Advances in Water Resources 186 (2024) 104677

Contents lists available at ScienceDirect

Advances in Water Resources

journal homepage: www.elsevier.com/locate/advwatres

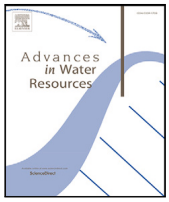

Polynomial chaos enhanced by dynamic mode decomposition for order-reduction of dynamic models

G. Libero [a], D.M. Tartakovsky [b], V. Ciriello [a,*]

[a] Department of Civil, Chemical, Environmental and Materials Engineering, University of Bologna, Italy
[b] Department of Energy Science and Engineering, Stanford University, USA

ARTICLE INFO

Keywords:
Polynomial chaos expansion
Dynamic mode decomposition
Reduced-order model
Hydrology
Dynamic model

ABSTRACT

Thanks to their low computational cost, reduced-order models (ROMs) are indispensable in ensemble-based simulations used, e.g., for uncertainty quantification, inverse modeling, and optimization. Since data used to train a ROM are typically obtained by running a high-fidelity model (HFM) multiple times, a ROM's efficiency rests on the computational cost associated with the data generation and training phase. One such ROM, a polynomial chaos expansion (PCE), often provides a robust description of an HFM's response surface in the space of model parameters. To reduce the data-generation cost, we propose to train a PCE on multi-fidelity data, part of which come from the dynamic HFM and the remainder from dynamic mode decomposition (DMD); the latter is used to interpolate the HFM data in time. Our numerical experiments demonstrate the accuracy of the proposed method and provide guidelines for the optimal use of DMD for interpolation purposes.

1. Introduction

Rapid advances in software and hardware development have led to the proliferation of high-fidelity models (HFMs), which provide invaluable insights into hydrological processes. The growth in model complexity roughly matches that in computational power, so that one's ability to perform multiple runs of HFMs remains practically unchanged. That creates a computational bottleneck because most applications of subsurface flow and transport require predictions of a system's temporal evolution for different values of the system parameters. This requirement places particular premium on reducing the (often prohibitive) cost of HFMs (e.g., Ciriello et al., 2013, 2017).

Reduced-order models (ROMs) provide computationally-efficient representations of key features of the underlying dynamics of a complex system directly from (observational or HFM-simulated) data (e.g., Oladyshkin et al., 2012; Ciriello et al., 2019; Tartakovsky et al., 2020; He and Tartakovsky, 2021; Kang et al., 2022; Zhan et al., 2022). The amount of data required to train a ROM is a measure of its efficiency. ROMs yield a simple mathematical relationship between the relevant input parameters (or features) and the model response (output or target). In doing so, ROMs enable one to perform a large number of simulations, for multiple points in the parameter space, while preserving the interpretation provided by a HFM (e.g., Ciriello et al., 2013; Oladyshkin et al., 2012; Focaccia et al., 2021; Marzadri et al., 2024).

Polynomial chaos expansions (PCEs) are a representative class of ROMs, which approximates the input–output relation encoded in the HFM via orthogonal polynomials (Xiu and Karniadakis, 2002; Sudret, 2008). The probabilistic collocation method (PCM) (Webster et al., 1996; Sudret, 2008) provides a non-intrusive approach for the computation of the PCE coefficients. Unlike its intrusive counterparts, which require one to derive and solve a set of differential equations for the PCE coefficients, the PCM amounts to a regression between points on the model response surface computed as solutions of the HFM for optimally selected combinations of the model inputs (collocation points in the parameter space). A PCE-based ROM can dramatically accelerate ensemble-based computations, such as uncertainty quantification for flow and transport problems (Lin et al., 2010; Oladyshkin et al., 2012; Mohammadi et al., 2018; Ciriello and de Barros, 2020; Meles et al., 2022), the task that could be computationally prohibitive if performed via Monte Carlo simulations of the HFM.

Several factors compromise the efficiency of PCEs, occasionally causing the computational cost of a PCE to exceed that of standard Monte Carlo. They scale poorly with the number of random inputs, a phenomenon referred to as the curse of dimensionality (Barajas-Solano and Tartakovsky, 2016). The PCM performance depends on the regularity of the response surface, which decreases with both the variance of the model parameters (Barajas-Solano and Tartakovsky, 2016) and the degree of the model nonlinearity (Meles et al., 2023).

* Corresponding author.
E-mail address: v.ciriello@unibo.it (V. Ciriello).

https://doi.org/10.1016/j.advwatres.2024.104677
Received 25 December 2023; Received in revised form 27 February 2024; Accepted 11 March 2024
Available online 14 March 2024

For time-dependent problems, the cost of the training-data generation often dwarfs the cost of the PCE construction. A common strategy for the reduction of the computational burden of solving the HFM is to use as large a time step as the accuracy and stability constraints of a numerical scheme would allow. This time step might be too big for an optimal PCE approximation. The main idea of our approach is to use dynamic mode decomposition (DMD) to interpolate between these times in order to derive a virtually continuous in time PCE approximation of the HFM response surface.

DMD is a data-driven approach based on a factorization and dimensionality reduction technique for data sequences associated with complex dynamic systems (e.g., Schmid, 2010). Grounded in singular value decomposition, DMD yields a best-fit linear operator to approximate the relationship between time-shifted snapshots of the state variable for assigned values of the parameters. It also identifies spatiotemporal structures that are dominant in the data and reconstructs the underlying processes from these structures. In the present analysis, we use the standard DMD (sDMD) algorithm (Tu et al., 2014; Kutz et al., 2016) and xDMD (Lu and Tartakovsky, 2021), a DMD variant that enables one to handle inhomogeneous partial differential equations and inhomogeneous boundary conditions.

We posit that a fusion of PCE and DMD yields a robust model-reduction framework for dynamic models. This approach takes advantage of the ability of PCE to provide an approximation of the response in the parameter space (thus enabling to perform stochastic analysis and scenario development) and the ability of DMD to interpolate high-fidelity data in time. DMD is also usable for temporal extrapolation, with a general loss in accuracy (Lu and Tartakovsky, 2020), and provides a means for data compression and memory-storage reduction. Our DMD-enhanced PCE (in the following, DMD-PCE) uses DMD to replace a HFM during the PCE training, such that the PCE is trained on multi-fidelity data. We test the accuracy of our approach on two-dimensional multiphase (nonlinear) flow in heterogeneous media (Song and Tartakovsky, 2021).

The paper is organized as follows. Section 2 is devoted to a presentation of the proposed methodological framework. Section 3 describes the selected case study and the implementation of the methodological framework to this problem. In Section 4, the results are presented and analyzed, while Section 5 provides final remarks.

2. Materials and methods

2.1. Polynomial chaos expansion

Let ω denote an output response provided by a HFM, $f(\mathbf{p})$, where $\mathbf{p}$ is the vector of system parameters (inputs). Their variability and uncertainty are captured by modeling them as independent random variables. The assumption of independence also holds for (i) cross-correlated parameters, for which singular-value decomposition techniques, such as a truncated Karhunen-Loève transformation, can be used to approximate $\mathbf{p}$ with a set of mutually uncorrelated identically distributed random variables; and (ii) correlated random variables, for which the Rosenblatt transform can be applied (e.g. Um et al., 2019) to map $\mathbf{p}$ onto a set of independent and identically distributed random variables.

Once the probabilistic behavior of the parameters $\mathbf{p}$ is characterized, and provided the response surface has finite variance σ_ω^2, PCE produces an approximation of the response surface, $\hat{\omega}$, in the random input space in terms of a polynomial series (e.g., Xiu and Karniadakis, 2002),

$$\omega = f(\mathbf{p}) \rightarrow \qquad \hat{\omega} = \sum_{j=0}^{P-1} a_j \Psi_j(\mathbf{p}), \qquad P = \frac{(N_{\text{par}}+q)!}{N_{\text{par}}!\,q!}. \tag{1}$$

Here, $N_{\text{par}} = \dim(\mathbf{p})$, Ψ_j are multivariate polynomials of degree not exceeding q that constitute an orthonormal basis with respect to the joint probability density function (PDF) of $\mathbf{p}$, q is the maximum degree of the expansion and coefficients a_j are the deterministic coordinates of the spectral decomposition. For Gaussian parameters, the Hermite polynomials serve as a basis; different types of polynomials are required for optimal convergence rate in the case of non-Gaussian parameters (Xiu and Karniadakis, 2002). The P-term truncation of an infinite-term PCE introduces an approximation error. This error has been the subject of many theoretical and numerical investigations (Ghanem and Spanos, 1991; Shi and Tartakovsky, 2022).

To compute the PCE coefficients, a non-intrusive regression-based approach is used. It comprises the minimization of the variance of the residual $\varepsilon = |\hat{\omega} - \omega|$ with respect to the PCE coefficients $\mathbf{a} = (a_0, \dots, a_{P-1})$ (Sudret, 2008),

$$\mathbf{a} = \operatorname{argmin} \frac{1}{N_{\text{rp}}} \sum_{i=1}^{N_{\text{rp}}} \left\{ \omega - \sum_{j=0}^{P-1} a_j \Psi_j(\mathbf{p}) \right\}^2, \tag{2}$$

where N_{rp} is the number of regression points. This is done on a dataset generated by the HFM for values of $\mathbf{p}$ provided by PCM (Webster et al., 1996). These values are termed collocation points ($N_{\text{rp}} = N_{\text{cp}}$ is their number) and they are selected based on the same arguments adopted for integral estimation through Gaussian quadrature. Specifically, PCM employs the roots of the polynomial of one order higher than q to assure proper sampling of the region associated with the largest probability in the distributions of the input parameters (Webster et al., 1996). The fewer the number of collocation points, the more effective PCE is. Note that $N_{\text{cp}} \geq P$, i.e., it increases with N_{par} and q (Ciriello et al., 2013).

If the HFM describes the behavior of a spatially-distributed time-variant response, i.e., $f(\mathbf{x}, t, \mathbf{p})$, the space–time dependence of the response ω is embedded in the PCE coefficients $\mathbf{a}$:

$$\hat{\omega}(\mathbf{x}, t) = \sum_{j=0}^{P-1} a_j(\mathbf{x}, t)\, \Psi_j(\mathbf{p}). \tag{3}$$

This means that to obtain the PCE approximation at each $(\mathbf{x}, t)$ of interest, ω has to be computed at the same space–time locations for each value of $\mathbf{p}$ required by the PCM to solve (2). If the HFM is solved numerically on a grid of N elements for M time points, then we need to compute ω at $N \cdot M$ locations N_{cp} times (e.g., Ciriello and de Barros, 2020).

2.2. Dynamic mode decomposition

Consider a set of $(M+1)$ snapshots of the output response $\boldsymbol{\omega}$, namely $\boldsymbol{\omega}_k$ with $t_{k+1} = t_k + \Delta t$ and $k = 0, \dots, M$. Let $\mathbf{X}, \mathbf{X}' \in \mathbb{R}^{N \times M}$ denote a couple of matrices whose columns are the vectors $\boldsymbol{\omega}_0, \dots, \boldsymbol{\omega}_{M-1}$ and $\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_M$, respectively. In the standard DMD (sDMD) formulation, the temporal evolution of $\boldsymbol{\omega}(t)$ is approximated by a linear model

$$\boldsymbol{\omega}_{k+1} \approx \mathbf{A}\boldsymbol{\omega}_k, \qquad \mathbf{A} = \mathbf{X}'\mathbf{X}^\dagger \in \mathbb{R}^{N \times N}. \tag{4}$$

In a typical application, $M \ll N$ so that the rank of $\mathbf{A}$ is at most M. Even though, computing $\mathbf{A}$ (or its spectral decomposition) is generally onerous. Instead, the truncated SVD of $\mathbf{X} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$, with rank $r < M$, is used:

$$\mathbf{A} \approx \mathbf{X}'\mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^\top, \tag{5}$$

where $\mathbf{U} \in \mathbb{R}^{N \times r}$, $\boldsymbol{\Sigma} \in \mathbb{R}^{r \times r}$, $\mathbf{V} \in \mathbb{R}^{M \times r}$. If r is smaller than the number of nonzero singular values (i.e., the rank of $\mathbf{X}$), then the truncated SVD is a proxy of $\mathbf{X}$.

In case of problem's inhomogeneity, the generalized DMD algorithm adds a bias term $\mathbf{b}_g \in \mathbb{R}^N$ to the preceding formulation,

$$\boldsymbol{\omega}_{k+1} \approx \mathbf{A}_g\boldsymbol{\omega}_k + \mathbf{b}_g. \tag{6}$$

Here, $[\mathbf{A}_g \;\; \mathbf{b}_g] = \mathbf{X}'\tilde{\mathbf{X}}^\dagger \in \mathbb{R}^{N \times N+1}$, where $\tilde{\mathbf{X}}^\top = [\mathbf{X} \;\; \mathbf{1}]$ and $\tilde{\mathbf{X}} \in \mathbb{R}^{N+1 \times M}$. The computational cost is reduced by obtaining the best-fit linear operator through the SVD of the matrix $\tilde{\mathbf{X}} \approx \mathbf{U}_g\boldsymbol{\Sigma}_g\mathbf{V}_g^\top$, such that

$$[\mathbf{A}_g \;\; \mathbf{b}_g] \approx \mathbf{X}'\mathbf{V}_g\boldsymbol{\Sigma}_g^{-1}\mathbf{U}_g^\top, \tag{7}$$

where $\mathbf{U}_g \in \mathbb{R}^{N+1\times r}, \mathbf{\Sigma}_g \in \mathbb{R}^{r\times r}, \mathbf{V}_g \in \mathbb{R}^{M\times r}$. By construction, the error of this gDMD method is equal to or smaller than the standard DMD's (Lu and Tartakovsky, 2021).

The xDMD algorithm introduced by Lu and Tartakovsky (2021) endows gDMD with a residual-learning idea by approximating the relationship between $\mathbf{Y} = \mathbf{X}' - \mathbf{X}$ and $\mathbf{X}$,

$$\mathbf{y}_{k+1} = \mathbf{B}_x \boldsymbol{\omega}_k + \mathbf{b}_x. \tag{8}$$

Here, $[\mathbf{B}_x \ \ \mathbf{b}_x] = \mathbf{Y}\tilde{\mathbf{X}}^\dagger \in \mathbb{R}^{N\times N+1}$, and $\tilde{\mathbf{X}}$ is defined as before. For computational saving, the best-fit linear operator is obtained through the SVD of the matrix $\tilde{\mathbf{X}}$ as

$$[\mathbf{B}_x \ \mathbf{b}_x] \approx \mathbf{Y}\mathbf{V}_g\mathbf{\Sigma}_g^{-1}\mathbf{U}_g^\top. \tag{9}$$

The error of xDMD equals to or is smaller than that of the residual DMD without bias (Lu and Tartakovsky, 2021).

Since n is large in any application of practical significance, an efficient computational strategy to derive prediction in (8) is as follows

$$\mathbf{y}_{k+1} = \mathbf{C}_x \mathbf{d}_x, \tag{10}$$

where $\mathbf{C}_x = \mathbf{Y}\mathbf{V}_g\mathbf{\Sigma}_g^{-1} \in \mathbb{R}^{N\times r}$, and $\mathbf{d}_x = \mathbf{U}_g^\top \tilde{\boldsymbol{\omega}}_k \in \mathbb{R}^{r\times 1}$. A similar computational approach can be conveniently adopted to derive prediction in (6).

DMD algorithms can be used for nonlinear PDEs, whose solution is confined in $\mathcal{H} \subseteq \mathbb{R}^N$ (to satisfy the assumptions in Lemma 2.1 in Lu and Tartakovsky, 2021). Let $\boldsymbol{\omega}_k^{\mathcal{L}}$ be the DMD approximation of $\boldsymbol{\omega}_k$. We assess the performance of either sDMD or xDMD in terms of the relative error

$$\varepsilon_{\mathcal{L}}^k = \frac{\|\boldsymbol{\omega}_k^{\mathcal{L}} - \boldsymbol{\omega}_k\|^2}{\|\boldsymbol{\omega}_k\|^2}. \tag{11}$$

As a truncation rank of sDMD and xDMD, we use either $r = r_{90}$—the number of diagonal elements of Σ accounting for 90% of the cumulative energy in the SVD of $\tilde{\mathbf{X}}$—or $r = r^*$ with

$$r^* = \min(i) : \sigma_i \le 10^{-5} \sum_{k=0}^{M-1} \sigma_k. \tag{12}$$

2.3. Integrated model reduction framework

The regression given by (2) has to be solved at each space–time location where we want to compute $\hat{\omega}$. Suppose the space–time locations are defined by indices $i = 1, \dots, N$, and $k = 1, \dots, M$. High-fidelity simulations need to be run at each (i, k), N_{cp} times. This traditional approach is indicated here as HF-PCE. The computational cost associated with this step is the only relevant cost of the PCE framework. Let η denote a fraction of temporal snapshots estimated via DMD, $\boldsymbol{\omega}_{k^\dagger}^{\mathcal{L}}$, for the use in PCE training. We perform the regression on high-fidelity data at $(1-\eta)$% of the time points t_{k_η}, i.e., on data $\boldsymbol{\omega}_{k_\eta} = \boldsymbol{\omega}(t_{k_\eta})$. Here, $k_\eta = 1, \dots, M$ with time step $1/(1-\eta)$ forms a subset $\mathcal{M}_\eta$ of the total number of snapshots, M, and $k^\dagger \neq k_\eta$ belongs to its complement $\mathcal{M}_\eta^\dagger$, such that $\|\mathcal{M}_\eta\| + \|\mathcal{M}_\eta^\dagger\| = M$. This approach is indicated in the following as DMD-PCE. The HF-PCE approach corresponds to $\eta = 0$, in which case the PCE coefficients are computed in all the space–time locations based on high-fidelity data.

We indicate the PCE coefficients computed at a generic (i, k) based on high-fidelity data with $\mathbf{a}^{\text{HF}}$, and those computed at a generic (i, k) based on DMD interpolated data with $\mathbf{a}^{\text{DMD}}$. A first assessment of the accuracy of the method is performed through comparison at a given $t_{k^\dagger}$ between $\mathbf{a}^{\text{DMD}}$, as computed through (2) for a given η, and $\mathbf{a}^{\text{HF}}$ as computed through the same equation in case of $\eta = 0$. The metric used for each coefficient is

$$\varepsilon_{a_j}^{k^\dagger} = \frac{\|\mathbf{a}_{j,k^\dagger}^{\text{DMD}} - \mathbf{a}_{j,k^\dagger}^{\text{HF}}\|^2}{\|\mathbf{a}_{j,k^\dagger}^{\text{HF}}\|^2}, \tag{13}$$

where $j = 0, \dots, P-1$ indicates the PCE coefficient, while $\mathbf{a}_{j,k^\dagger}^{\text{DMD}}$ and $\mathbf{a}_{j,k^\dagger}^{\text{HF}}$ are the vectors of values for the $(j+1)$th coefficient at all the space locations, i.e. for all i and for a specific $k^\dagger$, computed with the DMD-interpolated data and high-fidelity data, respectively.

Next, we directly compare at each space location and $t_{k^\dagger}$ the PCE approximations of the response provided by the DMD-PCE approach for a given η, i.e. $\hat{\boldsymbol{\omega}}_{i,k^\dagger}^{\text{DMD}}$, and by the HF-PCE approach (case of $\eta = 0$), i.e. $\hat{\boldsymbol{\omega}}_{i,k^\dagger}^{\text{HF}}$. In both cases, the quality of the proposed approach reflects the accuracy of DMD interpolation.

3. Application

3.1. Physics-based HFM

Models of multiphase flow play a notable role in many phenomena, ranging from contaminant migration and carbon sequestration to geothermal energy. The need to solve strongly nonlinear governing PDEs results in a high, often prohibitive, computational cost, which limits uncertainty quantification and any detailed study requiring a large number of simulations in the parameter space (Song and Tartakovsky, 2021).

We test our methodology by considering a two-dimensional flow of two incompressible and immiscible fluids, with viscosities μ_1 and μ_2, in a heterogeneous porous medium. The latter is incompressible, isotropic, and characterized by porosity ϕ and intrinsic permeability $k(\mathbf{x})$ (Taverniers et al., 2020; Song and Tartakovsky, 2021). The propagation of each phase $(l = 1, 2)$ is described by the combination of a mass conservation equation

$$\phi \frac{\partial S_l}{\partial t} + \nabla \cdot \mathbf{v}_l + q_l = 0, \qquad \mathbf{x} \equiv (x; y)^\top \in D, \quad t \in [0; T] \tag{14}$$

and the generalized Darcy's law

$$\mathbf{v}_l = -k \frac{k_{rl}}{\mu_l} \nabla P_l. \tag{15}$$

Here, $S_l(\mathbf{x}; t)$ is the saturation of the lth phase, constrained by $S_1 + S_2 = 1$; and q_l is the source/sink term. The macroscopic flow velocity $\mathbf{v}_l(\mathbf{x}; t)$ in (15) depends on the relative permeability of the lth phase, $k_{rl} = k_{rl}(S_l)$, which varies with the corresponding saturation in accordance with the Brooks-Corey constitutive model (Corey, 1954). The pressure is assumed to be equal within the two phases $P = P_1 = P_2 \equiv P(\mathbf{x}, t)$, hence the capillarity forces are ignored (Taverniers et al., 2020).

Similarly to Taverniers et al. (2020), the flow is studied on a square simulation domain D of size 150 m × 150 m, with impermeable bottom (Γ_b) and top Γ_t boundaries, and Dirichlet conditions imposed along the left Γ_l and right Γ_r boundaries:

$$\frac{\partial P}{\partial y} = 0, \ \mathbf{x} \in \Gamma_b \cup \Gamma_t; \quad P = 10 + \Delta p \ \wedge \ S_1 = 1.0, \ \mathbf{x} \in \Gamma_l; \quad P = 10, \ \mathbf{x} \in \Gamma_r. \tag{16}$$

Pressure P is expressed in MPa. Initial conditions are

$$P(\mathbf{x}, 0) = 10.1, \ S_1(\mathbf{x}, 0) = 0, \ \mathbf{x} \in D. \tag{17}$$

To account for heterogeneity, two different intrinsic permeability scenarios (SC1, SC2) are considered. Two permeability maps are generated as a second-order stationary random field, such that $Y(\mathbf{x}) = \ln k(\mathbf{x})$ is multivariate Gaussian with an exponential covariance function $C_Y(\mathbf{x}, \mathbf{y}) = \sigma_Y^2 \exp(-|\mathbf{x} - \mathbf{y}|/\lambda_Y)$. We consider two sets of values for the mean μ_Y, variance σ_Y^2, and correlation length λ_Y (in m): $\{\mu_Y, \sigma_Y^2, \lambda_Y\} = \{0, 2, 19\}$ and $\{0, 3, 10\}$. A truncated Karhunen–Loève expansion, counting 31 terms, chosen to capture 95% of the energy of the field Y, is employed to approximate $Y(\mathbf{x})$ (Taverniers et al., 2020). These KL coefficients do not serve as PCE inputs, i.e., the two randomly generated fields are considered as deterministic scenarios.

Eqs. (14) and (15) are discretized using a finite volume scheme in space and an implicit Euler scheme in time, yielding a highly non-linear

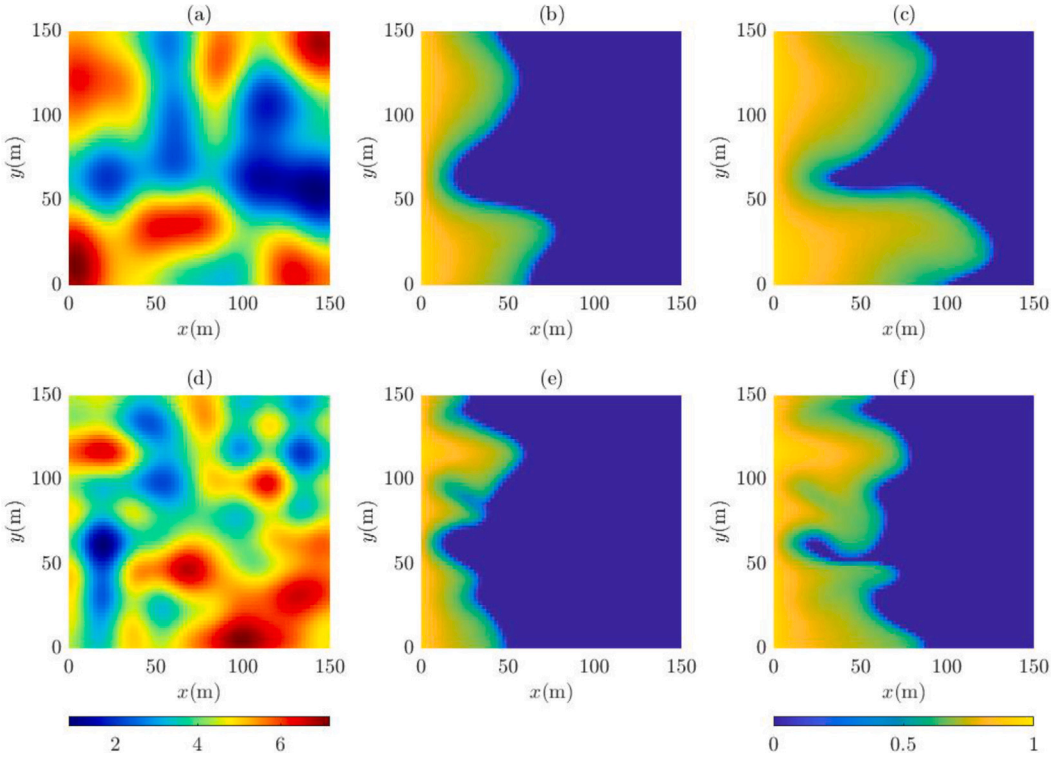

Fig. 1. (a) and (d) show the two scenarios (SC1 and SC2) of log-permeability fields $Y(\mathbf{x}) = \ln k(\mathbf{x})$ (with $k(\mathbf{x})$ is expressed in mDarcy); for each scenario (b)–(c) and (e)–(f) show the saturation maps at $t_{M/2}$ and t_M respectively.

system of equations, which is solved iteratively at each time step using the Newton–Raphson method, with modified Appleyard saturation update damping (Appleyard et al., 1981) to improve convergence. Three criteria are specified to ensure convergence of both flow and transport solutions: normalized residual norm, maximum saturation update, and maximum relative pressure update. Their respective tolerances are set to $\varepsilon_1 = 10^{-6}$, $\varepsilon_2 = 10^{-2}$, and $\varepsilon_3 = 10^{-3}$ (Taverniers et al., 2020; Song and Tartakovsky, 2021).

3.2. Experimental design

Among the input parameters (Taverniers et al., 2020), we focus on uncertainty related to the porosity ϕ and the external pressure difference Δp, i.e., $N_{\text{par}} = 2$. Following the PCE approach, we model the two inputs as independent random variables. Specifically, we assume two normal distributions $\phi \simeq \mathcal{N}(m_\phi = 0.25; \sigma_\phi = 0.05)$, $\Delta p \simeq \mathcal{N}(m_{\Delta p} = 0.20$ MPa$; \sigma_{\Delta p} = 0.04$ MPa) and, consequently, employ the Hermite polynomials in (1) (Xiu and Karniadakis, 2002). We use a second-order PCE approximation, i.e., $q = 2$, so that $P = 6$ and

$$\hat{\mathbf{S}}_1(\mathbf{x}, t) = a_0 + a_1\phi_n + a_2\Delta p_n + a_3(\phi_n^2 - 1) + a_4\phi_n \cdot \Delta p_n + a_5(\Delta p_n^2 - 1), \quad (18)$$

where the PCE coefficients $a_i = a_i(\mathbf{x}, t)$ vary in space and time, and ϕ_n and Δp_n are standard normal random variables obtained using an isoprobabilistic transform, i.e. the normal random variables ϕ and Δp are turned into standard ones to make the polynomial basis orthogonal with respect to the joint PDF of the parameters (Sudret, 2008). (After rescaling, (18) can also be written in terms of orthonormal polynomial basis (e.g., Oladyshkin and Nowak, 2018) that is useful to employ the analytical properties of PCE as in the case of global sensitivity analysis (Ciriello et al., 2013).) Under these conditions, PCM returns $N_{\text{cp}} = P = 6$ optimal collocation points in the random parameter space (Table 1), i.e., six combinations of values for the two uncertain inputs, ϕ and Δp, to determine the coefficients a_i by solving the regression expressed by (2) at each $(\mathbf{x}, t)$.

Table 1
Collocation points in the random parameter space, with $j = 1, \ldots, N_{\text{cp}} = P = 6$.

j	ϕ (-)	Δp (MPa)
1	2.50E−01	2.00E−01
2	3.37E−01	2.00E−01
3	2.50E−01	2.69E−01
4	1.63E−01	2.00E−01
5	2.50E−01	1.31E−01
6	3.37E−01	2.69E−01

For each collocation point in Table 1, high-fidelity simulations are performed for the permeability scenarios SC1 and SC2. Each high-fidelity simulation involves the computation of the response, which is S_1, on a 64×64 grid ($N = 4096$), and for $M = 1820$ ($k = 1, \ldots, 1820$) time instants in the range $t = 8, \ldots, 1827$ days with $\Delta t = 1$ day. Fig. 1 shows the two permeability fields that determine SC1 and SC2 and, for each, two snapshots of the response (saturation map) for $k = M/2$ ($t_{M/2} = 910$ days) and $k = M$ ($t_M = 1820$ days).

Once the entire set of high-fidelity simulations is available (i.e., for all the collocation points), the PCE (2) can be computed at all the space–time locations of interest considering $\eta = 0$ (HF-PCE). Instead, we use a limited number of high-fidelity data to calibrate the DMD approximation of the response in time and then use DMD to interpolate and fill the missing high-fidelity data, i.e., to use a mix of high-fidelity and DMD interpolated data to compute the PCE (2) for a given η. The DMD-PCE is then compared to the HF-PCE in terms of accuracy. Fig. 2 illustrates the DMD-PCE framework and its application to the case study; the DMD and PCE algorithms are used sequentially, with DMD deployed to generate training data, which then inform the PCE.

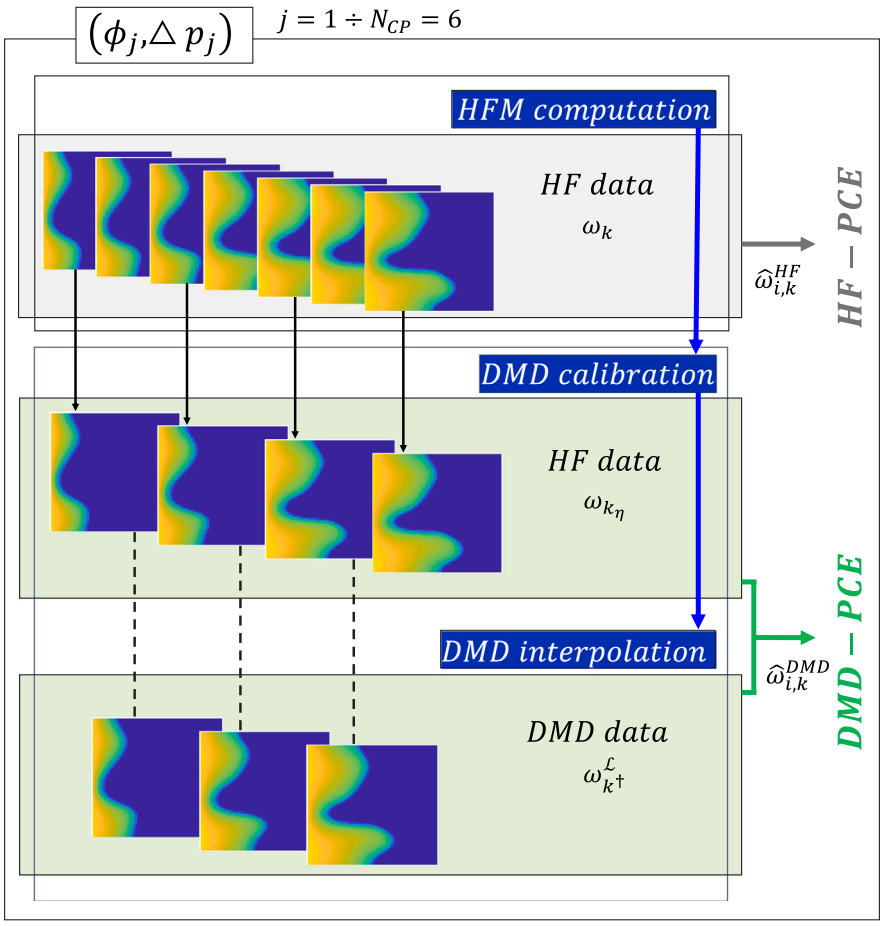

Fig. 2. Methodological framework of the DMD-PCE approach.

Table 2
Cases analyzed for the two scenarios of the permeability field (SC1, SC2).

	Case 1: $\eta = 0.5$	Case 2: $\eta = 0.75$	Case 3: $\eta = 0.9$
representation	$k_\eta = 1, 3, \dots, M-1$	$k_\eta = 1, 5, \dots, M-3$	$k_\eta = 1, 11, \dots, M-9$
interpolation	$k^\dagger = 2, 4, \dots, M$	$k^\dagger = 2, 3, 4, 6, 7, 8, \dots, M$	$k^\dagger = 2 \div 10, 12 \div 20, \dots, M$
SC1: rank(X)	890–909	454	181
SC1: r^*	75–129	75–129	74-128
SC1: r_{90}	11–16	10–16	10-16
SC2: rank(X)	765–908	454	181
SC2: r^*	61–105	61–104	60-104
SC2: r_{90}	8–13	8–13	8-13

We consider $\eta = 0.5, 0.75, 0.9$. For each value of η, a different DMD approximation is computed with both the sDMD and xDMD algorithms and different truncation ranks $r = (\text{rank(X)}, r^*, r_{90})$. The summary of the cases analyzed in this study is reported in Table 2.

4. Results and discussion

4.1. DMD representation error

As mentioned in Section 3.2, both sDMD and xDMD are trained based on high-fidelity data provided by the numerical model described in Section 3.1. The performance in reproducing the same data used for training (representation error) is compared between the two DMD algorithms, and evaluated for different dimensions of the training set. Specifically, for each permeability scenario (SC1/SC2, see Section 3.1) and collocation point (Table 1), three different subsets counting 50% (case 1, $\eta = 0.5$), 25% (case 2, $\eta = 0.75$), and 10% (case 3, $\eta = 0.90$) of high-fidelity data are used for training (see Table 2). Based on these subsets, we derive the DMD for different truncation ranks ($r = \text{rank(X)}$, r^*, r_{90}, see Section 2.2) to explore the accuracy of the two algorithms, sDMD and xDMD.

Figs. 3 and 4 show the results for permeability scenarios SC1 and SC2, respectively. In both figures, columns correspond to cases 1-3 in Table 2 (i.e., different dimensions of the training set). The first row of both the Figures depicts the singular values associated with the SVD of $\mathbf{X}$ for the three cases; while the second and third rows report the representation error (computed with Eq. (11) and averaged over all the time instants), $\varepsilon_{\mathcal{L}}$, of the sDMD and xDMD, respectively, as function of the truncation rank r of the SVD. In all the plots, different lines correspond to high-fidelity simulations produced for the six collocation points (different values of the random model parameters, ϕ and Δp collected in Table 1), while the ranks corresponding to different truncation

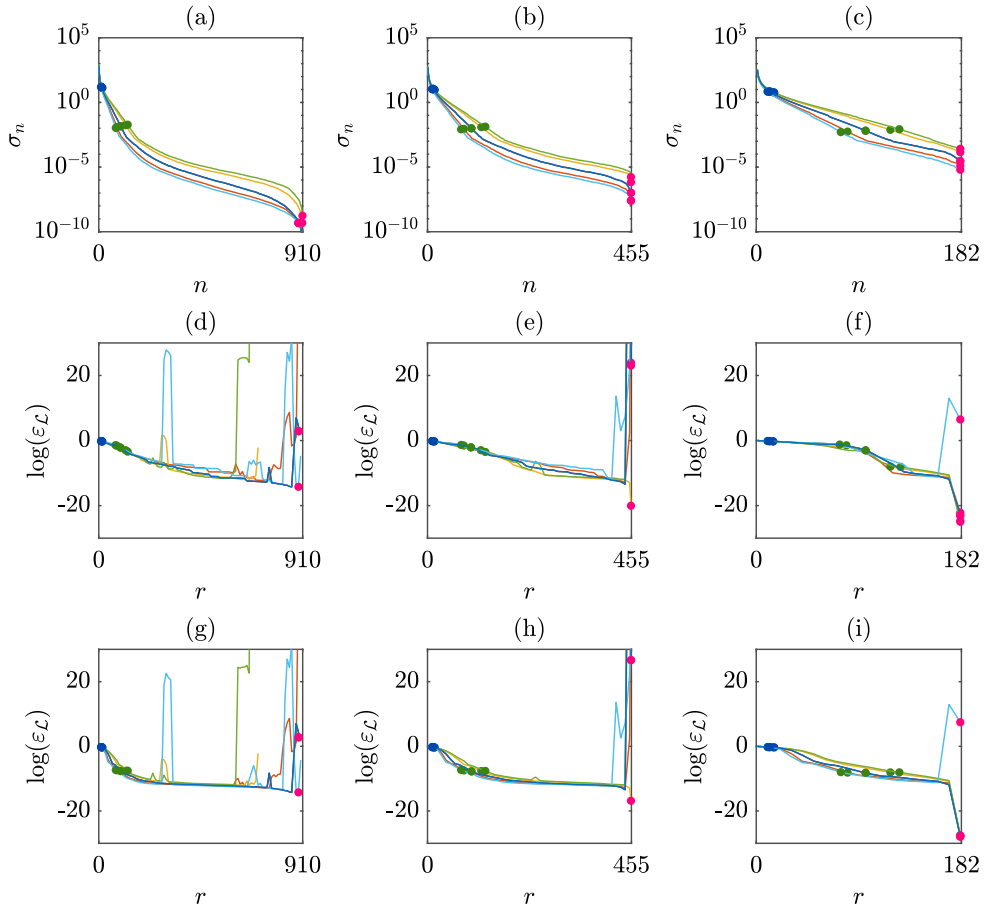

Fig. 3. For the first scenario of the permeability field (SC1): (a)–(c) represent the singular values associated with the SVD of X for cases 1-3 in Table 2, respectively. For the same cases (d)–(f) and (g)–(h) show the representation error (averaged over the time instants) of the sDMD and xDMD respectively as function of the truncation rank r of the SVD. Different lines in each panel correspond to the different dataset produced with the HFM for each collocation point. In each panel dots correspond to rank(X) (in red), r^* (in green), r_{90} (in blue) (the numerical values are in Table 2).

criteria are also reported: pointed in red is $r = \text{rank}(\mathbf{X})$, in green $r = r^*$, and blue $r = r_{90}$.

Figs. 3 and 4(a)–(c) show that, for both SC1 and SC2, values of r_{90} are similar for the three cases, and the same holds for r^* (see also Table 2 for the exact values). On the other hand, rank($\mathbf{X}$) changes from case to case, being equal (cases 2-3) or very close (case 1) to the number of training snapshots. In addition, most of the information is carried by a limited number of singular values which is captured by r^*, and that does not change increasing the dimension of the training set. As shown in Figs. 3 and 4(d)–(i), this is reflected in the representation error, $\varepsilon_{\mathcal{L}}$ (averaged over all the time instants), which becomes unstable after r^* for some of the high-fidelity simulations associated with different collocation points (in particular for SC2), suggesting all the exceeding features (correspondent to $r > r^*$) can be interpreted as noise or negligible details. For this reason, in the subsequent analyses, related to the use of DMD in interpolation, we truncate the DMD models at $r = r^*$.

Table 3 collects the values of the (time-averaged) representation error for both the algorithms, sDMD and xDMD, in case of SC1/SC2 and $r = r^*$, r_{90}, averaged for the six collocation points. These values complement what observed in Figs. 3 and 4(d)–(i). In particular, the performance of sDMD and xDMD is comparable in the case of $r = r_{90}$, with xDMD returning a slightly lower error in both SC1 and SC2. On the contrary, in the case of $r = r^*$, if for the sDMD the accuracy improves by about one order of magnitude, for the xDMD the error decreases by seven orders with respect to the value associated with $r = r_{90}$ (i.e. the representation error decreases rapidly from r_{90} to r^* for the xDMD, see also Figs. 3 and 4(g)–(i)). We also note that the error increases slowly with η, except in the case of the xDMD truncated at $r = r^*$. Furthermore, the variation is small passing from SC1 to SC2, denoting the robustness of the DMD approach under different conditions and dimensions of the training set.

Table 3
Representation error, $\varepsilon_{\mathcal{L}}$, of the DMD algorithms truncated at r^* (green dots in Figs. 3–4 (d)–(i)) and r_{90} (blue dots in Figs. 3–4 (d)–(i)), for the different scenarios of the permeability field (SC1, SC2), averaged for the six collocation points.

DMD	Case1: $\eta = 0.5$	Case2: $\eta = 0.75$	Case3: $\eta = 0.9$
SC1, sDMD, r^*	1.32 E−02	1.55 E−02	1.61 E−02
SC1, sDMD, r_{90}	6.16 E−01	6.24 E−01	6.36 E−01
SC1, xDMD, r^*	3.15 E−08	2.78 E−08	7.94 E−09
SC1, xDMD, r_{90}	5.32 E−01	5.37 E−01	5.50 E−01
SC2, sDMD, r^*	4.29 E−03	5.13 E−03	3.28 E−03
SC2, sDMD, r_{90}	6.64 E−01	6.67 E−01	6.76 E−01
SC2, xDMD, r^*	2.08 E−08	1.76 E−08	7.75 E−09
SC2, xDMD, r_{90}	5.63 E−01	5.67 E−01	5.78 E−01

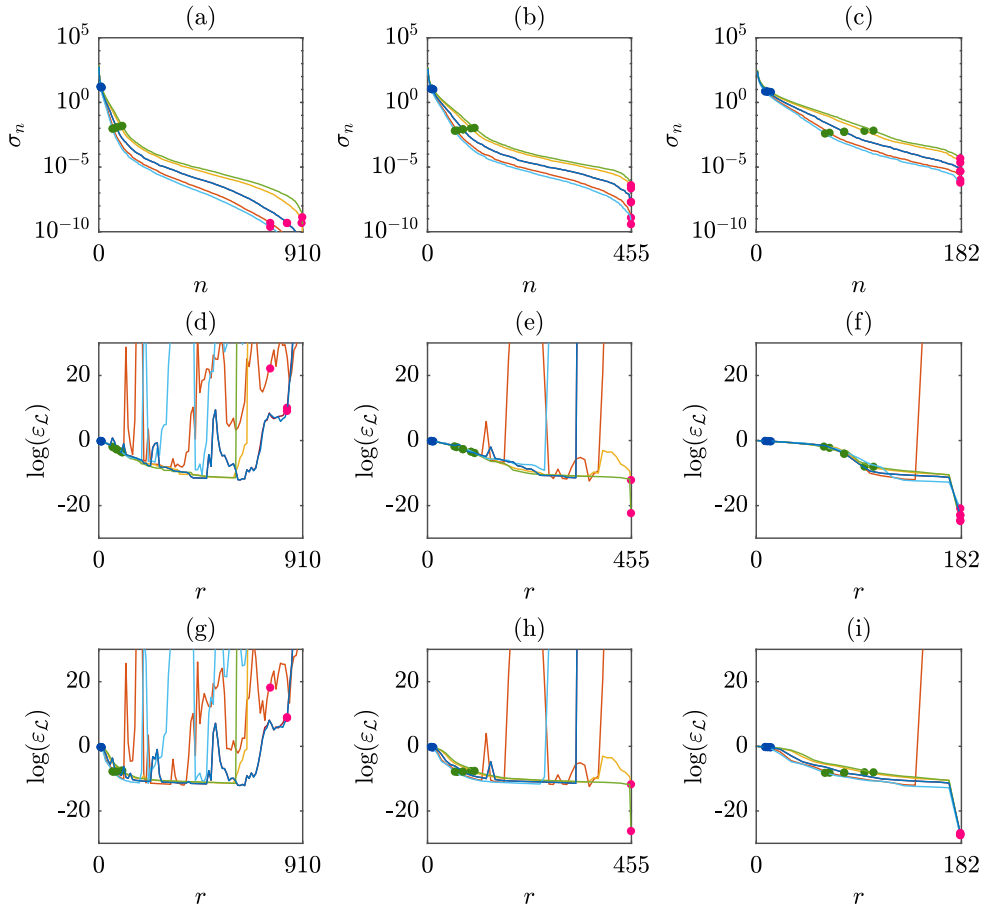

Fig. 4. For the second scenario of the permeability field (SC2): (a)–(c) represent the singular values associated with the SVD of X for cases 1-3 in Table 2, respectively. For the same cases (d)–(f) and (g)–(h) show the representation error (averaged over the time instants) of the sDMD and xDMD respectively as function of the truncation rank r of the SVD. Different lines in each panel correspond to the different dataset produced with the HFM for each collocation point. In each panel dots correspond to rank(X) (in red), r^* (in green), r_{90} (in blue) (the numerical values are in Table 2).

4.2. DMD interpolation error

Once calibrated on high-fidelity data, sDMD and xDMD are employed to interpolate the model response. We consider the DMD models trained in the previous Section, i.e., on 50%, 25%, and 10% of the overall set of high-fidelity data, and truncated at $r = r^*$, to reconstruct saturation maps for the missing time steps, i.e. thus corresponding to $\eta = 0.5, 0.75, 0.9$, respectively (see Table 2 cases 1-3).

Figs. 5 and 6(a)–(f) show the interpolation error in time, $\epsilon^k_{\mathcal{L}}$, for permeability scenarios SC1 and SC2, respectively. In both the Figures, columns correspond to cases 1-3 (Table 2). Specifically, the performance of the sDMD is represented in panels (a)–(c), while the xDMD is in panels (d)–(f). Also in this case, the interpolation error is computed for each collocation point (Table 1) through Eq. (11). The logarithm of the error of the sDMD is in the range $(-6, -1)$ for all the cases, showing an increasing variability for all the collocation points when the dimension of the training set decreases. Conversely, for the xDMD the range is sensibly lower within $(-9, -3)$, and the variability of the error is less remarkable; also the error generally decreases rapidly with time for the different cases and collocation points, denoting a more stable behavior of this algorithm. This holds for both SC1 and SC2. We can also observe that the sDMD algorithm is more sensitive than the xDMD to the HFM parametric variability: the interpolation error displays different trends and values for the different collocation points (i.e. values of ϕ and Δp) and permeability fields (SC1/SC2).

Given the higher accuracy of the xDMD, the difference between the high-fidelity saturation map and the corresponding map obtained by xDMD interpolation at $k = M - 1$ is represented in Figs. 5 and 6(g)–(i), for cases 1-3 and for SC1 and SC2, respectively. The difference is maximum for case 3 though in the range $(-5 \cdot 10^{-5}, 5 \cdot 10^{-5})$. These interpolation errors propagate in the DMD-PCE approach inducing the approximation discussed below.

4.3. Assessment of DMD-PCE accuracy

We represent the error of the proposed DMD-enhanced PCE approach (DMD-PCE) relative to the PCE obtained only with high-fidelity data (HF-PCE). As shown in Fig. 2, according to the DMD-PCE method, the snapshots produced by interpolation with either the sDMD or xDMD (Section 4.2) are mixed with the high-fidelity snapshots employed for the DMD training (Section 4.1) with a 50/50% ($\eta = 0.5$, case 1), 75/25% ($\eta = 0.75$, case 2), 90/10% ($\eta = 0.9$, case 3), as summarized in Table 2. These three mixed datasets are associated with a decreasing computational cost for their generation. The PCE is described by (18) and the PCE coefficients are computed with (2) at all the space–time locations of interest (see Section 3.2) based on each one of these mixed datasets as well as only on high-fidelity data for comparison.

To assess the DMD-PCE accuracy, we first analyze the error in the estimate of the PCE coefficients. Specifically we compute the logarithm of $\epsilon^{k^\dagger}_{a_j}$ (see Eq. (13)) for the selected $\eta = 0.5, 0.75, 0.9$. For SC1, Fig. 7

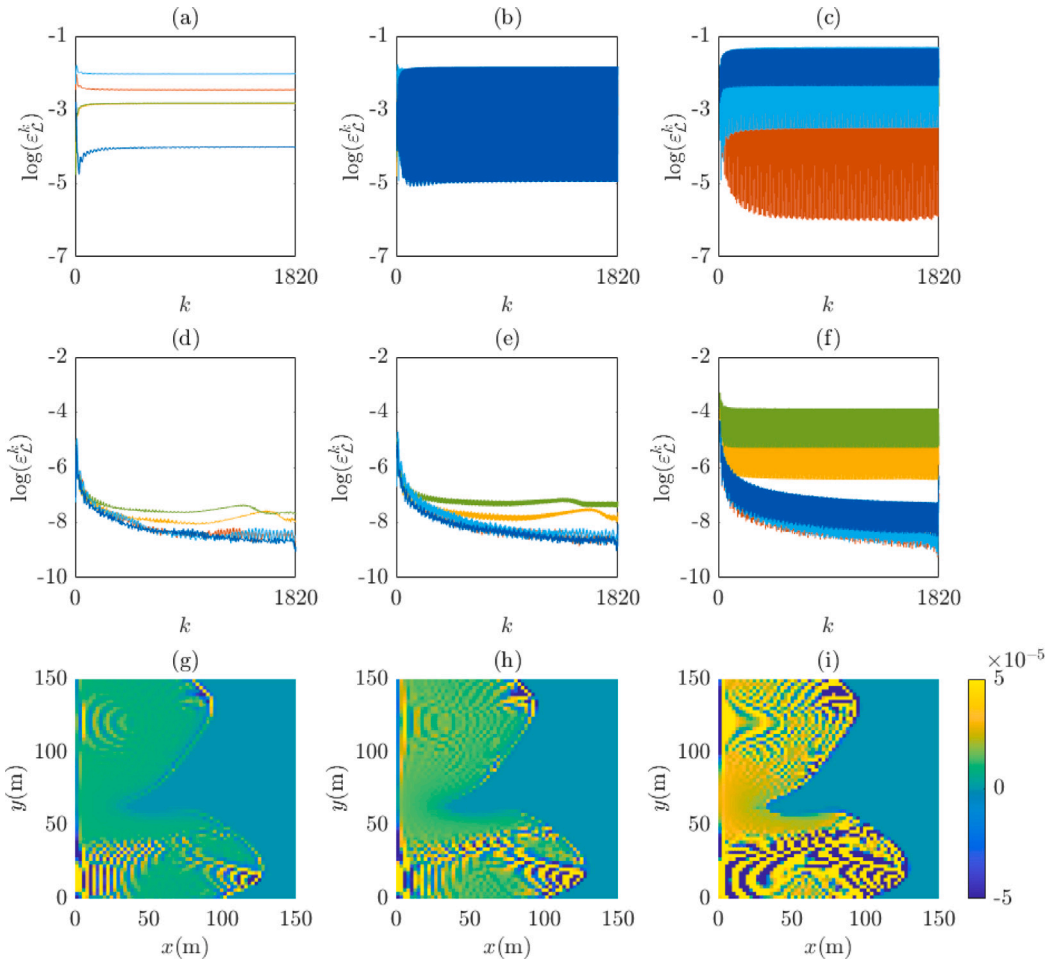

Fig. 5. For the first scenario of the permeability field (SC1): (a)–(c) represent the interpolation error in time for the sDMD for cases 1-3 in Table 2 (different lines in each panel correspond to the different collocation point); (d)–(f) show the same results for the xDMD; (g)–(h) show the difference between high-fidelity data and the predictions produced by the xDMD in interpolation at $k = M - 1$ for cases 1-3.

represents in each panel the box plot of this error (at the different space–time locations) associated with each coefficient a_j, with $j = 0, \ldots, 5$, in the first row for the sDMD and in the second row for the xDMD. Fig. 8 shows the same for SC2. We may observe that the error generally increases with η, and that the xDMD returns a sensibly more accurate approximation than the sDMD. In the case of $\eta = 0.9$ the error increases more significantly but still the approximation is acceptable for the xDMD thus denoting the robustness of this interpolation algorithm for the selected case study despite the nonlinearity of the phenomena and the different scenarios of heterogeneity.

The accuracy of the approach is finally assessed by comparing the prediction of the response at each space location and $t_{k^\dagger}$ provided by the DMD-PCE for either $\eta = 0.5$, 0.75 or 0.9, i.e. $\hat{\omega}^{\mathrm{DMD}}_{i,k^\dagger}$, with the prediction at the same space–time locations provided by the HF-PCE ($\eta = 0$), i.e. $\hat{\omega}^{\mathrm{HF}}_{i,k^\dagger}$ (see Section 2.3). Figs. 9 and 10 represent the regression plot to compare these predictions for SC1 and SC2, respectively. In these Figures the performance of the sDMD and xDMD are shown on the first and second row, respectively; while the three columns correspond to cases 1-3 for $\eta = 0.5$, 0.75 or 0.9. We observe that for SC1 and $\eta = 0.5$ (Fig. 9(a),(d)) the PCE informed by both the DMD algorithms returns a prediction sensibly equal to the one of the HF-PCE; while moving to SC2 (Fig. 10(a),(d)), a decrease in the accuracy of the PCE informed by the sDMD is already detectable. The remarkable accuracy of the PCE based on the xDMD interpolation holds for increasing η in both SC1 (Fig. 9(e)–(f)) and SC2 (Fig. 10(e)–(f)). Conversely, the accuracy of the PCE fed by the sDMD rapidly decreases for increasing η in the two scenarios (Figs. 9 and 10(b)–(c)).

5. Conclusion

The proposed DMD-PCE is robust and accurate for multiphase flow in heterogeneous media selected to exemplify and assess the method. In particular, the PCE informed by interpolated snapshots produced with the xDMD algorithm shows remarkable accuracy and a negligible sensitivity towards the parametric variability, mimicking different flow conditions. This is observed for $\eta = 0.5, 0.75$. However, even if for $\eta = 0.9$ the interpolation error of the xDMD becomes more sensitive to the model parameters, the accuracy of the resulting DMD-PCE is sensibly the same, thus unraveling huge potential of the proposed method to drastically reduce the computational cost associated with the calculation of the model response at each space–time location of interest multiple times equal to the number of collocation points. The maintenance of high accuracy in the DMD-PCE prediction for increasing interpolation rate is understandable looking at the error in the estimate of the PCE coefficients that increases by about two orders of magnitude in the case of 90% but still below 10^{-4} for all the coefficients with the xDMD. The same does not hold when the PCE is informed by the sDMD. In that case, only for a $\eta = 0.5$ the accuracy is comparable to the HF-PCE. In conclusion, our results on the proposed DMD-PCE method, based on the xDMD algorithm introduced by Lu and Tartakovsky (2021), suggest it may have critical implications in model reduction of hydrological models with time-variant response allowing, among others, for uncertainty quantification and scenarios development at an even lower computational cost with respect to the traditional PCE method relying only on high-fidelity simulations.

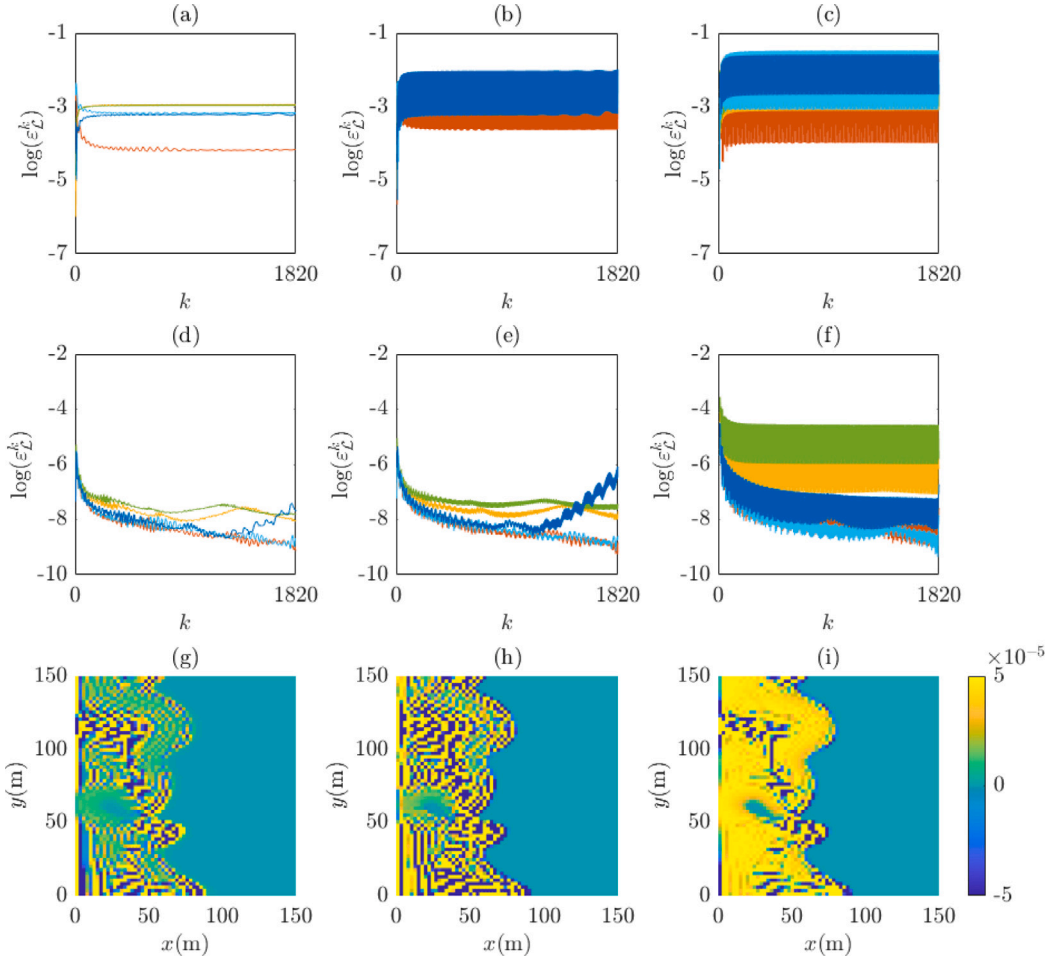

Fig. 6. For the second scenario of the permeability field (SC2): (a)–(c) represent the interpolation error in time for the sDMD for cases 1-3 in Table 2 (different lines in each panel correspond to the different collocation point); (d)–(f) show the same results for the xDMD; (g)–(h) show the difference between high-fidelity data and the predictions produced by the xDMD in interpolation at $k = M - 1$ for cases 1-3.

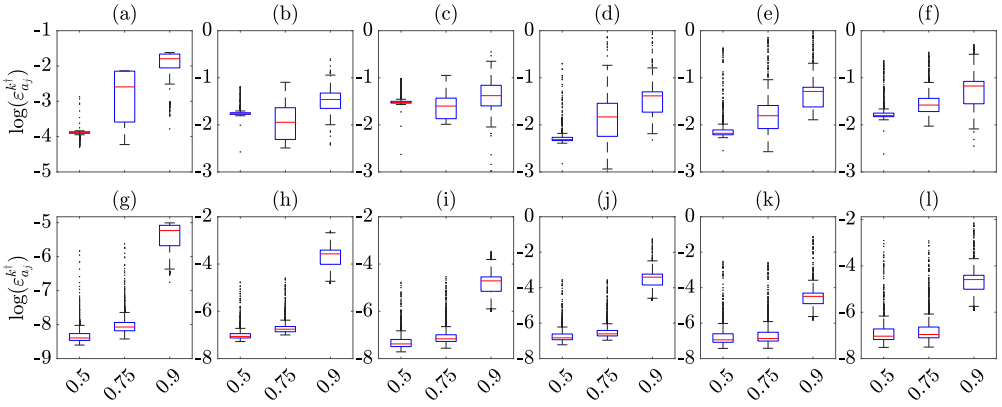

Fig. 7. For the first scenario of the permeability field (SC1): error in the estimate of PCE coefficients when using DMD interpolated data in place of the high-fidelity data for $\eta = 0.5, 0.75, 0.9$. The first row shows results for the sDMD and the second one for the xDMD.

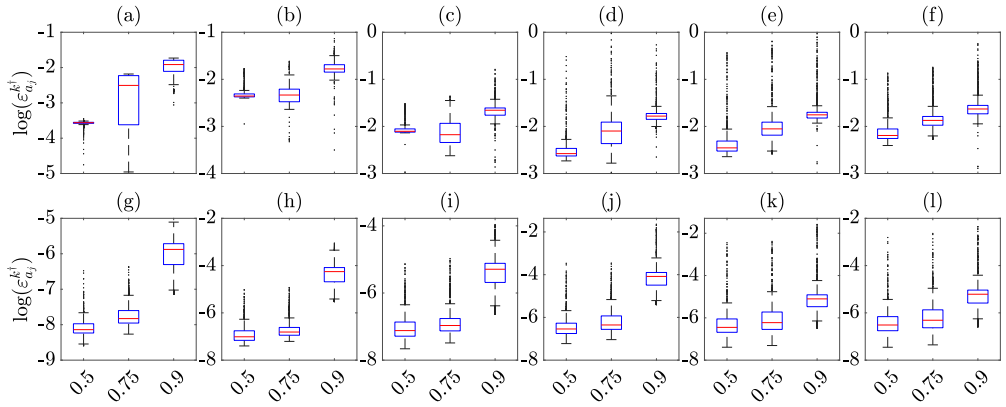

Fig. 8. For the second scenario of the permeability field (SC2): error in the estimate of PCE coefficients when using DMD interpolated data in place of the high-fidelity data for $\eta = 0.5, 0.75, 0.9$. The first row shows results for the sDMD and the second one for the xDMD.

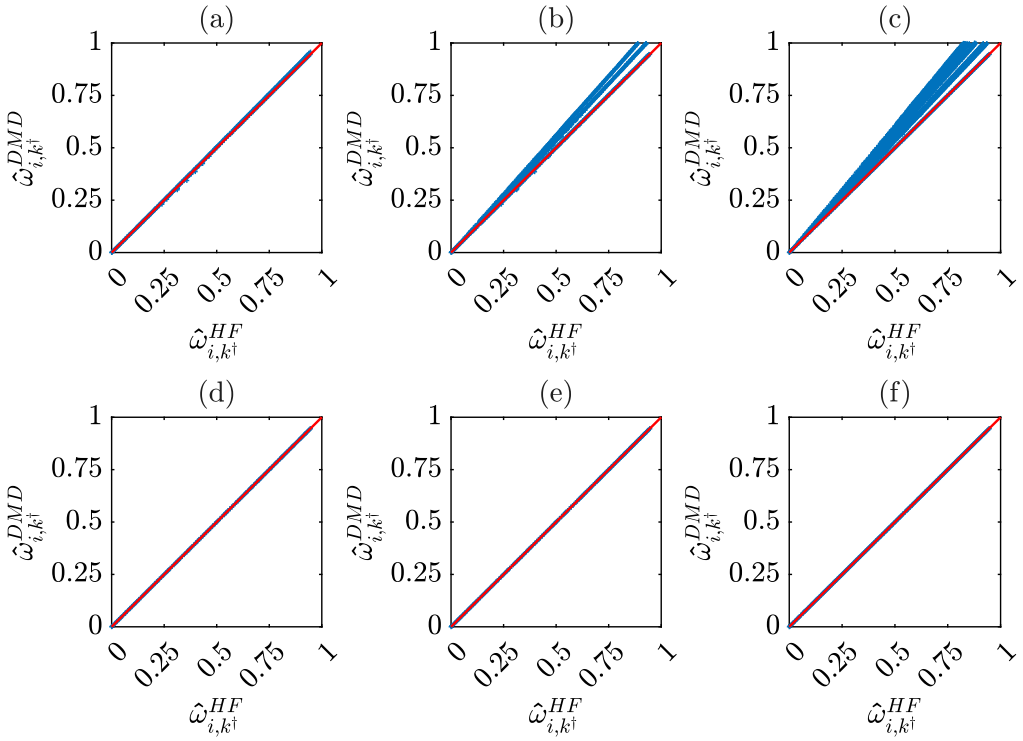

Fig. 9. For the first scenario of the permeability field (SC1): (a)–(c) show, for the sDMD, the DMD-PCE predictions associated with $\eta = 0.5, 0.75, 0.9$ respectively vs the HF-PCE predictions (obtained for $\eta = 0$); (d)–(f) show the same results for the xDMD.

CRediT authorship contribution statement

G. Libero: Writing – original draft, Software, Investigation, Methodology, Funding acquisition, Data curation, Conceptualization. **D.M. Tartakovsky:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization. **V. Ciriello:** Writing – review & editing, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

There are no data sharing issues since all of the numerical information is provided in the figures produced by solving the equations in the paper.

Acknowledgments

We thank Niklas Linde (University of Lausanne), Sergey Oladyshkin (University of Stuttgart) and the third anonymous reviewer for their constructive comments. GL acknowledges support from the Marco Polo mobility scholarship for research 2021 of the University of Bologna. VC acknowledges support from University of Bologna RFO (Ricerca Fondamentale Orientata) 2020–2021. DT was funded, in part, by the Air Force Office of Scientific Research, United States under grant FA9550-21-1-0381, by the Office of Advanced Scientific Computing Research (ASCR), United States within the Department of Energy Office

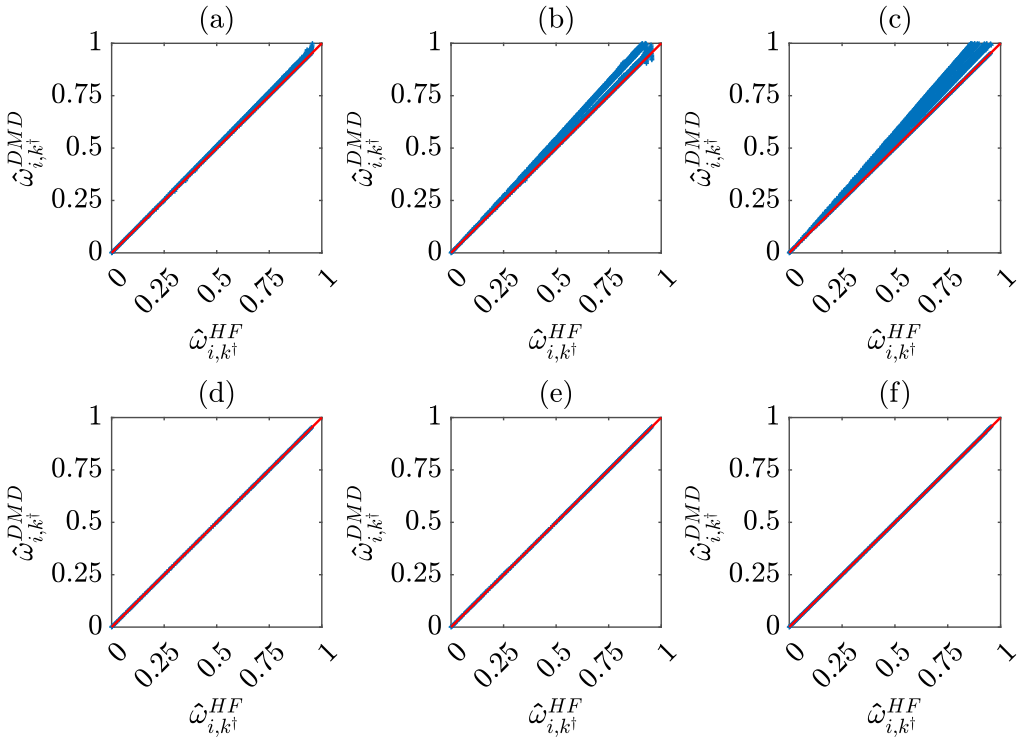

Fig. 10. For the second scenario of the permeability field (SC2): (a)–(c) show, for the sDMD, the DMD-PCE predictions associated with $\eta = 0.5, 0.75, 0.9$ respectively vs the HF-PCE predictions (obtained for $\eta = 0$); (d)–(f) show the same results for the xDMD.

of Science under award number DE-SC0023163, and by the Strategic Environmental Research and Development Program (SERDP), United States of the Department of Defense under award RC22-3278.

References

Appleyard, J.R., Cheshire, I.M., Pollard, R.K., 1981. Special techniques for fully implicit simulators. In: Proc. European Symp. on Enhanced Oil Recovery, Bournemouth, UK. pp. 395–408.

Barajas-Solano, D.A., Tartakovsky, D.M., 2016. Stochastic collocation methods for nonlinear parabolic equations with random coefficients. SIAM/ASA J. Uncert. Quant. 4 (1), 475–494. http://dx.doi.org/10.1137/130930108.

Ciriello, V., de Barros, F.P.J., 2020. Characterizing the influence of multiple uncertainties on predictions of contaminant discharge in groundwater within a Lagrangian stochastic formulation. Water Resour. Res. 56 (10), http://dx.doi.org/10.1029/2020wr027867.

Ciriello, V., Di Federico, V., Riva, M., Cadini, F., Sanctis, J.D., Zio, E., Guadagnini, A., 2013. Polynomial chaos expansion for global sensitivity analysis applied to a model of radionuclide migration in a randomly heterogeneous aquifer. Stoch. Environ. Res. Risk Assess. 27 (4), 945–954. http://dx.doi.org/10.1007/s00477-012-0616-7.

Ciriello, V., Lauriola, I., Bonvicini, S., Cozzani, V., Di Federico, V., Tartakovsky, D.M., 2017. Impact of hydrogeological uncertainty on estimation of environmental risks posed by hydrocarbon transportation networks. Water Resour. Res. 53 (11), 8686–8697. http://dx.doi.org/10.1002/2017wr021368.

Ciriello, V., Lauriola, I., Tartakovsky, D.M., 2019. Distribution-based global sensitivity analysis in hydrology. Water Resour. Res. 55 (11), 8708–8720. http://dx.doi.org/10.1029/2019wr025844.

Corey, A.T., 1954. The interrelation between gas and oil relative permeabilities. Producers Monthly 19, 38–41.

Focaccia, S., Panini, G., Pedrazzoli, P., Ciriello, V., 2021. A meta-modeling approach for hydrological forecasting under uncertainty: Application to groundwater nitrate response to climate change. J. Hydrol. 603, 127173. http://dx.doi.org/10.1016/j.jhydrol.2021.127173.

Ghanem, R.G., Spanos, P.D., 1991. Stochastic Finite Elements: A Spectral Approach. Springer, New York, http://dx.doi.org/10.1007/978-1-4612-3094-6.

He, Q., Tartakovsky, A.M., 2021. Physics-informed neural network method for forward and backward advection-dispersion equations. Water Resour. Res. 57, http://dx.doi.org/10.1029/2020wr029479.

Kang, X., Kokkinaki, A., Shi, X., Yoon, H., Lee, J., Kitanidis, P.K., Wu, J., 2022. Integration of deep learning-based inversion and upscaled mass-transfer model for DNAPL mass-discharge estimation and uncertainty assessment. Water Resour. Res. 58 (10), http://dx.doi.org/10.1029/2022WR033277, e2022WR033277.

Kutz, J.N., Brunton, S.L., Brunton, B.W., Proctor, J.L., 2016. Dynamic Mode Decomposition. Society for Industrial and Applied Mathematics, http://dx.doi.org/10.1137/1.9781611974508.

Lin, G., Tartakovsky, A.M., Tartakovsky, D.M., 2010. Uncertainty quantification via random domain decomposition and probabilistic collocation on sparse grids. J. Comput. Phys. 229 (19), 6995–7012. http://dx.doi.org/10.1016/j.jcp.2010.05.036.

Lu, H., Tartakovsky, D.M., 2020. Prediction accuracy of dynamic mode decomposition. SIAM J. Sci. Comput. 42 (3), A1639–A1662. http://dx.doi.org/10.1137/19M1259948.

Lu, H., Tartakovsky, D.M., 2021. Extended dynamic mode decomposition for inhomogeneous problems. J. Comput. Phys. 444, 110550. http://dx.doi.org/10.1016/j.jcp.2021.110550.

Marzadri, A., Ciriello, V., de Barros, F.P.J., 2024. Hyporheic flows in stratified sediments: Implications on residence time distributions. Water Resour. Res. 60 (1), http://dx.doi.org/10.1029/2023WR035625, e2023WR035625.

Meles, G.A., Amaya, M., Levy, S., Marelli, S., Linde, N., 2023. Bayesian tomography using polynomial chaos expansion and deep generative networks. Geophys. J. Int. 237 (1), 31–48. http://dx.doi.org/10.1093/gji/ggae026.

Meles, G.A., Linde, N., Marelli, S., 2022. Bayesian tomography with prior-knowledge-based parametrization and surrogate modelling. Geophys. J. Int. 231 (1), 673–691. http://dx.doi.org/10.1093/gji/ggac214.

Mohammadi, F., Kopmann, R., Guthke, A., Oladyshkin, S., Nowak, W., 2018. Bayesian selection of hydro-morphodynamic models under computational time constraints. Adv. Water Resour. 117, 53–64. http://dx.doi.org/10.1016/j.advwatres.2018.05.007.

Oladyshkin, S., de Barros, F., Nowak, W., 2012. Global sensitivity analysis: A flexible and efficient framework with an example from stochastic hydrogeology. Adv. Water Resour. 37, 10–22. http://dx.doi.org/10.1016/j.advwatres.2011.11.001.

Oladyshkin, S., Nowak, W., 2018. Incomplete statistical information limits the utility of high-order polynomial chaos expansions. Reliab. Eng. Syst. Saf. 169, 137–148. http://dx.doi.org/10.1016/j.ress.2017.08.010.

Schmid, P.J., 2010. Dynamic mode decomposition of numerical and experimental data. J. Fluid Mech. 656, 5–28. http://dx.doi.org/10.1017/s0022112010001217.

Shi, W., Tartakovsky, D.M., 2022. Polynomial chaos expansions for stiff random ODEs. SIAM J. Sci. Comput. 44 (3), A1021–A1046. http://dx.doi.org/10.1137/21M1432545.

Song, D.H., Tartakovsky, D.M., 2021. Transfer learning on multi-fidelity data. J. Mach. Learn. Model. Comput. 3 (1), 31–47. http://dx.doi.org/10.1615/JMachLearnModelComput.2021038925.

Sudret, B., 2008. Global sensitivity analysis using polynomial chaos expansions. Reliab. Eng. Syst. Saf. 93 (7), 964–979. http://dx.doi.org/10.1016/j.ress.2007.04.002.

Tartakovsky, A.M., Marrero, C.O., Perdikaris, P., Tartakovsky, G., Barajas-Solano, D., 2020. Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems. Water Resour. Res. 56 (5), http://dx.doi.org/10.1029/2019wr026731.

Taverniers, S., Bosma, S.B.M., Tartakovsky, D.M., 2020. Accelerated multilevel Monte Carlo with kernel-based smoothing and Latinized stratification. Water Resour. Res. 56 (9), http://dx.doi.org/10.1029/2019wr026984.

Tu, J.H., Rowley, C.W., Luchtenburg, D.M., Brunton, S.L., Kutz, J.N., 2014. On dynamic mode decomposition: Theory and applications. J. Comput. Dynam. 1 (2), 391–421. http://dx.doi.org/10.3934/jcd.2014.1.391.

Um, K., Hall, E.J., Katsoulakis, M.A., Tartakovsky, D.M., 2019. Causality and Bayesian Network PDEs for multiscale representations of porous media. J. Comput. Phys. 394, 658–678. http://dx.doi.org/10.1016/j.jcp.2019.06.007.

Webster, M., Tatang, M.A., McRae, G.J., 1996. Application of the Probabilistic Collocation Method for an Uncertainty Analysis of a Simple Ocean Model. MIT Joint Program on the Science and Policy of Global Change Reports Series No. 4, MIT, Cambridge, MA.

Xiu, D., Karniadakis, G.E., 2002. The Wiener-Askey polynomial chaos for stochastic differential equations. SIAM J. Sci. Comput. 24 (2), 619–644. http://dx.doi.org/10.1137/s1064827501387826.

Zhan, C., Dai, Z., Soltanian, M.R., de Barros, F.P.J., 2022. Data-worth analysis for heterogeneous subsurface structure identification with a stochastic deep learning framework. Water Resour. Res. 58 (11), http://dx.doi.org/10.1029/2022WR033241, e2022WR033241.