Conditioning of hybrid variational data assimilation

Article

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: http://dx.doi.org/10.1002/nla.2534

www.reading.ac.uk/centaur

CentAUR

RESEARCH ARTICLE

WILEY

Conditioning of hybrid variational data assimilation

Shaerdan Shataer[1] | **Amos S. Lawless[1,2]** | **Nancy K. Nichols[1,2]**

[1]School of Mathematical, Physical and Computational Sciences, University of Reading, Reading, UK

[2]National Centre for Earth Observation, Reading, UK

Correspondence
Shaerdan Shataer, School of Mathematical, Physical and Computational Sciences, University of Reading, Reading, UK.
Email: s.shataer@reading.ac.uk

Abstract

In variational assimilation, the most probable state of a dynamical system under Gaussian assumptions for the prior and likelihood can be found by solving a least-squares minimization problem. In recent years, we have seen the popularity of hybrid variational data assimilation methods for Numerical Weather Prediction. In these methods, the prior error covariance matrix is a weighted sum of a climatological part and a flow-dependent ensemble part, the latter being rank deficient. The nonlinear least squares problem of variational data assimilation is solved using iterative numerical methods, and the condition number of the Hessian is a good proxy for the convergence behavior of such methods. In this article, we study the conditioning of the least squares problem in a hybrid four-dimensional variational data assimilation (Hybrid 4D-Var) scheme by establishing bounds on the condition number of the Hessian. In particular, we consider the effect of the ensemble component of the prior covariance on the conditioning of the system. Numerical experiments show that the bounds obtained can be useful in predicting the behavior of the true condition number and the convergence speed of an iterative algorithm

KEYWORDS

condition number, data assimilation, Hessian matrix, Hybrid 4D-Var, spectral theory

1 | INTRODUCTION

In weather forecasting, we use mathematical and numerical models to describe the system dynamics of the ocean and atmosphere. These models are highly nonlinear and often sensitive to noise in initial conditions. Because of the nonlinearity and instability, random perturbations in the initial data and errors in the model amplify rapidly through time, producing unreliable predictions.[1,2] In variational data assimilation, the goal is to find the maximum Bayesian a posteriori estimate of the system state from which to initialize the model. Major operational centers worldwide have adopted the four-dimensional variational assimilation scheme (4D-Var) for environmental forecasting in recent years. Similar applications arise in other fields such as physics, biology, and economics.[1,3,4]

In 4D-Var, we aim to obtain an optimal initial state variable (conventionally called the "analysis") by solving a nonlinear least squares problem, in which we try to find the best fit between a set of observations over a time window and an a priori estimate of the state at the start of the window, known as the background. In the case where observations are only given at one time, the method becomes three-dimensional variational assimilation, or 3D-Var. We assume that

the background state and observations have Gaussian, unbiased errors, with covariance matrices B and R respectively. Traditional variational data assimilation methods have used a climatological estimate[4] for the background error covariance matrix B, where flow-dependent ensemble information is not incorporated into the system.[5,6] Recent developments utilize a hybrid approach, in which an ensemble background error covariance matrix is estimated with ensemble members and then combined with the climatological part. This has a clear advantage of bringing in the variability of the system and updates the statistics in each prediction window, giving flow-dependent information that can improve the accuracy of predictions.[4] However, due to computational restrictions, the affordable number of ensemble members is normally small, which leads to a rank deficient ensemble background error covariance matrix. The method of combining the ensemble parts with the conventional 4D-Var is called Hybrid 4D-Var. In this method, B is given by $B = (1 - \beta)B_0 + \beta P_f$. Here, B_0, P_f are the climatological background error covariance and the ensemble error covariance matrix, and $(1 - \beta)$, β are their weights, where β is a scalar. For the details of the hybrid method, we refer readers to the review paper of Bannister.[4]

In this article, we are especially interested in establishing the relationship between the conditioning of Hybrid 4D-Var and the weight β on P_f. The 4D-Var problem is usually solved using iterative gradient methods, such as conjugate gradient or Quasi-Newton methods.[1,4] The condition number of the Hessian can be used to estimate the number of iterations required for convergence.[7,8] In addition, the condition number also reveals the sensitivity of the minimization problem with respect to random noise.[7] Here we establish that, in Hybrid 4D-Var, a transition point exists where the condition number of B sharply increases with the weight on P_f. Since P_f is rank deficient,[1,4-6] adding the ensemble background error covariance matrix may cause difficulties for solving the nonlinear least-squares minimization problem[1,4-6] and an adequate preconditioning scheme becomes desirable. This is typically achieved through a Control Variable Transformation (CVT).[1,4] CVT uses a decomposition of B_0 and P_f to transform the state variable such that the conditioning of the Hessian matrix is improved.[1,4] Implementation details of CVT are frequently described in previous works on 4D-Var[3,5,6,9] and Hybrid 4D-Var.[4] In terms of practical applications of Hybrid 4D-Var, we note that the nonlinear least-squares problem is often linearized and solved as a sequence of linear least-squares minimizations. This is known as the incremental 4D-Var method[1,3,4] and is equivalent to an approximate Gauss–Newton method.[10]

In practice, it is useful to understand the contribution of each component of DA, in such a way that the impact on the conditioning can be predicted when these components change. This then motivates a comprehensive theory that can predict the conditioning while separating the contribution of each error component (B_0, P_f, R), and relating it to the parameters that characterize these components. We note that both preconditioned and unpreconditioned 4D-Var are implemented by major operational centers (such as the UK Met Office). Such theories have previously been constructed for conventional 4D-Var. Haben et al.[3] established a theory to estimate the conditioning of a preconditioned 4D-Var system. This work is later developed by Tabeart et al.[8,9] for both unpreconditioned 3D-Var and preconditioned 4D-Var with Control Variable Transformation (CVT). In Tabeart et al.'s studies, a bound estimation is proved for the conditioning of the system, in which the contributions of the background error covariance matrix and observation error covariance matrix are separated. The impact of each component is then associated with its characterizing parameters. The research of Haben et al.[3] and Tabeart et al.[8,9] are tested and analyzed using small scale examples. However, they are extrapolated to justify observations in large scale real-life applications. To give a few examples, Mercier et al.[11] used the analysis of Haben et al.[3] to explain the convergence behavior of a block Krylov method for a 3D-Var application;[11] Desroziers et al.[12] cited the same result of Haben et al.[3] to guide their design of a preconditioned Lanczos/Conjugate Gradient algorithm. Hatfield et al.[13] cited Haben et al.'s analysis[3] to explain the effect of an increasing model error on the convergence of an incremental 4D-Var; Aabaribaoune et al.[14] used the result of Tabeart et al.[8] to analyze the convergence speed of a BFGS algorithm for solving a 3D-Var system, in the application of ozone profiling.

In this article, we aim to extend previous studies of Haben et al.[3] and Tabeart et al.[8,9] to a hybrid system. In particular, we study the impact of P_f on the condition number of the Hessian matrix, for both unpreconditioned cases and preconditioned cases with CVT.

We outline this article as follows. In Section 2, we briefly formulate the problem of Hybrid 4D-Var and introduce the Hessian matrix; in Section 3, we establish the theory of conditioning for unpreconditioned Hybrid 4D-Var and preconditioned Hybrid 4D-Var with CVT; in Sections 4 to 6, we provide multiple numerical experiments to illustrate the theories and analyze the conditioning; in Section 7, we show how the behavior of the condition number of the Hessian predicted by our theory is reflected in the convergence speed of a conjugate gradient algorithm. Section 8 gives a general summary of the results.

2 | PROBLEM FORMULATION

In this section, we introduce the Hybrid 4D-Var method, the incremental method, the preconditioning technique of Control Variable Transformation, and the Hessian matrix associated with the unpreconditioned and preconditioned least-squares problems.

2.1 | A general formulation of the Hybrid 4D-Var

A general formulation of the 4D-Var is given by Problem 1.

Problem 1. Solve for the optimal initial state x_a by solving a minimization problem given by,

$$\begin{cases} x_a = \arg\min_{x_0 \in \mathbb{R}^n} \mathcal{J}(x_0), \\ \mathcal{J}(x_0) := \frac{1}{2}||x_0 - x^b||^2_{B^{-1}} + \frac{1}{2}\sum_{i=0}^{N}||y_i - \mathcal{H}_i(x_i)||^2_{R_i^{-1}}, \\ x_i \in \mathbb{R}^n, \ y_i \in \mathbb{R}^p, \ B \in \mathbb{R}^{n,n}, \ R_i \in \mathbb{R}^{p,p}, \ \mathcal{H}_i : \mathbb{R}^n \to \mathbb{R}^p, \end{cases} \tag{1}$$

where $\mathcal{H}_i$ is the nonlinear observational operator; y_i is the vector of observational measurements, taken at time t_i; x_i is the state variable at time t_i, given by,

$$x_i = \mathcal{M}_{i,i-1}(x_{i-1}) = \mathcal{M}_{i,0}(x_0), \tag{2}$$

$\mathcal{M}_{i,i-1} : \mathbb{R}^n \to \mathbb{R}^n$ is the nonlinear model, and $\mathcal{M}_{i,0} = \mathcal{M}_{i,i-1} \cdot \mathcal{M}_{i-1,i-2} \cdots \mathcal{M}_{1,0}$ is a direct product (composition) of them. The vector $x^b \in \mathbb{R}^n$ is the prior information given by the model at t_0, known as the background state. A simplified problem, given by 3D-Var, is a special case of the general 4D-Var problem where $N = 0$, meaning there are only observations at time t_0.

In Hybrid 4D-Var and 3D-Var, we follow the formulation given by (1), but replace the background error covariance matrix B with $B = (1 - \beta)B_0 + \beta P_f$. We note that the problem involves solving a nonlinear least squares problem, and it is difficult to implement directly when the problem is large scale. Instead, this is often replaced with a linearized incremental formulation, which we describe next.

2.2 | Incremental Hybrid 4D-Var, control variable transformation and the Hessian matrix

In practice, especially in NWP, Problem 1 is often solved using the incremental method. In incremental 4D-Var, the nonlinear least squares problem is replaced with a sequence of linear least squares problems with a cost function of

$$\mathcal{J}(\delta x_0^k) = \frac{1}{2}||\delta x_0^k - \delta x_b^k||^2_{B^{-1}} + \frac{1}{2}\sum_{i=0}^{N}||d_i^k - H_i\delta x_i^k||_{R_i^{-1}}, \tag{3}$$

where the vector d_i^k is known as the innovation, defined by $d_i^k = y_i - \mathcal{H}_i(x_i^k)$; the linear operator H_i is the Jacobian of $\mathcal{H}_i$, and the vector δx_b^k is the increment, given by $\delta x_b^k = x_0^k - x_b^k$; the vector δx_i^k is computed from $\delta x_i^k = M_{i,0}\delta x_0^k$. The model operator $M_{i,0}$ is given by $M_{i,0} = \prod_{j=1}^{i} M_j$, where M_j is the Jacobian of $\mathcal{M}_{j,j-1}$. At each outer iteration k (the outer loop), we minimize (3) to solve for the increment, and then x_0^k is updated for the next iteration.

We now introduce the Hessian matrix of the cost function $\mathcal{J}$ in Problem 1, which is given by,[3,9]

$$S_{4D} = \left((1 - \beta)B_0 + \beta P_f\right)^{-1} + \sum_{i=0}^{N}(H_i M_{i,0})^T R_i^{-1} H_i M_{i,0}.$$

The sum of the matrix product above can be written in a simple compact form.

Notation 1. The matrix $\hat{H}$ is the general observation operator, given by $\hat{H} := \left[H_0^T, (H_1 M_{1,0})^T, \ldots, (H_N M_{N,0})^T\right]^T \in \mathbb{R}^{p(N+1),n}$, where $\hat{R}$ is a block diagonal matrix with its ith diagonal block given by R_i.

Following Notation 1, we can rewrite the Hessian matrix as follows,[3,9]

$$S_{4D} = ((1-\beta)B_0 + \beta P_f)^{-1} + \hat{H}^T \hat{R}^{-1} \hat{H}. \tag{4}$$

In the case of Hybrid 3D-Var, we recall that it is a special case of 4D-Var with $N = 0$ and its Hessian is given by

$$S_{3D} = ((1-\beta)B_0 + \beta P_f)^{-1} + H_0^T R_0^{-1} H_0. \tag{5}$$

In applications, preconditioning techniques are often applied to improve the conditioning of the system. Control Variable Transformation (CVT) is one of the popular preconditioning techniques. The detail of CVT is described by Nichols,[1] Bannister,[4] and Buehner.[15] In this approach, we utilize factorizations of B_0 and P_f, given by

$$B_0 = U^T U, \; P_f = X_f^T X_f, \; \text{where } U = B_0^{1/2} \text{ and } X_f = \frac{1}{\sqrt{m-1}}\left[\boldsymbol{x}_1 - \bar{\boldsymbol{x}}, \boldsymbol{x}_2 - \bar{\boldsymbol{x}}, \ldots, \boldsymbol{x}_m - \bar{\boldsymbol{x}}\right], \tag{6}$$

and $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_m$ are ensemble members, m is the number of samples and $\bar{\boldsymbol{x}}$ is the ensemble mean. These matrices are then used to transform the state variable as follows,[4]

$$U_h \delta \boldsymbol{v} = \delta \boldsymbol{x}, \; \text{where } U_h = \left[\sqrt{(1-\beta)}U \; \sqrt{\beta}X_f\right] \in \mathbb{R}^{n,n+m}. \tag{7}$$

Here $\delta \boldsymbol{x} \in \mathbb{R}^n$ is the increment of the state variable and $\delta \boldsymbol{v}$ is the increment of the control variable. Applying this transform to (3) leads to a new cost function that reads

$$J\left(\delta \boldsymbol{v}_0^k\right) = \frac{1}{2}\left\|\delta \boldsymbol{v}_0^k - \delta \boldsymbol{v}_b^k\right\|_{I_{n+m}} + \frac{1}{2}\sum_{i=0}^{N}\left\|y_i - \mathcal{H}_i\left(\boldsymbol{x}^k\right) - H_i M_{i,0} U_h \delta \boldsymbol{v}_i^k\right\|_{R_i^{-1}}. \tag{8}$$

A direct calculation then yields the Hessian of $J\left(\delta \boldsymbol{v}_0^k\right)$ with respect to $\delta \boldsymbol{v}_0^k$ as

$$S_{P4D} = I_{n+m} + U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h. \tag{9}$$

In the following sections, we will demonstrate that CVT prevents the condition number from going to infinity when β (the weight of the ensemble part) approaches 1. In addition, we also note that the adjoint of U_h does not need to be computed explicitly. This is discussed in detail by Smith et al.[5] and Bannister.[4]

3 | THEORY OF THE CONDITIONING OF HYBRID 4D-VAR

To better analyze the convergence of the nonlinear least squares problem of Hybrid 4D-Var, we aim to establish a set of theories that can predict changes in the conditioning of the system prompted by varying parameters (such as correlation length scale, error variance, etc). Typically, we are also keen to understand the impact of the rank deficient ensemble part on the conditioning. These motivate us to develop an estimation of the condition number that is informative of the actual conditioning of the system. In order to achieve such a goal, we use spectral theories to establish bounds for the condition number of the Hessian matrix.

We outline the structure of this section as follows. We start by introducing some pre-established results, then extend them to the Hybrid method. We will discuss the unpreconditioned cases and the preconditioned cases separately.

3.1 | Eigenvalues and conditioning

We begin with a brief review of previous work by Tabeart et al.[8] and some fundamental eigenvalue inequalities.[16] We will extend these results to Hybrid 4D-Var. In the scope of this article, we always assume that B and R are symmetric.

Notation 2. Let $\lambda_k(A)$ be the kth largest eigenvalue of a matrix $A \in \mathbb{R}^{n,n}$, $\lambda_1(A)$, $\lambda_n(A)$ be the largest and smallest eigenvalues of A, and $\kappa(A)$ be the condition number of A.

Definition 1. For a symmetric positive definite matrix $A \in \mathbb{R}^{n,n}$, its condition number is defined by $\kappa(A) = \lambda_1(A)/\lambda_n(A)$.

For the eigenvalues of the sum of two Hermitian matrices, Weyl[17] proved the following theorem:

Theorem 1. *Let A_1, A_2 be two symmetric matrices. Then the eigenvalues of $A = A_1 + A_2$ satisfy the following*:

$$\lambda_k(A_1) + \lambda_n(A_2) \leq \lambda_k(A) \leq \lambda_k(A_1) + \lambda_1(A_2). \tag{10}$$

The inequality for products is given by Wang and Zhang as follows,[18]

Theorem 2. *Let $A_1, A_2 \in R^{n,n}$ be positive semidefinite Hermitian matrices. Then*

$$\max\left[\lambda_1(A_1)\lambda_n(A_2), \lambda_1(A_n)\lambda_1(A_2)\right] \leq \lambda_1(A_1 A_2) \leq \lambda_1(A_1)\lambda_1(A_2). \tag{11}$$

Applying these results to the Hessian matrix of 3D-Var, Tabeart et al.[8] established that,

Theorem 3. *Let $S_{3D} = B^{-1} + H_0^T R_0^{-1} H_0$ and given that B, R are symmetric positive definite, the condition number of S_{3D} is bounded as follows,*

$$\max\left[\frac{\kappa(B)}{1 + \lambda_1(B)\lambda_1(H_0^T R_0^{-1} H_0)}, \frac{1 + \lambda_1(B)\lambda_1(H_0^T R_0^{-1} H_0)}{\kappa(B)}\right] \leq \kappa(S_{3D}) \leq (1 + \lambda_n(B)\lambda_1(H_0^T R_0^{-1} H_0))\kappa(B). \tag{12}$$

3.2 | Conditioning of unpreconditioned Hybrid 4D-Var

In Hybrid 4D-Var, the background error covariance matrix B is replaced by a weighted sum of the static and ensemble background error covariance matrix, such that $B = (1 - \beta)B_0 + \beta P_f$. Although, we can still use (12) to estimate the condition number of the Hessian, we note that it does not have the capacity to separate the contribution of P_f and predict the condition number as the weight of P_f grows. This motivates us to establish theories that are specific to the hybrid case of 4D-Var.

Lemma 1. *Assume that P_f is rank deficient and let $B = (1 - \beta)B_0 + \beta P_f$. Then the extreme eigenvalues of B satisfy the following*:

$$\begin{cases} \max\left[(1 - \beta)\lambda_1(B_0), \ \beta\lambda_1(P_f) + (1 - \beta)\lambda_n(B_0)\right] \leq \lambda_1(B) \leq (1 - \beta)\lambda_1(B_0) + \beta\lambda_1(P_f), \\ (1 - \beta)\lambda_n(B_0) \leq \lambda_n(B) \leq \min\left[(1 - \beta)\lambda_1(B_0), \ \beta\lambda_1(P_f) + (1 - \beta)\lambda_n(B_0)\right]. \end{cases} \tag{13}$$

Proof. The conclusion follows from Theorem 1 and that $\lambda_n(P_f) = 0$ (since P_f is rank deficient). ∎

Lemma 2. *Given B_0, P_f are two symmetric matrices, let $B = (1 - \beta)B_0 + \beta P_f$, and assuming that P_f is rank deficient, the condition number of B is then bounded as follows,*

$$\max\left[\frac{1}{\kappa(B_0)} + \frac{\beta\lambda_1(P_f)}{(1 - \beta)\lambda_1(B_0)}, \left(\frac{1}{\kappa(B_0)} + \frac{\beta\lambda_1(P_f)}{(1 - \beta)\lambda_1(B_0)}\right)^{-1}\right] \leq \kappa(B) \leq \kappa(B_0)\left(1 + \frac{\beta\lambda_1(P_f)}{(1 - \beta)\lambda_1(B_0)}\right). \tag{14}$$

Proof. By the definition of $\kappa(B)$ and Lemma 1, we find that,

$$\kappa(B) = \frac{\lambda_1(B)}{\lambda_n(B)} \leq \frac{(1 - \beta)\lambda_1(B_0) + \beta\lambda_1(P_f)}{(1 - \beta)\lambda_n(B_0)} = \kappa(B_0) + \frac{\beta\lambda_1(P_f)}{(1 - \beta)\lambda_n(B_0)},$$

and similarly,

$$\kappa(B) \geq \max \left[\frac{(1-\beta)\lambda_n(B_0) + \beta\lambda_1\left(P_f\right)}{(1-\beta)\lambda_1(B_0)}, \frac{(1-\beta)\lambda_1(B_0)}{(1-\beta)\lambda_n(B_0) + \beta\lambda_1\left(P_f\right)} \right]$$

$$= \max \left[\frac{1}{\kappa(B_0)} + \frac{\beta\lambda_1\left(P_f\right)}{(1-\beta)\lambda_1(B_0)}, \left(\frac{1}{\kappa(B_0)} + \frac{\beta\lambda_1\left(P_f\right)}{(1-\beta)\lambda_1(B_0)} \right)^{-1} \right].$$

∎

We highlight that the condition number of B diverges to infinity as $\beta \to 1$. This is reflected by the lemma above as the lower bound of $\kappa(B)$ goes to infinity.

We now use these results to obtain a bound on the condition number in Theorem 4. We emphasize that the conclusion of these results cannot be directly derived from Theorem 3 by simply replacing B with $(1-\beta)B_0 + \beta P_f$ in the upper and lower bounds. In fact, such an effort leads to a less sharp bound that does not provide analytical or theoretical insight on how the bound changes with β. Moreover, this approach does not reveal how the bound responds to changes in physical parameters related to B_0 and P_f. Using different strategies leads to a bound that clearly separates the contribution of each component (B_0, P_f) and their weights, making it possible to analyze their impact on the condition number and their interaction as well.

Notation 3. For simplicity of presentation, we introduce γ_z, Γ_z to represent the lower and upper bound of any $z \in \mathbb{R}$.

Theorem 4. *Recall that $S_{4D} = B^{-1} + \hat{H}^T\hat{R}^{-1}\hat{H}$, $B = (1-\beta)B_0 + \beta P_f$, and let,*

$$\Gamma_{\lambda_n(B)} = \min((1-\beta)\lambda_n(B_0) + \beta\lambda_1\left(P_f\right), (1-\beta)\lambda_1(B_0)), \quad \gamma_{\kappa(B)}$$

$$= \max \left[\frac{1}{\kappa(B_0)} + \frac{\beta\lambda_1\left(P_f\right)}{(1-\beta)\lambda_1(B_0)}, \left(\frac{1}{\kappa(B_0)} + \frac{\beta\lambda_1\left(P_f\right)}{(1-\beta)\lambda_1(B_0)} \right)^{-1} \right],$$

$$\Gamma_{\kappa(B)} = \kappa(B_0)\left(1 + \frac{\beta\lambda_1\left(P_f\right)}{(1-\beta)\lambda_1(B_0)} \right).$$

Given that B, R are symmetric, the condition number of S_{4D} is then bounded as follows,

$$\begin{cases} \kappa(S_{4D}) \geq \max \left[\frac{1}{\Gamma_{\kappa(B)}} + (1-\beta)\lambda_n(B_0)\lambda_1\left(\hat{H}^T\hat{R}^{-1}\hat{H}\right), \left(\frac{1}{\gamma_{\kappa(B)}} + \Gamma_{\lambda_n(B)}\lambda_1\left(\hat{H}^T\hat{R}^{-1}\hat{H}\right) \right)^{-1}, 1 \right] \\ \kappa(S_{4D}) \leq \Gamma_{\kappa(B)} + \left((1-\beta)\lambda_1(B_0) + \beta\lambda_1\left(P_f\right) \right)\lambda_1\left(\hat{H}^T\hat{R}^{-1}\hat{H}\right). \end{cases} \quad (15)$$

Proof. In Theorem 3, replacing $H_0^T R^{-1} H_0$ with $\hat{H}\hat{R}^{-1}\hat{H}$, then the left hand side of Inequality (12) can be written

$$\max \left[\frac{1}{\kappa(B)} + \lambda_n(B)\lambda_1\left(\hat{H}^T\hat{R}^{-1}\hat{H}\right), \left(\frac{1}{\kappa(B)} + \lambda_n(B)\lambda_1\left(\hat{H}^T\hat{R}^{-1}\hat{H}\right) \right)^{-1} \right] \leq \kappa(S_{4D}).$$

For the upper bound, we note that the upper bound can be written as follows,

$$\kappa(S_{4D}) \leq \kappa(B) + \lambda_1(B)\lambda_1(\hat{H}^T\hat{R}^{-1}\hat{H}).$$

Substituting $\kappa(B)$ and $\lambda_1(B)$ with their upper bounds, this then produces the upper bound on $\kappa(S_{4D})$. ∎

We note that the same bound can be derived by using eigenvalue inequalities. Crucially, the upper bound in Theorem 4 does not require any explicit computation of S_{4D}. We only need to compute the largest eigenvalues of $\lambda_1(B_0)$, $\lambda_1\left(P_f\right)$ and the condition number of B_0.

In the case of a special observation operator we can simplify the Hessian in the 3D-Var case as follows.

Corollary 1. *Assuming that each row of H_0 has one unit entry, with all other entries being zero, and that R_0 is diagonal such that $R_0 = \sigma_{R_0}^2 I_p$, where $I_p \in \mathbb{R}^{p,p}$ is the identity matrix, then the bounds in Theorem 4 in the case*

of 3D-Var simplify to:

$$\max\left[\frac{1}{\Gamma_{\kappa(B)}} + \frac{(1-\beta)\lambda_n(B_0)}{\sigma_{R_0}^2}, \left(\frac{1}{\gamma_{\kappa(B)}} + \frac{\Gamma_{\lambda_n(B)}}{\sigma_{R_0}^2}\right)^{-1}\right] \leq \kappa(S_{3D}) \leq \Gamma_{\kappa(B)} + \frac{(1-\beta)\lambda_1(B_0) + \beta\lambda_1(P_f)}{\sigma_{R_0}^2}. \quad (16)$$

Proof. Given that $R_0 = \sigma_{R_0}^2 I_p$, and with the assumption on H, the largest eigenvalue of $H_0^T R_0^{-1} H_0$ is always $\lambda_1(H_0^T R_0^{-1} H_0) = \sigma_{R_0}^{-2}$. Replacing relevant terms in Bounds (15), we then reach the conclusion. ∎

We emphasize that the upper bound on $\kappa(S_{4D})$ diverges to infinity as $\beta \to 1$. In such a case, the background error covariance matrix is dominated by P_f. However, the theory does not predict the same behavior for the lower bound. In terms of of the observation part, we observe that the number p of observation points does not alter the bounds; thus the theory cannot predict the behavior of $\kappa(S_{4D})$ when p varies.

Although unpreconditioned 4D-Var is still in use for real-life applications,[8] it is now a common practice to implement CVT for Hybrid 4D-Var. In the next section, we focus on developing similar theories for such cases.

3.3 | Conditioning of preconditioned Hybrid 4D-Var method with CVT

The preconditioning of Hybrid 4D-Var is broadly adopted in major operational centers, and the impact of different error components on the conditioning is important to determine the convergence of numerical schemes for incremental Hybrid 4D-Var, or more generally, for solving the nonlinear least squares problem in Hybrid 4D-Var. In this section, we prove two versions of the bounds for $\kappa(S_{4D})$ with different strategies. First, we remind readers some useful notation as follows: let $K := \hat{H}^T \hat{R} \hat{H}$, $U := B_0^{1/2}$ and $X_f := \frac{1}{\sqrt{m-1}}[x_1 - \overline{x}, x_2 - \overline{x}, \dots, x_m - \overline{x}]$, where $x_1, x_2, \dots, x_m$ is the set of ensemble members and $\overline{x}$ is the ensemble mean.

We recall that the motivation behind our approach to deriving the bound is to separate the contribution of each component in the Hessian, particularly the components of B_0 and P_f and their weights. We are especially keen to find out how the ratio of $(1-\beta)||B_0||_2/(\beta||P_f||_2)$ influences the bound, as this ratio is linked to the balancing of the climatological part and the ensemble part. Driven by this motivation, we discovered a decomposition of the matrix product $U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h$ that yields a useful max function which is directly related to this ratio. The detail is given by the proof of Theorem 5.

Theorem 5. *The condition number of the Hessian matrix of the Hybrid 4D-Var method satisfies*

$$\begin{cases} \kappa(S_{P4D}) \leq 1 + \sqrt{(\beta - \beta^2)\lambda_1(B_0)\lambda_1(P_f)\lambda_1(K^2)} + \max\left[(1-\beta)\lambda_1(B_0)\lambda_1(K), \beta\lambda_1(P_f)\lambda_1(K)\right], \\ \kappa(S_{P4D}) \geq 1 + \max\left[(1-\beta)\lambda_1(K)\lambda_n(B_0), \sqrt{(\beta - \beta^2)\lambda_n(B_0)\lambda_1(P_f)\lambda_1(K^2)}\right]. \end{cases} \quad (17)$$

Proof. We know that the Hessian matrix of the Hybrid 4D-Var can be expressed as follows,

$$\begin{cases} S_{P4D} = I_{n+m} + U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h, \\ U_h = \left[\sqrt{1-\beta}U \quad \sqrt{\beta}X_f\right]. \end{cases} \quad (18)$$

Given that $U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h$ is rank deficient, it is immediate that

$$\kappa(S_{P4D}) = 1 + \lambda_1\left(U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h\right). \quad (19)$$

Substituting U_h in S_{P4D}, we derive

$$S_{P4D} = I_{n+m} + \begin{pmatrix} \sqrt{1-\beta}U^T \\ \sqrt{\beta}X_f^T \end{pmatrix} K \begin{pmatrix} \sqrt{1-\beta}U & \sqrt{\beta}X_f \end{pmatrix} = I_{n+m} + \begin{pmatrix} (1-\beta)U^T KU & \sqrt{\beta-\beta^2}U^T KX_f \\ \sqrt{\beta-\beta^2}X_f^T KU & \beta X_f^T KX_f \end{pmatrix} \quad (20)$$

$$= I_{n+m} + \begin{pmatrix} (1-\beta)U^T KU & 0 \\ 0 & \beta X_f^T KX_f \end{pmatrix} + \begin{pmatrix} 0 & \sqrt{\beta-\beta^2}U^T KX_f \\ \sqrt{\beta-\beta^2}X_f^T KU & 0 \end{pmatrix} := I_{n+m} + A_1 + A_2. \quad (21)$$

Since $U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h = A_1 + A_2$, using Theorem 1, we get

$$\max\left[\lambda_1(A_1) + \lambda_{n+m}(A_2), \ \lambda_{n+m}(A_1) + \lambda_1(A_2)\right] \le \lambda_1(U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h) \le \lambda_1(A_1) + \lambda_1(A_2). \tag{22}$$

We note that

$$\lambda_1(U^T K U) = \lambda_1(U^T U K) = \lambda_1(B_0 K), \ \lambda_1(X_f^T K X_f) = \lambda_1(X_f^T X_f K) = \lambda_1(P_f K), \ \lambda_1(U^T K X_f X_f^T K U) = \lambda_1(B_0 P_f K^2). \tag{23}$$

Then, applying Theorem 2 to the matrix products, we derive that

$$\lambda_1(A_1) = \max\left[(1 - \beta)\lambda_1(U^T K U), \beta\lambda_1(X_f^T K X_f)\right] \le \max\left[(1 - \beta)\lambda_1(B_0)\lambda_1(K), \ \beta\lambda_1(P_f)\lambda_1(K)\right], \tag{24}$$

$$\lambda_1(A_2) = \sqrt{\lambda_1\left[(\beta - \beta^2)U^T K X_f X_f^T K U\right]} \le \sqrt{(\beta - \beta^2)\lambda_1(B_0)\lambda_1(P_f)\lambda_1(K^2)}. \tag{25}$$

So

$$\kappa(S_{P4D}) \le 1 + \max\left[(1 - \beta)\lambda_1(B_0)\lambda_1(K), \ \beta\lambda_1(P_f)\lambda_1(K)\right] + \sqrt{(\beta - \beta^2)\lambda_1(B_0)\lambda_1(P_f)\lambda_1(K^2)}. \tag{26}$$

In the lower bound, we use that

$$\max\left[\lambda_1(A_1), \lambda_1(A_2)\right] \le \max\left[\lambda_1(A_1) + \lambda_{n+m}(A_2), \ \lambda_{n+m}(A_1) + \lambda_1(A_2)\right] \le \lambda_1(U_h^T \hat{H}^T \hat{R}^{-1} \hat{H} U_h). \tag{27}$$

Applying Theorem 2 then yields (we note that K is rank deficient, such that $\lambda_n(K) = 0$)

$$\lambda_1(U^T K U) \ge \max\left[\lambda_n(B_0)\lambda_1(K), \lambda_1(B_0)\lambda_n(K)\right] = \lambda_n(B_0)\lambda_1(K), \tag{28}$$

$$\lambda_1(X_f^T K X_f) \ge \max\left[\lambda_1(P_f)\lambda_n(K), \ \lambda_n(P_f)\lambda_1(K)\right] = 0, \tag{29}$$

$$\lambda_1(U^T K X_f X_f^T K U) \ge \max\left[\lambda_1(B_0)\lambda_1(P_f)\lambda_n(K^2), \lambda_1(B_0)\lambda_n(P_f)\lambda_1(K^2), \lambda_n(B_0)\lambda_1(P_f)\lambda_1(K^2)\right] = \lambda_n(B_0)\lambda_1(P_f)\lambda_1(K^2) \tag{30}$$

and giving us

$$\lambda_1(A_1) \ge (1 - \beta)\lambda_1(K)\lambda_n(B_0), \ \lambda_1(A_2) \ge \sqrt{(\beta - \beta^2)\lambda_n(B_0)\lambda_1(P_f)\lambda_1(K^2)}. \tag{31}$$

By substituting $\max\left[\lambda_1(A_1), \lambda_1(A_2)\right]$ into (27) we conclude that

$$\kappa(S_{P4D}) \ge 1 + \max\left[(1 - \beta)\lambda_1(K)\lambda_n(B_0), \sqrt{(\beta - \beta^2)\lambda_n(B_0)\lambda_1(P_f)\lambda_1(K^2)}\right]. \tag{32}$$

∎

Another version of the bounds can be derived by directly applying Theorem 2 to the product of $U_h K U_h$ instead of using the decomposition such as that in Theorem 5.

Theorem 6. *Let S_{P4D} be the Hessian matrix of Hybrid 4D-Var with CVT, then the condition number of S_{P4D} satisfies*

$$1 \le \kappa(S_{P4D}) \le 1 + \left[(1 - \beta)\lambda_1(B_0) + \beta\lambda_1(P_f)\right]\lambda_1(K). \tag{33}$$

Proof. We note that

$$\lambda \left(U_h K U_h^T \right) = \lambda \left(U_h^T U_h K \right) = \lambda \left(B^T K \right). \tag{34}$$

Applying the eigenvalue inequality of the product, and using the symmetry of B, we find that

$$\lambda_1 \left(B^T K \right) \leq \lambda_1(B)\lambda_1(K). \tag{35}$$

Furthermore, the eigenvalue inequality of the sum yields

$$\lambda_1(B) \leq (1 - \beta)\lambda_1(B_0) + \beta\lambda_1\left(P_f\right). \tag{36}$$

Applying (35) and (36) directly to (19), we obtain that

$$\kappa(S_{P4D}) \leq 1 + \left[(1 - \beta)\lambda_1(B_0) + \beta\lambda_1\left(P_f\right)\right]\lambda_1(K). \tag{37}$$

For the lower bound, we note that applying Theorem 2 to $\lambda_1\left(B^T K\right)$ only yields

$$\lambda_1\left(B^T K\right) \geq \max\left[\lambda_1\left(B^T\right)\lambda_n(K), \lambda_n\left(B^T\right)\lambda_1(K)\right] \geq 0.$$

∎

In terms of the effectiveness of the preconditioning, we note that the upper bounds given by Theorems 5 and 6 are not controlled by the condition number of B, but by the largest eigenvalues of B_0, P_f, K^2 only. This is different from the upper bound of the unpreconditioned case. In addition, for the unpreconditioned case, $B = P_f$ at $\beta = 1$, indicating that B becomes closer to a singular matrix as β approaches 1. In such a case, the theory predicts that the condition number of S_{4D} diverges to infinity as $\beta \to 1$. However, by implementing CVT, this divergence is eliminated.

On the other hand, we note that the upper bound given by Theorem 5 is distinctively different from all the other versions, including the ones derived for the unpreconditioned cases. Namely, the upper bound of Theorem 5 contains a max function that selects the larger term of $(1 - \beta)\lambda_1(B_0)$ and $\beta\lambda_1\left(P_f\right)$. We can thus anticipate a switching point of the upper bound as a function of β, which does not exist in other versions of the upper bound. In fact, this switching point provides useful information about the behavior of $\kappa(S_{P4D})$ with respect to varying β. We will illustrate this with numerical examples in Section 4. Similar to the unpreconditioned scenario, we find that these bounds do not reflect any contribution of the number p of observation points, which means that they cannot predict the trend of $\kappa(S_{4D})$ with respect to a changing p.

4 | EXPERIMENTAL DESIGN

Recalling the motivation of the theoretical studies in Section 3, we note that the purpose of the bounds is to provide useful information on the actual condition number of the Hessian without having to compute the Hessian explicitly. However, for the bounds to be informative, we require two properties of them: the estimation given by these bounds should reflect the condition number; the bounds change with β similarly to the condition number. The first property ensures that we can use the bound directly to estimate the condition number in practice; the second property guarantees that the bound is useful in analyzing the impact of the ensemble part on the condition number. To demonstrate the theories we use the special case of 3DVar.

4.1 | Computing error covariance matrices and the observational matrix

In this section, we give details of computing the error covariance matrices of B_0, P_f, R and the observational matrix H on a one-dimensional grid.

Notation 4. Let L_0, L_{ens} denote the correlation length scale associated with B_0, P_f, and $\sigma_{R_0}^2, \sigma_{B_0}^2, \sigma_{P_f}^2$ be the variances of R, B_0, P_f. Let r denote the radius of a two dimensional disk $\mathcal{B}^2(0, r)$, θ be the angular location of a point on the boundary of $\mathcal{B}^2(0, r)$ and $\theta_{i,j}$ be the angular difference between two grid point on the boundary of $\mathcal{B}^2(0, r)$.

To simulate real-life applications, we use the Second-order Auto-regressive Correlation (SOAR) function to generate B_0, with a correlation length scale of L_0. The SOAR function is used by the UK Met Office for Numerical Weather Prediction (NWP).[8,19] It is also frequently used in other DA applications.[20,21] Following the formulation discussed in Tabeart et al.,[8] we can compute the SOAR matrix from the following formula,[8,22]

$$D_L(i,j) = \left[1 + 2rL^{-1}\sin\left(\frac{\theta_{i,j}}{2}\right)\right]\exp\left[1 + 2rL^{-1}\sin\left(\frac{\theta_{i,j}}{2}\right)\right], \tag{38}$$

where L is a correlation length scale.

Applying (38), we can obtain B_0 directly by replacing L with L_0, and $\theta_{i,j}$ with $2\pi/n$. Here we assume that the grid discretizing the boundary of $\mathcal{B}^2(0,r)$ is uniform and we set $r = 1$. The resulting background error covariance matrix is then given by,

$$B_0 = \sigma_{B_0}^2 D_{L_0}. \tag{39}$$

To produce the ensemble part P_f, we sample from a different SOAR matrix B_1 associated with a length scale of L_{ens}, such that,

$$B_1 = \sigma_{P_f}^2 D_{L_{ens}}. \tag{40}$$

To obtain the covariance matrix P_f sampling from B_1, we generate a set of random vectors $\boldsymbol{w}_k \sim \mathcal{N}(0,1)$, then compute the sample covariance matrix of the set $\left\{B_1^{1/2}\boldsymbol{w}_1, B_1^{1/2}\boldsymbol{w}_2, \ldots, B_1^{1/2}\boldsymbol{w}_m\right\}$. As the size of the set—denoted by m—is restricted by $m < n$, P_f is then guaranteed to be rank deficient.

For the observational error covariance matrix R, we simply choose that,

$$R = \sigma_{R_0}^2 I_p, \tag{41}$$

where $I_p \in \mathbb{R}^{p,p}$ is the identity matrix.

In terms of the observational operator, we consider four different types in our experiments. They are given as follows,

$$H^{(1)}(i,j) = \begin{cases} 1, & i = j, \text{ for } i = 1 \to p, \\ 0, & i \neq j. \end{cases} \quad ; \quad H^{(3)}(i,j) = \begin{cases} 1/5, & j = \frac{n}{p}i - 2 \to \frac{n}{p}i + 2, \ (mod(n)), \text{ for } i = 1 \to p, \\ 0. \end{cases}$$

$$H^{(2)}(i,j) = \begin{cases} 1, & j = \frac{n}{p}i, \text{ for } i = 1 \to p, \\ 0, & else. \end{cases} \quad ; \quad H^{(4)}(i,j) = \begin{cases} 1, & i,j, \text{ are chosen randomly and non-repeating}, \\ 0. \end{cases} \tag{42}$$

The choice of $H^{(1)}$ corresponds to observing the first p points of the domain, $H^{(2)}$ observes every n/p points, $H^{(3)}$ is an observation which is a weighted sum over 5 grid points and $H^{(4)}$ chooses p observations at random locations on the grid. We note that $H^{(1)}, H^{(2)}$ and $H^{(4)}$ satisfy the condition of Corollary 1, while $H^{(3)}$ does not. In all of the experiments, we fix the value of n to be 500 such that the computation is small scale and similar to previous studies,[3,8] and therefore the results can be compared.

4.2 | A comparison of different bounds

We note that four different versions of the bound have been presented in this article, that is, Theorems 3 and 4 for the unpreconditioned cases, Theorems 5 and 6 for the preconditioned cases with CVT. Before further investigating the details of these results, we want to determine which versions of these are the most effective (in terms of revealing the trend of the condition number itself and estimating it with a small error). Considering the unpreconditioned cases, Theorem 4 has an apparent advantage of separating the impact of P_f. However, through repeated experiments we observe that in most cases the lower bound in Theorem 4 remains close to 1; hence it not effective. On the other hand, the lower bound provided by Theorem 3 produces a better result. Still, it is not close enough to the actual condition number to be a good estimate.

In Figure 1, we show typical examples of the variation of these bounds with β. Figure 1a shows that, for the unpreconditioned case, as β increases toward 1, the upper bound captures the shape of $\kappa(S_{3D})$. The divergence of $\kappa(S_{3D})$ to infinity
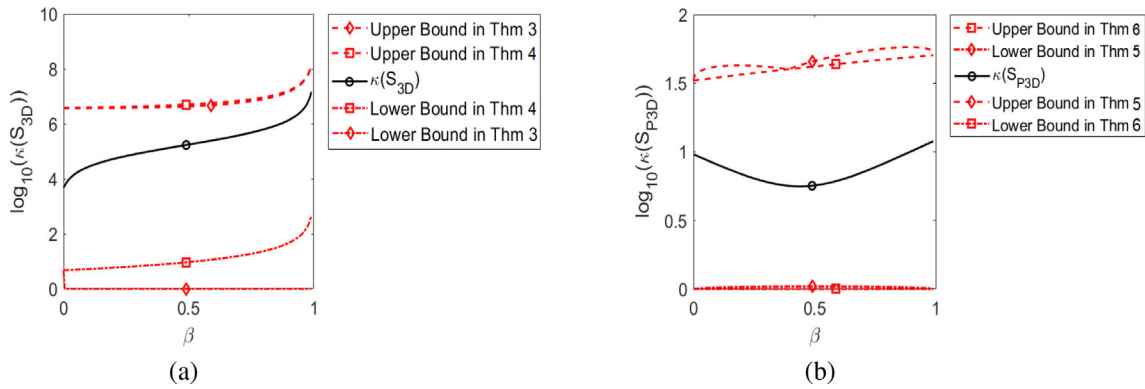
FIGURE 1 A comparison of the variation of the bounds with β for (a) $\kappa(S_{3D})$ and (b) $\kappa(S_{P3D})$. The parameters are as follows, $L_{ens} = 0.05, L_0 = 0.2, \sigma_{B_0} = \sigma_{P_f} = \sigma_{R_0} = 1, m = 50, p = 100$. The linear observation operator is chosen to be $H_0 = H^{(4)}$.

is captured by both versions of the upper bounds. The upper bounds computed from Theorems 1 and 3 do not differ significantly. This is a general result for all choices of parameters we have used to conduct numerical case studies. Based on these observations, we decide to focus only on the upper bound from here onward.

In the preconditioned case, Figure 1b, we note that the condition number is drastically reduced by the CVT compared to the unpreconditioned case. We observe an inflection point in the upper bound of Theorem 5. This coincides with the transition point of the condition number (from decreasing to increasing with β) and so predicts the minimum of $\kappa(S_{3D})$. The bound produced by Theorem 6 does not capture such behavior, while not being significantly closer to the actual condition number either.

These results are valid for all different choices of matrices and parameters. Therefore we focus our study on the bound produced by Theorems 4 and 5 only, because of their superiority in separating the impact of different matrices and capturing the shape of the actual condition number (see Figure 1b). We note that the condition numbers are plotted in $\log_{10}$ scale in this paper.

5 | CASE STUDIES FOR UNPRECONDITIONED 3D-VAR

In this section, we illustrate the upper bound given by Theorem 4 with multiple case studies. In these case studies we change parameters associated with $B_0, P_f, \hat{H}, \hat{R}$ and observe the responses in the condition number of S and its upper bound. In particular, we are especially interested in the relationship between the conditioning of the system and the weight of the ensemble part (i.e., β). We note again that the case studies in the following sections use 3D-Var as a special case of 4D-Var to demonstrate the theories. Therefore, $\hat{H}$ and $\hat{R}$ are replaced with H_0 and R_0, respectively. Additionally, we will denote the Hessian as S_{3D} and consider it as a special case of S_{4D}.

As the ensemble error covariance matrix is rank deficient, we can then expect that the balance of B_0 and P_f is particularly important in the conditioning of the system. For example, when P_f is the dominant component of the Hessian matrix S_{3D}, the condition number of S_{3D} is likely to be large and the system ill-conditioned. In addition, as $\beta \rightarrow 1$, the Hessian matrix $S_{3D} \rightarrow P_f^{-1} + H_0^T R_0^{-1} H_0$, the limit is clearly ill-defined as P_f^{-1} does not exist. Meanwhile, the balance of B_0 and P_f is also controlled by the relative sizes of physical length scales L_0, L_{ens} and the variances $\sigma_{B_0}, \sigma_{P_f}$.

Recalling the upper bound in Theorem 4, it is predominantly controlled by the largest eigenvalues of B_0 and P_f. Meanwhile, the largest eigenvalues of B_0 and P_f vary with their correlation length scales.[3] As Figure 2 shows, the largest eigenvalue of B_0 increases with the correlation length scale; the general trend is similar for P_f except for random fluctuations caused by the sampling noise. Consequently, we expect $\kappa(S_{3D})$ and its bound to change when the physical length scales (L_0, L_{ens}) alter.

As an important justification of the effectiveness of the upper bound, we observe that the shape of the upper bound is similar to the condition number in all four case studies presented in Figure 3. We also note that in this set of experiments, we observe that the impact of L_{ens} on the conditioning and the bounds is less significant than L_0 (compare Figure 3a,b). In Figure 3c, we observe that the trend of the condition number with respect to σ_{P_f} is well captured by the upper bounds. Also, in Figure 3d, we observe that upper bound does reveal a general trend of the condition number with respect to
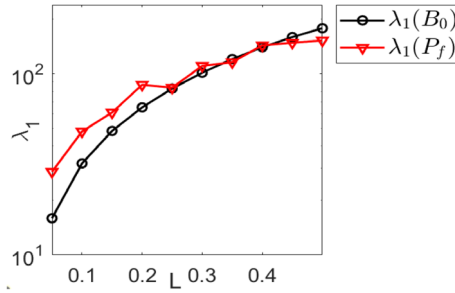
FIGURE 2 The largest eigenvalues of B_0 and P_f as functions of correlation length scale.
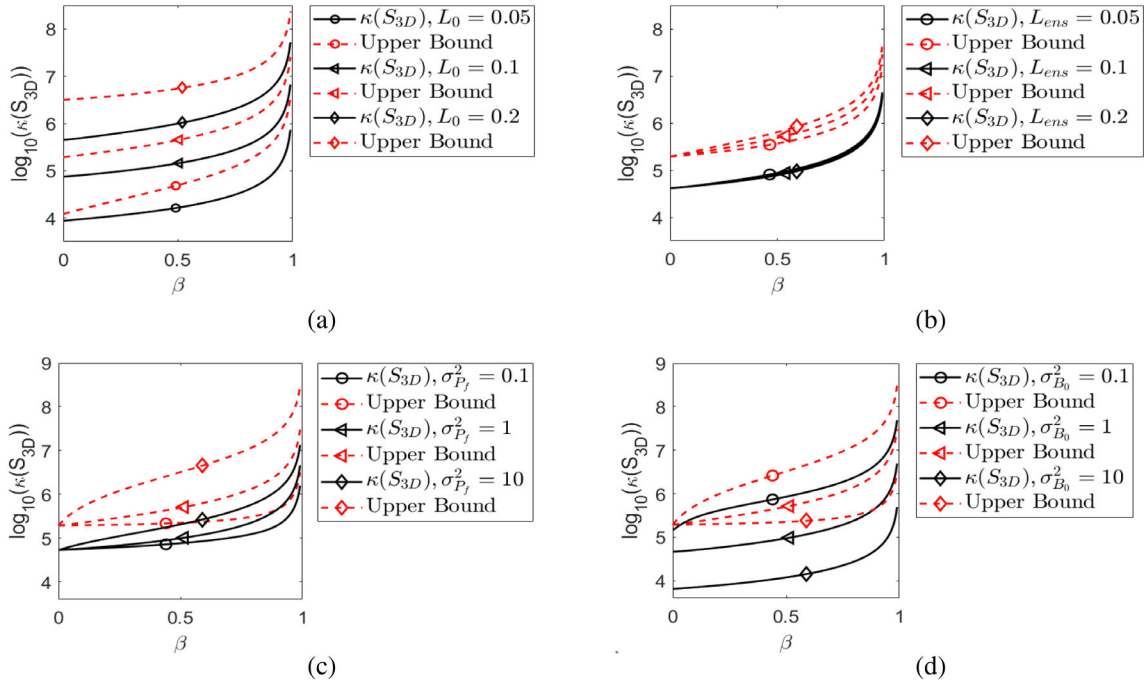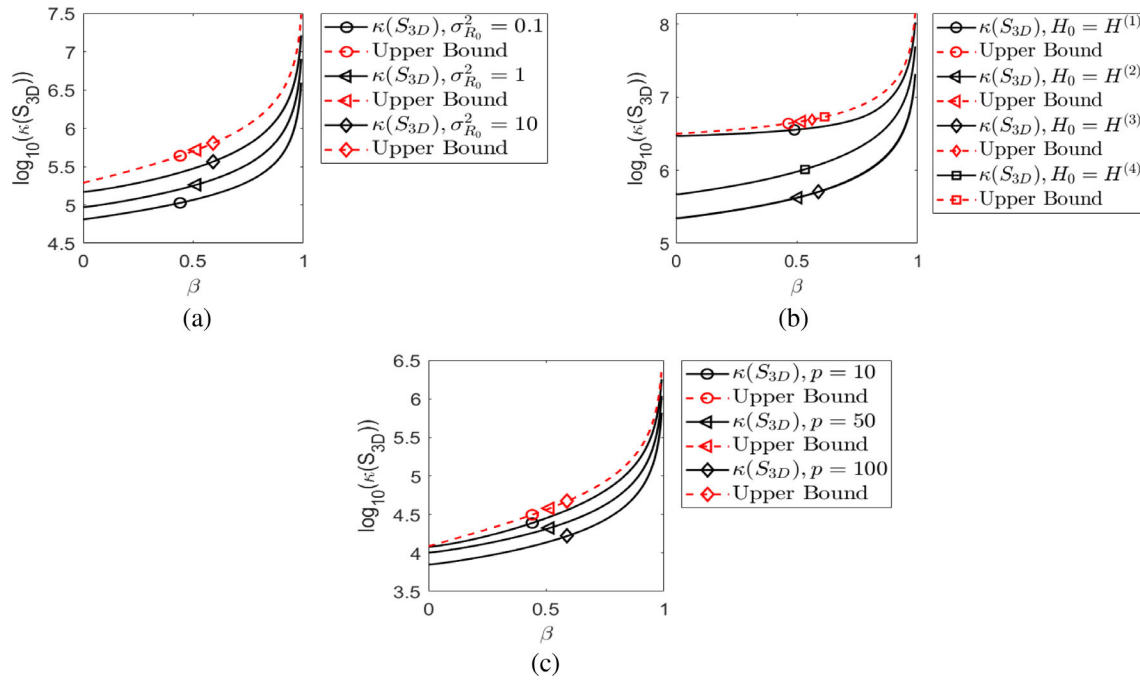

FIGURE 3 The variation of the condition number and the upper bound with β for different values of (a) L_0, (b) L_{ens}, (c) σ_{B_0}, and (d) σ_{P_f}, for the unpreconditioned case with $m = 100, p = 100, \sigma_{P_f}^2 = \sigma_{B_0}^2 = \sigma_{R_0}^2 = 1$ and $H_0 = H^{(4)}$. (a) Showcases the effect of L_0 (fixing $L_{ens} = 0.1$). (b) Demonstrates the effect of L_{ens} (fixing $L_0 = 0.1$). (c) Demonstrates the effect of $\sigma_{P_f}^2$ (fixing $\sigma_{B_0}^2 = 1$). (d) Illustrates the effect of $\sigma_{B_0}^2$ (fixing $\sigma_{P_f}^2 = 1$).

σ_{B_0}, but the upper bound does not reflect a significant reduction of the condition number at $\beta \sim 0$. The reason for this discrepancy is unknown.

It is an important finding that the upper bound starts to sharply increase and diverge to infinity, and this transition occurs at about the same β as the condition number of the Hessian. This then indicates that the bound is informative from a numerical point of view, for example, it informs a constraint on β if one seeks to avoid a sudden deterioration of the conditioning of the system in Hybrid 4D-Var. On the other hand, we also note that in the three cases of changing L_0, σ_{B_0} and σ_{P_f}, the upper bounds reflect the direction of change in the condition number of the Hessian. Thus it shows promising results that the upper bound can be used to qualitatively analyze the conditioning of the system when these parameters alter with time. Furthermore, the upper bound is two orders of magnitude above the condition number. We find similar conclusions for the cases of changing σ_{R_0}, the observation operator H_0 and the number of observations p (see Figure 4).

We note that the upper bound for a fixed B_0 and P_f does not change insofar as the largest eigenvalue of $H_0^T R_0 H_0$ stays the same (see Theorem 4). Comparing $\lambda_1(H_0^T R_0 H_0)$ for different versions of H_0, we find that $H^{(1)}, H^{(2)}$, and $H^{(4)}$ have the same largest eigenvalue. It is slightly smaller for $H^{(3)}$ but not significantly. We also point out that previous investigations in

FIGURE 4 The variation of the condition number and the upper bound with β for different (a) values of $\sigma_{R_0}^2$ (fixing $\sigma_{B_0}^2 = \sigma_{P_f}^2 = 1, p = 100,\ H_0 = H^{(4)}$), (b) observation operators and (c) number of observations (fixing $\sigma_{R_0}^2 = \sigma_{B_0}^2 = \sigma_{P_f}^2 = 1, H_0 = H^{(4)}$). This group of figures showcase the unpreconditioned case with parameters of $m = 100, L_{ens} = L_0 = 0.1$.

4D-Var[3,8] also indicate that a small change in $\sigma_{R_0}^2$ does not change the bound noticeably. Thus we anticipate that the upper bound remains similar for different versions of H_0. This is confirmed by cases in Figure 4b. Meanwhile, the condition number of the Hessian is the largest for the choice of $H_0 = H^{(1)}$, while choosing $H_0 = H^{(2)}$ or $H_0 = H^{(3)}$ produces a similar result and they yield the best conditioning of the Hessian; the condition number associated with the randomized version $H_0 = H^{(4)}$ lies between those of $H^{(1)}$ and $H^{(2)}, H^{(3)}$. Thus the observation indicates that evenly distributed observation points produce better conditioning of the system, while partial observations concentrated in a small region lead to the opposite.

The upper bound in (16) also implies that changing the number of observations does not change the upper bound (as the case study in Figure 4c confirms). Last, we also verified that the upper bound also remains effective when the number of ensemble members changes. We tested cases where the number of ensemble members increases from 50 to 400. In all these cases, we find that the upper bounds have similar shapes to the condition number, and they provide a good estimation of the value of the condition number.

6 | NUMERICAL EXPERIMENTS FOR THE PRECONDITIONED HYBRID 3D-VAR WITH CVT

Following the case studies for the unpreconditioned cases, we now conduct similar experiments to illustrate Theorem 5 and examine the predictions using Theorem 5 for preconditioned cases with CVT. In the experiments, we focus on three major issues: first, we validate the correctness of the bound; second, we illustrate the bound reflects the behavior of the conditioning in general (with a few exceptions); last, we make a comparison to the unpreconditioned cases. We note that each error covariance matrix and observation operator is given by Section 3.

An immediate observation from Theorem 5 is that the switching point in the upper bound shifts with the relative sizes of $\lambda_1(B_0)$ and $\lambda_1(P_f)$. More specifically, the max function in the upper bound in (17) switches at a larger value of β when $\lambda_1(B_0)$ increases, and the opposite occurs when $\lambda_1(P_f)$ increases. Furthermore, we know that $\lambda_1(B_0)$ and $\lambda_1(P_f)$ become larger as L_0 and L_{ens} increase (see Figure 2), and the same is true for increasing $\sigma_{B_0}, \sigma_{P_f}$. As Figure 5 shows, the results of these four case studies justifies the observations implied by Theorem 5, and the switching point of the upper bound
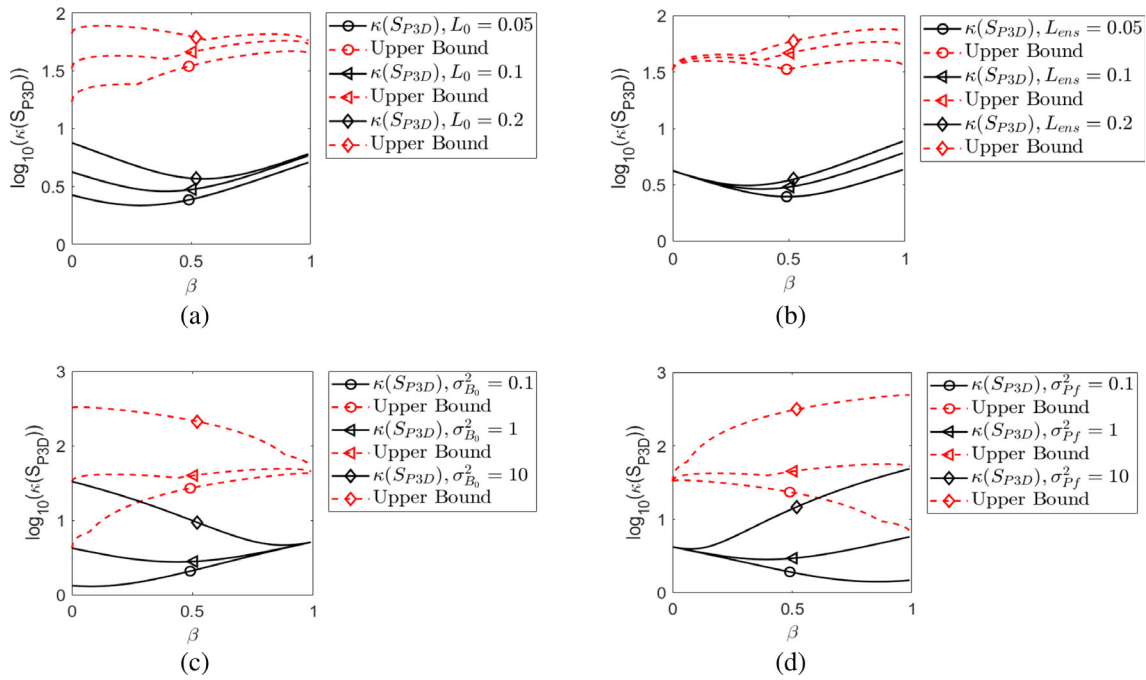
FIGURE 5 Same settings as Figure 3, but for the CVT preconditioned cases.

also predicts the minimum of the condition number. On the other hand, as $\lambda_1(B_0)$, $\lambda_1(P_f)$ increase with $L_0, L_{ens}, \sigma_{B_0}, \sigma_{P_f}$, we anticipate that the upper bound would increase with these parameters. This is confirmed by the case studies (see Figure 5). Crucially, we find that the upper bound predicts the trend of the condition number of Hessian with respect to β in all four cases. For example, in Figure 5a, we observe that the inflection points of the upper bounds predict the minima of the condition numbers, and they both move rightward when L_0 increases. In Figure 5b, we find that the inflection points of the upper bounds and the minima of the condition numbers move leftwards when L_{ens} increases. In Figure 5c, the inflection points of the upper bounds and the minima of the condition numbers move rightwards when σ_{B_0} increases, and in Figure 5d we find that the trend of the condition number reacting to σ_{P_f} is well captured by the upper bounds, the inflection points of the upper bounds and the minima of the condition numbers move leftwards when σ_{P_f} increases. However, we want to point out that the changes in the conditioning is very limited in these preconditioned cases and are normally less than one order of magnitude.

Furthermore, compared to the unpreconditioned case, these results show a clear improvement of the conditioning in Hybrid 4D-Var with CVT (we note that this is something that is also found in standard 4D-Var in previous work[9]). We observe that CVT leads to a maximum reduction of six magnitudes in the condition number of Hessian. We also note that in all cases presented in the preconditioned Hybrid 4D-Var, the condition number of Hessian does not diverge to infinity at $\beta = 1$, and this is a crucial difference from the unpreconditioned Hybrid 4D-Var.

Theorem 5 also shows that the upper bound takes a larger value when the largest eigenvalue of $K = H_0^T R_0 H_0$ is bigger. Furthermore, since $\sigma_{R_0}^{-2}$ is a scaling factor of K, we can then anticipate that the upper bound grows with a decreasing $\sigma_{R_0}^2$. The case presented in Figure 6a not only justifies this observation but also shows that the condition number itself follows the same trend.

On the impact of choosing different versions of H_0, the trend is similar to the unpreconditioned case; we find that evenly spreading out the observation across the whole domain leads to better conditioning and skewing observations into a local region result in worse conditioning. We note that the upper bound remains the same for $H_0 = H^{(1)}, H^{(2)}$ and $H^{(4)}$ (see Figure 6b,c). This is because $K = H_0^T R_0 H_0$ shares the same largest eigenvalue for these three versions, whereas $H^{(1)}$ produces a smaller maximum eigenvalue for K, therefore the upper bound is smaller. On the other hand, the impact of the number p of observations on the condition number is opposite to that of the unpreconditioned case (this is similar to previous reports of 4D-Var[8,9]), but the bound estimation does not reflect this trend, as the largest eigenvalue of K does not change with p. We note that in both unpreconditioned and preconditioned cases, the upper bounds cannot predict
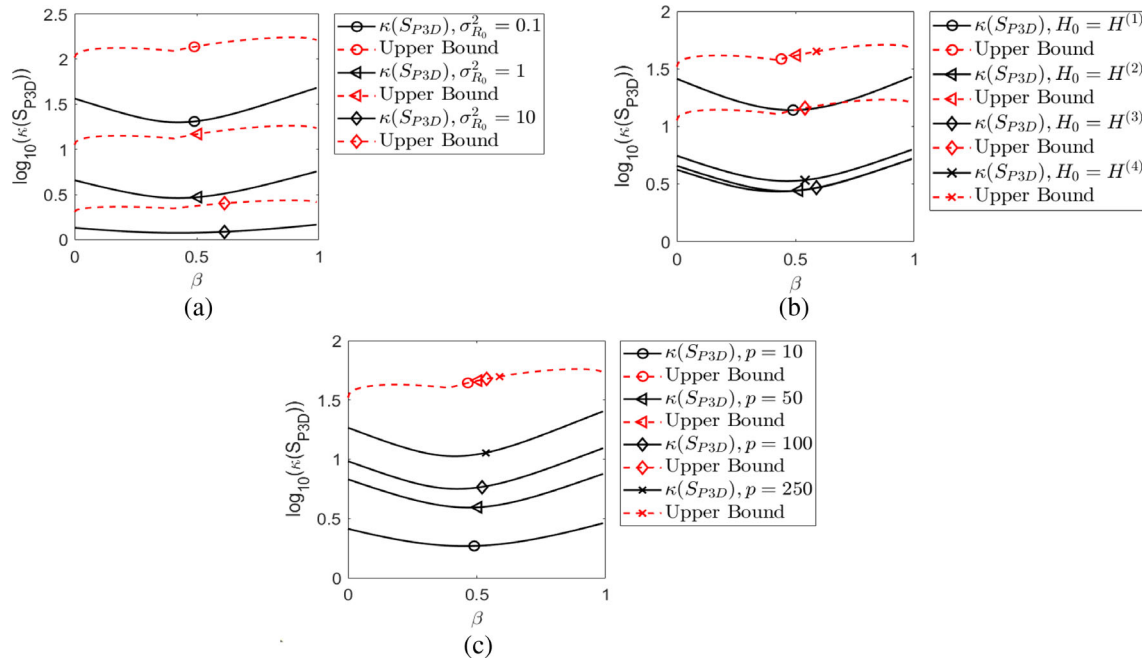
FIGURE 6   As settings as Figure 4, but for the CVT preconditioned cases.

the impact of p on the conditioning of the system. The effect of sampling noise in P_f is similar to the unpreconditioned case, which is that the impact on the conditioning or the upper bound is insignificant. We find that this is true even with a small sample size.

As an important justification of the effectiveness of the upper bound given by Theorem 5, we observe the changes of the upper bound with respect to β provide valuable information about the transition of the condition number (from decreasing to increasing); the inflection point of the upper bound predicts the minimum of the condition number. This is particularly important because it informs an optimal choice of β from a numerical perspective of obtaining the best conditioning of the system. We therefore can conclude that the upper bound is useful for providing qualitative information about the actual conditioning of the system.

## 7 | CONVERGENCE TEST OF A CONJUGATE GRADIENT ROUTINE

We note that there are well-known situations where the condition number provides a pessimistic indication of convergence speed (e.g., in the case of repeated or clustered eigenvalues).[9] Here we use hybrid 3D-Var as a special case of Hybrid 4D-Var to illustrate that for hybrid variational assimilation the convergence speed follows a similar trend to the condition number as the weight of the ensemble part increases.

Following a similar method to section 5.3.2. of Tabeart et al.,[8] we study how the speed of convergence of a conjugate gradient method applied to the linear system $S_{3D}\boldsymbol{x} = \boldsymbol{b}$ changes with the weight β of the ensemble part. In the first test, the matrix S_{3D} is given by the Hessian of the unpreconditioned 3D-Var (Section 5), and the vector $\boldsymbol{b}$ is given by Haben[3] in Section 3.2 ($\boldsymbol{b} = B^{-1}(\boldsymbol{x}_b - \boldsymbol{x}_0) - H_0^T\boldsymbol{d}$), where the vectors $\boldsymbol{x}_b - \boldsymbol{x}_0, \boldsymbol{d}$ are chosen to be random at the beginning of the trial). For the computation of S_{3D}, we choose the parameters as follows, $L_0 = 0.1, L_{ens} = 0.05, \sigma_{B_0} = \sigma_{P_f} = \sigma_{R_0} = 1, p = 100$ and $H_0 = H^{(4)}$, which are in line with our previous case studies.

As Figure 7 shows, the condition number of the Hessian shows the same trend as the number of iterations performed to reach the tolerance threshold. We can observe that the conditioning is a good proxy to study the convergence speed in this case. For the preconditioned case with CVT, shown in Figure 8, we find that the difference in the condition number as β varies is marginal relative to the unpreconditioned case. One of the reason for this is that the Hessian has more eigenvalues clustered around 1, which is a result of CVT. However, the number of iterations taken to converge still follows the general pattern of the condition number, with higher values when β is close to 0 or 1.
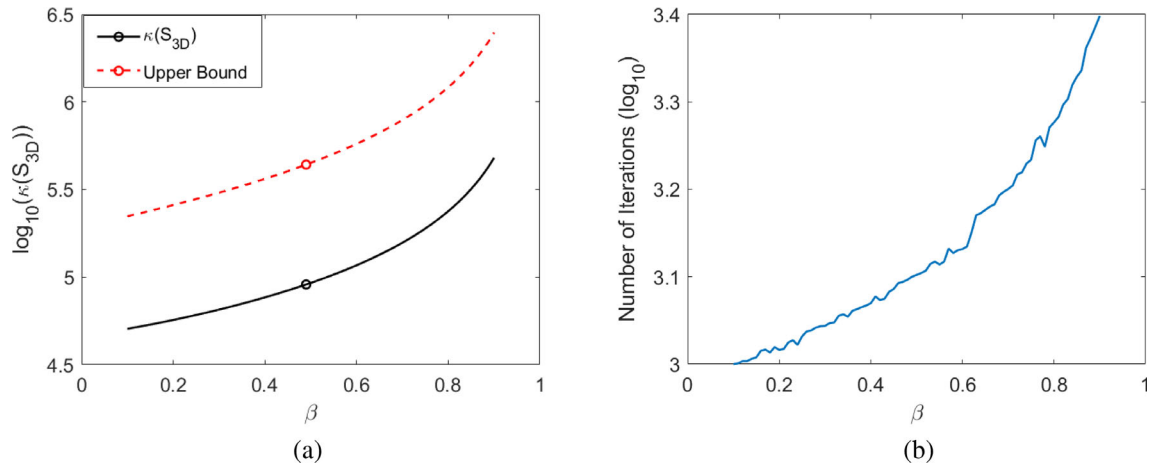
FIGURE 7 Convergence test for an unpreconditioned case. (a) displays the condition number of S_{3D} and its upper bound. The parameters used to compute S_{3D} are $L_0 = 0.1, L_{ens} = 0.05, \sigma_{B_0} = \sigma_{P_f} = \sigma_{R_0} = 1, p = 100$ and we choose $H_0 = H^{(4)}$. (b) displays the number of iterations.
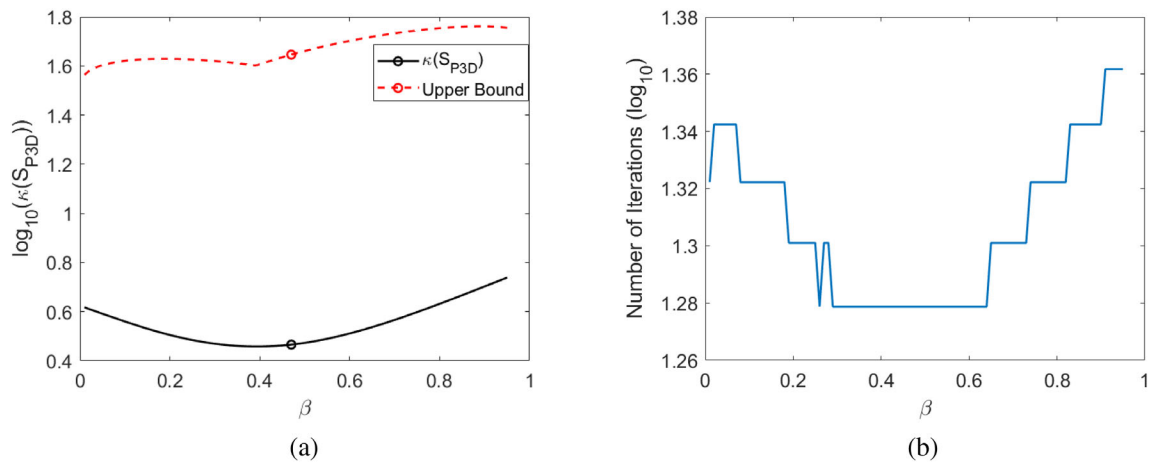
FIGURE 8 Convergence test for a case preconditioned with CVT. (a) Displays the condition number of S_{P3D} (as a special case of S_{P4D}) and its upper bound. The parameters used to compute S_{P3D} are $L_0 = 0.1, L_{ens} = 0.05, \sigma_{B_0} = \sigma_{P_f} = \sigma_{R_0} = 1, p = 100$ and we choose $H_0 = H^{(4)}$. (b) Displays the number of iterations.

Finally, we conducted experiments with stopping criteria of different orders of magnitude to investigate the impact on the convergence speed trend. For $\epsilon > 10^{-6}$ there is very little variation in the number of iterations as β changes. For values of ϵ smaller than 10^{-6}, the number of iterations increases slightly, but the trend in Figure 8b remains consistent.

8 | SUMMARY

In this article, we established a set of theories for the conditioning of Hybrid 4D-Var. These theories provide effective upper bounds for the condition number of the Hessian. These theoretical results are illustrated by numerical case studies using the special case of Hybrid 3D-Var. In numerical experiments we tested that the upper bounds have similar shapes to the condition number with respect to the weight of the ensemble part (i.e., β). Thus they can provide a useful estimation of the behavior of the condition number of the Hessian. In addition, the upper bound enabled us to study the condition number through parameters associated with different components and observe their interactions. We conclude that the upper bound is effective in explaining the conditioning of the Hessian matrix in general. However, the lower bound does not provide useful information about the condition number in all cases studied. We summarize our findings as follows,

- For the unpreconditioned cases, the condition number of the Hessian increases gradually at first with β then quickly diverges to infinity as $\beta \to 1$. This transition is predicted by the upper bounds.
- For the preconditioned cases with CVT, a general trend is that the condition number reduces at first and then increases as β increases. There is an optimal β at which the condition number is at its minimum. The upper bound predicts the optimal β effectively.
- The preconditioning with CVT improves the conditioning drastically and eliminates the divergence of the condition number of Hessian at $\beta = 1$.
- For preconditioned cases, the upper bound changes in the same direction as the condition number of the Hessian with respect to changes in $L_0, L_{ens}, \sigma_{B_0}, \sigma_{P_f}$, and σ_{R_0}. In unpreconditioned cases, we have similar conclusion for L_0, σ_{B_0}, and σ_{P_f}.

In preconditioned cases, we find that the upper bound reveals the trend of the condition number with respect to changing parameters such as the correlation length scale and the variance. However, the theories in Section 3 cannot explain the impact of the four different choices of H_0 that we presented in this paper. Furthermore, the bounds do not change with the number of observations p and the sample size m, although p does directly influence the actual condition number.

Meanwhile, the tests show that these theories do provide useful predictions on the impact of the balancing of variances and correlation length scales for the preconditioned cases. In unpreconditioned cases, the upper bound can predict well the influence of L_0 and σ_{P_f}. The bounds can inform the impact of these components on the convergence of iterative numerical algorithms.

It is well-known that the condition number of the Hessian matrix is a useful proxy to study the convergence speed of the least-squares minimization of Hybrid 4D-Var. The results presented in this paper could then inform applications in terms of the restriction of the weight of the ensemble part (for the unpreconditioned cases), such that extreme ill-conditioning can be avoided. For the preconditioned Hybrid 4D-Var with CVT, we established a theory that effectively predicts an optimal weight of the ensemble part such that the conditioning is optimal.

ACKNOWLEDGMENTS

CONFLICT OF INTEREST STATEMENT
The authors declare that there is no conflict of interest regarding the publication of this paper.

DATA AVAILABILITY STATEMENT
The data that support the findings of this study are openly available in my github page at https://github.com/Shaerdan.

ORCID
Shaerdan Shataer https://orcid.org/0000-0001-6022-8953

REFERENCES
1. Nichols NK. Mathematical concepts of data assimilation. Data Assimilation. New York: Springer; 2010. p. 13–39.
2. Gratton S, Lawless AS, Nichols NK. Approximate Gauss–Newton methods for nonlinear least squares problems. SIAM J Optim. 2007;18(1):106–32.
3. Haben SA, Lawless AS, Nichols NK. Conditioning of incremental variational data assimilation, with application to the Met Office system. Tellus A Dyn Meteorol Oceanography. 2011;63(4):782–92.
4. Bannister RN. A review of operational methods of variational and ensemble-variational data assimilation. Q J R Meteorol Soc. 2017;143(703):607–33.
5. Smith PJ, Lawless AS, Nichols NK. Estimating forecast error covariances for strongly coupled atmosphere–ocean 4D-Var data assimilation. Mon Weather Rev. 2017;145(10):4011–35.
6. Smith PJ, Lawless AS, Nichols NK. Treating sample covariances for use in strongly coupled atmosphere-ocean data assimilation. Geophys Res Lett. 2018;45(1):445–54.
7. Golub GH, Van Loan CF. Matrix Computations. Baltimore, Maryland: JHU Press; 2013.

8. Tabeart JM, Dance SL, Haben SA, Lawless AS, Nichols NK, Waller JA. The conditioning of least-squares problems in variational data assimilation. Numer Linear Algebra Appl. 2018;25(5):e2165.

9. Tabeart JM, Dance SL, Lawless AS, Nichols NK, Waller JA. New bounds on the condition number of the Hessian of the preconditioned variational data assimilation problem. Numer Linear Algebra Appl. 2022;29(1):e2405.

10. Lawless AS, Gratton S, Nichols NK. An investigation of incremental 4D-Var using non-tangent linear models. Quart J Royal Meteorol Soc A J Atmos Sci Appl Meteorol Phys Oceanography. 2005;131(606):459–76.

11. Mercier F, Michel Y, Montmerle T, Jolivet P, Gürol S. Speeding up the ensemble data assimilation system of the limited-area model of Météo-France using a block Krylov algorithm. Q J R Meteorol Soc. 2019;145(720):910–29.

12. Desroziers G, Berre L. Accelerating and parallelizing minimizations in ensemble and deterministic variational assimilations. Q J R Meteorol Soc. 2012;138(667):1599–610.

13. Hatfield S, McRae A, Palmer T, Düben P. Single-precision in the tangent-linear and adjoint models of incremental 4D-var. Mon Weather Rev. 2020;148(4):1541–52.

14. Aabaribaoune ME, Emili E, Guidard V. Estimation of the error covariance matrix for IASI radiances and its impact on ozone analyses. Atmos Measurement Tech Discussions. 2020;2020:1–26.

15. Buehner M. Ensemble-derived stationary and flow-dependent background-error covariances: Evaluation in a quasi-operational NWP setting. Q J R Meteorol Soc A J Atmos Sci Appl Meteorol Phys Oceanography. 2005;131(607):1013–43.

16. Wilkinson J. The algebraic eigenvalue problem. Handbook for Automatic Computation. Linear Algebra. Volume II. New York: Springer-Verlag; 1971.

17. Weyl H. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). Math Ann. 1912;71(4):441–79.

18. Wang B, Zhang F. Some inequalities for the eigenvalues of the product of positive semidefinite Hermitian matrices. Linear Algebra Appl. 1992;160:113–8.

19. Simonin D, Ballard S, Li Z. Doppler radar radial wind assimilation using an hourly cycling 3D-Var with a 1.5 km resolution version of the Met Office Unified Model for nowcasting. Q J R Meteorol Soc. 2014;140(684):2298–314.

20. Gong H, Yu Y, Li Q, Quan C. An inverse-distance-based fitting term for 3D-Var data assimilation in nuclear core simulation. Ann Nucl Energy. 2020;141:107346.

21. Cheng S, Qiu M. Observation error covariance specification in dynamical systems for data assimilation using recurrent neural networks. Neural Comput & Applic. 2021;34:1–19.

22. Waller JA, Dance SL, Nichols NK. Theoretical insight into diagnosing observation error correlations using observation-minus-background and observation-minus-analysis statistics. Q J R Meteorol Soc. 2016;142(694):418–31.