



A reduced basis ensemble Kalman method

Francesco A. B. Silva¹ · Cecilia Pagliantini² · Martin Grepl³ · Karen Veroy¹

Received: 1 November 2022 / Accepted: 25 July 2023 / Published online: 1 September 2023

© The Author(s) 2023

Abstract

In the process of reproducing the state dynamics of parameter dependent distributed systems, data from physical measurements can be incorporated into the mathematical model to reduce the parameter uncertainty and, consequently, improve the state prediction. Such a data assimilation process must deal with the data and model misfit arising from experimental noise as well as model inaccuracies and uncertainties. In this work, we focus on the ensemble Kalman method (EnKM), a particle-based iterative regularization method designed for *a posteriori* analysis of time series. The method is gradient free and, like the ensemble Kalman filter (EnKF), relies on a sample of parameters or particle ensemble to identify the state that better reproduces the physical observations, while preserving the physics of the system as described by the best knowledge model. We consider systems described by parameterized parabolic partial differential equations and employ model order reduction techniques to generate surrogate models of different accuracy with uncertain parameters. Their use in combination with the EnKM involves the introduction of the model bias which constitutes a new source of systematic error. To mitigate its impact, an algorithm adjustment is proposed accounting for a prior estimation of the bias in the data. The resulting RB-EnKM is tested in different conditions, including different ensemble sizes and increasing levels

✉ Francesco A. B. Silva

f.a.b.silva@tue.nl

Cecilia Pagliantini

cecilia.pagliantini@unipi.it

Martin Grepl

grepl@igpm.rwth-aachen.de

Karen Veroy

k.p.veroy@tue.nl

¹ Department of Mathematics and Computer Science, Eindhoven University of Technology, Eindhoven, The Netherlands

² Department of Mathematics, University of Pisa, Pisa, Italy

³ Institute of Geometry and Practical Mathematics, RWTH Aachen University, 52056 Aachen, Germany

of experimental noise. The results are compared to those obtained with the standard EnKF and with the unadjusted algorithm.

Keywords Inverse problems · Ensemble Kalman method · Model order reduction · Representation error

Mathematics Subject Classification 65M32 · 76M21 · 62M20 · 35Q86

1 Introduction

The problem of estimating model parameters of static and dynamical systems is encountered in many applications from earth sciences to engineering. In this work we focus on the parameter estimation of dynamical systems described by parameterized parabolic partial differential equations (pPDEs). Here, we assume that a limited and polluted knowledge of the solution is available at multiple time instances through noisy local measurements.

For solving this kind of inverse problem, countless deterministic and stochastic methods have been proposed. Among them, a widely used technique is the so-called ensemble Kalman filter (Evensen 2003), a recursive filter employing a series of measurements to obtain improved estimates of the variables involved in the process. The idea of using the EnKF for reconstructing the parameters of dynamical systems traces back to Anderson (2001) and Lorentzen et al. (2001), in which trivial artificial dynamics for the parameters was assumed to make the estimation possible. This was naturally accompanied by efforts for improving the performance of the method in terms of stability, by introducing covariance inflation (Hamill et al. 2001; Anderson and Anderson 1999) and localization (Hamill et al. 2001; Houtekamer and Mitchell 2001), and in terms of computational cost. Relevant to the latter have been the development of multi-level (Hoel et al. 2016) and multi-fidelity Popov et al. (2021) and Donoghue and Yano (2022) methods, the use of model order reduction (MOR) techniques with offline (Pagani et al. 2017; da Silva and Colonius 2018) and with on-the-fly (Donoghue and Yano 2022) training, as well as the introduction of further surrogate modeling techniques (Popov and Sandu 2022). The use of approximated models inevitably led to the study of the impact of model error on the EnKF (Mitchell et al. 2002; Mitchell and Carrassi 2015), alongside with other data assimilation methods (Calvetti et al. 2018; Huttunen and Kaipio 2007).

Although ensemble Kalman methods were originally meant for sequential data assimilation, i.e., for real-time applications, they proved to be reliable also for asynchronous data assimilation (Sakov et al. 2010). The first paper proposing to adapt the EnKF to a retrospective data analysis was (Skjervheim et al. 2007). For analysis, the data are employed all at once at the end of an assimilation window, which is in common with a series of methods, e.g., variational methods (Li and Navon 2001) such as 4D-VAR (Thepaut and Courtier 1991) and other smoothers (Anderson and Moore 1979). Compared to those approaches, the EnKF is particularly appealing since it does not require the computation of Fréchet derivatives, a major complication for data assimilation algorithms.

In Iglesias et al. (2013), Iglesias et al. introduced what they called the *ensemble Kalman method*, an EnKF-based asynchronous data assimilation algorithm. Depending on the design of the algorithm, this method has connections to Bayesian data assimilation (Schillings and Stuart 2018) and to maximum likelihood estimation (Chen and Oliver 2012). In particular, in the latter case, the method constitutes an ensemble-based implementation of so-called *iterative regularization methods* (Kaltenbacher et al. 2008). In the case of perfect models, the EnKM has already been analyzed in depth in Schillings and Stuart (2018) and Evensen (2018) and convergence and identifiability enhancements have been proposed in Wu et al. (2019) and Iglesias (2016). Due to the iterative nature of the EnKM, dealing with high-dimensional parametric problems is often computationally challenging. In Gao and Wang (2021) a multi-level strategy has been proposed to improve the computational performance of the method.

In this work we propose an algorithm, called *reduced basis ensemble Kalman method* (RB-EnKM), that leverages the computational efficiency of surrogate models obtained with MOR techniques to solve asynchronous data assimilation problems via ensemble Kalman methods. The use of the EnKM allows us to avoid adjoint problems that are often difficult to reduce and intrinsically depend on the choice of measurement positions. Model order reduction, already employed in other data assimilation problems (Gong et al. 2019; Nadal et al. 2015), is used as a key tool for accelerating the method. However, the use of approximate models within the EnKM introduces a model error that could hinder the convergence of the method. In this work, we propose to deal with this error by including a prior estimation of the bias in the data. Specifically, we incorporate empirical estimates of the mean and covariance of the bias in the Kalman gain. In some instances, those quantities can be computed at a negligible cost by employing the same training set used for the construction of the reduced model.

The paper is structured as follows: in Sect. 2 we introduce the asynchronous data assimilation problem together with the standard ensemble Kalman method (Algorithm 1). Subsequently, in Sect. 3.1, we present an overview on reduced basis (RB) methods and describe how to use them in combination with the ensemble Kalman method to derive the RB-EnKM (Algorithm 2). In Sect. 4, we test the new method on two numerical examples. In the first example, we estimate the diffusivity in a linear advection-dispersion problem in 2D (Sect. 4.1), while in the second, we estimate the hydraulic log-conductivity in a non-linear hydrological problem (Sect. 4.2). In both cases, we compare the behavior of the full order and reduced order models in different conditions. Section 5 provides conclusions and considerations on the proposed method and on its numerical performance.

2 Problem formulation

Let $(\mathcal{U}, \mathcal{H})$ be a suitable pair of function spaces and let $\mathcal{P} \subset \mathbb{R}^{N_p}$, with $N_p \in \mathbb{N}^+$, be a set of model parameters. We consider the pPDE: for any parameter $\boldsymbol{\mu} \in \mathcal{P}$, find $u(\boldsymbol{\mu}) \in \mathcal{U}$ such that $\partial_t u(\boldsymbol{\mu}) = \mathcal{F}_{\boldsymbol{\mu}} u(\boldsymbol{\mu})$, $u(0, \boldsymbol{\mu}) = u_0(\boldsymbol{\mu})$. Here $\mathcal{F}_{\boldsymbol{\mu}}$ is a generic parameterized differential operator, ∂_t is the first order partial time derivative and $u_0(\boldsymbol{\mu}) \in \mathcal{H}$ is a parameterized initial condition. This pPDE provides the constraint

to the inverse problem of estimating the unknown parameter $\mu^* \in \mathcal{P}$ from data or observations given by

$$\begin{aligned} \mathbf{y}(\mu^*, \eta) &= \mathcal{L}u(\mu^*) + \eta \quad s.t. \\ \partial_t u(\mu^*) &= \mathcal{F}_{\mu^*} u(\mu^*), \quad u(0, \mu^*) = u_0(\mu^*). \end{aligned} \tag{1}$$

Here, $\mathcal{L} : \mathcal{U} \rightarrow \mathbb{R}^{N_m}$, with $N_m \in \mathbb{N}^+$, maps the space of the solutions to the space of the measurements, simulating the observation process, and η is an unknown realization of a Gaussian random variable with zero mean and given covariance, $\Sigma \in \mathbb{R}^{N_m \times N_m}$. Note that both the observed data \mathbf{y} and the additive experimental noise η are N_m -dimensional vector-valued quantities and that Σ is a symmetric positive-definite matrix defining the inner product $\|\cdot\|_{\Sigma^{-1}}^2 := \|\Sigma^{-1/2} \cdot\|_2^2$ on \mathbb{R}^{N_m} , where $\|\cdot\|_2$ is the Euclidean norm.

To solve this inverse problem, we must explicitly solve the pPDE (1). This is done using a suitable discretization, in space and time, of the differential operators \mathcal{F}_μ and ∂_t . To this end, we introduce the approximation spaces $\mathcal{V}_h \subset \mathcal{U}$ and $\mathcal{H}_h \subset \mathcal{H}$ and the discretized initial condition $u_{h,0}(\mu) \in \mathcal{H}_h$ so that the approximate problem reads:

$$\text{find } u_h(\mu) \in \mathcal{V}_h \quad s.t. \quad \partial_t u_h(\mu) = \mathcal{F}_\mu^h u_h(\mu), \quad u_h(0, \mu) = u_{h,0}(\mu). \tag{2}$$

The discretization of the pPDE can be chosen according to the specific problem of interest. In all numerical examples proposed in this work, we employ a space-time Petrov–Galerkin discretization of (1) with piecewise polynomial trial and test spaces, as described in Sect. 4, and we assume (2) to be sufficiently accurate such that we can take $\mathbf{y}(\mu^*, \eta) = \mathcal{L}u_h(\mu^*) + \eta$.

To characterize the observation of the solution, we introduce the *forward response* map $\mathcal{G} : \mathcal{P} \rightarrow \mathbb{R}^{N_m}$ defined as $\mathcal{G}(\mu) := \mathcal{L}u_h(\mu)$ for any solution of the pPDE (2). Although the use of the map \mathcal{G} results in a more compact notation, omitting its dependence on the solution of the pPDE conceals a key aspect of the method, i.e., the mapping from the parameter vector to the corresponding space-time pPDE solution. For this reason, and because it makes it harder to introduce the problem discretization, it will be used with caution.

2.1 The ensemble Kalman method

The data assimilation problem presented above can be recast as a minimization problem for the cost functional, $\Phi(\mu | \mathbf{y}) := \|\mathbf{y}(\mu^*, \eta) - \mathcal{L}u_h(\mu)\|_{\Sigma^{-1}}^2$, representing the misfit between the experimental data, $\mathbf{y}(\mu^*, \eta)$, and the forward response. The optimal parameter estimate $\mu_{\text{opt}}(\mathbf{y})$ is thus given by

$$\begin{aligned} \mu_{\text{opt}}(\mathbf{y}) &= \arg \min_{\mu \in \mathcal{P}} \Phi(\mu | \mathbf{y}) \quad s.t. \\ \partial_t u_h(\mu) &= \mathcal{F}_\mu^h u_h(\mu), \quad u_h(0, \mu) = u_{h,0}(\mu). \end{aligned} \tag{3}$$

This is equivalent to a maximum likelihood estimation, given the *likelihood* function, $l(\mu | \mathbf{y}) = \exp\{-\frac{1}{2}\Phi(\mu | \mathbf{y})\}$, associated with the probability density function of the data, $\mathbf{y}|\mu$, i.e., the probability of observing \mathbf{y} if μ is the parametric state. The shape

of the function follows from the probability density function of the Gaussian noise realization.

Among various methods proposed to solve this optimization problem, the EnKM relies on a sequence of parameter ensembles \mathcal{E}_n , with $n \in \mathbb{N}^+$, to estimate the minimum of the cost functional. Each ensemble consists of a collection $\{\boldsymbol{\mu}_n^{(j)}\}_{j=1}^J$ of $J \in \mathbb{N}^+$ parameter vectors $\boldsymbol{\mu}_n^{(j)}$, hereby named ensemble members or particles, whose interaction, guided by the experimental measurements, causes them to cluster around the solution of the problem as iterations proceed. At the beginning of each iteration, the solution of the pPDE and its observations are computed for each $j \in \{1, \dots, J\}$. Subsequently, the ensemble is updated based on the empirical correlation among parameters and between parameters and measurements, as well as on the misfits between the experimental measurements $\mathbf{y}(\boldsymbol{\mu}^*, \boldsymbol{\eta})$ and the particle measurements $\mathcal{L}u_h(\boldsymbol{\mu}_n^{(j)})$. A single iteration, equivalent to the one in Iglesias et al. (2013), is formalized in the following pseudo algorithm:

Algorithm 1 Iterative ensemble method for inverse problems.

Input. Let \mathcal{E}_0 be the initial ensemble with elements $\{\boldsymbol{\mu}_0^{(j)}\}_{j=1}^J$ sampled from a given distribution $\Pi_0(\boldsymbol{\mu})$. Let $\boldsymbol{\Sigma}$ be the *a priori* known noise covariance and \mathbf{y} the vector of noisy measurements collected from the physical system. Let $\tau \ll 1$ be the termination tolerance.

For $n = 0, 1, \dots$

- (i) **Prediction step.** Compute the synthetic measurements of the solution over a time interval \mathcal{I} for each particle in the last updated ensemble:

$$\begin{aligned} \mathcal{G}(\boldsymbol{\mu}_n^{(j)}) &= \mathcal{L}u_h(\boldsymbol{\mu}_n^{(j)}) \quad \text{for all } j \in \{1, \dots, J\} \quad \text{s.t.} \\ \partial_t u_h(\boldsymbol{\mu}_n^{(j)}) &= \mathcal{F}_{\boldsymbol{\mu}_n^{(j)}}^h u_h(\boldsymbol{\mu}_n^{(j)}), \quad u_h(0, \boldsymbol{\mu}_n^{(j)}) = u_{h,0}(\boldsymbol{\mu}_n^{(j)}). \end{aligned} \tag{4}$$

- (ii) **Intermediate step.** From the last updated ensemble measurements and parameters, compute the sample means and covariances:

$$\mathbf{P}_n = \frac{1}{J} \sum_{j=1}^J \mathcal{G}(\boldsymbol{\mu}_n^{(j)}) \mathcal{G}(\boldsymbol{\mu}_n^{(j)})^\top - \bar{\mathcal{G}}_n \bar{\mathcal{G}}_n^\top \quad \text{with} \quad \bar{\mathcal{G}}_n = \frac{1}{J} \sum_{j=1}^J \mathcal{G}(\boldsymbol{\mu}_n^{(j)}) \tag{5}$$

$$\mathbf{Q}_n = \frac{1}{J} \sum_{j=1}^J \boldsymbol{\mu}_n^{(j)} \mathcal{G}(\boldsymbol{\mu}_n^{(j)})^\top - \bar{\boldsymbol{\mu}}_n \bar{\mathcal{G}}_n^\top \quad \text{with} \quad \bar{\boldsymbol{\mu}}_n = \frac{1}{J} \sum_{j=1}^J \boldsymbol{\mu}_n^{(j)}. \tag{6}$$

- (iii) **Analysis step.** Update each particle in the ensemble: for all $j \in \{1, \dots, J\}$

$$\boldsymbol{\gamma}_n^{(j)} \sim \mathcal{N}(0, \boldsymbol{\Sigma}), \tag{7}$$

$$\boldsymbol{\mu}_{n+1}^{(j)} = \boldsymbol{\mu}_n^{(j)} + \mathbf{Q}_n (\mathbf{P}_n + \boldsymbol{\Sigma})^{-1} (\mathbf{y} - \mathcal{G}(\boldsymbol{\mu}_n^{(j)}) - \boldsymbol{\gamma}_n^{(j)}). \tag{8}$$

- (iv) **Termination step.** Stop the algorithm when the termination criterion is satisfied. Here, we terminate when the relative change in the mean parameter is less than

the tolerance:

$$\|\bar{\mu}_{n+1} - \bar{\mu}_n\|_2 \leq \tau \|\bar{\mu}_{n+1}\|_2 \quad \text{with} \quad \bar{\mu}_{n+1} = \frac{1}{J} \sum_{j=1}^J \mu_{n+1}^{(j)}. \tag{9}$$

In the last step of the algorithm, the cross correlation matrices \mathbf{P}_n and \mathbf{Q}_n are used to compute the Kalman gain $\mathbf{K}_n := \mathbf{Q}_n(\mathbf{P}_n + \Sigma)^{-1}$. This modulates the extent of the correction: a low-gain corresponds to conservative behavior, i.e., small changes in the particle positions, while a high-gain involves a larger correction. Note that the experimental data are perturbed with artificial noise sampled from the same distribution assumed for the experimental noise η . This leads to an improved estimate over the unperturbed case.

A termination criterion for the algorithm is essential for the proper implementation of the method. The one presented in Iglesias et al. (2013) is based on the discrepancy principle and consists in stopping the algorithm when the error between the experimental data and the measurements is comparable to the experimental noise, that is, when $\|\mathbf{y} - \mathcal{G}(\bar{\mu}_n)\|_{\Sigma^{-1}}^2 \leq \sigma \|\eta\|_{\Sigma^{-1}}^2$ for some $\sigma \geq 1$. An alternative approach is to set a threshold for the norm of the parameter update, i.e., to terminate the algorithm when $\|\bar{\mu}_{n+1} - \bar{\mu}_n\|_2 \leq \tau \|\bar{\mu}_{n+1}\|_2$ for some $\tau \ll 1$. The latter criterion is more robust to model errors and is therefore used in our numerical experiments.

Equally important for the method is the choice of the distribution Π_0 from which the initial ensemble (or first guess) \mathcal{E}_0 is sampled. In most of the cases, including those considered in our numerical experiments, the distribution Π_0 comes from an *a priori* knowledge of the range of admissible parameters. In other scenarios, e.g., when the parameters live in an infinite-dimensional space, it may be necessary to define additional criteria on how to treat the parameter space. The initial ensemble plays a fundamental role in stabilizing the inverse problem. Indeed, it has been shown in Iglesias et al. (2013) that all the ensembles generated by Algorithm 1 are contained in the space spanned by the initial ensemble, that is

$$\mathcal{E}_n \in \mathcal{A} := \text{span} \{ \mu_0^{(j)} \}_{j=1}^J \quad \text{for all } n \in \mathbb{N}^+. \tag{10}$$

Furthermore, in the mean-field limit, i.e., in the case of infinite particles, and assuming an affine relationship between parameters and synthetic measurements, the distribution Π_0 plays the same role as the Tikhonov regularization in variational data assimilation, see (Asch et al. 2016). In particular, the stabilization term is given by $-\log_e \Pi_0(\mu)$.

The main sources of error of the EnKM are associated with the ensemble size and with the evaluation of $\mathcal{G}(\mu_n^{(j)}) = \mathcal{L}u_h(\mu_n^{(j)})$. Indeed, while the observation of the solution is accurate and computationally cheap to evaluate, due to the linearity of the operator, the accuracy in the computation of the pPDE solution intrinsically depends on the quality of the numerical discretization. High order numerical discretizations might require prohibitively large computational costs, especially if the pPDE (2) is solved for many values of the parameter and over long temporal intervals.

The other steps of Algorithm 1 involve the following operations: (i) the assembly of \mathbf{P}_n and \mathbf{Q}_n in (5)-(6), with computational complexity of order $\mathcal{O}(JN_m^2)$, and (ii) the

inversion of the matrix $\mathbf{P}_n + \Sigma$ in the *analysis step* (8) with complexity $\mathcal{O}(N_m^3)$. The solution of the pPDE (2), for all $j \in \{1, \dots, J\}$, in the prediction step of Algorithm 1 is thus the computational bottleneck of the EnKM algorithm.

3 Surrogate models

3.1 Reduced basis methods

Given the need to solve the pPDE (2) for several instances of the parameter, the use of MOR techniques appears an ideal choice. Model order reduction has allowed exceptional computational speed-ups in settings that require repeated model evaluations, such as multi-query simulations. In MOR the high-dimensional problem is replaced with a surrogate model of reduced dimensionality that still possesses optimal or near-optimal approximation properties but that can be solved at a considerably reduced computational cost. In this work, we focus on a particular class of MOR techniques, known as reduced basis methods (Prud'homme et al. 2002).

The reduced basis method typically consists of two phases: an offline phase and an online phase. In the computationally expensive offline phase a low-dimensional approximation of the solution space, namely the *reduced space*, is constructed and a surrogate model is derived via projection of the full order model onto the reduced space. Then, the resulting low-dimensional reduced model can be solved in the online phase for many instances of the parameter at a computational cost independent of the size of the full order model.

To be more precise, let $\mathcal{M} := \{u_h(t, \boldsymbol{\mu}) \in \mathcal{V}_h \mid \partial_t u_h(\boldsymbol{\mu}) = \mathcal{F}_\mu^h u_h(\boldsymbol{\mu}), u_h(0, \boldsymbol{\mu}) = u_{h,0}(\boldsymbol{\mu}) \text{ for all } \boldsymbol{\mu} \in \mathcal{P}, t \in \mathcal{I}\}$ be the solution set which collects the solution of the discretized pPDE (2) evaluated at times $t \in \mathcal{I} := (0, T]$, with $T \in \mathbb{R}^+$, for a set of parameters $\boldsymbol{\mu} \in \mathcal{P}$. The parametric problem (2) is said to be reducible if the solution set \mathcal{M} can be well approximated by a low-dimensional linear subspace. In this case, such a subspace is obtained as the span of a problem-dependent basis derived from a collection of full order solutions or snapshots, $\{u_h(t_n, \boldsymbol{\mu}_s)\}_{s,n=1}^{S,R}$, with $S, R \in \mathbb{N}^+$, at sampled values, $\{\boldsymbol{\mu}_s\}_{s=1}^S$, of the parameter, and at discrete times, $\{t_n\}_{n=1}^R$. The set $\mathcal{P}_{\text{TRAIN}} := \{\boldsymbol{\mu}_s\}_{s=1}^S \subset \mathcal{P}$ of training parameters is a sufficiently rich subset of the parameter space that can be obtained by drawing random samples from a uniform distribution in \mathcal{P} or with other sampling techniques, such as statistical methods and sparse grids, see (Quarteroni et al. 2015, Chapter 6) and references therein. The extraction of the basis functions from the snapshots is usually performed using SVD-type algorithms such as the proper orthogonal decomposition (POD) (Berkooz et al. 1993) or greedy algorithms. For problems that depend on both time and parameters, the so-called POD-Greedy method (Greppl and Patera 2005; Haasdonk and Ohlberger 2008) combines a greedy algorithm in parameter space with the proper orthogonal decomposition in time at a given parameter. In the numerical tests of this work, we rely on the Weak-POD-Greedy algorithm, which is the preferred method whenever a rigorous error bound can be derived, while the POD and the Strong-POD-Greedy are often used when a bound is unavailable, i.e., for most non-linear problems. Note

that, under the same choice of training parameters, the latter are more accurate but computationally less efficient.

Once an N_ε -dimensional set of spatial reduced basis functions $\{\psi_i\}_{i=1}^{N_\varepsilon}$ is obtained and a set of time basis functions $\{v_n\}_{n=1}^{N_t}$ is selected, the reduced spaces $\mathcal{H}_\varepsilon = \text{span}\{\psi_i\}_{i=1}^{N_\varepsilon} \subset \mathcal{H}_h$ and $\mathcal{V}_\varepsilon = \text{span}\{v_n \otimes \psi_i\}_{i,n=1}^{N_\varepsilon, N_t} \subset \mathcal{V}_h$ are constructed. The full model solution $u_h(\boldsymbol{\mu})$ and the initial condition $u_{h,0}(\boldsymbol{\mu})$, for a given $\boldsymbol{\mu}$, are approximated by the functions $u_\varepsilon(\boldsymbol{\mu})$ in \mathcal{V}_ε and $u_{\varepsilon,0}(\boldsymbol{\mu})$ in \mathcal{H}_ε ,

$$u_\varepsilon(\boldsymbol{\mu}) = \sum_{i,n=1}^{N_\varepsilon, N_t} u_{i,n}(\boldsymbol{\mu}) v_n \psi_i, \quad u_{\varepsilon,0}(\boldsymbol{\mu}) = \sum_{i=1}^{N_\varepsilon} u_{i,0}(\boldsymbol{\mu}) \psi_i, \quad \boldsymbol{\mu} \in \mathcal{P},$$

where $(u_{1,1}(\boldsymbol{\mu}), \dots, u_{N_\varepsilon, N_t}(\boldsymbol{\mu}))^\top \in \mathbb{R}^{N_\varepsilon N_t}$ and $(u_{1,0}(\boldsymbol{\mu}), \dots, u_{N_\varepsilon, 0}(\boldsymbol{\mu}))^\top \in \mathbb{R}^{N_\varepsilon}$ denote the vectors of expansion coefficients in the reduced basis. The reduced model thus reads:

$$\text{find } u_\varepsilon(\boldsymbol{\mu}) \in \mathcal{V}_\varepsilon \quad \text{s.t.} \quad \partial_t u_\varepsilon(\boldsymbol{\mu}) = \mathcal{F}_\mu^\varepsilon u_\varepsilon(\boldsymbol{\mu}), \quad u_\varepsilon(0, \boldsymbol{\mu}) = u_{\varepsilon,0}(\boldsymbol{\mu}), \quad (11)$$

where the operator $\mathcal{F}_\mu^\varepsilon$ is obtained by projecting the full order operator \mathcal{F}_μ^h onto the reduced space \mathcal{V}_ε . Note that we set $N_t = R$ in the sequel since we do not consider a temporal compression. However, choosing $R < N_t$ is also possible.

The computational gain derived from solving problem (11) instead of the full order model (2) hinges on the feasibility of a complete decoupling of the offline and online phases. A computational complexity of the online phase independent of the size of the full order problem can be achieved under the assumption of linearity and parameter-separability of the operator \mathcal{F}_μ^h . To deal with general non-linear operators, hyper-reduction techniques are required. These include methods for approximating the high-dimensional non-linear term \mathcal{F}_μ^h with an empirical affine decomposition, such as the EIM (Barrault et al. 2004), and methods for reducing the cost of evaluating the non-linear term, such as linear program empirical quadrature (Yano and Patera 2019) and empirical cubature (Hernández et al. 2017).

3.2 A reduced basis ensemble Kalman method

In this section, we discuss the implications of replacing the high-fidelity model in the prediction step of the EnKM by a surrogate model derived via model order reduction, as described in Sect. 3.1. The use of MOR for particle-based methods is particularly desirable in multi-query contexts since it allows us to significantly reduce the computational cost of solving the inverse problem. However, the approximation introduced by the model order reduction inevitably produces (small) deviations of the reduced solution from the full order one. This constitutes a problem for data assimilation algorithms, as already documented and investigated in Calvetti et al. (2018) and in other works. Indeed, the error in the solution results in discrepancies between approximated and exact measurements. Although we can expect the mismatch $\delta_\varepsilon(\boldsymbol{\mu}) := \mathcal{L}u_h(\boldsymbol{\mu}) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu})$ to decrease with the approximation error of $u_\varepsilon(\boldsymbol{\mu})$, this

bias will inevitably entail a distortion of the loss functional obtained by simple model substitution, i.e.,

$$\tilde{\Phi}(\boldsymbol{\mu} | \mathbf{y}) := \|\mathbf{y}(\boldsymbol{\mu}^*, \boldsymbol{\eta}) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu})\|_{\Sigma^{-1}}^2. \tag{12}$$

Note that this cost function does not vanish in the parameter $\boldsymbol{\mu}$ we are trying to estimate, not even in noise free conditions. This systematic error, independent of the magnitude of the experimental noise, can be mitigated by modifying the cost function and consequently the EnKM. A modified algorithm, which we refer to as *adjusted RB-EnKM* is presented in the following sections. This algorithm is in contrast to what we refer to as the *biased RB-EnKM*, i.e., the algorithm obtained by the simple substitution of the full order model with the reduced order model in Algorithm 1, as presented in (12).

The modification of the algorithm can proceed in two ways. One possibility is to rewrite the exact cost function in terms of the surrogate model and the measurement bias, namely substituting $\mathcal{L}u_h(\boldsymbol{\mu}) = \mathcal{L}u_\varepsilon(\boldsymbol{\mu}) + \boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu})$ in the minimization problem (3) to obtain

$$\begin{aligned} \Phi_1(\boldsymbol{\mu} | \mathbf{y}) &:= \|\mathbf{y}(\boldsymbol{\mu}^*, \boldsymbol{\eta}) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu}) - \boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu})\|_{\Sigma^{-1}}^2 \\ &= \|\mathcal{L}u_h(\boldsymbol{\mu}^*) - \mathcal{L}u_h(\boldsymbol{\mu}) + \boldsymbol{\eta}\|_{\Sigma^{-1}}^2 \equiv \Phi(\boldsymbol{\mu} | \mathbf{y}). \end{aligned} \tag{13}$$

A second option is to correct the experimental data involved in the biased cost function (12) so that, at least in noise free conditions, its minimum coincides with the minimum of the exact cost function. This means subtracting $\boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu}^*)$ instead of $\boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu})$, and results in the new cost function

$$\begin{aligned} \Phi_2(\boldsymbol{\mu} | \mathbf{y}) &:= \|\mathbf{y}(\boldsymbol{\mu}^*, \boldsymbol{\eta}) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu}) - \boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu}^*)\|_{\Sigma^{-1}}^2 \\ &= \|\mathcal{L}u_\varepsilon(\boldsymbol{\mu}^*) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu}) + \boldsymbol{\eta}\|_{\Sigma^{-1}}^2 \not\equiv \Phi(\boldsymbol{\mu} | \mathbf{y}). \end{aligned} \tag{14}$$

In noise free conditions, i.e., if $\boldsymbol{\eta} = \mathbf{0}$, both cost functions vanish at the exact value $\boldsymbol{\mu}^*$. Since the cost functions Φ_1 and Φ_2 are non-negative, the minimum attained in $\boldsymbol{\mu}^*$ is necessarily also a global minimum.

In the following, we focus on the second approach. The reason is that the first approach requires the evaluation of the bias at all parameter values $\boldsymbol{\mu} \in \mathcal{P}$ which is too expensive to perform. Furthermore, at the algorithmic level, the substitution of the true model with the sum of the surrogate model and its bias would significantly change the computation of \mathbf{P}_n and \mathbf{Q}_n , and thus the algorithm structure. By contrast, the second approach is based on the assumption that the true model is incorrect and on the subsequent correction of the experimental data. This implies that $\boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu}^*)$ is the only bias involved and it requires just a single full order evaluation. However, since the argument $\boldsymbol{\mu}^*$ is unknown, this is clearly not possible, and we must instead exploit the prior epistemic uncertainty on $\boldsymbol{\mu}^*$, encoded in $\Pi_0(\boldsymbol{\mu})$, to modify the cost function.

If $\boldsymbol{\mu}^*$ is treated as a random variable with probability measure Π_0 , then the data bias $\boldsymbol{\delta}_\varepsilon^* = \boldsymbol{\delta}_\varepsilon(\boldsymbol{\mu}^*)$ is in turn a random variable with probability measure $\Pi_0 \circ \boldsymbol{\delta}_\varepsilon^{-1}$. The moments of this distribution, henceforth denoted by $\bar{\boldsymbol{\delta}}_\varepsilon$ and $\boldsymbol{\Gamma}_\varepsilon$, can be empirically

estimated via pointwise evaluations of the bias without further assumptions on the nature of the distribution itself. However, the assumption of Gaussianity, although improperly implying the linearity of $\delta_\varepsilon : \mathcal{P} \rightarrow \mathbb{R}^{N_m}$, is consistent with the other assumptions of Gaussianity and linearity required for the derivation of the EnKF (Evensen 2003). Furthermore, it allows us to obtain closed-form results, as shown in the next paragraphs.

In view of the fact that δ_ε^* is considered as a random variable, we change (14) to make the dependence of the cost function Φ_2 on δ_ε^* explicit; i.e.

$$\Phi_\varepsilon(\boldsymbol{\mu} \mid \mathbf{y}, \delta_\varepsilon^*) := \|\mathbf{y}(\boldsymbol{\mu}^*, \boldsymbol{\eta}) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu}) - \delta_\varepsilon^*\|_{\boldsymbol{\Sigma}}^2. \tag{15}$$

In order to make the estimate of $\boldsymbol{\mu}$ dependent only on the experimental data, we must remove the conditioning on δ_ε^* , i.e., marginalize out the random variable. The easiest way to do this is by employing Bayesian statistics and particularly recovering the same marginal distribution $\mathbf{y} \mid \boldsymbol{\mu}$ mentioned at the beginning of Sect. 2.1. To this end, we consider the likelihood function $l(\boldsymbol{\mu} \mid \mathbf{y}, \delta_\varepsilon^*) := \exp\{-\frac{1}{2}\Phi_\varepsilon(\boldsymbol{\mu} \mid \mathbf{y}, \delta_\varepsilon^*)\}$, proportional to the density of $(\mathbf{y} \mid \boldsymbol{\mu}, \delta_\varepsilon^*) \sim \mathcal{N}(\delta_\varepsilon^* + \mathcal{L}u_\varepsilon(\boldsymbol{\mu}), \boldsymbol{\Sigma})$. Employing (Särkkä 2013, Lemma 1.A), concerning the mean and covariance of the joint distribution of Gaussian variables, it can be easily proven that, if $\delta_\varepsilon^* \sim \mathcal{N}(\bar{\delta}_\varepsilon, \Gamma_\varepsilon)$, then $\mathbf{y} \mid \boldsymbol{\mu} \sim \mathcal{N}(\bar{\delta}_\varepsilon + \mathcal{L}u_\varepsilon(\boldsymbol{\mu}), \boldsymbol{\Sigma} + \Gamma_\varepsilon)$ and consequently we derive the marginalized cost functional

$$\Phi_\varepsilon(\boldsymbol{\mu} \mid \mathbf{y}) := \|\mathbf{y} - \mathcal{L}u_\varepsilon(\boldsymbol{\mu}) - \bar{\delta}_\varepsilon\|_{(\boldsymbol{\Sigma} + \Gamma_\varepsilon)}^2. \tag{16}$$

Hence, by analogy with Sect. 2.1, we can adapt the EnKM to optimize the new cost function under the surrogate model constraint (11). The resulting adjusted RB-EnKM is summarized in Algorithm 2. Unlike the reference EnKM, we distinguish between an offline and an online phase. In the offline phase, the training set of full order solutions is generated and used both to construct the surrogate model and to estimate the moments of δ_ε^* . In the online phase, the actual optimization is performed.

Algorithm 2 Iterative ensemble method with reduced basis surrogate models and accounting for the associated measurements bias.

Offline:

Input. Let $\mathcal{P}_{\text{TRAIN}}$ be a set of S parameters $\{\boldsymbol{\mu}_0^{(s)}\}_{s=1}^S$ sampled from a given probability distribution $\Pi_0(\boldsymbol{\mu})$ and let $\{u_h(\boldsymbol{\mu}_s)\}_{s=1}^S$ be the associated training set of full order solutions. Let $\varepsilon \in \mathbb{R}^+$ be a prescribed tolerance.

- (i) **Model order reduction.** Relying on the training set $\{u_h(\boldsymbol{\mu}_s)\}_{s=1}^S$, construct a surrogate model of accuracy ε as explained in Sect. 3.1 and compute the set of reduced basis solutions $\{u_\varepsilon(\boldsymbol{\mu}_s)\}_{s=1}^S$.
- (ii) **Data bias estimation.** Define the training biases as

$$\delta_\varepsilon(\boldsymbol{\mu}^{(s)}) = \mathcal{L}u_h(\boldsymbol{\mu}^{(s)}) - \mathcal{L}u_\varepsilon(\boldsymbol{\mu}^{(s)}) \quad \text{for all } s \in \{1, \dots, S\} \tag{17}$$

and the associated empirical moments

$$\Gamma_\varepsilon = \frac{1}{S} \sum_{s=1}^S \delta_\varepsilon(\boldsymbol{\mu}^{(s)}) \delta_\varepsilon(\boldsymbol{\mu}^{(s)})^\top - \bar{\delta}_\varepsilon \bar{\delta}_\varepsilon^\top \quad \text{with} \quad \bar{\delta}_\varepsilon = \frac{1}{S} \sum_{s=1}^S \delta_\varepsilon(\boldsymbol{\mu}^{(s)}). \quad (18)$$

Online:

Input. Let \mathcal{E}_0 be the initial ensemble with elements $\{\boldsymbol{\mu}_0^{(j)}\}_{j=1}^J$ sampled from a given distribution $\Pi_0(\boldsymbol{\mu})$. Let $\boldsymbol{\Sigma}$ be the *a priori* known noise covariance and \mathbf{y} the vector of noisy measurements collected from the physical system. Let $\tau \ll 1$ be the termination parameter.

For $n = 0, 1, \dots$

- (i) **Prediction step.** Compute the biased measurements of the approximated solution over a time interval \mathcal{I} for each particle in the last updated ensemble:

$$\begin{aligned} \mathcal{G}_\varepsilon(\boldsymbol{\mu}_n^{(j)}) &= \mathcal{L}u_\varepsilon(\boldsymbol{\mu}_n^{(j)}) \quad \text{for all } j \in \{1, \dots, J\} \quad \text{s.t.} \\ \partial_t u_\varepsilon(\boldsymbol{\mu}_n^{(j)}) &= \mathcal{F}_{\boldsymbol{\mu}_n^{(j)}}^\varepsilon u_\varepsilon(\boldsymbol{\mu}_n^{(j)}), \quad u_\varepsilon(0, \boldsymbol{\mu}_n^{(j)}) = u_{\varepsilon,0}(\boldsymbol{\mu}_n^{(j)}). \end{aligned} \quad (19)$$

- (ii) **Intermediate step.** From the last updated ensemble measurements and parameters, define the sample means and covariances:

$$\mathbf{P}_{n,\varepsilon} = \frac{1}{J} \sum_{j=1}^J \mathcal{G}_\varepsilon(\boldsymbol{\mu}_n^{(j)}) \mathcal{G}_\varepsilon(\boldsymbol{\mu}_n^{(j)})^\top - \bar{\mathcal{G}}_{n,\varepsilon} \bar{\mathcal{G}}_{n,\varepsilon}^\top \quad \text{with} \quad \bar{\mathcal{G}}_{n,\varepsilon} = \frac{1}{J} \sum_{j=1}^J \mathcal{G}_\varepsilon(\boldsymbol{\mu}_n^{(j)}), \quad (20)$$

$$\mathbf{Q}_{n,\varepsilon} = \frac{1}{J} \sum_{j=1}^J \boldsymbol{\mu}_n^{(j)} \mathcal{G}_\varepsilon(\boldsymbol{\mu}_n^{(j)})^\top - \bar{\boldsymbol{\mu}}_n \bar{\mathcal{G}}_{n,\varepsilon}^\top \quad \text{with} \quad \bar{\boldsymbol{\mu}}_n = \frac{1}{J} \sum_{j=1}^J \boldsymbol{\mu}_n^{(j)}. \quad (21)$$

- (iii) **Analysis step.** Update each particle in the ensemble: for all $j \in \{1, \dots, J\}$

$$\boldsymbol{\gamma}_n^{(j)} \sim \mathcal{N}(\bar{\delta}_\varepsilon, \boldsymbol{\Sigma} + \Gamma_\varepsilon), \quad (22)$$

$$\boldsymbol{\mu}_{n+1}^{(j)} = \boldsymbol{\mu}_n^{(j)} + \mathbf{Q}_{n,\varepsilon} (\mathbf{P}_{n,\varepsilon} + \Gamma_\varepsilon + \boldsymbol{\Sigma})^{-1} (\mathbf{y} - \mathcal{G}_\varepsilon(\boldsymbol{\mu}_n^{(j)}) - \boldsymbol{\gamma}_n^{(j)}). \quad (23)$$

- (iv) **Termination step.** Stop the algorithm when the termination criterion is satisfied:

$$\|\bar{\boldsymbol{\mu}}_{n+1} - \bar{\boldsymbol{\mu}}_n\|_2 \leq \tau \|\bar{\boldsymbol{\mu}}_{n+1}\|_2 \quad \text{with} \quad \bar{\boldsymbol{\mu}}_{n+1} = \frac{1}{J} \sum_{j=1}^J \boldsymbol{\mu}_{n+1}^{(j)}. \quad (24)$$

By employing the same training set for constructing the surrogate model and for evaluating $\bar{\delta}_\varepsilon$ and Γ_ε , we provide the largest possible training set to the model order reduction algorithm for a fixed value of S . However, we also introduce a bias in the estimation of the moments of $\delta_\varepsilon(\boldsymbol{\mu}^*)$ due to the underestimated values of $\delta_\varepsilon(\boldsymbol{\mu}^{(s)})$. The bias could be removed, e.g., by partitioning the training set $\{\boldsymbol{\mu}_0^{(s)}\}_{s=1}^S$ into two

sub-sets (or by introducing two independent sets with cardinality $S/2$), one for the construction of the surrogate model and one for the independent estimation of $\bar{\delta}_\varepsilon$ and Γ_ε . The disadvantage of this approach (for fixed S) would be a smaller training set for the surrogate model construction, and a poorer yet unbiased statistics for the moments' estimation.

We note from the offline part in Algorithm 2 in the data bias estimation that the RB approximation needs to approximate the full order model uniformly well globally, i.e., over the entire parameter domain. We can thus not replace the classical offline-online decomposition by an on-the-fly adaptation of the reduced basis (Donoghue and Yano 2022), which is beneficial if only local approximations, e.g. along the optimization path, are required.

In Algorithm 2, the prior probability $\Pi_0(\boldsymbol{\mu})$ used for the estimation of the moments of δ_ε^* could be substituted at every iteration by an updated probability measure of $\boldsymbol{\mu}$. However, the computation of the updated probability measure might compromise the computational gain obtained with the use of reduced models. One possibility to address this shortcoming is to use a Gaussian process regression of the initial ensemble biases to estimate the moments of δ_ε^* with respect to the new probability measure of $\boldsymbol{\mu}$. The development and study of this strategy together with its effect on the accuracy and performances of the RB-EnKM will be investigated in future studies.

4 Numerical experiments

In the following section, we consider two data assimilation problems for the estimation of model parameters in pPDEs. The first problem involves a linear advection dispersion problem with unknown Péclet number. The corresponding model is linear in the observed state $c(\boldsymbol{\mu})$, but it is non-linear in the parameter to estimate. The second problem concerns the transport of a contaminant in an unconfined aquifer with unknown hydraulic conductivity. It involves two coupled PDEs: a stationary non-linear equation which describes the pressure field induced by an external pumping force and a time-dependent linear equation describing the advection-dispersion of the contaminant in a medium whose properties depend non-linearly on the pressure field.

Both models describe 2D systems, and each exhibits ideal characteristics to test the proposed algorithms. The first, while leading to a non-linear inverse problem, is sufficiently simple to allow for a comparison between the adjusted and biased RB-EnKM and the reference full order EnKM. Moreover, its affine dependence on the parameter enables the use of error bounds for the efficient construction of the reduced space. The second problem, which is non-linear and non-affine in the six-dimensional parameter vector, is complex enough to serve as a non-trivial challenge for the proposed RB-EnKM algorithm, while the reference EnKM cannot even be tested due to the computational cost. From an *a priori* estimate, performing full order tests with the same statistical relevance as the reduced basis ones would have taken up to 20 days on our machine.

The two problems are presented in Sects. 4.1 and 4.2. We first introduce the pPDE, then present the full order discretization followed by the reduced basis approximation. The measurement operator is then introduced, and a first analysis of the inversion

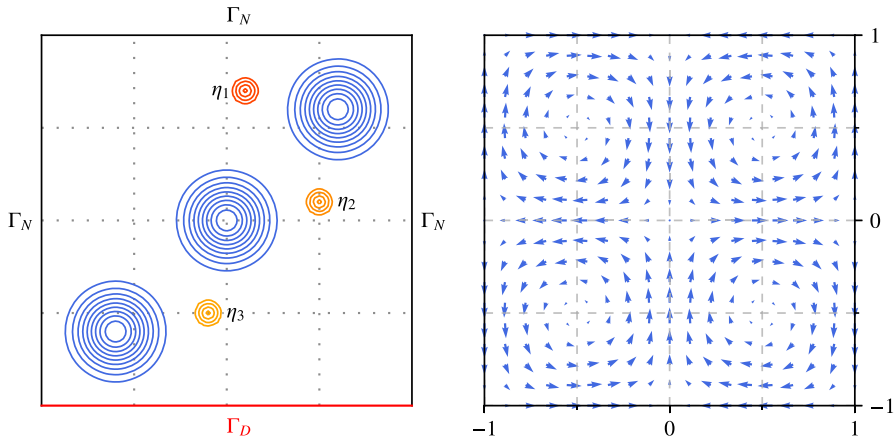


Fig. 1 Spatial domain of the Taylor–Green problem. On the left: the initial condition c_0 in blue, the sensor shape functions η_i , and the Neumann and Dirichlet boundaries, Γ_N , Γ_D . On the right, the velocity field, β , with four Taylor–Green vortices

method is carried out. Finally, we study the impact of the ensemble size, of the experimental noise magnitude, and of the error of the reduced model on the reconstruction error of the EnKM. All the computations are performed using Python on a computer with 2.20 GHz Intel Core i7-8750H processor and 32 GB of RAM.

4.1 Taylor–Green vortex problem

Let us consider the dispersion of a contaminant modeled by the 2D advection–diffusion equation with a Taylor–Green vortex velocity field (Kärcher et al. 2018). We introduce the spatial domain $\Omega = (-1, 1)^2$ with Dirichlet boundary $\Gamma_D := (-1, 1) \times \{-1\}$ and Neumann boundary $\Gamma_N := \partial\Omega \setminus \Gamma_D$, and the time domain $\mathcal{I} := (0, T]$ with $T = 2.5$. We consider the problem of estimating the inverse of the Péclet number $\mu = 1/Pe$ in the interval $\mathcal{P} := [1/50, 1/10]$. The governing pPDE is given by: find $c(\mu) : \Omega \times (0, T] \rightarrow \mathbb{R}$ such that

$$\begin{cases} \partial_t c - \mu \Delta c + \beta \cdot \nabla c = 0, & \text{in } \Omega \quad \times \mathcal{I}, \\ \nabla c(\mathbf{x}, t; \mu) \cdot \mathbf{n} = 0, & \text{on } \Gamma_N \quad \times \mathcal{I}, \\ c(\mathbf{x}, t; \mu) = 0, & \text{on } \Gamma_D \quad \times \mathcal{I}, \\ c(\mathbf{x}, 0; \mu) = c_0(\mathbf{x}; \mu), & \text{in } \Omega. \end{cases} \tag{25}$$

Here, the velocity field $\beta := (\sin(\pi x_1) \cos(\pi x_2), -\cos(\pi x_1) \sin(\pi x_2))^T$, $\mathbf{x} = (x_1, x_2)$, is a solenoidal field, and the initial condition $c_0(\mu) : \Omega \rightarrow \mathbb{R}$ is given by the sum of three Wendland functions $\psi_{2,1}$ (Wendland 1995) of radius 0.4 and centers located at $(-0.6, -0.6)$, $(0, 0)$, and $(0.6, 0.6)$. The velocity field and the initial condition are shown in Fig. 1.

The full order model is obtained by a nodal finite element discretization of (25) using piecewise continuous polynomial functions, $\zeta_i : \Omega \rightarrow \mathbb{R}$, $i = 1, \dots, N_h$, of degree 2




Fig. 2 Solution of the advection–diffusion equation for three increasing values of Pe at four time instances t . Snapshots normalized to unitary $L^\infty(\Omega)$ norm

over a uniform Cartesian grid of width $h = 0.04$, for a total of $N_h = 10,100$ degrees of freedom. The resulting system of ordinary differential equations is integrated over time using a Crank–Nicolson scheme with uniform time step $\Delta t = 0.01$. As shown in (Thomé 2006, Chapter 12), this is equivalent to performing a Petrov–Galerkin projection of (25) with trial and test spaces defined as follows: we consider the partition of the temporal interval \mathcal{I} into the union of equispaced subintervals, $\mathcal{I}_n := (t_{n-1}, t_n]$, of length Δt with $n = 1, \dots, N_t$ and $N_t := T/\Delta t$. Let $\omega_n : \mathcal{I} \rightarrow \mathbb{R}$ be a piecewise constant function with support in \mathcal{I}_n , and let $v_n : \mathcal{I} \rightarrow \mathbb{R}$ be a hat function with support in $\mathcal{I}_n \cup \mathcal{I}_{n+1}$. We define the trial space $\mathcal{V}_h := \text{span}\{v_n \cdot \zeta_i\}_{i,n=1}^{N_h, N_t}$ and the test space $\mathcal{W}_h := \text{span}\{\omega_n \cdot \zeta_i\}_{i,n=1}^{N_h, N_t}$, respectively.

To solve the spatial problems arising at each time step, we use the sparse `splu` function implemented in the `scipy.sparse.linalg`¹ package. The computational time to obtain a single full order solution is on average 0.56s. Snapshots of the solution at times $t \in \{0.2, 0.8, 1.4, 2.0\}$ and for the three parameter values $\mu \in \{1/10, 1/30, 1/50\}$ are shown in Fig. 2.

The high-fidelity model is used in combination with the time-gradient error bound $\Delta_R^{pr}(\mu)$ introduced in Aretz (2021) to implement a Weak-POD-Greedy algorithm for the selection of the reduced basis functions. To this end, we consider the training set Ξ_{TRAIN}^μ with parameters $\mu^{(s)} = 1/(9.5 + 0.5s)$ for all $s \in \mathbb{N} \cap [1, S]$ of size $S = 81$. We

¹ <https://docs.scipy.org/doc/scipy/reference/sparse.linalg.html>.

prescribe a target accuracy of 10^{-2} for the maximum time-gradient relative error bound and we obtain an RB space of size 42. We can construct surrogate models of different accuracy by selecting $N_\varepsilon \in \mathbb{N}$ basis functions $\psi_i : \Omega \rightarrow \mathbb{R}$, for $i = 1, \dots, N_\varepsilon$, out of these 42. Each choice corresponds to a relative error for the model given by

$$\varepsilon_c := \sup_{\mu \in \mathcal{D}} \frac{\|c_h(\mu) - c_\varepsilon(\mu)\|_{L^2(\mathcal{I}, H^1(\Omega))}}{\|c_h(\mu)\|_{L^2(\mathcal{I}, H^1(\Omega))}}. \tag{26}$$

Once the reduced basis has been computed, we construct a reduced model via a Petrov–Galerkin projection of (25) in the same way as we did for the full order model. For this purpose, we define the trial space $\mathcal{V}_\varepsilon := \text{span}\{\nu_n \otimes \psi_i\}_{i,n=1}^{N_\varepsilon, N_t}$ and the test space $\mathcal{W}_\varepsilon := \text{span}\{\omega_n \otimes \psi_i\}_{i,n=1}^{N_\varepsilon, N_t}$.

We then look for a reduced solution of the form

$$c_\varepsilon(\mu) = \sum_{i=1}^{N_\varepsilon} \sum_{n=1}^{N_t} c_{n,i}(\mu) \nu_n \psi_i, \tag{27}$$

where the expansion coefficients $c_{0,i}$, for $i = 1, \dots, N_\varepsilon$, result from the projection of the initial condition onto \mathcal{V}_ε , while the remaining coefficients $c_{n,i}$, with $i = 1, \dots, N_\varepsilon$ and $n = 1, \dots, N_t$, satisfy the equation

$$\sum_{j=1}^{N_\varepsilon} \left(\mathbf{M}_{ij} + \frac{\Delta t}{2} (\mathbf{A}_{ij} + \mu \mathbf{K}_{ij}) \right) c_{n,j} = \sum_{j=1}^{N_\varepsilon} \left(\mathbf{M}_{ij} - \frac{\Delta t}{2} (\mathbf{A}_{ij} + \mu \mathbf{K}_{ij}) \right) c_{n-1,j}. \tag{28}$$

Here the matrices $\mathbf{M}, \mathbf{K}, \mathbf{A} \in \mathbb{R}^{N_\varepsilon \times N_\varepsilon}$ denote the mass, stiffness, and advection matrix, respectively, and are given by

$$\mathbf{M}_{ij} := \int_{\Omega} \psi_j \psi_i \, d\Omega, \quad \mathbf{K}_{ij} := \int_{\Omega} \nabla \psi_j \cdot \nabla \psi_i \, d\Omega, \quad \mathbf{A}_{ij} := \int_{\Omega} (\boldsymbol{\beta} \cdot \nabla \psi_j) \psi_i \, d\Omega. \tag{29}$$

The solution of the system of equations (28), equivalent to a Crank–Nicolson scheme, can be obtained iteratively solving N_t linear systems of size N_ε for an online complexity $\mathcal{O}(N_\varepsilon^3 + N_t N_\varepsilon^2)$. This complexity can be achieved due to the time-independence of the pPDE by performing the LU factorization of the left-hand side before entering the time integration loop. Employing all $N_\varepsilon = 42$ basis functions, the computational time for a reduced basis solution (online cost) is on average 5.4ms, significantly less than the approximately 0.56s required for a full order solution. The acceleration achieved is over 100, which justifies the 47s necessary for the construction of the RB model (offline cost), considering that the online phase requires computing up to 150 reduced basis solutions per iteration. Let us remark that such a cheap training phase is due to the low-dimensionality of the parameters space \mathcal{P} and the availability of a tight error bound for this class of linear problems.

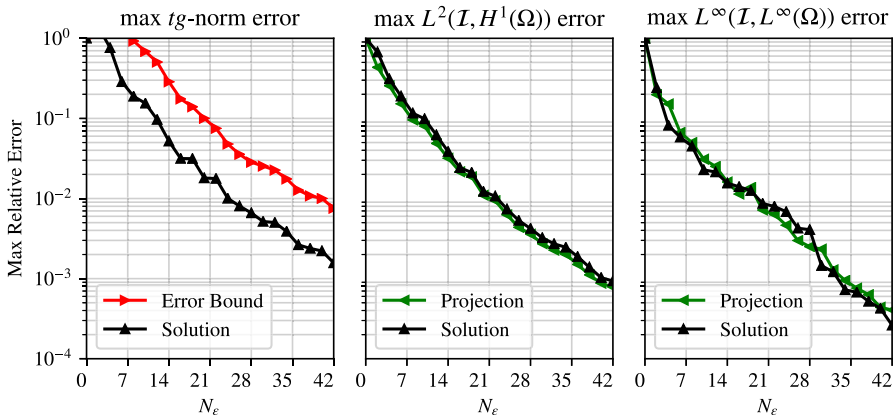


Fig. 3 Left: maximum relative time-gradient error and error bound of the advection–diffusion solution versus N_ϵ . Center: maximum $L^2(\mathcal{I}, H^1(\Omega))$ relative error of the projection and of the solution versus N_ϵ . Right: maximum $L^\infty(\mathcal{I}, L^\infty(\Omega))$ relative error of the projection and of the solution versus N_ϵ . Projections based on the $L^2(\Omega)$ inner product of the gradients

Note that both the online computational cost and the accuracy of the solution depend on Δt and on N_ϵ . The first is kept fixed, $\Delta t = 0.01$, while the latter varies in some of the experiments. In order to keep track of the error associated with different choices of N_ϵ , we proceed with the characterization of the error between the surrogate model solution $c_\epsilon(\mu)$ and the full model solution $c_h(\mu)$ for different values of N_ϵ . This analysis is provided in Fig. 3, depicting the maximum relative errors in $L^2(\mathcal{I}, H^1(\Omega))$, $L^\infty(\mathcal{I}, L^\infty(\Omega))$ and the time-gradient norm versus the reduced basis size. It shows a nearly exponential error decay as N_ϵ increases. The maxima are computed on an independent test set, $\Xi_{\text{TEST}}^\mu := \{1/(9.75 + 0.5s), \text{ for all } s \in \mathbb{N} \cap [1, 80]\}$. Furthermore, the reduced solutions do not appear to deviate significantly from the projection of their full order counterparts onto the associated RB space, and the error bound employed demonstrates a good effectivity.

For the implementation of the EnKM as presented in Sect. 2.1, it is necessary to provide a mathematical model for the measurement process. We take 40 measurements in time at the three sensor locations, $\eta_i, i \in \{1, 2, 3\}$, shown in Fig. 1. For this purpose, we introduce the measurement operator $\mathcal{L} : L^2(\mathcal{I}, L^2(\Omega)) \rightarrow \mathbb{R}^{120}$, which can be seen as a vector of linear functionals $\ell_k : L^2(\mathcal{I}, L^2(\Omega)) \rightarrow \mathbb{R}$ for all $k \in \mathbb{N} \cap [1, 120]$. Each of those linear functionals has a unique Riesz representer $\rho_k : \mathcal{I} \times \Omega \rightarrow \mathbb{R}$, with respect to the $L^2(\mathcal{I}, L^2(\Omega))$ norm, that can be written as

$$\rho_k = v_j \cdot \eta_i \quad \text{with } k = 3j + i \quad \text{for all } j \in \mathbb{N} \cap [1, 40], i \in \mathbb{N} \cap [1, 3],$$

where the spatial fields $\eta_i : \Omega \rightarrow \mathbb{R}$ are Wendland functions $\psi_{2,1}$ of radius 0.1 and center coordinates $(x_i, y_i) \in \{(0.1, 0.7), (-0.1, -0.5), (0.5, 0.1)\}$ (see Fig. 1), while, for each $j \in \mathbb{N} \cap [1, 40]$, $v_j : \mathcal{I} \rightarrow \mathbb{R}$ is a piecewise linear function supported over the interval $\mathcal{I}_j := [t_j - 2\Delta t, t_j + 2\Delta t]$, where $t_j := \Delta t(33 + 5j)$; v_j is assumed to be symmetric with respect to t_j and constant between $t_j - \Delta t$ and $t_j + \Delta t$.

Table 1 Comparison of reference FE (\cdot_h)—biased RB (\cdot_ε)—adjusted RB (\cdot_ε^*) EnKM in low-noise conditions $\sigma^2 = 10^{-6}$. The test was performed by averaging 25 estimations obtained employing ensembles of 150 particles and using reduced basis models of size $N_\varepsilon = 42$ ($\varepsilon_c \approx 0.001$). H refers to the mean of the estimation error, while S denotes the standard deviation of the estimation error. t.c. and o.c. indicate the total and online cost of one parameter estimation, respectively

Iter	H_h	S_h	H_ε	S_ε	H_ε^*	S_ε^*
0	$1.962e-2$	$2.168e-3$	$1.927e-2$	$1.740e-3$	$2.021e-2$	$1.931e-3$
1	$1.549e-6$	$1.039e-6$	$2.510e-4$	$2.584e-6$	$1.655e-5$	$1.737e-6$
2	$1.035e-7$	$6.405e-8$	$1.641e-5$	$2.336e-7$	$8.685e-7$	$1.335e-7$
3	$7.273e-8$	$5.152e-8$	$1.600e-5$	$1.317e-7$	$8.300e-7$	$6.766e-8$
4	$5.827e-8$	$3.835e-8$	$1.585e-5$	$1.048e-7$	$8.379e-7$	$6.469e-8$
5	$4.301e-8$	$2.840e-8$	$1.578e-5$	$8.631e-8$	$8.249e-7$	$5.630e-8$
t.c./o.c.	$6' 58''/6' 58''$		$55''/8''$		$1' 33''/8''$	

Given this description of the observation process and the surrogate model, we next test the data assimilation scheme. We start with the estimation of the unknown parameter $\mu^* = 0.04$ given the experimental measurements $\mathbf{y}(\mu^*, \boldsymbol{\eta}) \in \mathbb{R}^{120}$, with noise $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \Sigma)$. We compare the performances of the EnKM employing a full order model and a surrogate model of accuracy $\varepsilon_c = 10^{-3}$ with $N_\varepsilon = 42$. In order to obtain reliable statistics, we consider 25 ensembles \mathcal{E}_0 of size $J = 150$ with particles sampled from the uniform prior distribution, $\Pi_0(\mu) = U(0.02, 0.10)$. The results obtained for a fixed value of $\sigma^2 = 10^{-6}$, at different iterations of the algorithm, are shown in Table 1. We observe a quick stabilization of the error means H_h , H_ε and H_ε^* , and of the error covariances, S_h , S_ε and S_ε^* , after just a few steps. The full order algorithm performs significantly better than the biased reduced basis algorithm, while the adjusted version of the algorithm exhibits an excellent performance, very close to the full order one.

The comparison of the ensemble standard deviation, reported in Table 2, with the average error, reported in Table 1, shows a positive correlation between the two quantities in the reference and the adjusted case. Contrarily, it is clear a decorrelation between the two quantities in the biased case as the iteration index increases. From this observation, we infer that, at least in this case, the ensemble covariance can be used as an error indicator when the reference or the adjusted algorithm is employed.

We next investigate the sensitivity of the algorithm with respect to the accuracy of the reduced model, to the effect of the ensemble size, and to the noise magnitude. First, we repeat the estimation of the reference parameter $\mu^* = 0.04$ for different values of the ensemble size $J = 4k$, with $k \in \mathbb{N} \cap [1, 10]$. In this experiment, we employ the same surrogate model used before and consider the relative noise magnitude $\sigma / \|\mathcal{G}(\mu^*)\|_\infty = 10^{-3}$. The results, shown in Fig. 4, indicate a larger sensitivity to J for the full order algorithm than for the other two. It requires a larger number of particles before stabilizing on a large-ensemble asymptotic behavior (or mean-field behavior), while the reduced basis algorithms exhibit a much faster convergence, possibly as a consequence of a lower-dimensional state space. Among the three iterations

Table 2 Same experimental conditions as in Table 1. E refers to the ensemble mean, Σ to the ensemble standard deviation. Both quantities are computed as the average over 25 ensembles

Iter	E_h	Σ_h	E_ε	Σ_ε	E_ε^*	Σ_ε^*
0	0.0599106	$2.210e-2$	0.0592677	$2.339e-2$	0.0602145	$2.304e-2$
1	0.0399985	$1.740e-5$	0.0402510	$1.708e-5$	0.0400165	$2.076e-5$
2	0.0400001	$1.326e-6$	0.0400164	$1.332e-6$	0.0400009	$1.484e-6$
3	0.0400001	$9.354e-7$	0.0400160	$9.294e-7$	0.0400008	$1.051e-6$
4	0.0400000	$7.383e-7$	0.0400159	$7.571e-7$	0.0400008	$8.593e-7$
5	0.0400000	$6.357e-7$	0.0400158	$6.520e-7$	0.0400008	$7.472e-7$

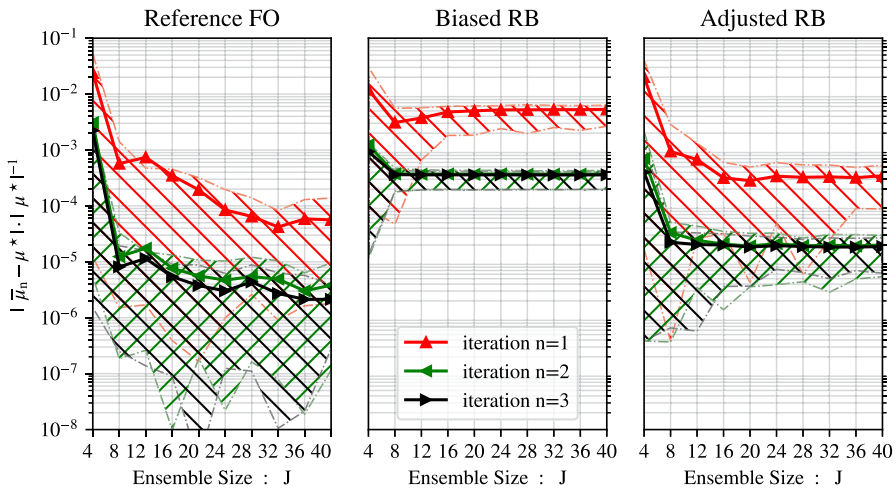


Fig. 4 Relative error in the parameter estimation versus ensemble size J for fixed noise magnitude, $\sigma = 10^{-3} \|\mathcal{G}(\mu^*)\|_\infty$. The standard full order EnKM is shown on the left, the biased RB-EnKM in the center, and the adjusted RB-EnKM on the right. The solid lines represent the average error over 64 ensembles, while the dashed lines correspond to the 10th and 90th percentiles

considered, the first appears to be the most affected, while, as the algorithm converges, the ensemble size seems to become less relevant.

In a second experiment, we consider the same parameter estimation, but we let the relative noise $\sigma / \|\mathcal{G}(\mu^*)\|_\infty$ take values 10^{-i} for $i \in \mathbb{N} \cap [2, 6]$. Moreover, we employ $J = 40$ particles per ensemble and the same reduced basis model as before. Each estimation is replicated 64 times for different noise realizations. The results are shown in Fig. 5: for the full order EnKM we observe a linear dependence between the reconstruction error and the experimental noise, while the results for the biased RB-EnKM show that an untreated model bias introduces a systematic error independent of the noise magnitude. The most important result is the one related to the adjusted RB-EnKM: the error behavior achieved with this algorithm is comparable with the one obtained using a full order model. This demonstrates the effectiveness of the proposed method in compensating for the bias introduced by the reduced basis model, at least in this case study.

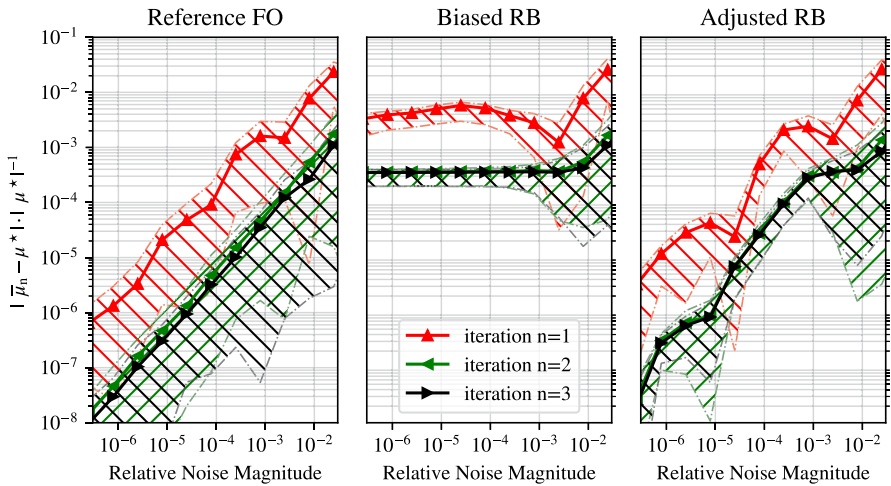


Fig. 5 Relative error in the parameter estimation versus relative noise magnitude $\sigma/\|\mathcal{G}(\mu^*)\|_\infty$ for fixed ensemble size $J = 40$. The standard full order EnKM is shown on the left, the biased RB-EnKM in the center, and the adjusted RB-EnKM on the right. The solid lines represent the average error over 64 ensembles, while the dashed lines correspond to the 10th and 90th percentiles

This conclusion is further confirmed by the last experiment, in which the performances of the biased and adjusted RB-EnKM are tested for all the parameters in the test set Ξ_{TEST}^μ already employed to test the reduced basis model. Each parameter in the set is estimated using surrogate models of increasing sizes. Each estimation is performed 64 times in very low-noise conditions, that is $\sigma/\|\mathcal{G}(\mu^*)\|_\infty = 10^{-5}$, employing $J = 40$ particles per ensemble. For each surrogate model employed, the results from the 64 ensembles are averaged and the maximum over the test set is computed. The results, shown in Fig. 6, demonstrate the ability of the correction to compensate for the presence of a model bias very well. As a consequence, the worst-case reconstruction error for the adjusted RB-EnKM barely depends on the reduced model size and it is always significantly lower than its biased counterpart. These results confirm the good performance of the adjusted RB-EnKM and its superiority over the biased RB-EnKM.

4.2 Tracer transport problem

We now consider the tracer transport problem from Conrad et al. (2018), describing the non-homogeneous and non-isotropic transport of a non-reactive tracer in an unconfined aquifer. We introduce the spatial domain $\Omega := (0, 1)^2$ divided into six sub-regions $\Omega = \bigcup_{r=1}^6 \Omega_r$ illustrated in Fig. 7 and defined as follows: $(x, y) \in \Omega$ is in Ω_r if the subscript r is the smallest integer for which $x_0^r < x < x_1^r$ and $y_0^r < y < y_1^r$ where the points $\{(x_0^r, y_0^r)\}_{r=1}^6$ and $\{(x_1^r, y_1^r)\}_{r=1}^6$ are defined in Table 3. We denote by $\partial\Omega$ the outer boundary of the domain and define the parallel walls $\Gamma_D := (0, 1) \times \{0, 1\}$ and $\Gamma_N := \partial\Omega \setminus \Gamma_D$. Based on this partition, we define the conductivity field as the piecewise constant function $k(\mu) : \Omega \rightarrow \mathbb{R}$ over the six sub-regions Ω_r . The




Fig. 6 Parameter error versus reduced basis size for the biased and the adjusted RB-EnKM

conductivity can be affinely decomposed employing the coefficient vector $\boldsymbol{\mu} \in \mathbb{R}^6$, with components μ_r , and the indicator functions $\eta_r : \Omega \rightarrow \mathbb{R}$

$$k(\mathbf{x}; \boldsymbol{\mu}) = \sum_{r=1}^6 e^{\mu_r} \eta_r(\mathbf{x}) \quad \text{with } \eta_r(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in \Omega_r, \\ 0 & \text{if } \mathbf{x} \in \Omega \setminus \Omega_r. \end{cases} \quad (30)$$

We can now estimate the hydraulic log-conductivity $\boldsymbol{\mu}$, restricted to the orthotope $\mathcal{D} := \times_{r=1}^6 (\mu_r^{\min}, \mu_r^{\max}) \subset \mathbb{R}^6$, relying on measurements of the tracer concentration $c(\boldsymbol{\mu})$ collected over the time interval $\mathcal{I} := (0, T]$, with $T = 0.5$. This field satisfies the pPDE: find $c(\boldsymbol{\mu}) : \Omega \times \mathcal{I} \rightarrow \mathbb{R}$ such that

$$\begin{cases} \partial_t c - \nabla \cdot ((d_m \mathbf{I} + d_l \boldsymbol{\beta} \boldsymbol{\beta}^\top(\boldsymbol{\mu})) \nabla c) + \boldsymbol{\beta}(\boldsymbol{\mu}) \cdot \nabla c = f_c, & \text{in } \Omega \times \mathcal{I}, \\ \nabla c(\mathbf{x}, t; \boldsymbol{\mu}) \cdot \mathbf{n} = 0, & \text{on } \partial\Omega \times \mathcal{I}, \\ c(\mathbf{x}, 0; \boldsymbol{\mu}) = 0, & \text{in } \Omega. \end{cases} \quad (31)$$

In this equation, the dispersion coefficients $d_l = d_m = 2.5 \cdot 10^{-3}$ correspond to the flow-dependent component of the dispersion tensor and to its residual component, respectively. The forcing term f_c is assumed to be of the form $f_c := \sum_{i=1}^4 f_{c,i}$ and it models the injection of different amounts of tracer in four wells located at $(a_i, b_i) \in \{0.15, 0.85\}^2$; each $f_{c,i}$ is a Gaussian function centered in (a_i, b_i) , with covariance $\Gamma_c = 0.005$ and multiplicative coefficient p_i where $(p_1, p_2, p_3, p_4) = (10, 5, 10, 5)$. The velocity field $\boldsymbol{\beta}(\boldsymbol{\mu}) : \Omega \rightarrow \mathbb{R}^2$ is linearly dependent on the hydraulic head $u(\boldsymbol{\mu}) : \Omega \rightarrow \mathbb{R}$ through the relation $\boldsymbol{\beta}(\boldsymbol{\mu}) = -k(\boldsymbol{\mu}) \nabla u$. The latter field must satisfy the second constraint of the inverse problem, i.e., under the Dupuit–Forchheimer approximation

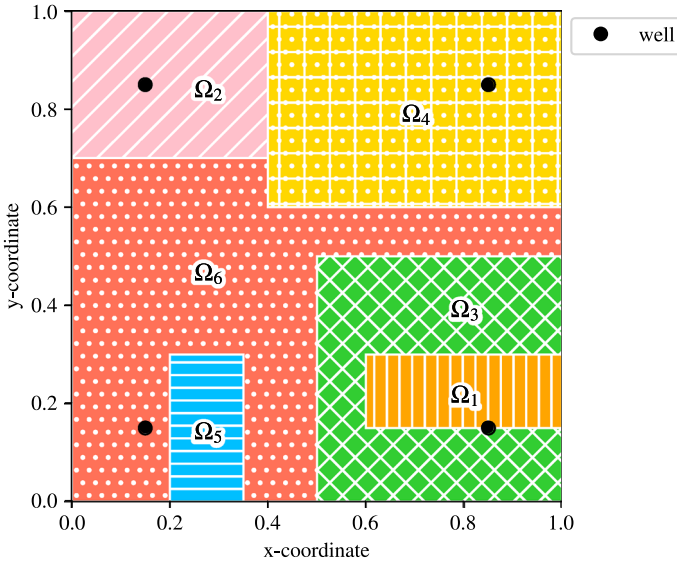


Fig. 7 Domain of the tracer transport problem and injection wells

Table 3 On the left: coordinates of the corners of the sub-regions Ω_r . On the right: true values of the parameters μ_r and boundaries of the uniform prior Π_0

Region	x_0^r	x_1^r	y_0^r	y_1^r	(*) Ref	Min	Max	μ_r
Ω_1	0.60	1.00	0.15	0.30	-0.75	-1.00	0.00	μ_1
Ω_2	0.00	0.40	0.70	0.10	-0.25	-1.00	1.00	μ_2
Ω_3	0.50	1.00	0.00	0.50	-0.50	-1.00	0.00	μ_3
Ω_4	0.40	1.00	0.60	1.00	1.00	0.00	2.00	μ_4
Ω_5	0.20	0.25	0.00	0.30	-0.25	-1.00	0.00	μ_5
Ω_6	0.00	1.00	0.00	1.00	3.00	2.00	5.00	μ_6

(Delleur 2016) it solves the non-linear elliptic pPDE: find $u(\boldsymbol{\mu}) : \Omega \rightarrow \mathbb{R}$ such that

$$\begin{cases} \nabla \cdot (k(\boldsymbol{\mu})u\nabla u) + f_u = 0, & \text{in } \Omega, \\ \nabla u(\mathbf{x}; \boldsymbol{\mu}) \cdot \mathbf{n} = 0, & \text{on } \Gamma_N, \\ u(\mathbf{x}; \boldsymbol{\mu}) = 0, & \text{on } \Gamma_D. \end{cases} \quad (32)$$

Here, the forcing term $f_u := \sum_{i=1}^4 f_{u,i}$ models an active pumping action at the four wells, each $f_{u,i}$ is a Gaussian function centered in (a_i, b_i) , of covariance $\Gamma_u = 0.02$ and coefficient q_i , where $(q_1, q_2, q_3, q_4) = (10, 50, 150, 50)$. Due to the combination of the quadratic dependence in u and the zero boundary conditions, the equation always admits pairs of opposite solutions u^+, u^- . However, in our study, we are only interested in the positive solution $u^+(\boldsymbol{\mu}) : \Omega \rightarrow \mathbb{R}^+$.

Full order solutions are obtained via a finite element approximation, employing piecewise linear functions, $\zeta_i : \Omega \rightarrow \mathbb{R}$, for $i = 1, \dots, N_h$, with $N_h = 44,972$




Fig. 8 Reference solutions at log-conductivity $\mu^* := [-0.75, -0.25, -0.5, 1, -0.25, 3]$. On the left: hydraulic head $u_h(\mu^*)$ and corresponding velocity field $\beta_h := -k(\mu^*)\nabla u_h$ (in red). On the right: tracer concentration $c_h(\mu^*)$, at time $t = 0.4$, and measurement wells (in red)

degrees of freedom (mesh size $h \approx 0.01$). The discretization of the elliptic equation (32) results in a discrete non-linear problem which is iteratively solved employing a Newton scheme with tolerance 10^{-6} . The approximate solution, u_h , is used to compute the velocity field, $\beta_h(\mu) := -k(\mu)\nabla u_h$, which is piecewise constant with $N_h - 1$ degrees of freedom. This is needed for the solution of the parabolic equation (31), whose discretization leads to a system of ordinary differential equations integrated over the time interval \mathcal{I} using the Crank–Nicolson scheme with uniform time step $\Delta t = 0.01$. This is equivalent to performing a Petrov–Galerkin projection of Equation (31), analogously to what have been shown for Equation (25) in Sect. 4.1.

Each full order simulation is obtained employing a FreeFEM++ solver (Hecht 2012) and takes roughly 2 min to be computed. Figure 8 shows the hydraulic head $u_h(\mu^*)$ and the relative velocity field $\beta_h(\mu^*)$ (on the left) and the tracer concentration field $c_h(0.4; \mu^*)$ (on the right), both associated with the reference log-conductivity

$$\mu^* = [-0.75, -0.25, -0.50, 1.00, -0.25, 3.00]^\top. \tag{33}$$

The same reference log-conductivity is used as the true parameter for the data assimilation problem. Pointwise observations are collected at five successive times $t_m \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, in 25 spatial location $\mathbf{x}_{ij} = (x_i, y_j)$ such that $x_i = 0.1 + 0.2i$ and $y_j = 0.1 + 0.2j$ for $i, j \in \{0, \dots, 4\}$. This operation is encoded in the measurement operator $\mathcal{L} : H^1(\Omega) \rightarrow \mathbb{R}^{125}$. Each noise-free measurement is polluted with i.i.d. Gaussian noise with mean zero and covariance σ , resulting in a noise covariance matrix $\Sigma = \sigma^2 \mathbf{I}$.

In order to solve the inverse problem with surrogate models of different accuracy, various approximations of (32) and (31) must be produced. This requires the introduction of spatial basis functions $\psi_i, \varphi_j : \Omega \rightarrow \mathbb{R}, i \in \mathbb{N} \cap [1, N_\varepsilon], j \in \mathbb{N} \cap [1, M_\varepsilon]$, selected by applying the method of snapshots (POD) to the two sets of full order solutions,

$$\Theta_{\text{TRAIN}}^u := \{u_h(\boldsymbol{\mu}^{(s)})\}_{s=1}^S \quad \text{and} \quad \Theta_{\text{TRAIN}}^c := \{c_h(t^{(z)}; \boldsymbol{\mu}^{(s)})\}_{z,s=1}^{Z,S},$$

with snapshot parameters, $\boldsymbol{\mu}^{(s)} \sim \Pi_0 := \times_{r=1}^6 U(\mu_r^{\min}, \mu_r^{\max})$, for all $s \in \mathbb{N} \cap [1, S]$ and sampling times $t^{(z)} = 0.01z$ for all $z \in \mathbb{N} \cap [1, Z]$, where $S = 2,000$ and $Z = 50$. The number of basis functions considered, $N_\varepsilon, M_\varepsilon \in \mathbb{N}$, is the one required to approximate the hydraulic head and the tracer concentration with relative accuracy $\varepsilon_u, \varepsilon_c \in \mathbb{R}^+$, where

$$\varepsilon_u := \sup_{\boldsymbol{\mu} \in \mathcal{D}} \frac{\|u_h(\boldsymbol{\mu}) - u_\varepsilon(\boldsymbol{\mu})\|_{H^1(\Omega)}}{\|u_h(\boldsymbol{\mu})\|_{H^1(\Omega)}}, \tag{34}$$

$$\varepsilon_c := \sup_{\boldsymbol{\mu} \in \mathcal{D}} \frac{\|c_h(\boldsymbol{\mu}) - c_\varepsilon(\boldsymbol{\mu})\|_{L^2(\mathcal{I}, H^1(\Omega))}}{\|c_h(\boldsymbol{\mu})\|_{L^2(\mathcal{I}, H^1(\Omega))}}. \tag{35}$$

Based on the first set of basis functions, the approximation space for the Galerkin projection of (32) is defined as $\mathcal{U}_\varepsilon := \text{span}\{\psi_i\}_{i=1}^{N_\varepsilon}$. From the second set of basis functions, instead, the RB test space $\mathcal{W}_\varepsilon := \text{span}\{\omega_n \otimes \varphi_i\}_{i,n=1}^{M_\varepsilon, N_t}$ and RB trial space $\mathcal{V}_\varepsilon := \text{span}\{\nu_n \otimes \varphi_i\}_{i,n=1}^{M_\varepsilon, N_t}$ are defined for the Petrov–Galerkin projection of (31). We look at reduced solutions of the form

$$u_\varepsilon(\boldsymbol{\mu}) = \sum_{i=1}^{N_\varepsilon} u_i(\boldsymbol{\mu}) \psi_i \tag{36}$$

$$c_\varepsilon(\boldsymbol{\mu}) = \sum_{j=1}^{M_\varepsilon} \sum_{n=1}^{N_t} c_{n,j}(\boldsymbol{\mu}) \nu_n \varphi_j, \tag{37}$$

where the expansion coefficients $c_{n,j}$ and u_i , with $i \in \mathbb{N} \cap [1, N_\varepsilon]$, $n \in \mathbb{N} \cap [1, N_t]$ and $j \in \mathbb{N} \cap [1, M_\varepsilon]$, satisfy the systems of algebraic equations

$$\sum_{p,q=1}^{N_\varepsilon, N_\varepsilon} \mathbf{N}_{ipq}(\boldsymbol{\mu}) u_p u_q = f_i, \tag{38}$$

$$\sum_{k=1}^{M_\varepsilon} \left(\mathbf{M}_{jk} + \frac{\Delta t}{2} \mathbf{D}_{jk}(\mathbf{u}, \boldsymbol{\mu}) \right) c_{n+1,k} = \left(\mathbf{M}_{jk} - \frac{\Delta t}{2} \mathbf{D}_{jk}(\mathbf{u}, \boldsymbol{\mu}) \right) c_{n,k} + g_j, \tag{39}$$

given the initial conditions $c_{0,j} = 0$ for all $j \in \mathbb{N} \cap [1, M_\varepsilon]$. The scalar forcing terms f_i, g_j are obtained by integrating their full order counterparts versus the basis functions ψ_i and φ_j , for all $i \in \mathbb{N} \cap [1, N_\varepsilon]$, $j \in \mathbb{N} \cap [1, M_\varepsilon]$

$$f_i := \int_{\Omega} f_h \psi_i d\Omega, \quad g_j := \Delta t \int_{\Omega} f_c \varphi_j d\Omega. \tag{40}$$

The mass and stiffness matrices $\mathbf{M}, \mathbf{K} \in \mathbb{R}^{M_\varepsilon}$ are defined as in (29), while the parameter dependent tensors $\mathbf{D}(\mathbf{u}, \boldsymbol{\mu}) \in \mathbb{R}^{M_\varepsilon^2}$ and $\mathbf{N}(\boldsymbol{\mu}) \in \mathbb{R}^{M_\varepsilon^3}$ depend affinely on the

multidimensional arrays $\mathbf{A} \in \mathbb{R}^{6 \times N_\varepsilon^3}$, $\mathbf{B} \in \mathbb{R}^{6 \times N_\varepsilon^2 \times M_\varepsilon^2}$, and $\mathbf{C} \in \mathbb{R}^{6 \times N_\varepsilon \times M_\varepsilon^2}$ defined as

$$\mathbf{A}_{ipqr} := \int_{\Omega} \frac{\eta_r}{2} (\psi_p(\nabla \psi_q \cdot \nabla \psi_i) + \psi_q(\nabla \psi_p \cdot \nabla \psi_i)) d\Omega, \tag{41}$$

$$\mathbf{B}_{jkpqr} := \int_{\Omega} \eta_r (\nabla \varphi_j \cdot \nabla \psi_p)(\nabla \varphi_k \cdot \nabla \psi_q) d\Omega, \tag{42}$$

$$\mathbf{C}_{jksr} := \int_{\Omega} \eta_r (\nabla \varphi_j \cdot \nabla \psi_s) \varphi_k d\Omega. \tag{43}$$

For a fixed value of the log-conductivity, $\boldsymbol{\mu}$, the tensors $\mathbf{N}(\boldsymbol{\mu})$ and $\mathbf{D}(\mathbf{u}, \boldsymbol{\mu})$ can be assembled. The latter, however, requires the evaluation of the discrete hydraulic head \mathbf{u} . They are respectively defined as

$$\mathbf{N}_{ipq}(\boldsymbol{\mu}) := \sum_{r=1}^6 e^{\mu_r} \mathbf{A}_{ipqr}, \tag{44}$$

$$\mathbf{D}_{jk}(\mathbf{u}, \boldsymbol{\mu}) := d_m \mathbf{K}_{jk} + d_l \sum_{p,q,r=1}^{N_\varepsilon, N_\varepsilon, 6} e^{2\mu_r} \mathbf{B}_{jkpqr} u_p u_q + \sum_{s,r=1}^{N_\varepsilon, 6} e^{\mu_r} \mathbf{C}_{jksr} u_s. \tag{45}$$

We emphasize that the accuracy of the solutions of (38) and (39), with the latter equivalent to a Crank–Nicolson discretization, depends on the number of basis functions and on the time step Δt . In Fig. 9, on the left and on the right, we show the maximum relative errors of the surrogate model ($\varepsilon_u, \varepsilon_c$) as a function of N_ε and M_ε . In the center, we show the $L^\infty(\mathcal{I}; L^\infty(\Omega))$ relative error of the tracer concentration, bounding from above the error on synthetic measurements. We compute these maximum relative errors on a set of parameters $\Xi_{\text{TEST}}^\mu := \{\boldsymbol{\mu}^{(s)} \sim \Pi_0(\boldsymbol{\mu})\}_{s=1}^{500}$ independent of the ones used for the model training. It can be observed that, for small values of N_ε , the error in the concentration stagnates after a certain value of M_ε , suggesting that, in this region, the error is dominated by the approximation of the hydraulic head. However, for $N_\varepsilon = 40$, this effect is no longer present, at least for the values of M_ε considered, and the tracer error only depends on M_ε . This allows us to modify the accuracy of the model by varying the dimension of the reduced model.

The construction of the reduced model has an offline cost of about 75 h. This includes the time required for the construction of a training set of 2,000 full order solutions (61h 16' 40''), the time for the computation of the POD basis functions (23' 40''), and the time for assembling the RB model tensors (13h 32' 47''). This cost corresponds roughly to the computational cost of 2,500 finite element solutions, each of which takes approximately 110s. The surrogate model obtained employing $N_\varepsilon = 40$, $M_\varepsilon = 320$ basis functions produces a solution in only 1.25s (online cost), which is about 1/90 of its full order equivalent. The same training set used for the POD is employed to estimate, at negligible cost, the empirical moments of δ_ε^* , i.e., $\bar{\delta}_\varepsilon$ and Γ_ε .

Note that we admittedly used a ‘‘brute force’’ POD approach to generate the basis for this nonlinear problem. A more offline-efficient method, for example using a POD-Greedy algorithm, could have been used. However, this would have been more complex in terms of both theory and implementation of the reduced model, and is beyond the




Fig. 9 Left and center: maximum relative error of the solution and of the projection of the tracer concentration versus M_ϵ for different values of N_ϵ ; projection—in space—performed with respect to the $H^1(\Omega)$ inner product. Right: maximum relative error of the projection and of the solution of the hydraulic head versus N_ϵ . Error norm shown above each plot

scope of this paper. Our focus here is on the EnKM and its modification in settings where surrogate models are used.

We now turn our attention to the inverse problem, as discussed in Sect. 2.1. We start by considering the estimation of the reference parameter μ^* given the measurements $\mathbf{y}(\mu^*, \eta) \in \mathbb{R}^{125}$, polluted by experimental noise of magnitude σ . To have a reliable statistic, we consider 32 independent initial ensembles \mathcal{E}_0 of variable size, sampled from the same distribution Π_0 .

As a first experiment, we compare the performances of the two RB-EnKM employing $J = 160$ particles and a surrogate model with error tolerance $\epsilon_c \approx 0.02$ (obtained with $N_\epsilon = 40$ and $M_\epsilon = 320$). The first test relies on the biased version of the RB-EnKM, as presented in Sect. 2.1, while the second test corresponds to the adjusted algorithm. For both simulations, we consider low-amplitude experimental noise, i.e., negligible if compared to the model error, $\sup_\mu \|\mathcal{L}(c_h(\mu) - c_\epsilon(\mu))\|_\infty \approx 10^{-2} > 10^{-3} = \sigma$, and we separately pollute the measurements employed by the ensembles. In Table 4, we report the average properties of the ensembles after 4 iterations: columns E_ϵ and E_ϵ^* contain the mean parameter estimation, i.e., the particle mean, averaged over the 32 ensembles. Here, columns Σ_ϵ and Σ_ϵ^* contain the average standard deviation of the ensembles. We can observe that the correction term has the effect of significantly lowering the reconstruction error from $\|E_\epsilon - \mu^*\|_\infty = 6.437e-3$ to $\|E_\epsilon^* - \mu^*\|_\infty = 7.870e-4$. We also notice that the variability of the estimate increases consistently with the presence of an additional term in the Kalman gain. At least in this test, and contrary to the previous numerical experiment, the standard deviation of the ensemble shows a good correlation with the reconstruction error for both the biased and adjusted algorithm.

We note, for this example, that the total cost of estimating a single parameter with the biased or the adjusted RB-EnKM is higher than with the FO-EnKM (approx. 75h 28' for the former and 19h 37' for the latter). This is due to the six-dimensional

Table 4 Comparison of biased RB (\cdot_ϵ)—adjusted RB (\cdot_ϵ^*) EnKM in low-noise conditions $\sigma = 10^{-6}$. The test was performed by averaging 32 estimations obtained employing ensembles of 160 particles and 4 iterations, and using reduced basis models of size $N_\epsilon = 40$, $M_\epsilon = 320$ ($\epsilon_C \approx 0.02$). E refers to the average parameter estimation, while Σ denotes the average ensemble standard deviation, and H the average estimation error. t.c. and o.c. indicate respectively the total and online cost of one parameter estimation

	E_ϵ	Σ_ϵ	H_ϵ	E_ϵ^*	Σ_ϵ^*	H_ϵ^*	μ^*
μ_1	-0.75189	1.261e-3	1.970e-3	-0.75079	1.543e-3	1.385e-3	-0.75
μ_2	-0.25644	1.860e-3	6.437e-3	-0.25040	2.785e-3	2.251e-3	-0.25
μ_3	-0.50313	1.001e-3	3.134e-3	-0.50040	1.353e-3	1.130e-3	-0.50
μ_4	0.99904	0.701e-3	1.018e-3	0.99976	1.016e-3	0.836e-3	1.00
μ_5	-0.25111	0.898e-3	1.276e-3	-0.25033	1.221e-3	1.014e-3	-0.25
μ_6	3.00013	0.630e-3	0.516e-3	2.99967	0.904e-3	0.775e-3	3.00
t.c.	75h 27' 39''			75h 27' 36''			
o.c.	13' 32''			13' 29''			

parameter space and the fact that the parameters are not highly correlated. We thus require a fairly large training set of size 2,000 to obtain a sufficiently accurate reduced order model over the whole parameter space. In combination with the “brute force” POD approach—as mentioned above—the offline cost is thus considerable. However, if one is interested in *multiple* parameter estimations, e.g., due to new data being available, the RB-EnKM algorithm significantly outperforms the FO-EnKM in terms of computational runtime since the online phase requires only 13' 32''. For example, repeating the parameter estimation 32 times in order to obtain a better statistical characterization of the method takes about 82h using the reduced basis method, but it would require more than 627h using the FO-EnKM.

As an extension of the previous experiment, we estimate the reference parameter μ^* employing the same surrogate model, noise magnitude and number of ensembles as before, but using ensembles of variable size $J = 20k$, with $k \in \mathbb{N} \cap [2, 16]$. This allows us to study the effect of the ensemble size on the parameter estimation obtained with the biased and adjusted RB-EnKM algorithms. The results shown in Fig. 10 indicate that, for both algorithms, very small ensembles lead to large relative errors and entail a large variability among the different samples. This behavior seems to be relevant only for ensembles with less than 40 particles when the biased RB-EnKM is employed, and with less than 80 particles when the adjusted version is used. Larger ensembles do not exhibit relevant fluctuations; we can therefore assume an ensemble of size $J = 160$ to be sufficiently large to ensure the independence from this quantity of the results in the upcoming tests.

A key quantity determining the performances of the method is the noise magnitude. Its effect on the two reduced basis algorithms is investigated by looking at the variation of the relative estimation error of the reference parameter μ^* when the noise magnitude varies. To this end, we consider seven noise values, $\sigma^2 = 10^{-m}$ with $m \in \mathbb{N} \cap [1, 7]$. We employ the same RB-EnKM used before, with a fixed ensemble size $J = 160$, and we average the results over 32 independent ensembles. The results, shown in Fig. 11, reiterate the inadequacy of the biased method in dealing with the systematic bias




Fig. 10 Biased and adjusted RB-EnKM parameter estimation relative error versus ensemble size, for fixed noise magnitude $\sigma = 10^{-3}$. The solid lines represent the average error over 32 ensembles at different algorithm iterations. The dashed lines represent the 10th and the 90th percentiles




Fig. 11 Biased and adjusted RB-EnKM parameter estimation relative error versus absolute noise magnitude, for fixed ensemble size $J = 160$. The solid lines represent the average error over 32 ensembles at different algorithm iterations. The dashed lines represent the 10th and the 90th percentiles

introduced in the measurements by the surrogate model. In fact, the plot corresponding to the biased method shows error stagnation for low-noise. On the contrary, the plot corresponding to the adjusted method highlights a mitigation of this effect, with an estimation error that keeps decreasing in low-noise conditions, although at a lower rate than in high-noise conditions.

In our last experiment, we test the performances of the biased and adjusted RB-EnKM by employing surrogate models of increasing accuracy. We fix the size of the reduced space \mathcal{U}_ε to a sufficiently large value, $N_\varepsilon = 40$, and we vary the size of the approximation space associated with the concentration: $M_\varepsilon = 10k$, with $k \in \mathbb{N} \cap [2, 32]$. Employing the resulting approximated models, we estimate the reference parameter μ^* in low-noise conditions, $\sigma = 10^{-3}$, averaging the results obtained over




Fig. 12 Parameter error versus reduced basis size and maximum relative error of the solution. The solid lines represent the average error over 16 ensembles at different algorithm iterations. The dashed lines represent the 10th and the 90th percentiles

16 ensembles of 160 particles each. In Fig. 12, we show the final relative error (after three algorithm iterations) as M_ϵ and ϵ_c change, both for the biased and the adjusted RB-EnKM. For both, we observe that the relative estimation error decreases, almost linearly, with the error of the surrogate model. Moreover, we observe that, with few exceptions, the error of the adjusted algorithm is smaller than the error of the biased algorithm. The few points where the two errors are very close can be explained by a strongly unbalanced distribution of the measurement bias in a region away from the reference parameter. Future developments that take into account, in the execution of the algorithm, the parameter estimate to adjust the bias correction should dampen this effect.

5 Conclusions

We proposed an efficient, gradient-free iterative solution method for inverse problems that combines model order reduction techniques, via the reduced basis method, and the Kalman ensemble method introduced in Iglesias et al. (2013). The use of surrogate models allows a significant speed-up of the online computational cost, but it leads to a distortion in the cost function optimized by the inverse problem. This in turn introduces a systematic error in the approximate solution of the inverse problem. To overcome this limitation, we have proposed the adjusted RB-EnKM which corrects for this bias by systematically adjusting the cost function and thus retrieving good convergence.

Using a linear Taylor–Green vortex problem, the performance of the method is compared versus the full order model as well as to the biased RB-EnKM in which no adjustment was made. The numerical results show that the biased method fails to achieve the same accuracy as the full order method. Contrarily, the adjusted RB-

EnKM attains the same accuracy as its full order counterpart for a large range of noise magnitudes at a significantly lower computational cost, and even approaches the mean-field limit faster as the ensemble size is increased. Furthermore, the dependence on model accuracy of the reconstruction error is essentially removed over the range of model accuracy considered.

The method was then applied to a non-linear tracer transport problem. The results for this example show that, despite a decrease in the order of convergence at low-noise, the stagnation of the reconstruction error observed in the biased RB-EnKM can be removed by adjusting the algorithm. Regarding the model accuracy, a substantial improvement of the adjusted EnKM with respect to the biased EnKM was observed, although less pronounced than in the linear problem.

Overall, our numerical tests show that the proposed method allows for the use of inexpensive surrogate models while empirically ensuring that the predicted result of the inversion remains accurate with respect to the full order inversion.

Although the online computational cost is significantly lower than the reference full order method, we do note that—depending on the problem at hand and the implementation—the offline cost can be considerable. As a result, the overall cost (offline plus online parameter inversion) for solving a single inversion problem may not be competitive with a plain full order inversion as observed in the second case study. However, if we consider the solution of multiple inverse problems, either due to new data being analyzed or in order to obtain a better statistics, the method becomes competitive also for the second numerical experiment considered.

Funding The research leading to these results received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 818473).

Declarations

Conflict of interest The authors have no conflict of interest to declare.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson, J.L., Anderson, S.L.: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review* **127**(12), 2741–2758 (1999). [https://doi.org/10.1175/1520-0493\(1999\)127<2741:AMCIOT>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<2741:AMCIOT>2.0.CO;2)
- Anderson, J.L.: An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review* **129**(12), 2884–2903 (2001). [https://doi.org/10.1175/1520-0493\(2001\)129<2884:AEAKFF>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<2884:AEAKFF>2.0.CO;2)
- Anderson, B.D.O., Moore, J.B.: *Optimal Filtering*, 1st edn. Prentice-Hall, Englewood Cliffs (1979)

- Aretz, N.: Data assimilation and sensor selection for configurable forward models: Challenges and opportunities for model order reduction methods. PhD thesis, IRTG-2379, RWTH Aachen, Germany (2021)
- Asch, M., Bocquet, M., Nodet, M.: Data Assimilation. Society for Industrial and Applied Mathematics, Philadelphia (2016)
- Barraut, M., Maday, Y., Nguyen, N.C., Patera, A.T.: An ‘empirical interpolation method’: application to efficient reduced-basis discretization of partial differential equations. *C. R. Math.* **339**(9), 667–672 (2004). <https://doi.org/10.1016/j.crma.2004.08.006>
- Berkooz, G., Holmes, P., Lumley, J.L.: The proper orthogonal decomposition in the analysis of turbulent flows. *Ann. Rev. Fluid Mech.* **25**, 539–575 (1993). <https://doi.org/10.1146/annurev.fl.25.010193.002543>
- Calvetti, D., Dunlop, M., Somersalo, E., Stuart, A.: Iterative updating of model error for Bayesian inversion. *Inverse Prob.* **34**(2), 025008–38 (2018). <https://doi.org/10.1088/1361-6420/aaa34d>
- Chen, Y., Oliver, D.S.: Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Math. Geosci.* **44**(1), 1–26 (2012). <https://doi.org/10.1007/s11004-011-9376-z>
- Conrad, P.R., Davis, A.D., Marzouk, Y.M., Pillai, N.S., Smith, A.: Parallel local approximation MCMC for expensive models. *SIAM/ASA J. Uncertain. Quantif.* **6**(1), 339–373 (2018). <https://doi.org/10.1137/16M1084080>
- da Silva, A.F.C., Colonius, T.: Ensemble-based state estimator for aerodynamic flows. *AIAA J.* **56**(7), 2568–2578 (2018). <https://doi.org/10.2514/1.J056743>
- Delleur, J.W.: Elementary Groundwater Flow and Transport Processes, pp. 73–102. CRC Press, Boca Raton (2016)
- Donoghue, G., Yano, M.: A multi-fidelity ensemble kalman filter with hyperreduced reduced-order models. *Computer Methods in Applied Mechanics and Engineering* **398**, 115282 (2022). <https://doi.org/10.1016/j.cma.2022.115282>
- Evensen, G.: The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dyn.* **53**, 343–367 (2003). <https://doi.org/10.1007/s10236-003-0036-9>
- Evensen, G.: Analysis of iterative ensemble smoothers for solving inverse problems. *Comput. Geosci.* **22**(3), 885–908 (2018). <https://doi.org/10.1007/s10596-018-9731-y>
- Gao, H., Wang, J.-X.: A bi-fidelity ensemble Kalman method for PDE-constrained inverse problems in computational mechanics. *Comput. Mech.* **67**(4), 1115–1131 (2021). <https://doi.org/10.1007/s00466-021-01979-6>
- Gong, H., Maday, Y., Mula, O., Taddei, T.: PBDW method for state estimation: error analysis for noisy data and nonlinear formulation (2019). [arxiv:1906.00810](https://arxiv.org/abs/1906.00810)
- Grepl, M.A., Patera, A.T.: A posteriori error bounds for reduced-bias approximations of parametrized parabolic partial differential equations. *Math. Model. Numer. Anal.* **39**(1), 157–181 (2005). <https://doi.org/10.1051/m2an:2005006>
- Haasdonk, B., Ohlberger, M.: Reduced basis method for finite volume approximations of parametrized linear evolution equations. *Math. Model. Numer. Anal.* **42**(2), 277–302 (2008). <https://doi.org/10.1051/m2an:2008001>
- Hamill, T.M., Whitaker, J.S., Snyder, C.: Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Monthly Weather Review* **129**(11), 2776–2790 (2001). [https://doi.org/10.1175/1520-0493\(2001\)129<2776:DDFOBE>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<2776:DDFOBE>2.0.CO;2)
- Hecht, F.: New development in freefem++. *J. Numer. Math.* **20**(3–4), 251–265 (2012). <https://doi.org/10.1515/jnum-2012-0013>
- Hernández, J.A., Caicedo, M.A., Ferrer, A.: Dimensional hyper-reduction of nonlinear finite element models via empirical cubature. *Comput. Methods Appl. Mech. Eng.* **313**, 687–722 (2017). <https://doi.org/10.1016/j.cma.2016.10.022>
- Hoel, H.K., Law, K.J.H., Tempone, R.: Multilevel ensemble Kalman filtering. *SIAM J. Numer. Anal.* **54**(3), 1813–1839 (2016). <https://doi.org/10.1137/15M100955X>
- Houtekamer, P.L., Mitchell, H.L.: A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review* **129**(1), 123–137 (2001). [https://doi.org/10.1175/1520-0493\(2001\)129<0123:ASEKFF>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<0123:ASEKFF>2.0.CO;2)
- Huttunen, J.M.J., Kaipio, J.P.: Approximation error analysis in nonlinear state estimation with an application to state-space identification. *Inverse Prob.* **23**(5), 2141–2157 (2007). <https://doi.org/10.1088/0266-5611/23/5/019>
- Iglesias, M.A.: A regularizing iterative ensemble Kalman method for PDE-constrained inverse problems. *Inverse Prob.* **32**(2), 025002–45 (2016). <https://doi.org/10.1088/0266-5611/32/2/025002>

- Iglesias, M.A., Law, K.J.H., Stuart, A.M.: Ensemble Kalman methods for inverse problems. *Inverse Prob.* **29**(4), 045001–20 (2013). <https://doi.org/10.1088/0266-5611/29/4/045001>
- Kaltenbacher, B., Neubauer, A., Scherzer, O.: *Iterative Regularization Methods for Nonlinear Ill-posed Problems*. De Gruyter, Berlin (2008). <https://doi.org/10.1515/9783110208276>
- Kärcher, M., Boyaval, S., Grepl, M.A., Veroy, K.: Reduced basis approximation and a posteriori error bounds for 4D-Var data assimilation. *Optim. Eng.* **19**(3), 663–695 (2018). <https://doi.org/10.1007/s11081-018-9389-2>
- Li, Z., Navon, I.M.: Optimality of variational data assimilation and its relationship with the Kalman filter and smoother. *Q. J. R. Meteorol. Soc.* **127**(572), 661–683 (2001). <https://doi.org/10.1002/qj.49712757220>
- Lorentzen, R.J., Fjælde, K.K., Frøyen, J., Lage, A.C.V.M., Nævdal, G., Vefring, E.H.: Underbalanced and low-head drilling operations: Real time interpretation of measured data and operational support. *SPE Annual Technical Conference and Exhibition*, vol. All Days (2001). <https://doi.org/10.2118/71384-MS>
- Mitchell, L., Carrassi, A.: Accounting for model error due to unresolved scales within ensemble Kalman filtering. *Q. J. R. Meteorol. Soc.* **141**(689), 1417–1428 (2015). <https://doi.org/10.1002/qj.2451>
- Mitchell, H.L., Houtekamer, P.L., Pellerin, G.: Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Mon Weather Rev* **130**(11), 2791–2808 (2002). [https://doi.org/10.1175/1520-0493\(2002\)130<2791:ESBAME>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<2791:ESBAME>2.0.CO;2)
- Nadal, E., Chinesta, F., Díez, P., Fuenmayor, F.J., Denia, F.D.: Real time parameter identification and solution reconstruction from experimental data using the proper generalized decomposition. *Comput. Methods Appl. Mech. Engrg.* **296**, 113–128 (2015). <https://doi.org/10.1016/j.cma.2015.07.020>
- Pagani, S., Manzoni, A., Quarteroni, A.: Efficient state/parameter estimation in nonlinear unsteady PDEs by a reduced basis ensemble Kalman filter. *SIAM/ASA J. Uncertain. Quantif.* **5**(1), 890–921 (2017). <https://doi.org/10.1137/16M1078598>
- Popov, A.A., Sandu, A.: In: Park, S.K., Xu, L. (eds.) *Multifidelity data assimilation for physical systems*, pp. 43–67. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-77722-7_2
- Popov, A.A., Mou, C., Sandu, A., Iliescu, T.: A multifidelity ensemble kalman filter with reduced order control variates. *SIAM J. Sci. Comput.* **43**(2), 1134–1162 (2021). <https://doi.org/10.1137/20M1349965>
- Prud'homme, C., Rovas, D.V., Veroy, K., Patera, A.T.: A mathematical and computational framework for reliable real-time solution of parametrized partial differential equations. *Math. Model. Numer. Anal.* **36**(5), 747–771 (2002). <https://doi.org/10.1051/m2an:2002035>
- Quarteroni, A., Manzoni, A., Negri, F.: *Reduced basis methods for partial differential equations: An introduction*, p. 296. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-15431-2>
- Sakov, P., Evensen, G., Bertino, L.: Asynchronous data assimilation with the EnKF. *Tellus A: Dyn. Meteorol. Oceanogr* **62**(1), 24–29 (2010). <https://doi.org/10.1111/j.1600-0870.2009.00417.x>
- Särkkä, S.: *Bayesian filtering and smoothing*. Institute of Mathematical Statistics Textbooks, vol. 3, p. 232. Cambridge University Press, Cambridge (2013). <https://doi.org/10.1017/CBO9781139344203>
- Schillings, C., Stuart, A.M.: Convergence analysis of ensemble Kalman inversion: the linear, noisy case. *Appl. Anal.* **97**(1), 107–123 (2018). <https://doi.org/10.1080/00036811.2017.1386784>
- Skjervheim, J.-A., Evensen, G., Aanonsen, S.I., Ruud, B.O., Johansen, T.A.: Incorporating 4D seismic data in reservoir simulation models using ensemble Kalman filter. *SPE J.* **12**(03), 282–292 (2007). <https://doi.org/10.2118/95789-PA>
- Thepaut, J.-N., Courtier, P.: Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model. *Q. J. R. Meteorol. Soc.* **117**(502), 1225–1254 (1991). <https://doi.org/10.1002/qj.49711750206>
- Thomé V.: *Galerkin Finite Element Methods for Parabolic Problems* vol. 25. Springer, Berlin (2006). 10.1007/3-540-33122-0 12
- Wendland, H.: Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Adv. Comput. Math.* **4**(4), 389–396 (1995). <https://doi.org/10.1007/bf02123482>
- Wu, J., Wang, J.-X., Shadden, S.C.: Improving the convergence of the iterative ensemble Kalman filter by resampling (2019). [arxiv:1910.04247](https://arxiv.org/abs/1910.04247)
- Yano, M., Patera, A.T.: An LP empirical quadrature procedure for reduced basis treatment of parametrized nonlinear PDEs. *Comput. Methods Appl. Mech. Eng.* **344**, 1104–1123 (2019). <https://doi.org/10.1016/j.cma.2018.02.028>