mathbf{P}}^f + \mathbf{Q}. \quad (35)$$

Whitaker and Hamill [45] proposed to combine the additive and multiplicative inflations. Indeed, Anderson and Anderson [46] have shown that additive and multiplicative inflation account for distinct kind of errors. The multiplicative inflation tends to account for sampling errors due to the small ensemble size, while the additive inflation tends to correct intrinsic errors of the model. Multiplicative inflation only requires the tuning of one parameter. Adjusting additive inflation requires tuning as many parameters as the dimension of the state. This requires more work and a deeper understanding of the dynamical system. Raanes [38] showed that the additive inflation is a robust technique to account for model errors, more complex inflation schemes can be constructed to reduce further model errors which are not considered in this study [38, 45].

F. Flow estimation with an ensemble of aggregated vortex models

With the aggregated vortex model summarized in Section II and the details of the ensemble filtering methods discussed above, we are now in a position to summarize the overall method. A schematic of this method is presented in Fig. 2. In the forecast, an ensemble of aggregated vortex models are advanced by one time step. At the end of this forecast, multiplicative and additive covariance inflation are applied to each member of the ensemble. The analysis step is then carried out, either by applying Algorithm 1 for the sEnKF or Algorithm 3 for the ETKF.

It is useful to interpret the method physically. Before we discuss the physical roles of each step of the filter, it is important to understand the effect of the ensemble and its propagation on the vorticity field. To demonstrate this effect in the clearest manner, let us assume that the q vortex elements in each ensemble member are singular: point vortices rather than blobs. Their blob form is only used to regularize the Biot–Savart interactions between them, and our discussion here focuses only on the interpretation of the vorticity field itself.

We can then write the vorticity field at location \mathbf{r} and time step k as a function of the random state vector \mathbf{X}_k , whose components (aside from LESPC) constitute the strengths and positions of the singular elements:

$$\omega(\mathbf{r}, \mathbf{X}_k) = \sum_{j=1}^q \Gamma_k^j \delta(\mathbf{r} - \mathbf{r}_k^j), \quad (36)$$

where δ is the Dirac delta function. The expected value of the vorticity field at the end of time step k is given (ideally, for infinite ensemble) in terms of the posterior distribution function

$$\mathbb{E}[\omega(\mathbf{r}, \mathbf{X}_k)] = \int \omega(\mathbf{r}, \mathbf{x}_k) \pi_{\mathbf{X}_k}(\mathbf{x}_k) d\mathbf{x}_k. \quad (37)$$

(For simplicity of notation, we have omitted the fact that the distribution $\pi_{\mathbf{X}_k}$ is conditioned on the observations made thus far.) Under our Gaussian assumption, it is particularly easy to evaluate these integrals, and we arrive at

$$\mathbb{E}[\omega(\mathbf{r}, \mathbf{X}_k)] = \sum_{j=1}^q \frac{\bar{\Gamma}_k^j}{2\pi |\mathbf{P}_k^{r_j}|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{r} - \bar{\mathbf{r}}_k^j)^\top \mathbf{P}_k^{r_j^{-1}} (\mathbf{r} - \bar{\mathbf{r}}_k^j)\right), \quad (38)$$

where $|\cdot|$ denotes determinant; $\bar{\Gamma}_k^j$ and $\bar{\mathbf{r}}_k^j$ denote, respectively, the mean circulation and position of vortex element j at time step k ; and $\mathbf{P}_k^{r_j}$ is the 2×2 covariance submatrix associated with the position of vortex element j at step k .

In other words, we can interpret each vortex element’s uncertainty as defining an elliptically-shaped region, centered at its mean location and endowed with its mean strength. The behavior of this elliptical region’s shape over time is determined by the combined influences of three processes in the filter: the forecast, the inflation, and the analysis.

The role of the forecast step is straightforward: it constitutes an inviscid (i.e., advective) advancement of the vortex elements by one step and the creation of new vorticity to satisfy modeled edge conditions. The ensemble of such inviscid models establishes a set of slightly different displacements of each vortex element. Since we interpret this ensemble as approximating a Gaussian distribution both before and after the forecast, the set of displacements of each vortex define a constrained transformation of the vortex’s ellipse: a net advection of the center and a stretching and rotation of its shape.

The inflation step imposes two influences on each region’s shape. The multiplicative inflation stretches the ellipse uniformly by a small fraction in every direction. The additive inflation, on the other hand, comprises a single step of a random walk. It is well known that a random walk, applied over a large number of steps, approaches a Wiener process, and the associated probability distribution satisfies a linear diffusion equation [47]. This diffusion causes the elliptical region to spread, reminiscent of core spreading in vortex methods [48]. To simulate diffusion of viscosity ν , the random step is chosen from a normal distribution with standard deviation $\sqrt{2\nu\Delta t}$.

In the context here, the additive inflation occurs among other processes, so its interpretation is less clear. Chorin’s random vortex method [49], utilizing a large number of overlapping vortex blobs undergoing random walks, is known to converge to the solution of the Navier–Stokes equation as $q^{-1/2} \log q$, where q is the number of blobs [50]. In the EnKF context, it is possible to identify a stochastic differential equation that asymptotically describes the forecast and inflation steps in the limit of large ensemble [51]. We leave a rigorous interpretation of the ensemble of vortex models in this manner for future work. We merely observe that the variance of the distribution from which we select the additive inflation parameter loosely specifies the viscosity of a diffusion process. However, we do not attempt to match this effective viscosity to the actual viscosity in the truth system; rather, we tune the additive inflation to balance the trust between the forecast and analysis steps.

Finally, in order to reconcile the new observations made in this step—namely, the pressures measured on the surface of the wing—with the pressures predicted via the Bernoulli equation, the vortex elements’ positions and strengths (and LESPC) need to be adjusted. The EnKF analysis step provides this adjustment by assembling a minimum least-squares solution to the problem $\mathbf{h}(\bar{\mathbf{x}}^a) = \mathbf{z}^*$, subject to measurement noise \mathbf{V} ; the minimization is regularized by the forecast $\bar{\mathbf{x}}^f$, with its associated covariance $\bar{\mathbf{P}}^f$, ensuring that the new state vector does not stray far from its forecast. Mathematically, this problem is expressed as [14, 52]:

$$\bar{\mathbf{x}}^a = \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} (\mathbf{z}^* - \mathbf{h}(\mathbf{x}))^\top \mathbf{V}^{-1} (\mathbf{z}^* - \mathbf{h}(\mathbf{x})) + \frac{1}{2\beta} (\mathbf{x} - \bar{\mathbf{x}}^f)^\top \bar{\mathbf{P}}^f^{-1} (\mathbf{x} - \bar{\mathbf{x}}^f), \quad (39)$$

with β the multiplicative covariance inflation factor. The analysis step causes the covariance to shrink, according to (27). The elliptical region associated with each vortex element is translated, stretched, and rotated by the analysis step, but its area shrinks.

It should be noted that the dimension of the state vector, n , changes with each time step of this filtering process: it increases typically by six (three per newly released blob originating from each edge), and occasionally decreases as elements of zero strength are eliminated. There are no inherent restrictions in the EnKF on changes to the dimension of the state. However, in order to keep a consistent state dimension, i.e., the same number of blobs across the different ensemble members, the vortex elements whose circulations are aggregated into another element are not removed, but rather, are simply assigned zero circulation. Similarly, even if the LESPC does not exceed the current estimate of the LESPC, a new vortex with zero circulation is still introduced. If the blob has zero circulation across all ensemble members, then it is removed from the state vector.

We discretize the vortex model with the forward Euler time scheme with time step $\Delta t^* = 0.01$. The blob radius δ (normalized by c) is set to 5×10^{-3} and 9×10^{-3} for the sEnKF and the ETKF, respectively. The difference between these two blob sizes is of little physical relevance: each blob radius was chosen concomitantly with the inflation and noise parameters to reduce the time-averaged root mean-square error of the normal force prediction and the appearance of spurious pressures on the surface for the sEnKF and the ETKF. We use the same blob radius for the sEnKF as the previous study by Darakananda et al. [29]. For the aggregated vortex model, the vortex elements are mostly isolated and their interactions require little regularization. The non-zero blob radius primarily regularizes the interactions of the vortex elements soon after their release from the edges. Throughout this study, we use an ensemble of size $M = 50$. At the initial time, no vortices are present in the state vector of the vortex model. The ensemble is initiated with random samples for the LESPC drawn from $\mathcal{N}(0.5, 0.1)$, i.e., a normal distribution with mean 0.5 and covariance 0.1.

We use a different set of inflation parameters for the sEnKF and the ETKF:

- A multiplicative inflation $\beta = 1.01, 1.028$ for the sEnKF and ETKF, respectively.

- An additive inflation for the position of the vortices drawn from $\mathcal{N}(0, 1 \times 10^{-5})$ (normalized by c)
- An additive inflation for the strength of the vortices drawn from $\mathcal{N}(0, 1 \times 10^{-3} \Delta t^*)$ (normalized by Uc)
- An additive inflation for the LESPc drawn from $\mathcal{N}(0, 5 \times 10^{-5})$, $\mathcal{N}(0, 8.5 \times 10^{-5})$ for the sEnKF and ETKF, respectively.
- The measurement noise \mathbf{V}_k is drawn from $\mathcal{N}(0, 1 \times 10^{-8})$ (normalized by ρU^2)

These parameters and the blob radius have been manually determined, independently for the sEnKF and the ETKF, to minimize the time-averaged root-mean-squared error of the predicted normal force and the occurrence of spurious pressures on the surface. The sEnKF and ETKF share the same additive inflation parameters for the vortex properties. It is important to note that each of these parameters is chosen in order to balance the trust between the vortex model and the analysis step; lower values, for example, lead to a vortex model that is less responsive to measurement innovation. For example, the relatively larger additive inflation for the LESPc makes this parameter more responsive than the vortex parameters. The same parameters are used for all examples in the next section. These parameters are nondimensionalized by the undisturbed characteristic scales of the problem: the density of the fluid ρ , the chord length c and the translational velocity of the plate U . In our various numerical experiments, we have found that this choice of the parameters work well for perturbations of different strengths and shapes.

IV. RESULTS

In this section, we present the results of flow estimation for two strongly-perturbed flows about an infinitely thin plate at 20° . In the first case, a sequence of perturbations that mimic pulse actuation is applied near the leading edge of the plate. In the second case, the plate is subject to large scale and coherent perturbations created in the wake of an upstream cylinder. In the discussion of the results, positive and negative vortex elements will refer to vortex elements with positive (counter-clockwise) and negative (clockwise) circulation respectively.

In this study, the true pressure jump measurements are generated from simulations of a flat plate at Reynolds number $Re = 500$, carried out with a high-fidelity Navier–Stokes solver, based on the immersed boundary projection method with lattice Green’s function [15, 53]. It is important to emphasize that the flow perturbations are only present in the truth system, and their effect is only made available to the vortex model through the assimilation of the true pressure measurements.

Due to the stochastic nature of the filtering problem, results may vary from one simulation to another for the same filter. We estimate the uncertainty in the results in order to draw consistent conclusions about the performance of each filter. For assessment purposes, the results presented in this work have been obtained by running an ensemble of 100 realizations of the same filter on each flow configuration, with each realization consisting of an application of the EnKF (using an $M = 50$ ensemble of vortex models). From this ensemble of realizations, we construct the sample mean and standard deviation of the quantity of interest for comparison with the truth. If we assume that the results obtained over the different runs follow a Gaussian distribution, the 95% confidence interval can be estimated by considering the plus and minus deviation of twice the standard deviation from the mean. It should be noted that the number of vortex elements—and hence, the dimensionality of the state vector—varies from one realization of the filter to the next, so it is not possible to define a mean state from the ensemble of realizations.

As mentioned in Section II, the near encounters between the singular vortex elements and the plate can lead to spurious errors in the estimate of the pressure distribution, and hence, in the prediction of overall normal force on the plate. We use a median filter to remove these spurious spikes from our presented results. At each time step, the current value is replaced by the median value of the $l = 7$ previous time steps. The median filter is causal, so it is directly integrated into the estimation sequence.

Unfortunately, there is no simple and direct way to compare the discrete vorticity distribution generated by the state estimate with the true continuous vorticity distribution. Instead, we compare the normal force coefficient C_n (i.e., the integral of this pressure distribution on the plate rescaled by $\rho c U^2 / 2$) from the high-fidelity simulation and from its evaluation for the analysis ensemble obtained from either the sEnKF or the ETKF. The performance of each filter is assessed with the root-mean-squared error (RMSE), the standard deviation of the ensemble [36], and the interquantile range defined by the 2.5% and the 97.5% quantiles. These metrics are time-averaged. We define the RMSE between the true observation C_n^* (from the high-fidelity simulation) and the mean posterior normal force coefficient \bar{C}_n^a (i.e., the mean normal force coefficient for the different realizations computed from the posterior ensemble) as $\text{RMSE} = \|C_n^* - \bar{C}_n^a\|_2$. We quantify the dispersion of the ensemble with the sample standard deviation of the normal force coefficient $\sigma_{C_n^a}$. While the RMSE and the standard deviation of the ensemble can be affected by the spurious force events, the interquantile range defined between the 2.5% and the 97.5% quantiles is insensitive to outliers, and is used as a companion to the standard deviation to quantify the spread of the ensemble.




FIG. 3. (a): Schematic of the plate at 20° at $t^* = 0$ subject to actuation. Green dot depicts the location of the actuation. (b) Time history of the flow actuation.

A. Translating plate subjected to pulse actuation disturbances

In this section, we assess the data-assimilated vortex model on the response of an impulsively translating plate at 20° to disturbances applied near the leading edge. We apply the perturbations at a point $0.475c$ from the centroid along the plate and $0.1c$ above it, as shown in Fig. 3(a). The perturbations are introduced as a superposition of vertical body force of nominal strength $F_n = 0.03\rho U^2 c$ at $t^* = 2.5, 2.9, 3.0$ and 3.2 , distributed in Gaussian form in time and space, with temporal standard deviation of $t_{std}^* = 0.1$ and spatial standard deviation $0.1c$. The vertical body force $f(t^*)$ is given by:

$$\frac{f(t^*)}{F_n} = \mathcal{N}(t^*; 2.5, t_{std}^*) + \mathcal{N}(t^*; 2.9, t_{std}^*) + \mathcal{N}(t^*; 3.0, t_{std}^*) + \mathcal{N}(t^*; 3.2, t_{std}^*), \quad (40)$$

where $\mathcal{N}(t^*; \mu, \sigma)$ denotes the temporal Gaussian kernel, with mean μ and standard deviation σ , evaluated at t^* . Fig. 3(b) shows the time history of the force actuation. For reference, the gusts considered in Darakananda et al. [29] were weaker and non-overlapping, applied further upstream of the plate at $t^* = 3, 4$ with nominal amplitude $F_n = 0.01\rho U^2 c$. The disturbance considered here is more challenging due to the presence of strong and overlapping perturbations.

Fig. 4 compares the history of the surface pressure response from the truth with the mean of 100 realizations of the sEnKF and the ETKF. During the first two convective times, we observe the development of the leading-edge vortex. The successive flow actuation disturbances are easily detected by the distinct regions of strongly negative pressure. The first suction region is due to the disturbance applied at $t^* = 2.5$, while the second one is due to the superposed response to the flow disturbances centered at $t^* = 2.9, 3.0$, and 3.2 . Both filters match well with the true pressure distribution. The uncertainty of each filter is characterized by the sample standard deviation of the pressure, computed over the 100 runs; these are shown in the right column of Fig. 4. In both filters, two narrow bands of high variance, particularly strong near the leading edge, can be identified at $t^* = 2.5$ and 3.0 . These correspond to the instants of local maxima in the disturbance force. However, the sEnKF has a higher level of variability in the pressure distribution from one run to another. In particular, there is a significant band of high dispersion for the sEnKF at around $t^* = 3.8$, and additional smaller bands at other times. The ETKF does not exhibit such bands, a direct consequence of its statistically consistent analysis update.

The normal force coefficient is the integral of this pressure distribution on the plate. The top left panel in Fig. 5 compares the force coefficient obtained from the truth system with the mean of the sEnKF and the ETKF applications over 100 realizations. The mean force predicted by each filter agrees very well with the true force response. The peaks created by the actuation and the subsequent drop of force around $t^* = 4$ are also well predicted. However, consistent with our observations of the surface pressure data, both filters show variability near the peak disturbances; Fig. 5 also shows that the peaks are slightly underpredicted.

The sEnKF exhibits more variability than the ETKF at all times, and this higher variability is apparent in the wider uncertainty envelope for the normal force coefficient, particularly after $t^* > 3.5$. To better appreciate the raw behavior of each filter, the top right panel in Fig. 5 show the same results, but without the use of the median filter. The sEnKF, in Fig. 5(b), exhibits frequent spikes throughout the simulation, and particularly so after $t^* > 3.5$. These spikes are not eliminated by the ETKF, but they are much weaker and less frequent. Due to the presence of spikes, the 95% confidence interval based on the standard deviation can be corrupted. To better appreciate the raw performance of the two filters, we have also plotted the 2.5% and 97.5% quantiles for each filter in the lower panel in Fig. 5. Table




FIG. 4. Left column [(a), (b), (d)]: Spatiotemporal map of the pressure coefficient jump for an impulsively translating plate at 20° subjected to pulse actuation disturbance from (a) high-fidelity numerical simulation at Reynolds number 500, and mean over 100 realizations of an inviscid vortex model with (b) the sEnKF and (d) the ETKF. Right column [(c), (e)]: Spatiotemporal map of standard deviation of the pressure coefficient jump over 100 realizations for (c) the sEnKF and (e) the ETKF

I reports the time-averaged RMSE, the standard deviation and the interquartile range of the normal force coefficient for the sEnKF and the ETKF, without the application of the median filter. The ETKF outperforms the sEnKF with an average reduction of 29% of the RMSE, 49% of the standard deviation and 17% of the interquartile range. We should note that there is no downside to apply the median filter.

The origins of this high uncertainty are clear when we examine the flow behavior. In Fig. 6 we show the vorticity distribution from the truth system at three instants, $t^* = 3.0, 4.0,$ and 5.0 . Since the vorticity in the fluid is modeled with limited number of regularized point vortices, it is not possible to make a direct comparison; instead, we plot the locations of the elements for one realization of each filter, with the elements' signed circulations represented in colors consistent with the sign of vorticity in the truth data. Most of the large-scale vortex dynamics are captured by the data-assimilated vortex model. Over the first two convective times, the impulsive translation of the plate creates a leading-edge vortex that grows in size and strength. This feature is captured by the continuous release of new vortex elements in both filters. The disturbance applied at $t^* = 2.5$ creates a coherent structure that is rapidly advected along the plate and merges with the initial leading-edge vortex. At around $t^* = 3.0$, the resulting vortex lies slightly above the mid chord of the plate, evident in the left column of Fig. 6. Both the sEnKF and ETKF capture this feature with positive vortex elements clustered at a similar position. It is also important to note that the vorticity associated with the disturbance itself, clearly evident at $t^* = 3.0$, is *not* observed in the vortex model results. The absence of this disturbance vorticity is by design; only its effect on vortex elements shed from the edge of the plate is captured by the assimilation of pressure data.

Over the time interval $t^* \in [3.0, 4.0]$, the vortex models predict a separation of a cluster of positive vortex elements from the plate with dynamics similar to the truth flow. The middle column of Fig. 6 compares the results at $t^* = 4.0$. The large coherent structure with positive vorticity is well captured by the inviscid model with both filters, but is more tightly clustered in the case of the ETKF. From $t^* = 4.0$ to 5.0 , the vortex is shed into the wake and triggers a large flux of negative vorticity from the trailing edge. The right column of Fig. 6, for $t^* = 5.0$, shows a good visual agreement between the truth and the prediction of the ETKF. The spatial vortex distribution predicted by the sEnKF is less representative of the true vorticity distribution, however. Overall, the ETKF predicts a more structured and physically consistent spatial distribution of the vortices than the sEnKF. The coherence of these vortex element

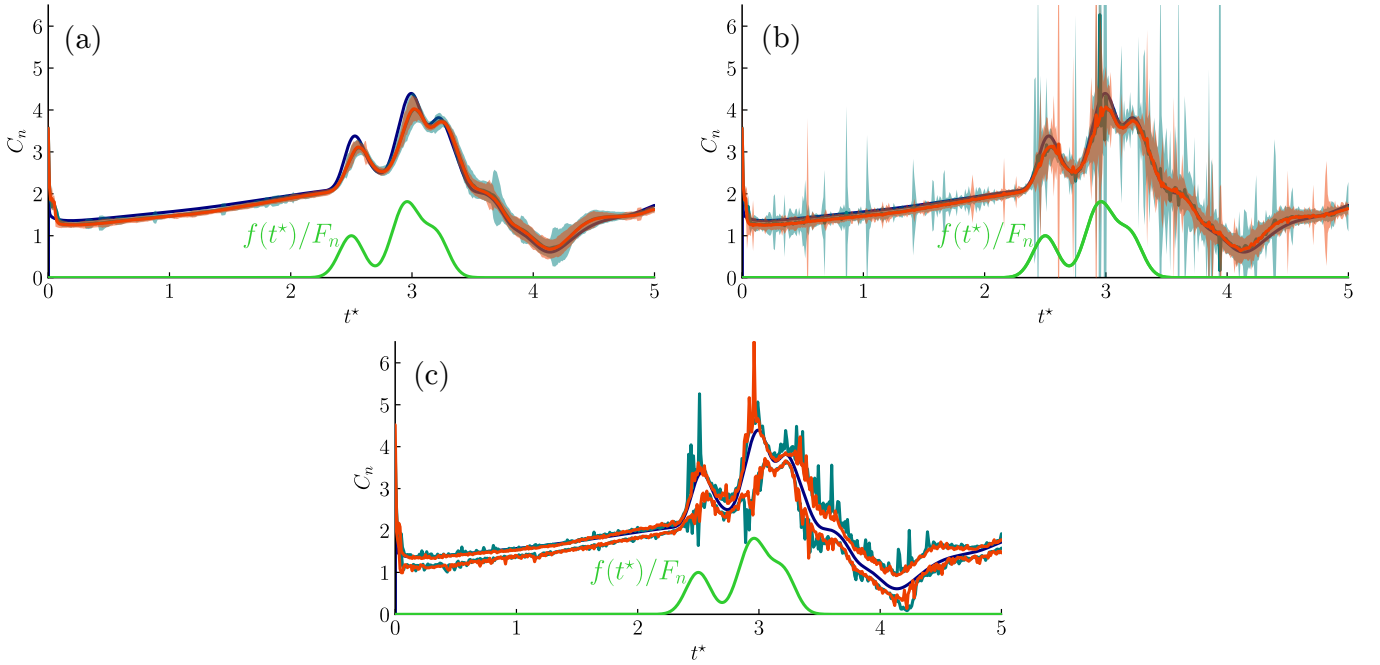


FIG. 5. Top left panel (a): Normal force coefficient of an impulsively translating plate at 20° subject to actuation, from high-fidelity numerical simulation at Reynolds number 500 (—■—), mean over 100 realizations of the inviscid vortex model with sEnKF (—■—), mean over 100 realizations of the inviscid vortex model with ETKF (—■—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF. Time history of the flow actuation (—■—). Top right panel (b): Comparison with the same normal force coefficient from high-fidelity simulation, but without application of the median filter for the sEnKF (—■—) and ETKF (—■—). Lower panel (c): Comparison of the same normal force coefficient from high-fidelity simulation, with the 2.5% (lower curve) and 97.5% (upper curve) quantiles of the normal force coefficient over 100 realizations of the inviscid vortex model for the sEnKF (—■—) and ETKF (—■—) without application of the median filter.

clusters is consistent with the narrower uncertainty in the normal force coefficient, and demonstrates a clear advantage of the ETKF over the sEnKF for modeling the flow response to unknown flow perturbations.

In Fig. 7 we compare the time histories of the LESPc estimated by the sEnKF and the ETKF. The LESPc is constrained to remain positive, and the constraint reverts to the Kutta condition if LESPc becomes zero. Before the flow is disturbed, the mean estimate of the LESPc stays on a plateau about 0.5, the mean value in the initial ensemble. This behavior supports the hypothesis of Ramesh et al. [24, 31]: the LESPc remains constant for a given Reynolds number and airfoil section. The time variation of the imposed disturbance is reflected in a similar variation of the LESPc. Indeed, the application of an actuation-like disturbance near the leading edge directly controls the vorticity flux about this edge. Large values of LESPc lead to weaker vorticity, temporarily suppressing the flux into the shear layer. The small decay of LESPc after the first disturbance increases the vorticity flux, triggering the creation of a new coherent structure that merges with the initial leading-edge vortex about the mid-chord. This leading-edge development is then halted after the next disturbance peak, and the leading-edge vortex is shed. The uncertainty envelopes of the two filters are very similar and tend to grow over time. The width of these envelopes is large, reflecting significant variation in the estimated values of LESPc from one realization to the next. This variation indicates a weak physical correlation between this threshold value and the pressure on the plate: this threshold's effect on pressure is only exerted indirectly, through the subsequent release of vorticity. (The LESP itself, in contrast, is more strongly correlated, since shed vortex elements contribute to this value [21, 24].)

The population histories of vortex elements is depicted in Fig. 8. The histories are nearly identical for each filter. In the initial time steps, the number of vortex elements increases linearly from 0 to 20; the model prevents aggregation of elements during this interval. Subsequently, the population grows slowly up to 35 ± 7 at $t^* = 5$. It should be noted that, without aggregation, this population would be approximately 1000 vortex elements (500 time steps, with two elements shed per step). The variation in population among the different realizations of each filter is attributable to the variation in LESPc, which sets the initial strengths of the elements, which in turn affects their later aggregation. It is interesting to note that this variation of vortex element populations is proportionally larger than the variation in the pressure and normal force, indicating that there is some non-uniqueness in the mapping from surface pressures to vortex element dynamics.

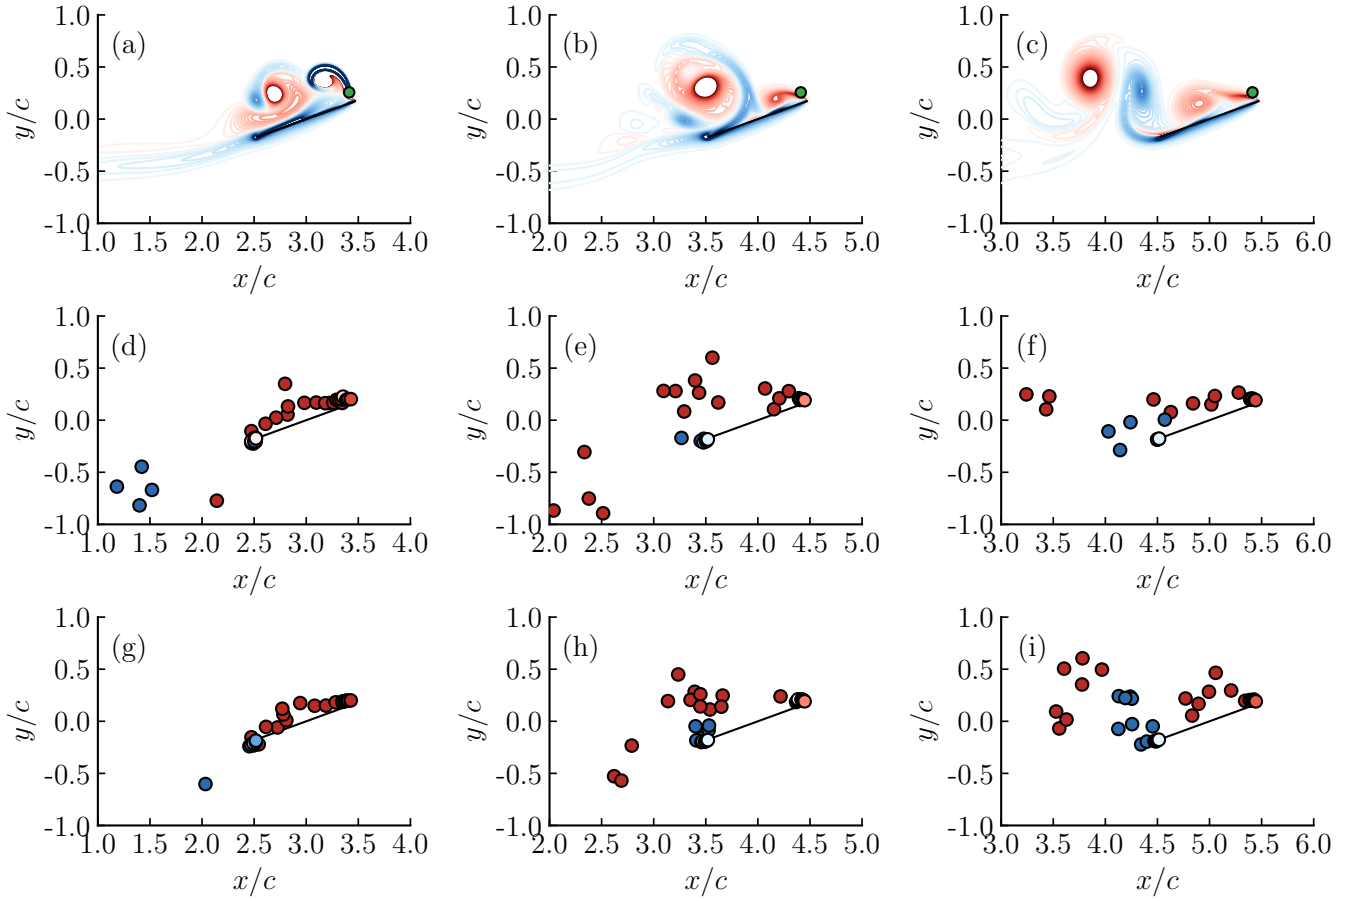


FIG. 6. Snapshots of the vorticity distribution at $t^* = 3.0$ (left column), $t^* = 4.0$ (middle column) and $t^* = 5.0$ (right column) for an impulsively translating plate at 20° subject to actuation, predicted from [(a)-(c)] high-fidelity numerical simulation at Reynolds number 500, [(d)-(f)] inviscid vortex model with sEnKF, and [(g)-(i)] inviscid vortex model with ETKF.




FIG. 7. Time history of the ensemble mean value of the LESPC of an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

Fig. 9 depicts the ensemble variances of the positions and strengths of the vortex elements and the LESPC estimate, averaged over the 100 realizations of each filter. Each variance is lower-bounded by the additive covariance inflation to avoid filter divergence. The variances of the strengths of the vortices and LESPC are fairly constant (and near the values set by the inflation parameter) while those of the x and y positions are more variable. These values are larger than the variance set by the inflation, likely due to additional error incurred by aggregation.




FIG. 8. Time history of the particle count for an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

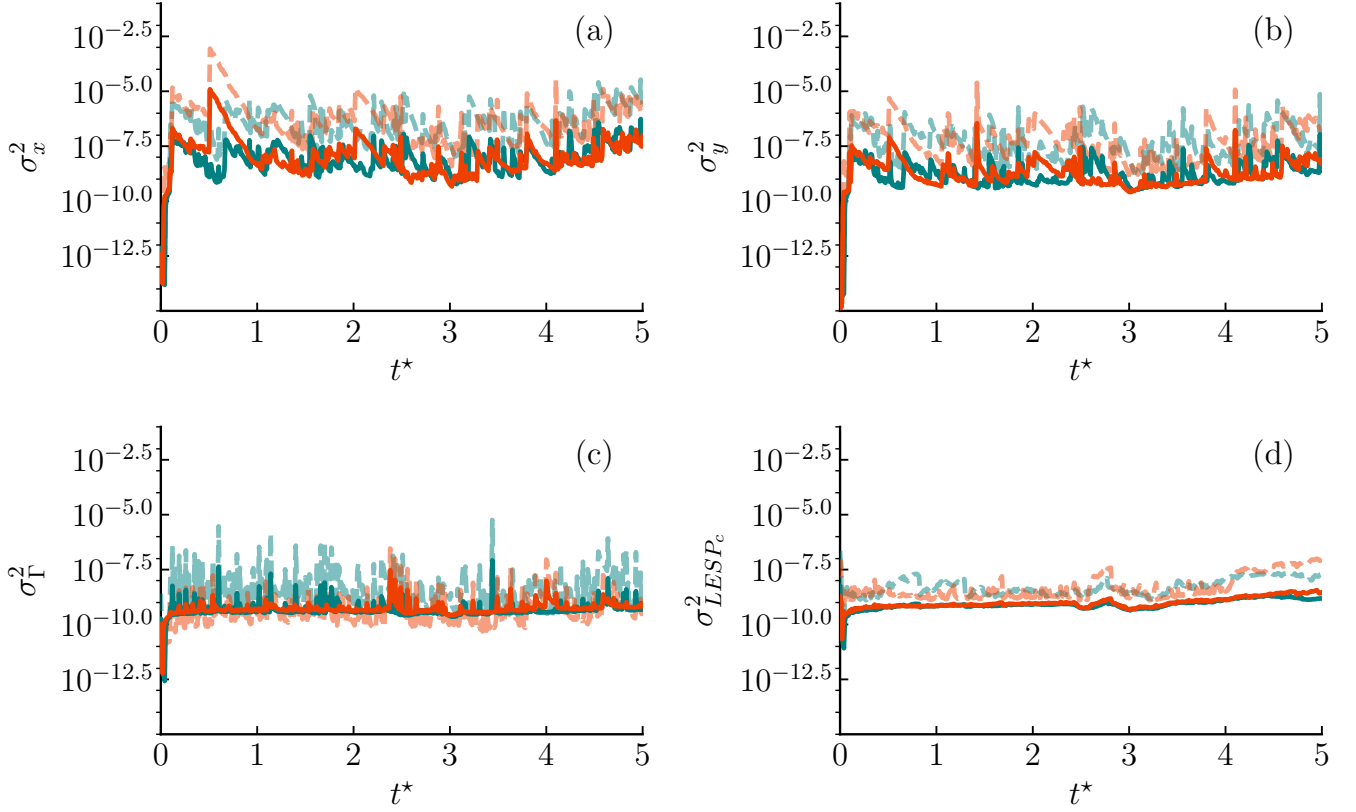


FIG. 9. Time history of the ensemble variances of an impulsively translating plate at 20° subject to actuation, averaged over 100 realizations. Mean variances for (a) the x coordinate of the blobs, (b) the y coordinate of the blobs, (c) the circulation of the blobs, and (d) the LESPC from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Fainter dashed lines show the standard deviation of the different variances over the 100 realizations for the inviscid vortex model with sEnKF and ETKF.

	RMSE	Standard deviation	Interquantile range
sEnKF	3.67	3.06×10^{-1}	4.71×10^{-1}
ETKF	2.61	1.56×10^{-1}	3.92×10^{-1}

TABLE I. Time-averaged RMSE, standard deviation and interquantile metrics of the normal force of the impulsively translating plate subject to actuation with the sEnKF and the ETKF, without application of the median filter. The lowest RMSE, the standard deviation and the interquantile range among the two filters is highlighted in bold.




FIG. 10. Schematic of a plate at 20° behind a cylinder. The plate and cylinder translate to the right at uniform speed.

B. Plate in the wake of a cylinder

In this part, we assess our flow estimator when applied to a plate at 20° angle of attack in the wake of a cylinder. In the truth system, high fidelity simulations are conducted of a cylinder of diameter $0.16c$ is centered 2 chord lengths upstream and 0.3 chord lengths above the plate's centroid, as shown in Fig. 10. The plate and cylinder are both impulsively set in motion at $t^* = 0$ at speed U , so that their relative configuration remains fixed for all time. The Reynolds number based on the cylinder diameter is 80, sufficiently large that the cylinder's wake exhibits a von Kármán vortex street. The presence of the plate triggers the wake to break symmetry and achieve this vortex street. This flow configuration distills the main features of a vehicle flying through the wake of a structure, e.g., buildings in an urban environment or other flying vehicles.

It is important to recall that we do not include the cylinder and its wake in our vortex model flow estimator; its effects are only felt through the pressure jump measurements obtained from surface sensors along the plate. Fig. 11 depicts the history of the true surface pressure distribution and those estimated by the sEnKF and ETKF; the associated standard deviations are shown on the right. Both filters estimate pressure fields that agree very well with the true distribution. As expected, it takes about 2 convective times for the cylinder wake to reach and be sensed by the plate. Over this time window, the pressure field is essentially disturbance-free. The encounter of the vortex structures shed by the cylinder leaves successive short-lived pressure disturbances.

The true vorticity distribution and mean sets of vortex elements are compared in Fig. 12 at $t^* = 3.5$ (soon after the wake has reached the plate), at $t^* = 4.2$ (when the cylinder wake transitions to a von Kármán vortex street) and at $t^* = 10$ (long after the flow has achieved periodic vortex shedding). As expected, the vortex models do not attempt to represent the cylinder wake with vortex elements; rather, the filter accommodates the influence of the wake vorticity by modifying the behavior of the vortex elements shed from the plate. Large scale structures of the flow around the plate are captured and match visually with the true vorticity field.

Concomitantly, the normal force estimate agrees well with the truth, as seen in Fig. 13. As in the first example, we have plotted the normal force estimate of the two filters, without the application of the median filter, and the associated uncertainty between different realizations based on the standard deviation (Fig. 13 (b)) and the interquantile range (Fig. 13 (c)). The transition from a symmetric cylinder wake to a periodic vortex shedding causes some challenge to the discrete vortex models, as highlighted by the temporary growth of the uncertainty envelope around $t^* = 4.0$. The envelope remains small for both filters after the periodic wake behavior has been established. Table II reports the time-averaged RMSE and the standard deviation and the interquantile range of the normal force coefficient for the sEnKF and the ETKF. The ETKF reduces on average the standard deviation by 16% and the interquantile range by 16%. The RMSE performance of the two filters are almost identical. The history of the LESPC estimate in Fig. 14 reveals the same growth and decay of this critical value observed in the previous example; here, they represent a response to individual cylinder wake vortices passing the leading edge of the plate. These variations are essential to our inviscid framework to control the leading-edge vortex shedding in the presence of flow perturbations. As in the previous example, the weak correlation between the LESPC and measured pressure causes significant volatility in the estimated LESPC among the different realizations. The vortex element population remains small, $O(60)$, after 12

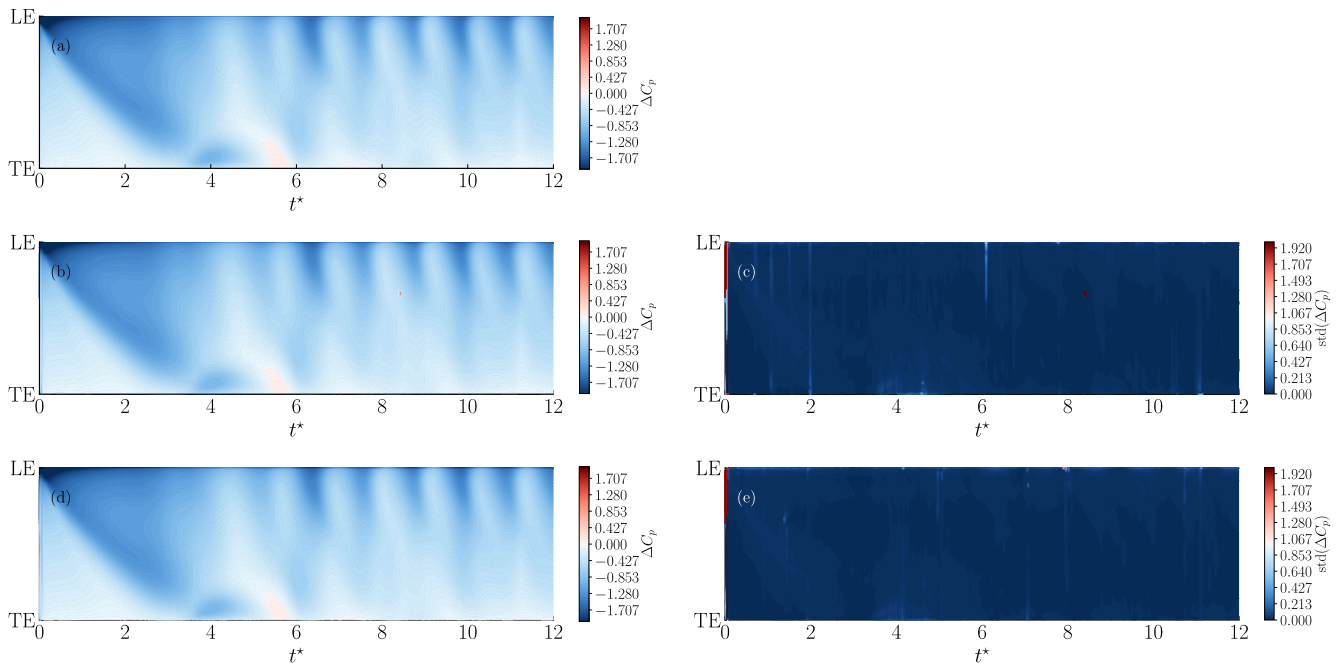


FIG. 11. Left column [(a), (b), (d)]: Spatiotemporal map of the pressure coefficient jump for an impulsively translating plate at 20° in a cylinder wake (a) high-fidelity numerical simulation at Reynolds number 500, mean over 100 realizations of an inviscid vortex model with the sEnKF (b) and the ETKF (d). Right column [(c), (e)]: Spatiotemporal map of standard deviation of the pressure coefficient jump over 100 realizations for the sEnKF(c) and ETKF (e)

	RMSE	Standard deviation	Interquantile range
sEnKF	2.63	1.63×10^{-1}	3.18×10^{-1}
ETKF	2.66	1.41×10^{-1}	2.69×10^{-1}

TABLE II. Time-averaged RMSE, standard deviation and interquantile metrics of the normal force of the impulsively translating plate subject in a cylinder wake with the sEnKF and the ETKF without the median filter. The lowest RMSE, the standard deviation and the interquantile range among the two filters is highlighted in bold.

convective time units. Without aggregation, we would have to track 2400 vortex elements.

C. Dissection of the EnKF analysis

The results shown thus far depict the overall evolution of the state with time, but do not reveal the manner in which individual components of the state—vortex element positions and strengths, and LESPC—are modified by the analysis step to account for pressure mismatch. Here, we examine this more closely, and particularly, ask whether there is some notion of *locality* inherent in the radius of action of a given pressure measurement on the correction of the state components. This question motivates our dissection of the input-output structure of the Kalman gain matrix $\mathbf{K} \in \mathbb{R}^{n,d}$ (20), the linear map from the d sensors (measurement innovation) to the n -dimensional state correction. The component K_{ij} of the Kalman gain tells us how much the j th measurement updates the i th component of the state; a low value (in magnitude) indicates a low impact. As we recall from (32), the Kalman gain is derived from the ensemble’s forecast anomalies.

Figs. 16, 17 examine how the ETKF algorithm modifies the state during the analysis step at $t^* = 1$ and 10, respectively. We use a different visualization convention for the vortex element distribution in Fig. 16(a) and 17(a). The size of each element inversely indicates its index in the collection of vortices. Since newly-added elements are assigned indices at the end of the list, larger circles correspond to “older” vortices. Colors still indicate vorticity sign and strength. With this convention, it is easier to see which vortex elements are updated by the filter. At $t^* = 1$, the state vector has dimension $n = 67$, corresponding to 22 vortex elements and the LESPC estimate. Fig. 16(b) shows a heat map of the magnitudes of the Kalman gain elements. Low measurement index corresponds to a pressure sensor




FIG. 12. Snapshots of the vorticity distribution at $t^* = 3.5$ (left column), $t^* = 4.2$ (middle column) and $t^* = 10.0$ (right column) for an impulsively translating plate at 20° in a cylinder wake, predicted from [(a)-(c)] high-fidelity numerical simulation at Reynolds number 500, [(d)-(f)] inviscid vortex model with sEnKF, and [(g)-(i)] inviscid vortex model with ETKF.

near the trailing edge. The LESPC estimate is always the last component of the state vector. Overall, the magnitudes of the components of the Kalman gain are small. The effect of these small gain values is highlighted in Fig. 16(e), which shows the mean analysis updates to the non-dimensional x and y positions and circulations of vortex elements. These updates are comparable to or smaller than the time step size 0.01, and thus, of similar order of magnitude to their forecast updates.

It is useful to recall the relationship between pressure and vorticity in incompressible flow. Though we formally compute the pressure in the vortex model with the unsteady Bernoulli equation, it can also be obtained through the following Poisson equation:

$$\nabla^2 \left(p + \frac{1}{2} \rho |\mathbf{u}|^2 \right) = \rho \nabla \cdot (\mathbf{u} \times \omega \mathbf{e}_3), \quad (41)$$

where \mathbf{u} denotes the velocity field. This velocity depends in part on the contributions from all other vortices (along with the body's motion and/or uniform flow). Thus, this equation clearly illustrates that pressure depends non-linearly on the vorticity field. But even if the model was made linearly dependent on vorticity, e.g., by linearizing about some base flow, the pressure remains non-linearly dependent on the positions of the vortices, due to the ellipticity of relationship (41). Therefore, the linear analysis step in the Kalman gain might be inadequate to estimate the true modification of a vortex due to a pressure mismatch. Also, the EnKF is a sequential filter, which means that it performs iterated state updates based on pressure difference. But even if each analysis step only slightly modifies the state, the long-term impact of these linear updates is important.

Some interesting physical features can be inferred by a closer study of Fig. 16(b). The top line of the heat map shows the impact of pressure measurements on the LESPC estimate. Only a few pressure sensors near the leading edge have a high impact on the LESPC. This is not surprising, since LESPC a measure of the integrated pressure about the leading edge, and LESPC is a threshold on this value. The heat map also shows that pressure sensors near the trailing edge do not impact the state substantially. For a given vortex element, the most impactful sensors are those that are closest. Most of the state update are due to sensors between the leading edge and mid-chord, consistent with the position of the leading-edge vortex about the mid-chord at $t^* = 1$. The newly released vortex elements are updated by the pressure sensors in the immediate vicinity of the leading-edge.

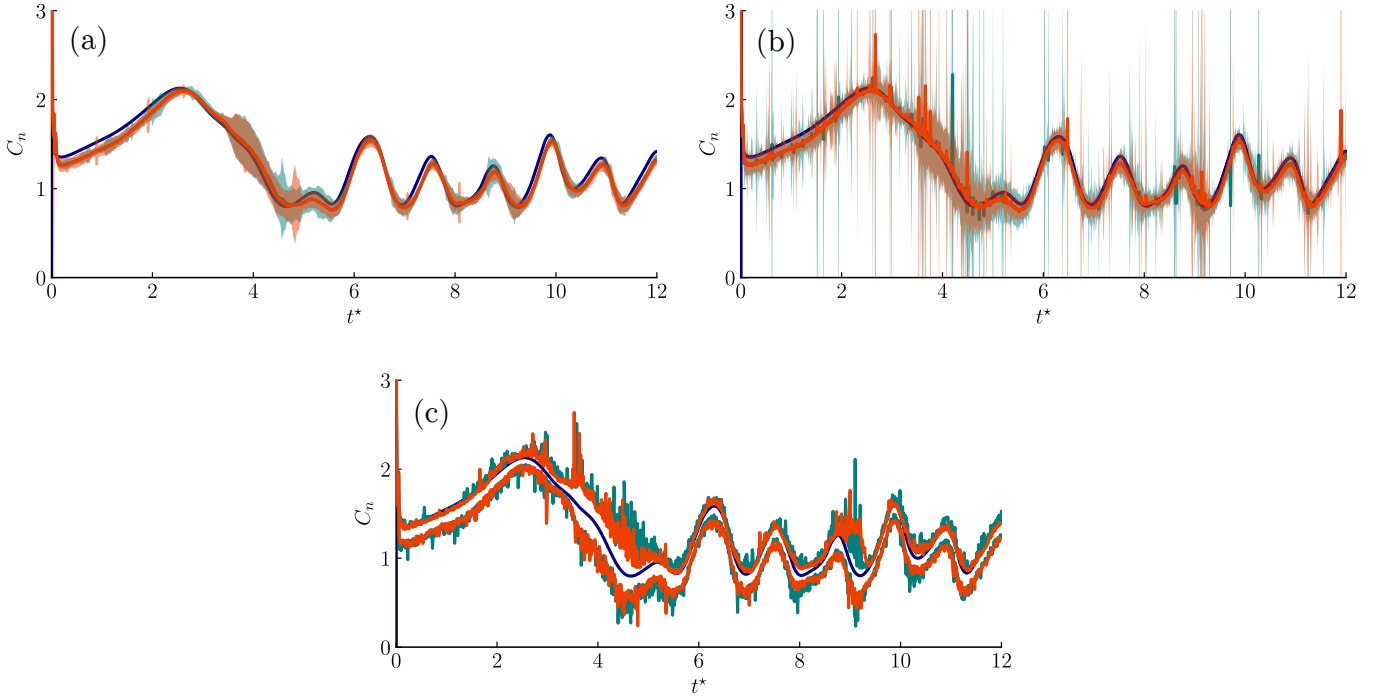


FIG. 13. Top left panel (a): Normal force coefficient of an impulsively translating plate at 20° in a cylinder wake, from high-fidelity numerical simulation at Reynolds number 500 (—), mean over 100 realizations of the inviscid vortex model with sEnKF (—), mean over 100 realizations of the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF. Top right panel (b): Comparison with the same normal force coefficient from high-fidelity simulation, but without application of the median filter for the sEnKF (—) and ETKF (—). Lower panel (c): Comparison of the same normal force coefficient from high-fidelity simulation, with the 2.5% (lower curve) and 97.5% (upper curve) quantiles of the normal force coefficient over 100 realizations of the inviscid vortex model for the sEnKF (—) and ETKF (—) without application of the median filter.




FIG. 14. Time history of the ensemble mean value of the LESPC of an impulsively translating plate at 20° in a cylinder wake, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

To explore the update process more deeply, we perform a singular value decomposition of the Kalman gain matrix:

$$\mathbf{K} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top, \quad (42)$$

where $\mathbf{K} \in \mathbb{R}^{n,d}$, $\mathbf{U} \in \mathbb{R}^{n,d}$, $\mathbf{\Sigma} \in \mathbb{R}^{d,d}$, $\mathbf{V} \in \mathbb{R}^{d,d}$. The columns of \mathbf{U} and \mathbf{V} are called the left and right singular vectors, and can also be interpreted here as modes of the response (of the state) and forcing (from the measurements), respectively. The response-forcing pairs are ranked by decreasing gain, given by the diagonal entries of the diagonal matrix $\mathbf{\Sigma}$, the singular values; these are shown in Fig. 16(c) for the Kalman gain at $t^* = 1$. Each measurement mode is transformed into a response mode, amplified by their corresponding singular value. The decay rate of the singular values is a useful indication of the effective number of pressure measurements used in the analysis step. Here, fourteen modes are necessary to capture 99% of the energy (sum of the square of the singular values). The response modes tell us which state indices are the most impacted. Fig. 16(d) depicts the first eight response modes that account for




FIG. 15. Time history of the particle count for an impulsively translating plate at 20° in a cylinder wake, averaged over 100 realizations, from the inviscid vortex model with sEnKF (—) and the inviscid vortex model with ETKF (—). Shaded areas show the 95% confidence interval for the inviscid vortex model with sEnKF and ETKF.

90% percent of the energy. The largest updates are applied to the LESpc estimate (the largest state index) and to the oldest vortices with large circulation. Though the results in Fig. 16 correspond to the ETKF, we have also found similar results for the sEnKF.

Fig. 17 shows equivalent plots for $t^* = 10$. The update of the LESpc is still driven by the sensors near the leading edge and a few about the mid-chord. The sensors near the leading edge still have a stronger impact on the state correction. The SVD of the Kalman gain reveals in Fig. 17(d) that the LESpc and a particular vortex element are the most affected components of the state. From Fig. 17(e), this particular element's update is displaced by $0.06c$, $0.02c$ in the x and y directions, respectively. The element is located five chord lengths away from the plate, an unlikely recipient of update from any surface measurements, particularly to sensors near the leading edge as panel (b) indicates. Similarly, we do not expect that pressure sensors about the mid-chord will update the LESpc value. These spurious updates reveal a flaw of using a finite ensemble size. Indeed, the Kalman gain that maps measurement discrepancies to state updates is estimated from the finite forecast ensemble. However, one cannot hope to recover the exact covariance matrices from a limited number of samples. The empirical covariance matrices suffer from sampling errors, leading to rank-deficiency and possible spurious long-range correlations. These errors in the estimated Kalman gain might lead, as in this case, to an inadequate state update.

In future work, we will explore the concept of *localization* to remove these artificial correlations by performing a local state update [11, 51, 54]. In other words, pressure measurements will be assimilated one at a time and only affect state components that are within a certain distance from that specific pressure sensor location. Intuitively, we want to embed into our estimation framework that vortices in the far wake of the plate should not be modified by pressure discrepancies on the plate. Localization is another form of regularization, complementary to the covariance inflation that helps overcome the rank deficiency of the covariance matrices. By performing local analyses, each reduced inverse problem can be made full rank and improves the overall performance of the filter [11].

V. CONCLUSIONS

This paper presents several refinements of the framework introduced by Darakananda et al. [29] for the estimation of two-dimensional unsteady aerodynamic flows. As in [29], the dynamical model is a discrete vortex model that advects a collection of regularized vortices, successively released from the edges of an infinitely thin plate and aggregated to maintain a modest population. The unrepresented physical features of the vortex model, e.g., effects of viscosity, are accounted for by the assimilation of pressure measurements obtained from the true physical process. Vorticity flux from the leading edge is obtained by a critical leading-edge suction parameter whose value is also obtained from the measurements.

The data assimilation has been carried out with the ensemble Kalman filter (EnKF), which relies on a finite-sized ensemble of vortex models to forecast the state and uncertainty of the system. We have interpreted this ensemble forecast physically as an effective translation, stretching, and rotation of the elliptically shaped regions defining each vortex element's uncertainty region. We have explored two different versions of the EnKF: the stochastic ensemble Kalman filter (sEnKF), which was also used in [29], and the ensemble transform Kalman filter (ETKF). The sEnKF introduces additional numerical noise by not exactly obtaining the correct covariance propagation during the analysis step. In the current vortex estimation procedure, this noise manifests itself as spurious pressures exerted on the surface of the wing. The ETKF [30], in contrast, reproduces the exact covariance relation even for finite sized ensembles. We have found that this latter filter greatly mitigates the appearance of spurious pressures on the surface, and clusters of vortex elements match better with the true coherent structures. As in [29], multiplicative and additive covariance inflation are used to prevent filter divergence. Overall, we have found that the ETKF has better performance and

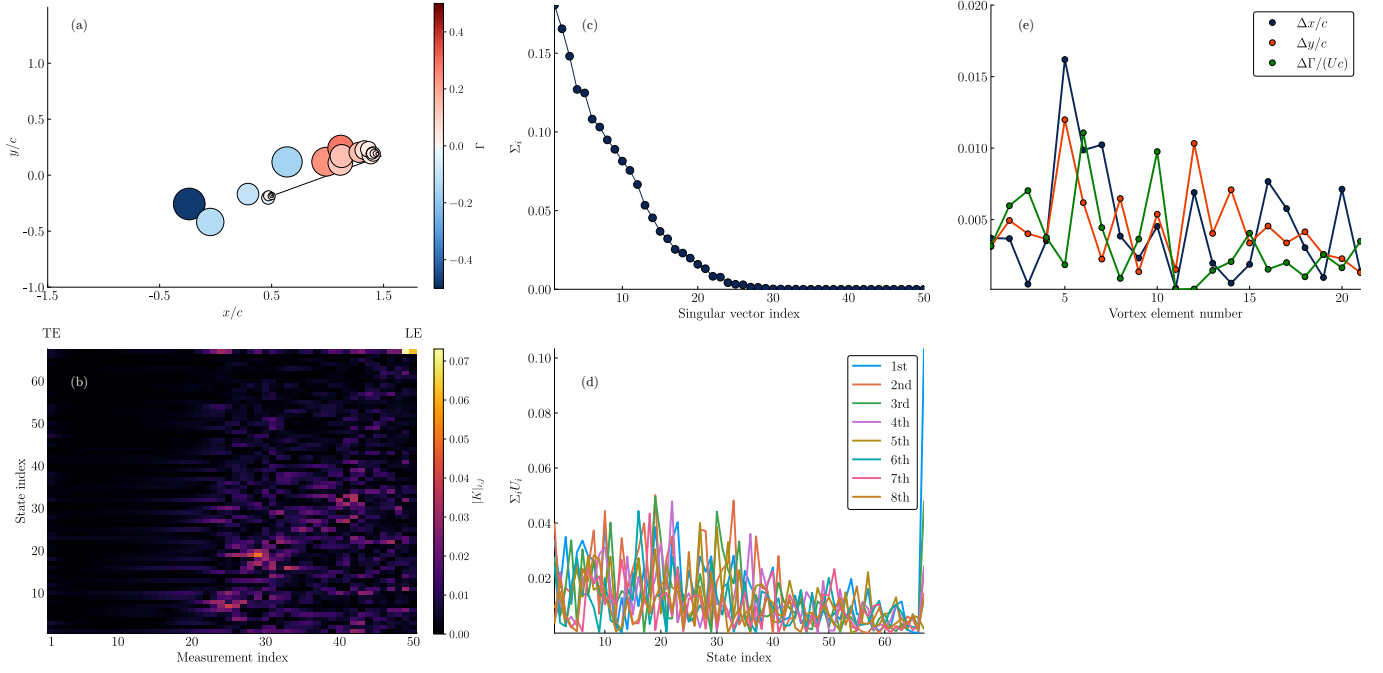


FIG. 16. Study of the analysis step at $t^* = 1$ for the impulsively translating plate in the cylinder wake. (a): Vorticity distribution from the discrete vortex model with ETKF. Sizes indicate "age" of the vortices (large means old). Colors indicate vorticity sign (blue is negative). (b): Element-wise magnitude of the Kalman gain. (c): Singular values of the Kalman gain. (d): First eight left singular vectors of the Kalman gain scaled by their associated singular value. (e): Mean difference of x, y position and circulation of vortices during the analysis step.

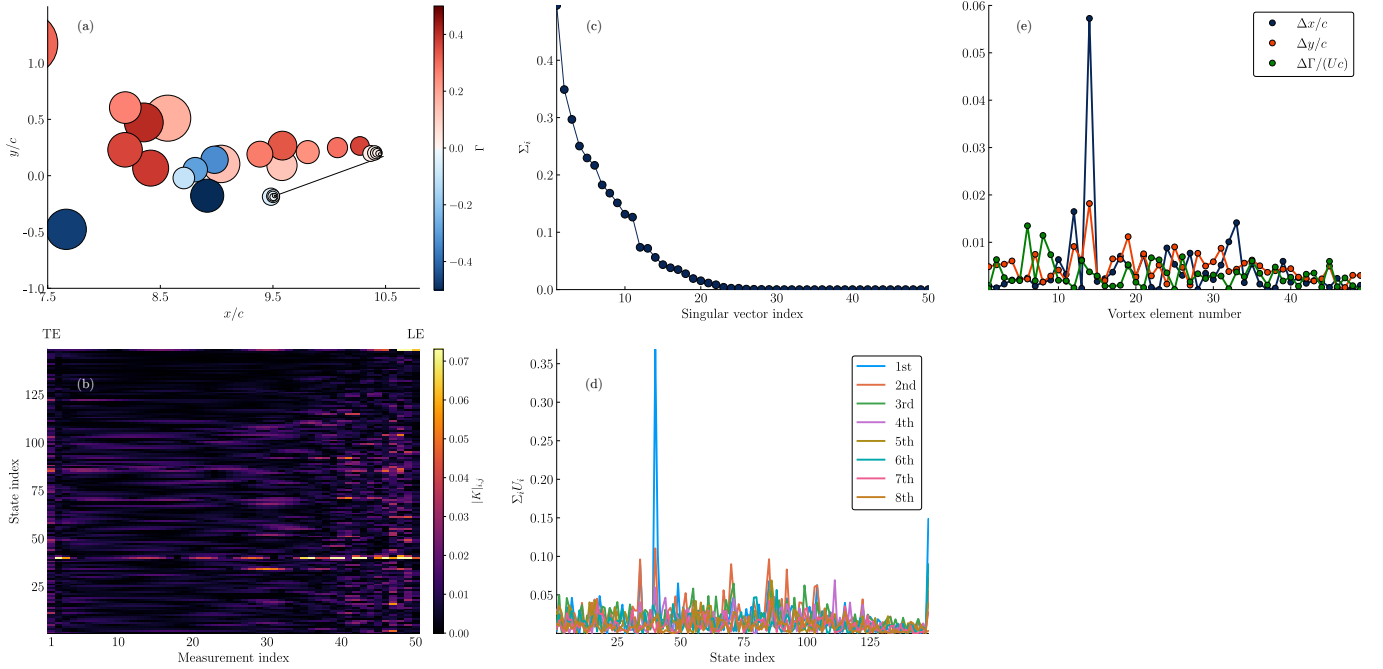


FIG. 17. Study of the analysis step at $t^* = 10$ for the impulsively translating plate in the cylinder wake. (a): Vorticity distribution from the discrete vortex model with ETKF. Sizes indicate "age" of the vortices (large means old). Colors indicate vorticity sign (blue is negative). (b): Element-wise magnitude of the Kalman gain. (c): Singular values of the Kalman gain. (d): First eight left singular vectors of the Kalman gain scaled by their associated singular value. (e): Mean difference of x, y position and circulation of vortices during the analysis step.

robustness than the sEnKF.

Two flow configurations have been considered to assess the performance of our flow estimator and the influence of the filtering algorithm. The true pressure measurements were obtained from high-fidelity numerical simulations of an impulsively translating plate at 20° and Reynolds number $Re = 500$, subject to flow disturbances. The first example examined the response to a sequence of strong body force pulses applied near the leading edge mimicking flow actuation. We demonstrate that the ensemble of data-assimilated vortex models accurately capture the non-linear flow interactions between the pulses' responses. It should be stressed that a linear convolutional model, with a convolution kernel based on a single pulse, would fail due to these non-linear interactions. The second example consisted of the response of a plate subject to the large scale and coherent perturbations created in the wake of an upstream cylinder. The normal force coefficient on the plate was accurately estimated over 12 convective time units. It should be stressed that the estimation framework does not require a model of the disturbance itself. The disturbance's effect on shed vorticity is accounted for by using measurements to update the state vector in the filter analysis step. This feature is particularly beneficial for a real-time aerodynamic flow estimator, in which the structure of incident flow disturbances such as gusts are unknown.

It is important to stress that our framework is independent of the source of the measurements, and our objective in ongoing work is to assimilate measurements from a small number of surface pressure transducers in a physical experiment or on board a flight vehicle. However, there are still challenges to address first before we can achieve this. In particular, we are exploring ways to further improve the analysis step. In this work, with the help of an SVD of the Kalman gain, we have illuminated the function of this step by examining how pressure discrepancies modify the state components. This examination has shown that, in most cases, the Kalman gain ensures that sensors primarily affect state components that are physically close, e.g., vortex elements near the plate or the value of LESPc. However, the finite ensemble size can lead to spurious long-range correlations and non-physical updates of the vortices in the far wake. To mitigate this issue, we are currently investigating localization, which aims to remove these long-range correlations by performing local analyses. We are also exploring approaches to systematically sparsify the sensors needed for analysis updates.

Given the nonlinearity of the state-space model (i.e., the Biot-Savart law and the unsteady Bernoulli equation), it is reasonable to expect the probability densities of the vortex states to be non-Gaussian. Therefore, the linear transformation of the analysis step of the EnKF can be inaccurate for estimating the true effect of a pressure discrepancy on a vortex. To address this limitation, in our current work we propose to use a generalization of the EnKF introduced by Spantini et al. [36], that relies on parsimonious and interpretable nonlinear transformations in the analysis step to reduce the bias of the EnKF but maintain its robustness. Moreover, real-world pressure data typically contains many outliers, and are not well represented by Gaussian distributions. To develop a flow estimator able to assimilate limited observations from a physical experiment, we are currently exploring the use of a Laplace density to replace the Gaussian density of the observational noise. A Laplace density has heavy tails that will assign higher probability to outliers, and be more representative of the statistics of real-world pressure data [55].

VI. ACKNOWLEDGMENTS

Support by the U.S. Air Force Office of Scientific Research FA9550-18-1-0440 is gratefully acknowledged. The authors also greatly benefited from discussions with Ricardo Baptista (MIT), Tim Colonius (Caltech), Youssef Marzouk (MIT), and David Williams (IIT).

Appendix A: Ensemble Kalman filter algorithms

-
- [1] Michael H Dickinson and Karl G Gotz. Unsteady aerodynamic performance of model wings at low reynolds numbers. *Journal of experimental biology*, 174(1):45–64, 1993.
 - [2] James M Birch and Michael H Dickinson. Spanwise flow and the attachment of the leading-edge vortex on insect wings. *Nature*, 412(6848):729–733, 2001.
 - [3] Michele Milano and Morteza Gharib. Uncovering the physics of flapping flat plates with artificial evolution. *Journal of Fluid Mechanics*, 534:403–409, 2005.
 - [4] Kunihiro Taira and TIM Colonius. Three-dimensional flows around low-aspect-ratio flat-plate wings at low reynolds numbers. *Journal of Fluid Mechanics*, 623:187–207, 2009.

Algorithm 1: Algorithm for the sEnKF, adapted from Asch et al. [11]

Initialize M ensemble member $(\mathbf{x}_0^{a,i}) \sim \mathcal{N}(\bar{\mathbf{x}}_0, \mathbf{P}_0)$
for $k = 1 : K$ **do**
 Compute the ensemble forecast:
 $\mathbf{x}_k^{f,i} = \mathbf{f}(\mathbf{x}_{k-1}^{a,i}) + \mathbf{w}_{k-1}^i$, for $i = 1, \dots, M$
 Sample from the measurement noise distribution $\pi_{\mathbf{v}_k}$:
 $\mathbf{v}_k^i \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_k)$, for $i = 1, \dots, M$
 Compute the sample means:
 $\bar{\mathbf{x}}_k^f = \frac{1}{M} \sum_{i=1}^M \mathbf{x}_k^{f,i}$, $\bar{\mathbf{v}}_k = \frac{1}{M} \sum_{i=1}^M \mathbf{v}_k^i$, $\bar{\mathbf{z}}_k^f = \frac{1}{M} \sum_{i=1}^M \mathbf{h}(\mathbf{x}_k^{f,i})$
 and the anomalies:
 $\mathbf{X}_k^{f,i} = (\mathbf{x}_k^{f,i} - \bar{\mathbf{x}}_k^f) / \sqrt{M-1}$, $\mathbf{Z}_k^{f,i} = (\mathbf{h}(\mathbf{x}_k^{f,i}) - \mathbf{v}_k^i - \bar{\mathbf{z}}_k^f + \bar{\mathbf{v}}_k) / \sqrt{M-1}$, for $i = 1, \dots, M$
 Invert these equations for (\mathbf{b}_k^i) :
 $\mathbf{Z}^{f'} \mathbf{Z}^{f'} \mathbf{b}_k^i = \mathbf{z}_k^* + \mathbf{v}_k^i - \mathbf{h}(\mathbf{x}_k^{f,i})$, for $i = 1, \dots, M$
 Update the ensemble:
 $\mathbf{x}_k^{a,i} = \mathbf{x}_k^{f,i} + \mathbf{X}_k^{f'} \mathbf{Z}_k^{f'} \mathbf{b}_k^i$
end

Algorithm 2: Algorithm to generate mean-preserving random rotations, adapted from [42, 43]

Input: integer M
Draw $(M-1) \times (M-1)$ samples from an standard normal distribution and store them in $\mathbf{\Omega} \in \mathbb{R}^{(M-1) \times (M-1)}$:
 $\mathbf{\Omega} = \text{randn}(M-1, M-1)$
SVD decomposition of $\mathbf{\Omega}$:
 $\mathbf{\Omega} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$
 $\mathbf{\Omega} = \mathbf{V}^\top$
Compute $\mathbf{b} = \mathbf{1} / \sqrt{M}$
Compute \mathbf{b}^{sign}
 $\mathbf{b}^{sign} = \mathbf{b}$
 $\mathbf{b}^{sign}[M] += \text{sign}(\mathbf{b}[M])$ # add $\text{sign}(\mathbf{b}[M])$ to the last entry
Initialize $\mathbf{B} \in \mathbb{R}^{M \times M}$
Fill columns of \mathbf{B} :
 $\mathbf{B}[:, 1] = \mathbf{b}$
 $\mathbf{B}[:, 2 : M] = -(\mathbf{I}_M + \mathbf{b}^{sign} \mathbf{b}^{sign \top} / (||\mathbf{b}|| + 1))[:, 1 : M-1]$
Initialize $\mathbf{\Lambda}_b \in \mathbb{R}^{M \times M}$
Fill $\mathbf{\Lambda}_b$
 $\mathbf{\Lambda}_b[1, 1] = 1$
 $\mathbf{\Lambda}_b[1 : M-1, 1 : M-1] = \mathbf{\Omega}$
return $\mathbf{U} = \mathbf{B} \mathbf{\Lambda}_b \mathbf{B}^\top$

- [5] Jeff D Eldredge and Anya R Jones. Leading-edge vortices: mechanics and modeling. *Annual Review of Fluid Mechanics*, 51:75–104, 2019.
- [6] Wesley Kerstens, Jens Pfeiffer, David Williams, Rudibert King, and Tim Colonius. Closed-loop control of lift for longitudinal gust suppression at low reynolds numbers. *AIAA journal*, 49(8):1721–1728, 2011.
- [7] Gino Perrotta and Anya R Jones. Unsteady forcing on a flat-plate wing in large transverse gusts. *Exp. Fluids*, 58(8):101, 2017.
- [8] Kai Fukami, Koji Fukagata, and Kunihiko Taira. Assessment of supervised machine learning methods for fluid flows. *Theoretical and Computational Fluid Dynamics*, pages 1–23, 2020.
- [9] Wei Hou, Darwin Darakananda, and Jeff D Eldredge. Machine-learning-based detection of aerodynamic disturbances using surface pressure measurements. *AIAA Journal*, 57(12):5079–5093, 2019.
- [10] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 1960.
- [11] Mark Asch, Marc Bocquet, and Maëlle Nodet. *Data assimilation: methods, algorithms, and applications*, volume 11. SIAM, 2016.
- [12] Geir Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*, 99(C5):10143–10162, 1994.
- [13] Andre FC da Silva and Tim Colonius. Ensemble-based state estimator for aerodynamic flows. *AIAA Journal*, 56(7):2568–2578, 2018.
- [14] Andre FC da Silva and Tim Colonius. Flow state estimation in the presence of discretization errors. *Journal of Fluid Mechanics*, 890, 2020.
- [15] Kunihiko Taira and Tim Colonius. The immersed boundary method: a projection approach. *Journal of Computational Physics*, 225(2):2118–2137, 2007.

Algorithm 3: Algorithm for the ETKF, adapted from Asch et al. [11]

U is a mean-preserving rotation generated from time to time, see Algorithm 2, otherwise set to I_M

Initialize M ensemble member $(\mathbf{x}_0^{a,i}) \sim \mathcal{N}(\mathbf{x}_0, \mathbf{P}_0)$

for $k = 1 : K$ **do**

 Compute the ensemble forecast:

$$\mathbf{x}_k^{f,i} = \mathbf{f}(\mathbf{x}_{k-1}^{a,i}) + \mathbf{w}_{k-1}^i \text{ for } i = 1, \dots, M$$

 Compute the prior mean and anomaly matrix:

$$\bar{\mathbf{x}}_k^f = \frac{1}{M} \sum_{i=1}^M \mathbf{x}_k^{f,i}, \quad \mathbf{X}_k^{f,i} = (\mathbf{x}_k^{f,i} - \bar{\mathbf{x}}_k^f) / \sqrt{M-1}, \text{ for } i = 1, \dots, M$$

 Compute the observations:

$$\mathbf{Z}_k^{f,i} = \mathbf{h}(\mathbf{x}_k^{f,i}), \text{ for } i = 1, \dots, M$$

 Compute mean of the observations $\bar{\mathbf{z}}_k^f$:

$$\bar{\mathbf{z}}_k^f = \frac{1}{M} \mathbf{Z}_k^f \mathbf{1}$$

 Construct the matrices \mathbf{S} and \mathbf{G} :

$$\mathbf{S} = \mathbf{V}_k^{-1/2} (\mathbf{Z}_k^f - \bar{\mathbf{z}}_k^f \mathbf{1}^\top) / \sqrt{M-1}$$

$$\mathbf{G} = (\mathbf{I}_M + \mathbf{S} \mathbf{S}^\top)^{-1}$$

 Construct the vectors $\boldsymbol{\delta}$, \mathbf{c} :

$$\boldsymbol{\delta} = \mathbf{V}_k^{-1/2} (\mathbf{z}_k^* - \bar{\mathbf{z}}_k^f)$$

$$\mathbf{c} = \mathbf{G} \mathbf{S}^\top \boldsymbol{\delta}$$

 Update the ensemble:

$$\mathbf{X}_k^a = \bar{\mathbf{x}}_k^f \mathbf{1}^\top + \mathbf{X}_k^{f,i} (\mathbf{c} \mathbf{1}^\top + \sqrt{M-1} \mathbf{G}^{1/2} \mathbf{U})$$

end

- [16] Kunihiro Taira, Steven L Brunton, Scott TM Dawson, Clarence W Rowley, Tim Colonius, Beverley J McKeon, Oliver T Schmidt, Stanislav Gordeyev, Vassilios Theofilis, and Lawrence S Ukeiley. Modal analysis of fluid flows: An overview. *Aiaa Journal*, pages 4013–4041, 2017.
- [17] T. von Kármán and W. R. Sears. Airfoil theory for non-uniform motion. *J. Aeronautical Sci.*, 5(10):379–390, 1938.
- [18] C. E. Brown and W. H. Michael. Effect of leading edge separation on the lift of a delta wing. *J. Aero. Sci.*, 21:690–694, 1954.
- [19] Turgut Sarpkaya. An inviscid model of two-dimensional vortex shedding for transient and asymptotically steady separated flow over an inclined plate. *Journal of Fluid Mechanics*, 68(1):109–128, 1975.
- [20] Georges-Henri Cottet and Petros D. Koumoutsakos. *Vortex methods: theory and practice*. Cambridge university press, 2000.
- [21] Jeff D. Eldredge. *Mathematical Modeling of Unsteady Inviscid Flows*. Springer International Publishing, 2019.
- [22] Chengjie Wang and Jeff D Eldredge. Low-order phenomenological modeling of leading-edge vortex formation. *Theoretical and Computational Fluid Dynamics*, 27(5):577–598, 2013.
- [23] X. Xia and K. Mohseni. Unsteady aerodynamics and vortex-sheet formation of a two-dimensional airfoil. *J. Fluid Mech.*, 830:439–478, 2017.
- [24] Kiran Ramesh, Ashok Gopalarathnam, Kenneth Granlund, Michael V Ol, and Jack R Edwards. Discrete-vortex method with novel shedding criterion for unsteady aerofoil flows with intermittent leading-edge vortex shedding. *Journal of Fluid Mechanics*, 751:500–538, 2014.
- [25] Darwin Darakananda and Jeff D Eldredge. A versatile taxonomy of low-dimensional vortex models for unsteady aerodynamics. *Journal of Fluid Mechanics*, 858:917–948, 2019.
- [26] Kiran Ramesh, Kenneth Granlund, Michael V Ol, Ashok Gopalarathnam, and Jack R Edwards. Leading-edge flow criticality as a governing factor in leading-edge vortex initiation in unsteady airfoil flows. *Theoretical and Computational Fluid Dynamics*, 32(2):109–136, 2018.
- [27] Philippe R Spalart. Vortex methods for separated flows. Von Karman Institute for Fluid Dynamics, 1988.
- [28] ArunVishnu SureshBabu, Kiran Ramesh, and Ashok Gopalarathnam. Model reduction in discrete-vortex methods for unsteady airfoil flows. *AIAA Journal*, 57(4):1409–1422, 2019.
- [29] Darwin Darakananda, André Fernando de Castro da Silva, Tim Colonius, and Jeff D Eldredge. Data-assimilated low-order vortex modeling of separated flows. *Physical Review Fluids*, 3(12):124701, 2018.
- [30] Craig H Bishop, Brian J Etherton, and Sharanya J Majumdar. Adaptive sampling with the ensemble transform kalman filter. part i: Theoretical aspects. *Monthly weather review*, 129(3):420–436, 2001.
- [31] Kiran Ramesh, Ashok Gopalarathnam, Jack R Edwards, Michael V Ol, and Kenneth Granlund. An unsteady airfoil theory applied to pitching motions validated against experiment and computation. *Theoretical and Computational Fluid Dynamics*, 27(6):843–864, 2013.
- [32] Darwin Darakananda. *Vortex Models for Data Assimilation*. PhD thesis, UCLA, 2017.
- [33] Alberto Carrassi, Marc Bocquet, Laurent Bertino, and Geir Evensen. Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *Wiley Interdisciplinary Reviews: Climate Change*, 9(5):e535, 2018.
- [34] Geir Evensen. The ensemble kalman filter for combined state and parameter estimation. *IEEE Control Systems Magazine*, 29(3):83–104, 2009.

- [35] Sanita Vetra-Carvalho, Peter Jan Van Leeuwen, Lars Nerger, Alexander Barth, M Umer Altaf, Pierre Brasseur, Paul Kirchgeßner, and Jean-Marie Beckers. State-of-the-art stochastic data assimilation methods for high-dimensional non-gaussian problems. *Tellus A: Dynamic Meteorology and Oceanography*, 70(1):1–43, 2018.
- [36] Alessio Spantini, Ricardo Baptista, and Youssef Marzouk. Coupling techniques for nonlinear ensemble filtering. *arXiv preprint arXiv:1907.00389*, 2019.
- [37] Jan Mandel, Loren Cobb, and Jonathan D Beezley. On the convergence of the ensemble kalman filter. *Applications of Mathematics*, 56(6):533–541, 2011.
- [38] Patrick N Raanes. *Improvements to ensemble methods for data assimilation in the geosciences*. PhD thesis, University of Oxford, 2015.
- [39] Gerrit Burgers, Peter Jan van Leeuwen, and Geir Evensen. Analysis scheme in the ensemble kalman filter. *Monthly weather review*, 126(6):1719–1724, 1998.
- [40] David M Livings, Sarah L Dance, and Nancy K Nichols. Unbiased ensemble square root filters. *Physica D: Nonlinear Phenomena*, 237(8):1021–1028, 2008.
- [41] Pavel Sakov and Peter R Oke. A deterministic formulation of the ensemble kalman filter: an alternative to ensemble square root filters. *Tellus A: Dynamic Meteorology and Oceanography*, 60(2):361–371, 2008.
- [42] Lars Nerger, Tijana Janjić, Jens Schröter, and Wolfgang Hiller. A unification of ensemble square root kalman filters. *Monthly Weather Review*, 140(7):2335–2345, 2012.
- [43] Julian Tödter and Bodo Ahrens. A second-order exact ensemble square root filter for nonlinear data assimilation. *Monthly Weather Review*, 143(4):1347–1367, 2015.
- [44] Matthias Katzfuss, Jonathan R Stroud, and Christopher K Wikle. Understanding the ensemble kalman filter. *The American Statistician*, 70(4):350–357, 2016.
- [45] Jeffrey S Whitaker and Thomas M Hamill. Evaluating methods to account for system errors in ensemble data assimilation. *Monthly Weather Review*, 140(9):3078–3089, 2012.
- [46] Jeffrey L Anderson and Stephen L Anderson. A monte carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127(12):2741–2758, 1999.
- [47] Grigorios A Pavliotis. *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*. Springer, 2014.
- [48] A. Leonard. Vortex methods for flow simulation. *J. Comput. Phys.*, 37(3):289–335, 1980.
- [49] Alexandre Joel Chorin. Numerical study of slightly viscous flow. *Journal of fluid mechanics*, 57(4):785–796, 1973.
- [50] Ding-Gwo Long. Convergence of the random vortex method in two dimensions. *Journal of the American Mathematical Society*, 1(4):779–804, 1988.
- [51] Geir Evensen. The Ensemble Kalman Filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343–367, 2003.
- [52] Kody Law, Andrew Stuart, and Konstantinos Zygalakis. *Discrete Time: Filtering Algorithms*, pages 79–114. Springer International Publishing, Cham, 2015.
- [53] Sebastian Liska and Tim Colonius. A fast immersed boundary method for external incompressible viscous flows using lattice green’s functions. *Journal of Computational Physics*, 331:257–279, 2017.
- [54] Peter L Houtekamer and Herschel L Mitchell. Data assimilation using an ensemble kalman filter technique. *Monthly Weather Review*, 126(3):796–811, 1998.
- [55] Donald Leslie McLeish. A robust alternative to the normal distribution. *The Canadian Journal of Statistics/La Revue Canadienne de Statistique*, pages 89–102, 1982.