



HAL
open science

Bayesian inference of spatially varying Manning's n coefficients in an idealized coastal ocean model using a generalized Karhunen-Loève expansion and polynomial chaos

Adil Siripatana, Olivier Le Maitre, Omar Knio, Clint Dawson, Ibrahim Hoteit

► To cite this version:

Adil Siripatana, Olivier Le Maitre, Omar Knio, Clint Dawson, Ibrahim Hoteit. Bayesian inference of spatially varying Manning's n coefficients in an idealized coastal ocean model using a generalized Karhunen-Loève expansion and polynomial chaos. *Ocean Dynamics*, 2020, 70 (8), pp.1103-1127. 10.1007/s10236-020-01382-4 . hal-03010151

HAL Id: hal-03010151

<https://hal.science/hal-03010151v1>

Submitted on 17 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bayesian Inference of Spatially-Varying Manning’s n Coefficients in an Idealized Coastal Ocean Model Using a Generalized Karhunen-Loève Expansion and Polynomial Chaos

Adil Siripatana^a, Olivier Le Maitre^b, Omar Knio^a, Clint Dawson^c, Ibrahim Hoteit^{a,*}

^aKing Abdullah University of Science and Technology, Thuwal, KSA, adil.siripatana@kaust.edu.sa, omar.knio@kaust.edu.sa and ibrahim.hoteit@kaust.edu.sa

^bCNRS, Centre de Mathématiques Appliquées, INRIA, École Polytechnique, IPP, Route de Saclay, 91128 Palaiseau, France olivier.le-maitre@polytechnique.edu

^cUniversity of Texas at Austin, Texas, USA, clint@oden.utexas.edu

Abstract

Bayesian inference with coordinate transformations and polynomial chaos for a Gaussian process with a parametrized prior covariance model was introduced in [61] to enable and infer uncertainties in a parameterized prior field. The feasibility of the method was successfully demonstrated on a simple transient diffusion equation. In this work, we adopt a similar approach to infer a spatially varying Manning’s n field in a coastal ocean model. The idea is to view the prior on the Manning’s n field as a stochastic Gaussian field, expressed through a covariance function with uncertain hyper-parameters. A generalized Karhunen-Loève (KL) expansion, which incorporates the construction of a reference basis of spatial modes and a coordinate transformation, is then applied to the prior field. To improve the computational efficiency of the method proposed in [61], we propose to use two polynomial chaos expansions to: (i) approximate the coordinate transformation, and (ii) build a cheap surrogate of the large-scale advanced circulation (ADCIRC) numerical model. These two surrogates are used to accelerate the Bayesian inference process using a Markov chain Monte Carlo algorithm. Water elevation data are inverted within an observing system simulation experiment framework, based on a realistic ADCIRC model, to infer the KL coordinates and hyper-parameters of a reference 2D Manning’s field. Our results demonstrate the efficiency of the proposed approach and suggest that including the hyper-parameter uncertainties greatly enhances the inferred Manning’s n field, compared to using a covariance with fixed hyper-parameters.

Keywords: Coastal ocean model, Manning’s n coefficients, Bayesian inference, Karhunen-Loève expansion, Polynomial Chaos, MCMC

1. Introduction

The global economy and security rely heavily on ocean productivity, with approximately 50% of the human population being located next to coastal and estuarine systems [75, 46]. Accurate forecasts of the dynamics of the coastal ocean are crucial to all aspects of marine activities. These include, but are not limited to, fisheries, nearshore construction, tourism, recreational activities, extreme event mitigation, contaminant transport management, and ocean conservation. In particular, the information gained from coastal ocean simulations and forecasts is critical for the authorities to make informed decisions and develop smart and sustainable policies to exploit the ocean resources.

Advanced coastal ocean models based on the shallow water equations (SWEs) have been widely utilized to simulate and predict shallow water circulation [29]. The SWEs are based on the simplified Navier-Stokes

*Corresponding author: 4700 King Abdullah University of Science and Technology, Thuwal 23955-6900, Kingdom of Saudi Arabia.

equations, which assume hydrostatic pressure and relatively large horizontal length scales with respect to vertical length scales [42]. Even when the model is assumed to be perfect in describing the underlying physical processes, the model solutions are always subject to uncertainties that emerge on account of errors originating from numerous sources [20], e.g., initial conditions, model parameters, forcing fields, and resolution of the implemented numerical methods. To provide robust and reliable forecasts of the ocean states, uncertainty quantification and reduction techniques have, in recent years, played a major role toward enhancing state-of-the-art coastal ocean simulation systems [20, 59, 28, 1, 46, 62, 57].

This work considers the uncertainty reduction of the model parameters in the coastal ocean, focusing on the inference of the Manning’s n coefficient of roughness, introduced in the SWEs through the bottom stress components in the momentum equation [29, 32]. The Manning’s n coefficients describe the bottom surface characteristics, which further define the resistance of water flow. Mayo et al. 2014 [46] demonstrated that water elevation is highly sensitive to the change in this parameter, particularly in shallow coastal areas. The Manning’s n coefficients are highly variable in space but cannot be measured directly [8], and as such often empirically derived. To infer such a parameter from the limited available observations, one may resort to Bayesian inversion [26, 6, 33, 60]. The Bayesian inference framework views the parameter estimation problem as the inference of a probability density function (pdf). A prior pdf is then updated into a posterior pdf by incorporating the likelihood of the observations [69].

Several well-established approaches to Bayesian inference are widely used in parameters inference problems, the most prominent of which is the Markov chain Monte Carlo (MCMC) method [25, 19, 49], which directly samples the posterior of the inferred parameters given the sample points of likelihood. A large number of model runs is required for MCMC to efficiently sample the posterior distribution [62, 23]. This renders the MCMC implementation computationally prohibitive for large-scale systems. To alleviate this hurdle, model reduction techniques have been exploited, both to build inexpensive surrogates of the forward models and to reduce the dimension of the parameter estimation problems. Of particular interest is the spectral projection method used with the polynomial chaos (PC) basis proposed by [45]. The PC method expresses the dependence of the model predictions on the model parameters using orthogonal stochastic polynomials. This spectral expansion enables the construction of extremely cost-effective surrogate models that can be used as substitutes of the full model predictions to efficiently sample the likelihood [20, 35]. The PC method has been successfully implemented with a wide scope of large-scale problems, including mechanical systems (e.g., [38, 21, 37]), chemical systems (e.g., [48, 39, 47]), and more recently, large-scale oceanic systems (e.g., [1, 62, 63, 58]). Nevertheless, the PC method suffers from the ‘curse of dimensionality’ [1], in which only a limited number of stochastic parameters can be accommodated to build the PC surrogate before it becomes computationally prohibitive. To alleviate this issue, one may implement a parameter search space reduction technique, namely, Karhunen-Loève (KL) expansion [40, 31]. This technique expresses stochastic processes in terms of an orthonormal set of eigenfunctions, which can be used to parametrize highly spatially variant parameters with only a few dominant modes. Marzouk *et al.* [44] demonstrated the efficiency of combining PC and KL schemes to infer the spatially varying diffusivity of a medium in the context of a transient diffusion equation.

Sequential Bayesian inference method, such as the ensemble Kalman filter (EnKF) [4, 7, 27], has also been intensively utilized for parameter estimation of ocean models (e.g., [2, 3, 5, 11, 55, 66, 70]). The EnKF has been found efficient, with advantages over the MCMC approach in accommodating large state-parameter vectors at reasonable computational cost [46, 57, 58]. Recently, the EnKF has been successfully used for the inference of a spatially varying Manning’s n field in an idealized coastal ocean framework, both in the full and reduced KL space [57]. The main merit in using an EnKF is its ability to accommodate larger member of parameters or KL modes. However, exploring the posteriors of these coordinates is limited to the directions defined by the subspace of the prior ensemble covariance model [17]. The prior covariance matrix describing the Gaussian process needs to therefore be well known *a priori*.

In both MCMC- and EnKF-based inference, presented in [44] and [57], respectively, the underlying Gaussian process is associated with prior covariance models that are often poorly known. This is because a prior covariance model that describes the stochastic parameter field is subjected to the uncertainty in hyper-parameters, e.g., prior variance and length scales. An early attempt to include the hyper-parameter uncertainty in Bayesian inference problems was proposed by Tagade and Choi [67]. In this study, a KL

expansion of a stochastic process is derived as a function of the hyper-parameters. The hyper-parameters are then cast as stochastic variables in the PC expansion, which can then be directly estimated alongside the KL coordinates via MCMC. However, this method can become computationally prohibitive due to the requirement in building the PC surrogate with increasing augmented stochastic dimensions.

In this work, we adopt the generalized KL expansion with the coordinate transformation approach introduced in [61] to account for the covariance hyper-parameters in a PC-based MCMC inference of spatially varying Manning’s n coefficients of a coastal ocean model. This approach only requires building the PC surrogate based on a KL basis decomposed from a reference covariance matrix. A change of basis methodology is then utilized to transform the KL coordinates based on one basis to a different basis. A similar coordinate transformation is finally applied to the KL expansion of the reference basis to obtain the PC expansion for any other basis required for the inference. Here, we avoid introducing the hyper-parameters as additional stochastic dimensions to the PC expansion. However, the computation of the coordinate transformation matrix proposed in [61] can be computationally quite demanding, since it requires decomposing of the prior covariance corresponding to each set of hyper-parameters. Therefore, we propose to introduce the construction of a PC surrogate of the coordinate transformation matrix. This present contribution is essential for the application of a generalized KL expansion to any large-scale fluid system, which involve very large prior covariance matrices.

The rest of this paper is organized as follows. Section 2 presents our Bayesian inference framework. Section 3 describes the change of coordinates methodology for incorporating hyper-parameters uncertainty to KL and PC expansions. The experimental setting for the evaluation of our proposed inference framework is outlined in Section 4. In Section 5, the results of the inference of the Manning’s n coefficients coastal ocean model are presented and discussed. Concluding remarks are finally offered in Section 6.

2. Bayesian inference framework

Our objective is to estimate spatially varying parameters of a coastal ocean model given a set of (indirect) observations $\mathbf{d} \in \mathbb{R}^{N_o}$, $N_o \geq 1$. The Manning’s n roughness parameters, for example, can be viewed as a deterministic field $m(\mathbf{x})$, $\mathbf{x} = (x, y) \in D$. We also assume that the observations can be predicted using the forward model (in our case the SWEs), which takes $m(\mathbf{x})$ as input and the model predictions of the observations as output. Bayes’ rule then updates the prior probability distribution function (pdf) of the unknown parameters m , conditioned on the discrepancy between the model prediction $\mathbf{u}(m)$ and observations \mathbf{d} as [51, 30]:

$$p(m, \sigma_o^2 \mid \mathbf{d}) \propto p(\mathbf{d} \mid m, \sigma_o^2) p_m(m) p_o(\sigma_o^2), \quad (1)$$

where $p(m, \sigma_o^2 \mid \mathbf{d})$ is the posterior pdf, $p(\mathbf{d} \mid m, \sigma_o^2)$ is the likelihood function of obtaining observations \mathbf{d} given the parameter field m and the observations error model hyper-parameter σ_o^2 , with priors $p_m(m)$ and $p_o(\sigma_o^2)$, respectively. In this work, we consider for simplicity an unbiased additive Gaussian error model, $\epsilon \doteq \mathbf{d} - \mathbf{u}(m)$, $\epsilon \sim \mathcal{N}(0, \sigma_o^2 \mathbf{I}_{N_o})$, where $\mathcal{N}(0, \sigma_o^2 \mathbf{I}_{N_o})$ denotes the multivariate Gaussian distribution with zero mean and diagonal covariance $\sigma_o^2 \mathbf{I}_{N_o}$.

We resort to the adaptive Metropolis-Hastings MCMC algorithm [24, 52] to efficiently sample the posterior pdf of the Manning’s n parameters m . Metropolis-Hastings algorithm iteratively generates a Markov chain based on the accept/reject methodology, in which the state of the chain at a given iteration depends only on the previous one. The high-dimensional nature of the Manning’s n field discretized on a computational mesh, however, renders the node-wise parameter inference using MCMC computationally prohibitive. A model reduction technique, which statistically describes the Manning’s n field as a stochastic process, parametrized using a few random variables, is thus required.

2.1. Karhunen-Loève (KL) expansion

The Manning’s n roughness parameters physically characterizing the seabed surface are generally continuous in space, with correlated neighboring values. It is therefore effective to describe a Manning’s n field as a stochastic process, $M(\mathbf{x}, \omega)$, where M is a function defined on the product space $\mathbf{D} \times \Omega$, $\mathbf{x} \in \mathbf{D}$

is the index set of the location of the physical domain, and $\omega \in \Omega$ is the sample space of the probability space, (Ω, \mathcal{F}, P) . Given a stochastic process $M(\mathbf{x}, \omega)$, with mean $\mu(\mathbf{x})$ and symmetric positive semi-definite covariance function $\mathbf{C}(\mathbf{x}, \mathbf{x}')$, the Mercer theorem [74] states that \mathbf{C} can be decomposed as:

$$\mathbf{C}(\mathbf{x}, \mathbf{x}') = \sum_{k=1}^{\infty} \lambda_k \phi_k(\mathbf{x}) \phi_k(\mathbf{x}'), \quad (2)$$

where $\lambda_k > 0$ are the eigenvalues of \mathbf{C} , and ϕ_k are the corresponding normalized eigenvectors. The eigenpairs satisfy the Fredholm equation of the second kind:

$$\int_{\Omega} \mathbf{C}(\mathbf{x}, \mathbf{x}') \phi_k(\mathbf{x}') d\mathbf{x}' = \lambda_k \phi_k(\mathbf{x}), \quad k = 1, 2, \dots \quad (3)$$

Sequencing the eigenvalues in decreasing order, any realization of M can then be written as [22]

$$M(\mathbf{x}, \omega) = \mu(\mathbf{x}) + \sum_{k=1}^{\infty} \sqrt{\lambda_k} \phi_k(\mathbf{x}) \eta_k(\omega). \quad (4)$$

The coordinate coefficients $\eta_k(\omega)$ can be deterministically computed by evaluating the integrals

$$\eta_k(\omega) = \int (M(\mathbf{x}, \omega) - \mu(\mathbf{x})) \phi_k(\mathbf{x}) d\mathbf{x}. \quad (5)$$

In case M is a Gaussian process (GP), $\eta_k(\omega)$ are Gaussian-independent identically distributed random variables with zero mean and unit variance, such that $\mathbb{E}[\eta_k \eta_{k'}'] = \delta_{kk'}$. The parameter field M is then fully characterized by a set of coordinates η_k 's given the eigenbasis is known. By truncating the series in Equation (4) to retain the first few K terms, the KL expansion efficiently approximates the high-dimensional parameter field using a few dominant modes. The size of the retained terms K depends on the desired energy percentage to be retained by the KL expansion, which is defined as $\sum_{k=1}^K \lambda_k / \sum_{k=1}^{\infty} \lambda_k$ [17]. In practice, it is preferable to choose K such that the truncated KL expansion encapsulates as much information as possible with respect to its infinite counterpart. Larger K are required for prior covariance functions with smaller correlation lengths for the KL expansion to capture a similar information percentage. There exist several strategies to set K in practice [40, 37]. For instance, for $\lambda_1 \geq \lambda_2 \geq \dots$, one can control the representation error on M by selecting K such that $\lambda_K < c\lambda_1$ or $\sum_{k=1}^K \lambda_k \leq c \sum_{k=1}^{\infty} \lambda_k$ for some truncation parameter $c \ll 1$.

The truncated KL expansion is optimal, in the mean square sense, in approximating stochastic processes [21]. It is used as the model reduction technique for Bayesian inversion. The problem of inferring the continuous field M from a few discretely collected observations is then reduced to the inference of the KL coordinates vector $\boldsymbol{\eta} \in \mathbb{R}^K$, written as

$$p(\boldsymbol{\eta}, \sigma_o^2 \mid \mathbf{d}) \propto p(\mathbf{d} \mid \boldsymbol{\eta}, \sigma_o^2) p_{\boldsymbol{\eta}}(\boldsymbol{\eta}) p_o(\sigma_o^2), \quad (6)$$

where $p_{\boldsymbol{\eta}}(\boldsymbol{\eta}) = \exp(-\boldsymbol{\eta}^T \boldsymbol{\eta} / 2) / (2\pi^{K/2})$ is the Gaussian prior of the KL coordinates. The unbiased additive Gaussian observational error model considered in this work leads to the likelihood function of the form

$$p(\mathbf{d} \mid \boldsymbol{\eta}, \sigma_o^2) = \prod_{i=1}^{N_o} p_{\epsilon}(d_i - u_i(\boldsymbol{\eta}), \sigma_o^2),$$

$$p_{\epsilon}(d_i - u_i(\boldsymbol{\eta}), \sigma_o^2) \doteq \frac{1}{\sqrt{2\pi\sigma_o^2}} \exp\left[-\frac{(d_i - u_i(\boldsymbol{\eta}))^2}{2\sigma_o^2}\right]. \quad (7)$$

2.2. Polynomial Chaos (PC) expansion

Even with a parameter reduction technique (e.g., KL expansion), MCMC sampling remains computationally prohibitive for realistic large-scale models. This is because the sampling of the posterior pdf requires a massive collection of likelihood evaluation, each one relying on a full model run. PC expansion is a suitable technique for building a model surrogate that helps accelerating this MCMC sampling process [35, 37]. The PC method expresses the dependency of model solution U on some uncertain model inputs, using a truncated spectral polynomial expansion [21, 37]. Generally, one has to approximate $U(\boldsymbol{\xi})$, where $\boldsymbol{\xi} = (\xi_1, \dots, \xi_s) \in \Omega^* \subseteq \mathbb{R}^s$ are independent second-order random variables, with given density $p_{\boldsymbol{\xi}} : \mathbb{R}^d \mapsto \mathbb{R}^+$, and Ω^* is the image of the sample space over $\boldsymbol{\xi}$. In the present work, the uncertain model inputs are the random variables η_k of the KL expansion such that $\boldsymbol{\xi}$ consists of the normalized vector $\boldsymbol{\eta}$ and $s = K$. If other independent uncertainty sources were present, $\boldsymbol{\xi}$ would consist of an extended vector with $s > K$.

Let $\{\Psi_k, k \geq 0\}$ be a set of s -variate orthogonal polynomials. The PC method expresses the model output U using a spectral decomposition of the form

$$U(\boldsymbol{\xi}) = \sum_{k=0}^{\infty} U_k \Psi_k(\boldsymbol{\xi}) \cong \sum_{k=0}^P U_k \Psi_k(\boldsymbol{\xi}), \quad (8)$$

where U_k are called the stochastic modes of U . Because we restrict our investigation to Gaussian processes, $\boldsymbol{\xi} = \boldsymbol{\eta}$ will be a vector of standard Gaussian random variables, and Ψ_k are the normalized multi-variate Hermite polynomials [72]. The number of terms P in (8) is defined by the polynomial order N_o and the stochastic dimension s , such that $P + 1 = (N_o + s)! / (N_o! s!)$. The terms in the PC expansion therefore increases exponentially with respect to an increase in both the polynomial order and the number of the stochastic inputs, which results in the so-called ‘curse of dimensionality’ for PC expansions [1].

Various projection methods can be used to determine the stochastic modes U_k of the PC expansion [68]. These methods can be categorized as the Galerkin and non-intrusive approaches. The Galerkin projection method [21, 37] defines the PC coefficients through the reformulation of the system equations, using a weak form with respect to the stochastic dimension. This requires modifying the model code, which makes the method only suitable for small problems. The non-intrusive approaches, as the name implies, do not require modifying the model code by defining the stochastic modes based on the ensemble of the deterministic evaluation of $U(\boldsymbol{\xi})$, using a suitable quadrature formula. In this study, we resort to the non-intrusive spectral projection (NISIP) method to build the PC surrogate.

From the infinite expansion (8), the orthogonality of the polynomials implies that

$$\langle U, \Psi_k \rangle = \sum_{l=0}^{\infty} U_l \langle \Psi_l, \Psi_k \rangle = U_k \langle \Psi_k, \Psi_k \rangle, \quad (9)$$

Where $\langle \cdot, \cdot \rangle$ denotes the stochastic inner product,

$$\langle u, v \rangle = \int_{\Omega^*} u(\mathbf{s})v(\mathbf{s})p_{\boldsymbol{\xi}}(\mathbf{s}). \quad (10)$$

In other words, the coefficient $U_k = \langle U, \Psi_k \rangle / \langle \Psi_k, \Psi_k \rangle$ is given by the correlation between $U(\boldsymbol{\xi})$ and the considered polynomial $\Psi_k(\boldsymbol{\xi})$, up to the squared norm of Ψ_k , which can be evaluated analytically [37]. The evaluation of the U_k , thus, involves computing a set of $P + 1$ integrals over Ω^* .

In NISIP, the correlation defining the projection coefficient is approximated by a finite sum, using an appropriate quadrature formula:

$$\int_{\Omega^*} U(\mathbf{s})\Psi_k(\mathbf{s})dp_{\boldsymbol{\xi}}(\mathbf{s}) \simeq \sum_{j=1}^{N_q} w_j U(\mathbf{s}_j)\Psi_k(\mathbf{s}_j). \quad (11)$$

Here, w_j denotes the weights associated to the node \mathbf{s}_j , and N_q is the number of quadrature nodes. The set of quadrature nodes is denoted by

$$\mathcal{S} = \{\mathbf{s}_j\}_{j=1}^{N_q} \subset \Omega^*. \quad (12)$$

To evaluate the summation (11), one needs to evaluate U for all $\mathbf{s}_q \in \mathcal{S}$. Let $\mathbf{\Pi} \in \mathbb{R}^{(P+1) \times N_q}$ denote the NISP projection matrix

$$\mathbf{\Pi}_{k,j} = \frac{w_j \Psi_k(\mathbf{s}_j)}{\langle \Psi_k, \Psi_k \rangle}, \quad k = 0, \dots, P, \quad j = 1, \dots, N_q,$$

and ζ be the vector with coordinates $\zeta_j = U(\mathbf{s}_j)$. Then, the vector of PC coefficients \mathbf{U} can be expressed as $\mathbf{\Pi}\zeta$, or in coordinate form,

$$U_k = \sum_{j=1}^{N_q} \mathbf{\Pi}_{kj} \zeta_j = \sum_{j=1}^{N_q} \mathbf{\Pi}_{kj} U(\mathbf{s}_j), \quad k = 0, \dots, P. \quad (13)$$

A more comprehensive derivation of the NISP method can be found in [1, 37]. The complexity in performing NISP scales with N_q and grows exponentially with stochastic dimension s . Consequently, the implementation of the NISP method to construct the PC surrogate is usually restricted to a limited small number of stochastic dimensions.

3. Change of coordinates for uncertain covariance functions

The inference framework described in the previous section incorporates the parameter reduction technique (i.e., KL expansion) and PC surrogate to enable the use of MCMC sampling of large dimensional spatially varying parameters, assuming the underlying covariance function is known. However, in real-world applications, this is rarely the case. The covariance functions are often characterized by some (uncertain) hyper-parameters, such as the variances and correlation length-scales.

The unknown hyper-parameters were traditionally estimated prior to inference using a Gaussian process regression (GPR) [50], where the marginal likelihood function is maximized with respect to the hyper-parameters via stochastic interpolation, using a set of noisy observations of the parameter field and data collected at a few locations. One may further use the model outputs observation to further reduce the hyper-parameter uncertainty via Bayesian inference [67]. This however requires building a more computationally demanding PC surrogate, compared to the case of fixed covariance model. Here, we resort to the ‘change of coordinates’ technique, proposed in [61] to simultaneously infer the covariance function hyper-parameters and the KL coordinates vector $\boldsymbol{\eta}$ without introducing the hyper-parameters as additional stochastic dimensions in the PC surrogate.

3.1. Generalized Karhunen-Loève expansion

Let \mathbf{q} be a random vector of covariance function hyper-parameters with joint density $p_{\mathbf{q}}$ (e.g. $\mathbf{q} = \{\sigma_f^2, l\}$, where σ_f^2 is the process variance, and l is the correlation length scale). The centered ($\mu(\mathbf{x}) = 0$) stochastic prior process M has a covariance function $\mathbf{C} = \mathbf{C}(\mathbf{x}, \mathbf{x}', \mathbf{q})$ parametrized by $\mathbf{q} \in \mathbb{R}^h$, where h is the number of hyper-parameters. The KL expansion in (4) can be truncated and modified to reflect the dependency of the covariance function on \mathbf{q} as

$$M_K(\mathbf{x}, \omega, \mathbf{q}) = \sum_{k=1}^K \sqrt{\lambda_k(\mathbf{q})} \phi_k(\mathbf{x}, \mathbf{q}) \eta_k(\omega). \quad (14)$$

Hereafter, we omit the \mathbf{x} dependence to simplify the notations and introduce the scaled eigenfunction $\Phi_k(\mathbf{q}) \doteq \sqrt{\lambda_k(\mathbf{q})} \phi_k(\mathbf{x}, \mathbf{q})$, which gives (14) the compact form

$$M_K(\omega, \mathbf{q}) = \sum_{k=1}^K \Phi_k(\mathbf{q}) \eta_k(\omega). \quad (15)$$

3.1.1. Coordinate transformation

To eliminate the \mathbf{q} -dependence of the eigenfunctions, we introduce ϕ_k^r , the ordered and normalized eigenvectors of reference covariance \mathbf{C}^r , and the projection coefficients $b_{kk'}(\mathbf{q})$. Since a set of reference eigenfunctions constitutes an orthogonal basis, a change of basis can be applied to express any scaled eigenfunction in the reference basis as

$$\Phi_k(\mathbf{q}) = \sum_{k'=1}^{\infty} b_{kk'}(\mathbf{q})\phi_{k'}^r, \quad b_{kk'}(\mathbf{q}) = (\phi_k^r, \Phi_{k'}(\mathbf{q}))_{\mathcal{D}} = \int_{\mathcal{D}} \phi_k^r(\mathbf{x})\Phi_{k'}(\mathbf{x}, \mathbf{q})d\mathbf{x}. \quad (16)$$

In practice, the latter expansion is truncated to the first K^r terms. Here, we shall assign $K^r = K$ to allow for convergence analysis with respect to the first few dominant KL modes. This leads to the expression of the generalized KL expansion of M as

$$M_K(\omega, \mathbf{q}) = \sum_{k=1}^K \Phi_k(\mathbf{q})\eta_k(\omega) \approx \widehat{M}_K(\omega, \mathbf{q}) \doteq \sum_{k=1}^K \left(\sum_{k'=1}^K b_{kk'}(\mathbf{q})\phi_{k'}^r \right) \eta_k(\omega). \quad (17)$$

To further simplify (17), $\widehat{M}_K(\omega, \mathbf{q})$ can be recast as

$$\widehat{M}_K(\omega, \mathbf{q}) = \sum_{k=1}^K \phi_k^r \widehat{\eta}_k(\omega, \mathbf{q}), \quad \widehat{\eta}_k(\omega, \mathbf{q}) = \sum_{k'=1}^K b_{kk'}(\mathbf{q})\eta_{k'}(\omega). \quad (18)$$

Since $\widehat{\eta}_k(\omega, \mathbf{q})$ is a linear combination of i.i.d standard random variables, this term can be cast in matrix as

$$\widehat{\boldsymbol{\eta}}(\omega, \mathbf{q}) = \mathbf{B}(\mathbf{q})\boldsymbol{\eta}(\omega), \quad (19)$$

where $\mathbf{B}(\mathbf{q})$ is called the coordinate transformation matrix.

Ultimately, Equation (19) is the result of recasting the \mathbf{q} -dependence of \mathbf{C} into the expansion of M_K in (15), using scaled eigenfunctions and the further approximation of M_K using a \mathbf{q} -dependent linear transformation $b_{kk'}(\mathbf{q})$. The above derivations require the continuity assumption of the scaled eigenfunctions $\Phi_k(\mathbf{q})$ with respect to \mathbf{q} (for more details, see [61] and [36]). To ensure such conditions, the orientation of the decomposed eigenfunctions needs to be consistent. Sraj *et al.* [61] achieved this by defining the orientation of the eigenfunctions, such that $(\Phi_k(\mathbf{q}), \phi_k^r)_{\mathcal{D}}$ has a consistent sign for all \mathbf{q} [53].

A KL decomposition of $\mathbf{C}(\mathbf{x}, \mathbf{x}', \mathbf{q})$ is required at the beginning of each MCMC iteration to compute the corresponding change of coordinates matrix $\mathbf{B}(\mathbf{q})$. This computation becomes costly for large-scale systems, and motivates us to build a PC surrogate of the coordinate transformation $\mathbf{B}(\mathbf{q})$ in order to bypass this expensive step. Hence, we introduce the PC expansion:

$$\mathbf{B}(\mathbf{q}) \approx \widetilde{\mathbf{B}}(\mathbf{q}) = \sum_{k=0}^P \mathbf{B}_k \Psi_k(\mathbf{q}), \quad (20)$$

where the Ψ_k are the orthogonal random polynomials in \mathbf{q} . Following the previous discussion on the NISP method, the PC modes of $\mathbf{B}(\mathbf{q})$ are defined by

$$\forall k, \quad \mathbf{B}_k \langle \Psi_k, \Psi_k \rangle = \langle \mathbf{B}, \Psi_k \rangle, \quad (21)$$

which right-hand-side can be estimated using a numerical quadrature rule. With this PC expansion of \mathbf{B} , the processes is approximated as

$$\widetilde{M}_K(x, \mathbf{q}) = \sum_{k=1}^K \Phi_k(x)\widetilde{\eta}_k(\mathbf{q}), \quad \widetilde{\boldsymbol{\eta}}_k(\mathbf{q}) \doteq \widetilde{\mathbf{B}}(\mathbf{q})\boldsymbol{\eta}. \quad (22)$$

3.1.2. Example

To illustrate the coordinate transformation and its PC approximation, we consider the case of a one-dimensional centered Gaussian process M defined on the unit interval $[0, 1]$ and with uncertain covariance function $C(x, x', \mathbf{q}) = \exp(-|x - x'|^\gamma / \gamma l^\gamma)$. The covariance has two independent hyper-parameters, being the correlation length l and the regularity parameter γ : $\mathbf{q} = (l, \gamma)$. The priors of the hyper-parameters are both uniform with ranges $l \sim U[0.25, 1]$, and $\gamma \sim U[1.5, 2]$. The integral equations defining the \mathbf{q} -dependent modes,

$$\int_0^1 C(x, x', \mathbf{q}) \phi(x', \mathbf{q}) dx' = \lambda(\mathbf{q}) \phi(x, \mathbf{q}),$$

is discretized using piecewise constant approximations on a mesh involving $N_{el} = 200$ elements with size $h = 1/N_{el}$. A Galerkin procedure is employed to obtain a linear system eigen-value problem for the value of the modes over each element. Note that the discrete problem admits at most N_{el} KL modes. Consistently with the notation above, we denote $M_K(x, \mathbf{q})$ the truncated version of $M(x, \mathbf{q})$ to the K -dominant modes,

$$M_K(x, \mathbf{q}) = \sum_{k=1}^K \sqrt{\lambda_k(\mathbf{q})} \phi_k(x, \mathbf{q}) \eta_k, \quad \eta_k = \int_0^1 M(x, \mathbf{q}) \phi_k(x, \mathbf{q}) dx,$$

$\widehat{M}_K(x, \mathbf{q})$ the exact projection of $M_K(x, \mathbf{q})$ on the reference space,

$$\widehat{M}_K(x, \mathbf{q}) = \sum_{k=1}^K \Phi_k(x) \widehat{\eta}_k(\mathbf{q}), \quad \widehat{\eta}_k(\mathbf{q}) = \int_0^1 M_K(x, \mathbf{q}) \Phi_k(x) dx,$$

and finally $\widetilde{M}_K(x, \mathbf{q})$ the approximation of M_K in the reference basis and using the PC approximation of the change of coordinate,

$$\widetilde{M}_K(x, \mathbf{q}) = \sum_{k=1}^K \Phi_k(x) \widetilde{\eta}_k(\mathbf{q})$$

with $\widetilde{\eta}_k(\mathbf{q}) = \widetilde{\mathbf{B}}(\mathbf{q}) \boldsymbol{\eta}$. As discussed, the reference basis is constructed on the averaged covariance,

$$\overline{C}(x, x') = \frac{1}{(1 - 0.25)(2 - 1.5)} \int_{0.25}^1 \int_{1.5}^2 C(x, x', l, \gamma) d\gamma dl,$$

while the PC expansion of $\mathbf{B}(\mathbf{q})$ uses a basis of Legendre polynomials (orthogonal with respect to the uniform measure [73]) and the Pseudo-Spectral Projection (PSP) method, see [13, 12], to compute the PC coefficients:

$$\mathbf{B}(\mathbf{q}) = \sum_{i=0}^P \mathbf{B}_i \Psi_i(\mathbf{q}), \quad \mathbf{B}_i = \sum_{j=1}^{N_{\text{PSP}}} \mathbf{B}(\mathbf{q}^{(j)}) W_{i,j},$$

where $W_{i,j}$ are the coefficients of the PSP matrix associated to the sparse-grid with points $\{\mathbf{q}^{(j)}, j = 1, \dots, N_{\text{PSP}}\}$. Overall, the PC projection requires the computation of the dominant bases for the N_{PSP} covariances given by $C(x, x', \mathbf{q}^{(j)})$, in addition of the reference basis.

Table 1 reports the resulting mean-square errors norm due to the representation in the reference basis, $M_K - \widehat{M}_K$, global error of the change of coordinate with PC approximation, $M_K - \widetilde{M}_K$, and error due to the PC approximation of $\mathbf{B}(\mathbf{q})$: $\widehat{M}_K - \widetilde{M}_K$. The errors are computed by squaring the differences, integrating for $x \in [0, 1]$, and averaging over \mathbf{q} and realizations of the process. Note that the mean-square norm of $M(x, \mathbf{q})$ is one, so the error reported in the Table 1 are normalized and, therefore, reported as percentages. The Table reports the errors for different dimensions K of the reference basis and the level ℓ of the PSP method. The level ℓ controls the polynomial degree of the PC expansion and the size of the PC basis. Table 1 shows that the error due to the representation in the reference basis decreases when K increases. The PC approximation of the change of coordinate introduces an additional error that becomes dominant

Error (%)	$K = 5 - \ell = 2$	$K = 10 - \ell = 2$	$K = 15 - \ell = 2$	$K = 10 - \ell = 3$	$K = 10 - \ell = 4$
$\ M_K - \widehat{M}_K\ $	0.88	0.23	0.10	0.26	0.27
$\ M_K - \widetilde{M}_K\ $	1.10	0.77	0.72	0.26	0.23
$\ \widehat{M}_K - \widetilde{M}_K\ $	0.66	0.74	0.71	0.13	0.03

Table 1: Error analysis of the Change of Coordinate. Errors are reported as percentages for different dimension K of the reference basis and different levels ℓ of the Pseudo-Spectral-Projection method.

for $K \geq 10$ when $\ell = 2$. The PC error quickly decays when ℓ increases from 2 to 3 and 4, such that the truncation error becomes dominant in $M_K - \widehat{M}_K$ with $K = 10$.

This example illustrates the possibility of approaching the change of coordinate by a PC expansion. However, the convergence of the PC expansion requires sufficient smoothness of the function $\mathbf{q} \mapsto \mathbf{B}(\mathbf{q})$. Also, using non-intrusive projection methods for the determination of the PC expansion requires a unique definition of \mathbf{B} at each quadrature node. In fact, the definition \mathbf{B} is not unique, as the definition of the eigenvectors $\phi_k(\cdot, \mathbf{q})$ can be ambiguous. In particular, it is well-known that the eigenvectors are defined up to a sign ($-\phi_k$ is also an eigenvector). This ambiguity can compromise the non-intrusive projection if one is not careful in setting the correct signs of the eigenvector at a given \mathbf{q} when constructing the PC expansion of $\mathbf{B}(\mathbf{q})$.

In this study, the coastal ocean domain is two-dimensional with several symmetries. As a result, eigenvalues with multiplicity ≥ 1 and the crossing of eigenbranches (with \mathbf{q}) emerge, making it more difficult to compute a numerical solution of smooth eigenfunctions in \mathbf{q} , and obtain smooth dependencies of the matrix $\mathbf{B}(\mathbf{q})$. One possibility that we started to explore is to rely on clustering algorithms (based on similarity measures) to the set of eigenvectors computed at several \mathbf{q} -values and obtain groups corresponding to distinct branches (\mathbf{q} -dependent modes). Although quite promising, these clustering strategies are beyond the scope of the present paper and will be presented elsewhere. To avoid the above complications in the experimental demonstrations of Section 5, we limit ourselves to low KL truncation orders K , ensuring sufficiently well-separated eigenvalues.

3.2. PC surrogate with coordinate transformation

Given that the hyper-parameters \mathbf{q} are considered to be uncertain and inferred alongside the KL coordinates $\boldsymbol{\eta}$, the model prediction is expressed as $U(\boldsymbol{\eta}, \mathbf{q})$. Intuitively, this suggests building a PC surrogate that includes \mathbf{q} as stochastic inputs, that is, $U(\boldsymbol{\xi}) = \sum_{k=0}^P U_k \Psi_k(\boldsymbol{\eta}, \mathbf{q})$. However, this approach drastically increases the computational requirement for evaluating the PC coefficients due to the exponential increase in the terms of the PC expansion. This is particularly problematic for a large-scale ocean model, which requires a large number of model integrations and a large polynomial order to obtain an accurate PC surrogate for such a system. Adopting a similar framework applied to KL expansion as presented in section 3.1 for the inference problem, we seek to replace the direct augmentation of \mathbf{q} to solve the PC problem with the change of coordinates method, based on a reference model problem dictated by a fixed covariance function.

An immediate approach would consist in seeking a PC expansion in terms of the reference coordinate $\widehat{\boldsymbol{\eta}}$ given by $\widehat{\boldsymbol{\eta}} = \mathbf{B}(\mathbf{q})\boldsymbol{\eta}$. However, depending on the selected reference basis, the actual distribution of $\widehat{\boldsymbol{\eta}}$ may not be easily derived. Instead, it is preferred to rely on a priori coordinates $\boldsymbol{\xi}$ with known prior distribution. A reasonable choice is to define the reference Gaussian process:

$$\widehat{M}_K^{\text{PC}}(\boldsymbol{\xi}) = \sum_{k=1}^K \sqrt{\lambda_k^r} \phi_k^T \boldsymbol{\xi}_k, \quad (23)$$

where the $\boldsymbol{\xi}_k$'s are independent standard Gaussian random variables. We now denote $\widetilde{U}(\boldsymbol{\xi})$ the reference PC surrogate, which approximates the predictions of the reference model, that is,

$$U(\boldsymbol{\eta}, \mathbf{q}) \approx \widetilde{U}(\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q})) = \sum_{k=0}^P U_k \Psi_k(\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q})). \quad (24)$$

Similar to Equation (19), which expresses the \mathbf{q} -dependent Gaussian process \widehat{M}_K through the mapping $(\boldsymbol{\eta}, \mathbf{q}) \mapsto \widehat{\boldsymbol{\eta}}$ using the transformation matrix $\mathbf{B}(\mathbf{q})$, we can apply the change of basis principle to express the surrogate \widetilde{U} through the mapping $(\boldsymbol{\eta}, \mathbf{q}) \mapsto \boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q})$ using another coordinate transformation matrix:

$$\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q}) = \widehat{\mathbf{B}}(\mathbf{q})\boldsymbol{\eta}, \quad (25)$$

where the \mathbf{q} -dependent $\widehat{\mathbf{B}}(\mathbf{q})$ denotes the change of coordinates, mapping the couples $(\boldsymbol{\eta}, \mathbf{q})$ to $\boldsymbol{\xi}$. Based on (23), it comes

$$\widehat{\mathbf{B}}_{kl}(\mathbf{q}) = \begin{cases} \frac{\widetilde{\mathbf{B}}_{kl}(\mathbf{q})}{\sqrt{\lambda_k^r}}, & \lambda_k^r/\lambda_k^l > \kappa, \\ 0, & \text{otherwise,} \end{cases} \quad (26)$$

where κ is a small constant related to the level of numerical accuracy introduced to discard the diminishingly small λ_k^r and avoid numerical overflows. Note that $\widehat{\mathbf{B}}$ is derived from \mathbf{B} , and therefore, the KL and PC expansions used for the Bayesian inference are said to have a similar coordinate transformation.

3.3. Bayesian inference with the change of coordinates (COC)

The Bayesian inference framework for uncertain hyper-parameters based on the change of coordinate (COC) methodology can be split into two main steps. The invariant quantities, which are used throughout the MCMC simulation, such as the reference basis and the PC coefficients, can be precomputed offline and stored for the later usage, while the calculation of $(\boldsymbol{\eta}, \mathbf{q})$ -dependent quantities, i.e., coordinate transformation matrix and approximated PC solution of U , are embedded inside each iteration of MCMC sampling, which is referred to as the online step.

3.3.1. Offline step

Given that the reference covariance function \mathbf{C}^r has been already selected, the KL decomposition is applied to \mathbf{C}^r according to (2). The eigenpairs $(\lambda_k^r, \phi_k^r)_{k=1,K}$ obtained from KL decomposition are then stored in memory. Subsequently, according to (23), the same eigenpairs are utilized to build $\widehat{M}_K^{\text{PC}}(\boldsymbol{\xi})$, which represents the reference Gaussian process that characterizes the stochastic parameter field of the reference mathematical model $\widehat{\mathcal{L}}(\boldsymbol{\xi})$. Finally, the coordinates $\boldsymbol{\xi}$, which specify the realization of $\widehat{\mathcal{L}}(\boldsymbol{\xi})$, are sampled using a suitable quadrature formula to construct the surrogate $\widetilde{U}(\boldsymbol{\xi})$ of U . The PC coefficients $\{U_k\}_{k=0}^P$ computed in the latter step are stored in memory, to be used later to evaluate the model solutions for likelihood sampling in the inference step. Additionally, in this work, we also construct the PC surrogate $\widetilde{\mathbf{B}}(\mathbf{q})$ of $\mathbf{B}(\mathbf{q})$ to be used for fast coordinate transformation in the online step.

3.3.2. Online step

Each iteration of MCMC embedded in the online step involves several substeps. First, \mathbf{q} and $\boldsymbol{\eta}$ are proposed. Then, the PC approximated coordinate transformation matrix $\widetilde{\mathbf{B}}(\mathbf{q})$ is used to compute the corresponding reference coordinate $\boldsymbol{\xi} = \widehat{\mathbf{B}}(\mathbf{q})\boldsymbol{\eta}$. Subsequently, we evaluate the PC surrogate of the reference model prediction $\widetilde{U}(\boldsymbol{\xi})$ for an infinitesimal CPU time. This PC approximation of $U(\boldsymbol{\eta}, \mathbf{q})$ is then used for the likelihood computation, generating the MCMC chains of the posterior of the couple $(\boldsymbol{\eta}, \mathbf{q})$, based on the accept/reject method imposed by the Metropolis-Hastings algorithm. The above processes are repeated until the MCMC simulation reaches its designated terminal iteration.

4. Experimental setup

The experimental setup for the problem of inferring the Manning's n coefficients (and any other spatially varying parameters) of the coastal ocean is presented in this section. Albeit rather idealized, the setup we propose for the experimentation represents well the dynamics of the coastal ocean.

4.1. The advanced circulation model (ADCIRC)

We employ the advanced circulation (ADCIRC) model [41] to test the proposed inference framework. ADCIRC has been widely implemented for simulating coastal and estuarine systems, as well as the analysis of tides and currents [42, 46]. It has also been validated intensively in many storm surge hindcast studies, such as those of hurricanes Betsy (1965), Ivan (2004), Dennis (2004), Katrina (2005), Rita (2005) [71, 9, 15], Gustav (2008) [16], and Ike (2008) [32].

The machinery behind ADCIRC solves a modified version of the SWEs, in which the continuity equation is replaced by the second-order, hyperbolic generalized wave continuity equation (GWCE) to prevent spurious oscillations that often arise from the numerical solution of the original form [43, 34]. Together with the momentum equation, these governing equations are solved on unstructured, triangular elements using a first-order continuous Galerkin finite element scheme. The time derivatives are approximated using centered finite differences in the GWCE and forward differences in the momentum equations. ADCIRC supports a wide range of domain specifications and complex bathymetry structures, including the range of scales necessary to represent the deep ocean basins, continental shelves, and coastal inland areas [10].

4.2. The model discretization

We consider an idealized inlet with an ebb shoal domain as depicted in Fig.1. The domain comprises a lagoon that is connected to the open ocean on the west through an inlet with twin jetties. The domain is 4500 m wide and 3000 m long, and is discretized to have 1518 grid nodes and 2828 elements. The bathymetry of the domain is designed to resemble the realistic vertical profile of a coastal ocean with the deepest depth in the open ocean area. The depth of the ocean floor increases linearly from 3.8 m at the left-most boundary to 1 m at the mouth of the inlet on the west side of the domain, leveling off at the landlocked area. The diameter of the ebb shoal is 750 m. On the open ocean boundary, ADCIRC is forced with the principal lunar semi-diurnal M_2 tidal constituent, with an amplitude of 0.25 m relative to the geoid. This domain is a simplified version of a real ebb shoal, which is a common natural feature of coastal inlets.

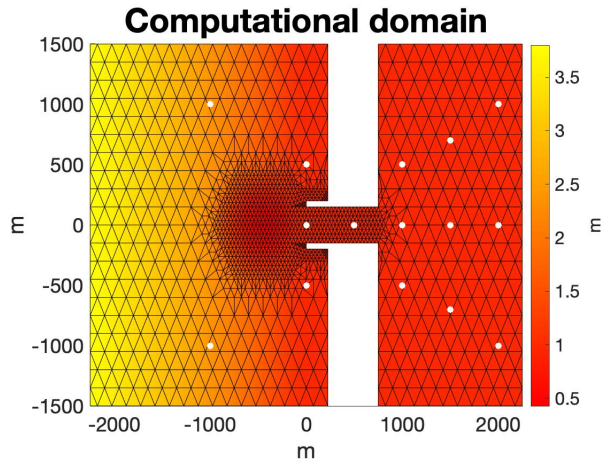


Figure 1: Idealized inlet with ebb shoal domain shown with the finite element discretization. Observation stations are marked with white dots. The color bar represents the bathymetry of the domain measuring down from the geoid (m).

4.3. Prior covariance function and the reference basis

We assume that the covariance function describing the Gaussian process underlying the uncertain Manning’s n field is isotropic. Specifically, we also note here that the field M is Gaussian, while the Manning’s n coefficients have a finite range $[0.005, 0.2]$. Thus, we use the normal cumulative distribution function (cdf) to map $\mathcal{N}(0, 1)$ to $\mathcal{U}[0, 1]$ and our own customized rescale function that maps $\mathcal{U}[0, 1]$ to $\mathcal{U}[0.005, 0.2]$, to transform between the realizations of a GP and the actual Manning’s n fields.

A popular class of covariance kernels employed in the field of geostatistics is the Matérn function [65],

$$\mathbf{C}(\mathbf{x}, \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}|\mathbf{x} - \mathbf{x}'|}{l} \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}|\mathbf{x} - \mathbf{x}'|}{l} \right). \quad (27)$$

Here, ν is a smoothness parameter, Γ is the gamma function, and K_ν is a modified Bessel function of the second kind. The roughness of the stochastic process has the inverse relation with the value of ν . As $\nu \rightarrow \infty$, this covariance kernel becomes infinitely differentiable, and the Matérn covariance function converges to the squared exponential covariance function, which can be written as

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', \mathbf{q}) = \sigma_f^2 \exp \left(-\frac{(\mathbf{x} - \mathbf{x}')^2}{2l^2} \right), \quad (28)$$

where $\mathbf{q} = \{\sigma_f^2, l\}$. We select this special case of the Matérn covariance to simulate the stochastic process; it is one of the most used covariance functions to simulate spatially varying parameter fields [64]. In our case, only the correlation length l affects the shape of the eigenfunctions, while the process variance σ_f^2 simply scales the eigenvalues. Therefore, we fix $\sigma_f^2 = 1$ for all experiments and only assume the uncertainty in l , that is, $\mathbf{q} = \{l\}$.

The choice of the reference basis is similar to that chosen in [61]. We compute the reference eigenbasis by solving (3), given the \mathbf{q} -averaged covariance function defined as

$$\mathbf{C}^r = \bar{\mathbf{C}} \doteq \int \mathbf{C}(\mathbf{q}) p_q(\mathbf{q}) d\mathbf{q}. \quad (29)$$

This choice of reference covariance function has been analyzed (see [61]) and contrasted with the choice $\mathbf{C}^r = \mathbf{C}(l^r)$ for several values of l^r . It has been shown to yield the lowest representation error (averaged relative error over \mathbf{q}) compared to selecting any reference covariance as $\mathbf{C}(l^r)$. This choice is optimal because it uses the eigenfunction spanning the optimal subspace to represent $M(\omega, \mathbf{q})$ when \mathbf{q} varies with probability law $p_q(\mathbf{q})$. Additionally, the error in approximating U with the PC surrogate \tilde{U} through coordinate transformation is consistent with the previous analysis, i.e., the error is minimized when $\bar{\mathbf{C}}$ is used for all truncation order K .

In practice, the explicit expression of the averaged covariance function in (29) is rarely available. It therefore needs to be approximated using a sampling-based approach. In our case, the Gauss quadrature rule was used to appropriately sample a set of $\mathbf{q} = \{l\}$ from $\mathcal{U}[l_{min}, l_{max}]$, where l_{min} and l_{max} are the minimum and maximum correlation length scales, respectively. In the present work, we take advantage of the moderate size of the mesh to create and store the 1518×1518 discrete covariance matrix whose entries are computed by quadrature. For each quadrature representing a certain l , the value of each element \mathbf{x}_{ij} in this matrix is simply $\mathbf{C}(\mathbf{x}_i, \mathbf{x}_j, l)$. We then perform the summation of this matrix over l using a for loop and take the average over the number of quadrature points, to finally obtain the approximated $\bar{\mathbf{C}}$ matrix. Note that this method only requires storing a single covariance matrix. For the problems with much larger spatial discretization in which computing a large covariance matrix is intractable, covariance approximation methods such as fixed rank kriging [14], covariance tapering [18], and the full-scale approximation [54] could be used.

4.4. Observation simulation system experiments (OSSEs)

To assess the performance of the proposed inference framework, we implement ‘observation simulation system experiments (OSSEs)’. In these experiments, the true Manning’s n field is chosen (assumed to be known). Then, this true field is used in the ADCIRC model, to compute the reference solution (i.e., water elevation), denoted by $u(x, t)$, where x is the location and t is time. Water elevation data are then extracted from the reference solution at 15 measurement stations every hour for 4.5 days, after the 12 hour ramp-up period. This yields a total $N_o = 15 \times 108 = 1620$ observations. These are then perturbed with i.i.d. additive noise $\epsilon_i \sim N(0, \sigma_\epsilon^2 = 0.01)$ and are used to infer the true Manning’s n field. We investigate two test cases:

- **Random profile:** $M^{ran}(\mathbf{x})$ drawn at random from a Gaussian process, $GP(0, \mathbf{C})$, where \mathbf{C} is the squared exponential covariance. To do so, we sample the quadruplet, $\{\eta_1, \eta_2, \eta_3, l\}$, that characterizes the random field $M_{K=3}$. We limit ourselves to only three modes ($K = 3$), on account of the limitation in the size of the stochastic dimension that can be accommodated in the PC construction scheme (i.e., NISP with full tensorization quadrature rules). This does not undermine the validity of our assessment, which is to identify the improvement of the inference using COC with respect to the inference without COC, using the same number of modes. Furthermore, one can alleviate this limitation of small stochastic dimensions due to the chosen PC scheme by using the sparse quadrature methods [1]. By constructing a true field from a known prior covariance, we can assess our method by comparing the posterior pdf and maximum a posteriori probability (MAP) estimate with the true field, in terms of $\boldsymbol{\eta}$ and \mathbf{q} directly (see in the next section). We show the image of the true field M^{ran} on the left of Fig.2. Here, the sampled quadruplet that generates this particular Manning’s n field realization is $\{\eta_1, \eta_2, \eta_3, l\} = \{1.73, 0.26, 0.04, 1255\}$.
- **Zonal profile:** the true Manning’s n profile $M^{zone}(\mathbf{x})$ is derived from the equation:

$$M^{zone}(\mathbf{x}) = \alpha M_1(\mathbf{x}) + \beta M_2(\mathbf{x}) + (\beta - \alpha) M_3(\mathbf{x}). \quad (30)$$

Here, $M_1(\mathbf{x})$, $M_2(\mathbf{x})$ and $M_3(\mathbf{x})$, are the parametrization coefficients, which take a value between 0 and 1. The Manning’s n coefficients in the open ocean area are equal to α , and β are the Manning’s n coefficients in the landlocked area. This parameterization is such that the east and the west sides of the inlet have constant Manning’s n values, and there is a linear increase (or decrease) in the Manning’s n coefficients inside the inlet, with the rate of change proportional to the difference between α and β . Therefore, we refer to this parameterization as the ‘zonal’ profile. The zonal profile has also been used in [46] to represent the simplistic approximation of the realistic Manning’s n field in the coastal domain, where the seabed roughness varies depending on the distance from the coastline, and the monotonic increase in the Manning’s n parameter within the inlet is caused by the sediment transport between the open ocean and the lagoon. By specifying $\alpha = 0.005$ and $\beta = 0.1$, we generate the true field M^{zone} , as shown in the right panel of Fig.2.

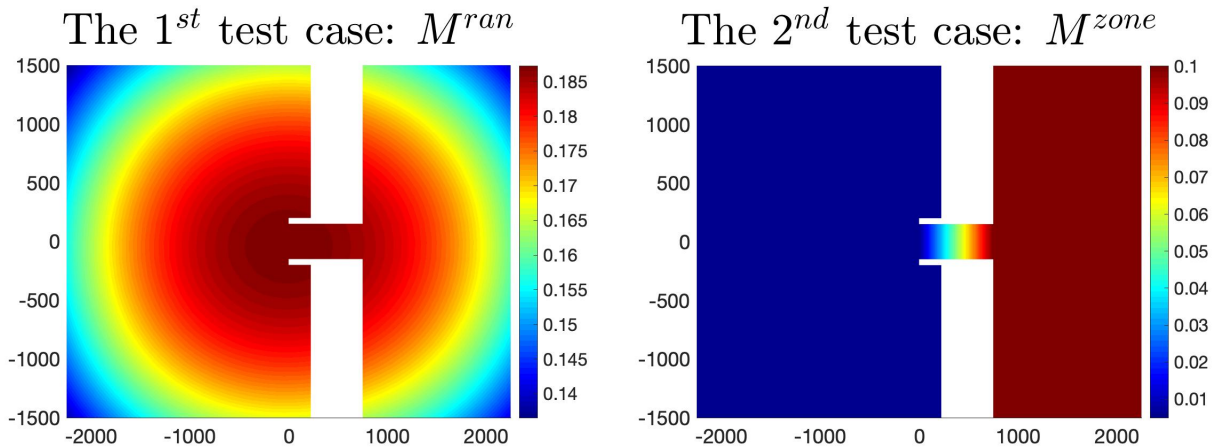


Figure 2: The true Manning’s n field of the two test cases. Truths were used to generate the synthetic water elevation observations for the inference. Left: M^{ran} , generated from sampling the KL coordinates and l of the squared exponential covariance model. Right: M^{zone} generated by setting $\alpha = 0.005$ and $\beta = 0.1$ in Equation (30).

For the inference, we consider the squared exponential prior covariance with $\mathbf{q} = \{l\}$ for both cases. The prior for l is uniform $\mathcal{U}[l_{min} = 1000 \text{ m}, l_{max} = 4000 \text{ m}]$. This range of the correlation length scale was chosen such that l_{max} does not exceed the domain length in any direction, and l_{min} is approximately 1/4

in diameter of the computational domain. This is to ensure that the prior covariance can be adequately represented by just a few eigenmodes. In addition, we do not want to assign an exceedingly small l_{min} to our inference problem because the sparse observation points in our domain cannot resolve the small-scale variability assumed by the prior covariance with small correlation length scale.

For the second test case, M^{zone} , we use a simplified 1D prior covariance in which we implicitly state that a priori, there is no dependence between two points in space on the transverse direction; the dependence between two points only exists in the horizontal (x) direction. To derive such a covariance function, we consider that the Gaussian kernel is separable, that is, the multidimensional kernel can be written as a product of one-dimensional kernels:

$$\mathbf{C}(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^{N_D} \mathbf{C}_i(x_i, x'_i), \quad (31)$$

where, N_D is the number of the dimension of \mathbf{D} . For 2D domain, we use

$$\mathbf{C}_1(x_1, x'_1) = \exp\left(-\frac{(x_1 - x'_1)^2}{2l^2}\right), \quad \mathbf{C}_2(x_2, x'_2) = 1, \quad (32)$$

that is an infinite correlation in the second dimension. The variability range for l is the same as previously.

5. Results and Discussions

5.1. Offline step constructions results

We first construct the reference basis $(\lambda_k^r, \phi_k^r)_{k=1, K}$ derived from the reference covariance $\mathbf{C}^r = \overline{\mathbf{C}}$ for both test cases. Fig.3 shows the first six eigenmodes (from left to right, top to bottom) and the spectrum (bottom) of the KL decomposition of \mathbf{C} of the first test case. We observe that, to retain about 90% of the energy, truncation order $K = 7$ is required. Similar plots for the second test case M^{zone} are shown in Fig.4. In this case, the reference covariance $\overline{\mathbf{C}}$ is simpler and $K = 4$ is sufficient to retain more than 90% of the energy.

To setup the MCMC inference with the change of coordinates method, we compute the coordinates transformation matrix $\mathbf{B}(\mathbf{q})$. For $M^{ran}(\mathbf{x})$, the dimension of $\mathbf{B}(\mathbf{q})$ is 3×3 . For the second test case, $M^{zone}(\mathbf{x})$, we increase K to 4; hence, $\mathbf{B}(\mathbf{q})$ is of the dimension 4×4 .

In a large-scale system, it is not practical to compute $\mathbf{B}(\mathbf{q})$, which requires performing KL decomposition each time \mathbf{q} changes at every MCMC iteration. Therefore, we build the PC surrogate $\tilde{\mathbf{B}}(\mathbf{q})$ of $\mathbf{B}(\mathbf{q})$ according to Equation (20) using an Pseudo-spectral projection (PSP) method [13, 12], a variant of the NISP which uses the nodes of a sparse grid in the domain Ω^* . PC construction using 15 Gaussian quadrature points (15 different realizations of l) based on Legendre polynomials with PC order $o = 5$ was found to be sufficient for approximating $\tilde{\mathbf{B}}(\mathbf{q})$ for both test cases. Fig. 5 shows the results of approximating the true field M^{ran} of the first test case, using the reference basis and the change of coordinates. The top left-plot is M^{ran} generated using Gaussian process with the true $\boldsymbol{\eta}$ and \mathbf{q} . The top-right plot is \widehat{M}_K , which is the approximation of M^{ran} , using the reference basis and transformation $\mathbf{B}(\mathbf{q})$. Finally, the bottom plot is \widehat{M}_K^{PC} , the approximation of M^{ran} with additional PC approximation of the coordinate transformation itself. \widehat{M}_K^{PC} exhibits a larger error in approximating the truth than \widehat{M}_K because it accumulates both the errors from the truncation in building the COC and the PC approximation of $\mathbf{B}(\mathbf{q})$. However, for this simple test case, the difference in error between \widehat{M}_K and \widehat{M}_K^{PC} is almost negligible because a sufficient number of quadratures and polynomial order were used to construct the PC surrogate. Nevertheless, the discrepancy in estimating the truth in the reference basis is noticeable. The circular gradient pattern of M^{ran} in the top-left plot is slightly compressed in the x-direction for both \widehat{M}_K and \widehat{M}_K^{PC} plots. Also, the range between the maximum and the minimum Manning's n coefficients is slightly underestimated. This is clearly due to the K -term truncation in Equation (17). However, overall, we have demonstrated that the COC method is efficient in approximating the Gaussian process \widehat{M}_K on the \mathbf{q} -averaged reference basis.

The second task in the offline step is to build the PC surrogate of the model solution $U(\boldsymbol{\eta}, \mathbf{q})$ following (24), based on the reference Gaussian process $\widehat{M}_K^{\text{PC}}(\boldsymbol{\xi})$ in (23). Following the former setup, $\widehat{M}_K^{\text{PC}}(\boldsymbol{\xi})$ is a stochastic process with \mathbf{q} -averaged covariance function $\widehat{\mathbf{C}}$, in which the KL decomposition is truncated to the first $K = 3$ dominant modes. Employing the NISP approach, we obtain the PC approximation $\widetilde{U}(\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q}))$ of the reference model solution. For the first test case, full tensorization of 1-D Gauss quadrature and Hermite polynomial order $o = 6$ were used to accurately emulate the response surface of the full solution $U(\boldsymbol{\eta}, \mathbf{q})$. With stochastic dimension $s = 3$ and 5 Gauss nodes in each direction, 125 ADCIRC model runs are required. After the surrogate is successfully built, we can evaluate the PC approximation of the ADCIRC forecasts, i.e., the water elevation at observation stations, denoted by $\tilde{u}(\boldsymbol{\xi})$. Given the location in space x_i and time t_i , the components of $\tilde{u}(\boldsymbol{\xi})$ are

$$\tilde{u}_i(\boldsymbol{\xi}) = \widetilde{U}(x_i, t_i, \boldsymbol{\xi}) = \sum_{k=0}^P U_k(x_i, t_i) \Psi_k(\boldsymbol{\xi}), \quad i = 1, \dots, 1620. \quad (33)$$

These model predictions then enter the likelihood computation in the MCMC inference. In Fig.6, we plot for the first test case, the time-evolution of the water elevation prediction, produced by the ADCIRC, contrasted with its surrogate counterpart at three observation stations, with each panel plotting a different realization of $\boldsymbol{\xi}$. We notice that the water elevation predictions are approximated well by the surrogate and, thus, the surrogate can be used in place of the full ADCIRC in MCMC.

Similar offline construction procedures were carried out for the second test case M^{zone} . However, we increase the KL truncation to $K = 4$, for the Gaussian process to capture better the small-scale variability in the true Manning's n field we attempt to infer. This requires 625 ADCIRC runs to build the surrogate, using the full tensorization quadrature method.

5.2. 1st case: inference of a random Manning's n field drawn from a Gaussian process

To perform the MCMC inversion, we start by specifying the priors. Since $\boldsymbol{\eta} = \{\eta_1, \eta_2, \eta_3\}$ are the standard normal random variables, $p_{\boldsymbol{\eta}}(\boldsymbol{\eta})$ is a centered Gaussian. As for the length scale l and the noise hyper-parameter σ_o^2 , we use the uninformative, improper, Jeffrey prior [61]. However, instead of inferring the actual correlation length scale, we seek to infer the normalized correlation length scale l^{norm} , which can be computed through the mapping $[1000, 4000] \mapsto [0, 1]$. For example, the $l^{\text{norm}} = 0.085$ is equivalent to $l = 1255$ m.

Having identified all priors, we can use the data predicted by the surrogate $\tilde{u}(\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q}))$ to compute the likelihood:

$$p(\mathbf{d} \mid \boldsymbol{\eta}, \mathbf{q}, \sigma_o^2) \approx \tilde{p}(\mathbf{d} \mid \boldsymbol{\eta}, \mathbf{q}, \sigma_o^2) = \prod_{i=1}^{N_o} p_{\epsilon}(d_i - u_i(\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q})), \sigma_o^2), \quad (34)$$

where $\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q})$ and p_{ϵ} is given in Equation (25) and Equation (6), respectively. To specify $\widehat{\mathbf{B}}$ in Equation (26), we set $\kappa = 0$ because $\lambda_{k \leq 3}^r / \lambda_k^1$ remains large enough for the current setting. With all components of the Bayes rule defined, we can proceed with generating the posterior $\tilde{p}(\mathbf{d} \mid \boldsymbol{\eta}, \mathbf{q}, \sigma_o^2)$ using an adaptive Metropolis-Hasting algorithm.

We first examine the inference results from using PC-MCMC with COC. A total of 3×10^5 MCMC iterations were found to be sufficient in obtaining the posterior with good resolution. We also observe well-mixed MCMC chains for all inferred parameters (not shown). The marginal posteriors are estimated using a standard kernel density estimation (KDE) method [56]. The posteriors of the inferred 3 first KL coordinates and the hyper-parameter l for the inference of M^{ran} are shown in Fig.7. We also infer σ_o^2 , which we found a MAP close to the value used to generate the data, i.e., $\sigma_{\epsilon}^2 = 0.01$. Similar findings are also the case for all later experiments (not shown for brevity). Concerning the posterior pdf of the KL coordinates, the information gained from the observations is found to be significant only for η_1 , which shows large divergence from the prior and significant decrease in the posterior pdf spread. An interesting result in these plots is in

the posterior pdf of the normalized length-scale l^{norm} as depicted on the right of the second panel. Here, the MAP estimate of the posterior of l^{norm} (presented with dashed blue) matches almost perfectly with the true normalized length-scale used to generate the field M^{ran} . This demonstrates the efficiency of PC-MCMC with COC to estimate the uncertain length-scale of the Gaussian process that best estimates the reference profile. However, Fig.7 is not informative in terms of validating the accuracy of the inferred KL coordinates with respect to the true field M^{ran} , since these plots only display the raw inference results, in which all parameters simultaneously change within the MCMC sampler. Each sample on the posterior is, in fact, defined on a different basis. For this reason, it is not possible to define "true" values of the coordinates, and assess their distance to the MAP, as the coordinate of M^{ran} depends on each sampled value of the hyper-parameters.

To assess the inference results against the truth in a fixed system of coordinates, one needs to express the sampled fields in the same basis. To this end, we express all sampled fields in the reference basis using the relation $\boldsymbol{\xi}(\boldsymbol{\eta}, \boldsymbol{q}) = \widehat{\boldsymbol{B}}(\boldsymbol{q})\boldsymbol{\eta}$ as given in (26). Fig.8 shows the posteriors of the KL coordinates $\boldsymbol{\xi}$ compared to the prior and the KL coordinates of the true field M^{ran} . Similar to the conclusion drawn from Fig.7, the clear information gained from the observations is only revealed by the improvement of the posterior of the 1st KL coordinate ξ_1 . The MAP of the posterior of ξ_1 clearly converges to the value of the first KL coordinate of the truth, and the posterior spread is greatly reduced compared to the prior. This is not the case for ξ_2 or ξ_3 , in which the MAP estimate accurately matches the truth, but shows no clear improvement in term of posterior spread. This is likely due to the fact that the mean of the priors was already close to the true values and because of the observation noise. Nevertheless, Fig.8 further emphasizes the strengths of the PC-MCMC with COC, which is efficient at inferring the hyper-parameters and concurrently improve the estimation of the KL coordinates.

After inferring $\boldsymbol{\xi}$, we can reconstruct the inferred Manning's n field using the KL expansion with known reference eigenbasis, and compare this field to the true field M^{ran} . We present in Fig.9 the comparison of the inferred Manning's n fields from different inference settings and methods: 1) the truth 2) the fixed l basis, 3) the reference basis without using the change of coordinates, and 4) PC-MCMC with COC. First, we observe the top right plot of Fig.9 where the correlation length-scale is fixed to $l^{norm} = 0.189$ or $l = 1567$ m. As expected, we obtain a poor estimate of the truth in this case, simply because the scheme is non-adaptive in \boldsymbol{q} . In addition, the size of fixed l^{norm} , in this case, is larger than the l^{norm} so it cannot resolve small-scale oscillations. As a result, the inference retrieves a MAP value that misses north-south features and appears more one-dimensional than the truth. We also plot in the bottom-left of Fig.9 the inferred Manning's n field using the basis derived from the \boldsymbol{q} -averaged prior covariance $\overline{\boldsymbol{C}}$. In this case, the pattern of the inferred field shows more resemblance to the truth, which can be explained by the better fit of the basis derived from $\overline{\boldsymbol{C}}$ compared to \boldsymbol{C}^r with fixed \boldsymbol{q} . It is nevertheless, visually, a poor estimate of the true field, with a pronounced asymmetry between the north and south parts of the domain. This asymmetry is present but much less significant in the truth. Finally, as expected, the PC-MCMC with COC case in the bottom-right of Fig.9 clearly suggests significant improvement in the spatial structure of the inferred field compared to the two former cases. It actually matches almost perfectly to the true field reconstruction based on the reference basis and COC present in Fig.5. This indicates that the COC approach is very robust, up to the limit of accuracy in the approximation of $\overline{\boldsymbol{C}}$ and $\boldsymbol{B}(\boldsymbol{\eta}, \boldsymbol{q})$. Further, this result underlines the potential of the proposed approach, which is able to produce much more satisfying estimates by encompassing a whole family of priors, and, in the present case, to handle satisfactorily the third coordinate associated with the third field of the reference basis shown in Fig.3.

We further investigate the accuracy of the estimated fields by calculating the percentage error $e_{\%}$ of each inferred field with respect to the truth (figures not shown), calculated node-wise as

$$e_{\%}(\boldsymbol{x}_i) = \left| \frac{M^{in}(\boldsymbol{x}_i) - M^{true}(\boldsymbol{x}_i)}{M^{true}(\boldsymbol{x}_i)} \right| \times 100, \quad i = 1, \dots, 1518, \quad (35)$$

where $M^{true} = M^{ran}$ in the first test case, M^{in} represents the inferred field, and \boldsymbol{x}_i represents the discretization node i in the domain. For the fixed l case, a large error is shown at the southeast corner of the land-locked area of the domain with the maximum error of about 150%. This clearly highlights the mismatch

between the truth and the inferred field when hyper-parameters uncertainty is not considered. Similar to the previous discussion, using the basis derived from $\overline{\mathbf{C}}$ yields very low error in most area. However, the exceedingly large error is still shown in both lower corners of the domain. In contrast, the PC-MCMC with COC yields relatively small $e\%$ throughout the domain, with the maximum relative error at about 60% in areas where the coefficient is the lowest. To complete the discussion on this test case, we mention that the results reported above are not sensitive to the particular realization of M^{ran} ; obviously, unless one uses systematically the correct (sampled) correlation length for each inference, the COC method provides better MAPs on average compared to using a fixed correlation length or the averaged prior.

In many scenarios, the ocean tides may not be very sensitive to the small change in the model parameters such as Manning’s n coefficients. We are also interested in examining the resulting water elevation data produced by ADCIRC run subjected to the inferred Manning’s n field we have obtained. Fig.10 show the time-evolution of the water elevation in the final 12 hours of the ADCIRC simulations. Each curve in each subplot is the water elevation produced with the different MAP values of the Manning’s n field. Here, we want to highlight the difference in the sensitivity of the water elevation to the inferred Manning’s n coefficients with respect to the observed location. Therefore, the three subplots represent the water elevation at different spatial locations in the domain (Fig.1) along the inlet cross-section: the open ocean ($x = -1000$ m, $y = 0$ m), the inlet ($x = 500$ m, $y = 0$ m), and the landlocked area ($x = 1500$ m, $y = 0$ m), respectively. In the open ocean area, the modelled water elevation is less sensitive to the difference in the field of the Manning’s n coefficient compared to the landlocked area. This is reasonable since the west side of the domain has deeper bathymetry, reducing the value of the drag law coefficient [46]. The differences are larger at stations where the model is more sensitive. Among the three inference results, the fixed hyper-parameter case clearly shows the largest difference in water elevation in the landlocked area with respect to the truth (RMSE = 0.0023), consistent with the corresponding inaccurately inferred Manning’s n field. For the landlocked area, using \mathbf{q} -averaged basis with and without COC produce accurate recovery of the water elevation. However, the case without COC shows the best estimate, with the smallest RMSE among all three cases at 0.00015. This can also be expected since, given the sufficiently well-approximated structure of the true field, it is often the case that water elevation in numerical models are not sensitive to its value at every point in the computational domain, but rather more sensitive to the range of the parameter in the domain itself. In our case, the range of Manning’s n coefficients given by the \mathbf{q} -averaged basis without COC is the closest to the range of the true Manning’s n field, while the case with COC slightly underestimates the maximum Manning’s n value. This may also suggest the need to better approximate $\mathbf{B}(\mathbf{q})$ by increasing the number of truncation terms K in (17), at the cost of increasing the computational load for building the PC surrogate of U .

5.3. 2nd case: Inference of a 2D parametrized zonal Manning’s n field

In the 2nd test case, we proceed our assessment of the PC-MCMC with COC method in similar manners as presented in the 1st test case. However, the true Manning’s n field M^{zone} used in this test case is derived independently from the Gaussian process assumption. As a result, this test case better simulates the inference problem in real-world settings. According to the observed spectrum in Fig.4, we set $K = 4$, which gives $\lambda_4/\lambda_1 = 0.101$. This ensures that the \mathbf{q} -averaged reference eigenbasis encapsulates roughly 90% of the information contained in $\overline{\mathbf{C}}$. Without loss of generality, we omit the inference results using fixed- \mathbf{q} basis on the experiments reported below, as we found in the previous section that the fixed- \mathbf{q} case is not very informative. We will instead focus on the comparison between the inference results using \mathbf{q} -averaged basis with and without COC. For simplicity, we shall also refer to the inference using \mathbf{q} -averaged basis without the change of coordinates as the baseline case, while the inference with the change of coordinates will be referred as the COC case.

Fig.11 outlines the inference of M^{zone} of the baseline case. The posteriors of the four coordinates are compared with their respective standard Gaussian priors. We notice the significant improvement of the posteriors distributions for the first two coordinates, while the other two only show the shift in the posterior with no improvement on the pdf spread. With COC, shown in Fig.12, better information gain has clearly been discerned for all four coordinates. Namely, the posteriors of the first three coordinates clearly exhibit smaller spread compared to their respective posteriors in the baseline case, and the posterior of the fourth

coordinate clearly show larger divergence from the prior than that of the baseline case, demonstrating the ability of the MCMC sampler to explore the posterior in a more suitable subspace. These improvements are indeed the result of including the uncertainty in l in the inference problem. Note also that, as in the previous case, it is not possible to provide true values for the coordinates of M^{zone} in the COC approach, as these vary with the sample value of the correlation length.

Next, we examine the fields of inferred Manning’s n coefficients constructed using the MAP estimates of η and l . The resulting Manning’s n fields and their respective absolute errors with respect to the truth M^{zone} are presented in Fig.13. Looking at the second line of Fig.13, the inferred Manning’s n fields from both the baseline and the COC case show resemblance in the overall spatial structure with respect to the truth. Both methods are observed to be equally successful in recovering the Manning’s n value in the open ocean area ($\alpha = 0.005$). However, overall, the COC case provides better results for the Manning’s n values in the inlet and the landlocked areas. The range of the Manning’s n values is overestimated in the two cases, but appears to be larger for the baseline case, as can be discerned by observing and comparing the first three subplots. In the bottom line of Fig.13, the percent error maps are presented. Overall, the percent error was found to be small over most of the area, except at the west side of the inlet, where the error reaches almost 390% for the baseline case and 360% for the COC case along the edge of the two barriers. This is due to a significant overestimation in areas where the coefficient is actually small. This behavior can be explained by the squared exponential prior covariance model considered to approximate the field, which is very smooth in space and therefore not well-suited for representing the discontinuous slopes of M^{zone} at the two inlet boundaries. The results then highlights that even, in a situation where the a priori is poor and the truth very unlikely, the introduction of the change of coordinate does not compromise the procedure but only slightly improve the MAP estimate.

To better discern the differences in the inferred field results between the baseline and the COC cases, we present in Fig.14 the cross-section plot of the parameter fields at $y = 0$ with varying Manning’s n value in the x -direction. Clearly, the truth M^{zone} is not a smooth function with sharp turns at the two entrances of the inlet, which is the result of the field parameterization in Equation (30). Similar to the conclusion we have drawn from the previous figure, both baseline and COC cases accurately estimate the parameter field in the open ocean with the exception of a small area around the western boundary. This is again the result of the attempt to fit the constant function in a basis of eigenmodes for squared exponential prior covariance. The consequence of this limitation is more pronounced when estimating large Manning’s n values in the even smaller landlocked area. Here, the results from both cases manifest significant discrepancy with respect to the truth. The shape of the inferred field subjected to a smooth Gaussian process with a few modes do not well-represent the constant profile in the landlocked area. That aside, the COC case demonstrates clear improvement compared to the baseline in recovering the M^{zone} profile in this area, with the maximum Manning’s n value of 0.1088; less than 10% larger than $\beta = 0.1$. The baseline, on the other hand, produces the maximum Manning’s n value of 0.1233; 20% larger than the true maximum. Furthermore, one can observe the clear benefit of including the uncertainty in l in the COC case, that is, it also attempts to recover the maximum Manning’s n coefficients of the eastern boundary of the landlocked area.

Finally, the sensitivity of the produced water elevation to the inferred field is investigated. The water elevation at three locations during the 12 final hours of ADCIRC simulation for different Manning’s n fields are illustrated in Fig.15. Similar to the 1st test case M^{ran} , the resulting water elevation in the open ocean of the baseline and the COC case is well estimated compared to the truth. This is both because the inferred Manning’s n fields from both cases in this area are well recovered and the water height is not very sensitive to the differences in Manning’s n values in the area with deep bathymetry. The distinction in the resulting water elevation between the three cases starts to be noticeable in the inlet and the landlocked region. The baseline produces the larger error compared to the COC case in both areas, mainly, due to the exaggeration in the range of the estimated Manning’s n values with the maxima exceeding more than 20% of that of the true field. The improvement in the resulting water elevation produced by the COC case over the baseline case is clearly outlined in the landlocked area.

To summarize, this section has demonstrated the efficiency of the change of coordinates approach for inference in a more realistic coastal ocean scenario compared to the first test case, with significant improvements in the inference results.

6. Conclusions

This work presented an extension of the Bayesian inference framework proposed in [61] and tested its application to a 2D spatially varying parameter field from a prior defined by a Gaussian process in the context of coastal ocean modeling. The approach is designed to handle hyper-parameters \mathbf{q} in the prior covariance kernel. To infer the spatial field, we relied on a generalized Karhunen-Loève (KL) expansion involving a reference (spatial) basis, independent of the hyper-parameters, while the effects of changing the hyper-parameters is accounted for by the random coordinates of the expansion. Specifically, the reference coordinates are defined by a Change Of Coordinates that maps the original KL coordinates $\boldsymbol{\eta}$, in the \mathbf{q} -dependent basis, into \mathbf{q} -dependent coordinates $\boldsymbol{\xi}(\boldsymbol{\eta}, \mathbf{q})$ in the reference basis. The reference covariance was chosen as the \mathbf{q} -averaged covariance, which has been shown to minimize the resulting representation error of the Gaussian process [61]. The change of coordinates is also exploited to construct a unique surrogate model of the forward model, which is then used to speed up the posterior sampling in a MCMC algorithm.

The Bayesian inference with COC can be summarized as follows. In the offline step, we defined and decomposed the reference covariance function from which the reference eigenbasis is computed through KL decomposition. Then, we construct the PC surrogate of the forward model that accounts for the dependence of the model prediction on the KL coordinates defined on the reference basis. Next, in the online step, every MCMC iteration samples $\boldsymbol{\eta}$ and \mathbf{q} are mapped to coordinates $\boldsymbol{\xi}$ in the reference basis, which is the result of transforming $\boldsymbol{\eta}$ through the relation $\boldsymbol{\xi} = \mathbf{B}(\mathbf{q})\boldsymbol{\eta}$. Using these reference coordinates, we evaluate the surrogate model of the predictions required in the likelihood evaluation to explore the Bayesian posterior. The most demanding part of the online step is the determination of the COC matrix $\mathbf{B}(\mathbf{q})$ for each sampled \mathbf{q} , as it requires the decomposition of the corresponding prior covariance. Therefore, in this work, we proposed and analyzed the possibility to approximate $\mathbf{B}(\mathbf{q})$ with an auxiliary PC expansion in the hyper-parameters \mathbf{q} . This approximation is constructed in the offline step and used in the online step to avoid multiple costly decompositions of the correlation kernel. This greatly accelerates the online step and allows the inference framework with COC to be applied to realistic ocean models.

We tested the efficiency the proposed inference framework in an idealized coastal ocean setting, by conducting OSSEs to infer the Manning’s n field from synthetic water elevation observations generated by running ADCIRC model with the assumed known true field. Two test cases with different reference Manning’s n profile were utilized for this analysis: 1) the true profile drawn randomly from a Gaussian process with a prescribed hyper-parameter value and 2) the zonal Manning’s n profile parametrized using a linear equation. While the first test case is quite informative in term of assessing the proposed framework by comparing directly the inferred $\boldsymbol{\eta}$ and \mathbf{q} to those used to generate the truth, the second test case represents a more realistic Bayesian inference scenario where the truth is not predetermined by a similar underlying stochastic process as the prior. In each test case, we compared the inference results using a fixed covariance function with those using the COC. We observed better information gain when the hyper-parameters were included in the inference, most notably in the first test case. The COC case also accurately recovered the true correlation length-scale in the first test case, resulting in a more accurate recovery of the true Manning’s n field. Similar results have been obtained in the second test case, that is, including the COC in the Bayesian inference helps producing a slightly better estimate of the true parameter profile compared to the baseline where the inference is performed without incorporating the change of coordinate technique.

Despite the efficiency of the Bayesian inference with COC, this approach calls for a few comments. Firstly, the COC method is very robust, but its accuracy in transforming the coordinates from one basis to another is limited by the approximation reference basis and the coordinate transformation matrix. Secondly, the COC method relies on the \mathbf{q} -continuity of $\mathbf{B}(\mathbf{q})$. In the large-scale problems where the domain has several spatial dimension, this assumption would not hold through the emergence of the multiplicity of eigenvalues and crossing of eigen-branches with $\mathbf{B}(\mathbf{q})$. This limits the scope of using COC to the few first dominant eigenmodes before the crossing/multiplicity issues become intractable. Motivated by these limitations, we plan in our future work to develop the inference method accounting for the uncertainty in the hyper-parameters that works on quantities that directly present smoothness and unambiguous dependences with \mathbf{q} , such as the prior covariance, thus alleviating the need to follow a particular eigen-mode as the hyper-parameter \mathbf{q} varies.

Acknowledgements

This work was supported by the Office of Sponsored Research (OSR) at King Abdullah University of Science and Technology (KAUST), Saudi Arabia (grant number CRG3-2016).

References

- [1] Alen Alexanderian, Justin Winokur, Ihab Sraj, Ashwanth Srinivasan, Mohamed Iskandarani, William C. Thacker, and Omar M. Knio. Global sensitivity analysis in an ocean general circulation model: a sparse spectral projection approach. *Computational Geosciences*, 16(3):757–778, Jun 2012.
- [2] MU Altaf, T Butler, X Luo, C Dawson, T Mayo, and Ibrahim Hoteit. Improving short-range ensemble kalman storm surge forecasting using robust adaptive inflation. *Monthly Weather Review*, 141(8):2705–2720, 2013.
- [3] MU Altaf, N Raboudi, ME Gharamti, C Dawson, MF McCabe, and I Hoteit. Hybrid vs adaptive ensemble kalman filtering for storm surge forecasting. *AGU Fall Meeting Abstracts*, 1:3352, 2014.
- [4] J.L. Anderson. An ensemble adjustment kalman filter for data assimilation. *Monthly weather review*, 129(12):2884–2903, 2001.
- [5] JD Annan, JC Hargreaves, NR Edwards, and R. Marsh. Parameter estimation in an intermediate complexity earth system model using an ensemble kalman filter. *Ocean Modelling*, 8(1):135–154, 2005.
- [6] J. Besag, P. Green, D. Higdon, , and K. Mengersen. Bayesian computation and stochastic systems. *Stat. Sci.*, 10:3–41, Mon. 1995.
- [7] Craig H. Bishop, Brian J. Etherton, and Sharanya J. Majumdar. Adaptive sampling with the ensemble transform kalman filter. part i: Theoretical aspects. *Monthly Weather Review*, 129(3):420–436, 2015/11/02 2001.
- [8] W. P. Budgell. Stochastic filtering of linear shallow water wave processes. *SIAM Journal on Scientific and Statistical Computing*, 8(2):152–170, 1987.
- [9] S. Bunya, J. Dietrich, J. Westerink, B. Ebersole, J. Smith, J. Atkinson, R. Jensen, D. Resio, R. Luettich, C. Dawson, and et.al. A high-resolution coupled riverine flow, tide, wind, wind wave, and storm surge model for southern louisiana and mississippi. part i: Model development and validation. *Monthly Weather Review*, 138(2):345–377, 2010.
- [10] T. Butler, M.U. Altaf, C. Dawson, I. Hoteit, X. Luo, and T. Mayo. Data assimilation within the advanced circulation (adcirc) modeling framework for hurricane storm surge forecasting. *Monthly Weather Review*, 140:2215–2231, 2012.
- [11] T Butler, MU Altaf, C Dawson, Ibrahim Hoteit, X Luo, and T Mayo. Data assimilation within the advanced circulation (adcirc) modeling framework for hurricane storm surge forecasting. *Monthly Weather Review*, 140(7):2215–2231, 2012.
- [12] P. R. Conrad and Y. M. Marzouk. Adaptive smolyak pseudospectral approximations. *SIAM J. Sci. Comp.*, 35(6):A2643–A2670, 2013.
- [13] P. G. Constantine, M. S. Eldred, and E. T. Phipps. Sparse pseudospectral approximation method. *Comp. Meth. App. Mech. and Eng.*, 229:1–12, 2012.
- [14] Noel Cressie and Gardar Johannesson. Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):209–226, 2008.
- [15] J. Dietrich, S. Bunya, J. Westerink, B. Ebersole, J. Smith, J. Atkinson, R. Jensen, D. Resio, R. Luettich, C. Dawson, and et al. A high-resolution coupled riverine flow, tide, wind, wind wave, and storm surge model for southern louisiana and mississippi. part ii: Synoptic description and analysis of hurricanes katrina and rita. *Monthly Weather Review*, 138(2):378–404, 2010.
- [16] J. Dietrich, J. Westerink, A. Kennedy, J. Smith, R. Jensen, M. Zijlema, L. Holthuijsen, C. Dawson, R. Luettich Jr, M. Powell, and et al. Hurricane gustav (2008) waves and storm surge: Hindcast, synoptic analysis, and validation in southern louisiana. *Monthly Weather Review*, 139(8):2488–2522, 2011.
- [17] A. H. ElSheikh, C. C. Pain, F. Fang, J. L. M. A. Gomes, and I. M. Navon. Parameter estimation of subsurface flow models using iterative regularized ensemble kalman filter. *Stochastic Environmental Research and Risk Assessment*, 27(4):877–897, 2013.
- [18] Reinhard Furrer, Marc G Genton, and Douglas Nychka. Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3):502–523, 2006.
- [19] D. Gamerman and H. Lopes. *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. Chapman and Hall/CRC, Boca Raton, 2006.
- [20] R. Ghanem and J. Red-Horse. Propagation of probabilistic uncertainty in complex physical systems using a stochastic finite element approach. *Fluid Dyn. Res.*, 133:137–144, 1999.
- [21] R. Ghanem and P. Spanos. : *Stochastic Finite Elements: A Spectral Approach*. Dover, New York, 2nd edition, 2002.
- [22] R.G. Ghanem and P.D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Civil, Mechanical and Other Engineering Series. Dover Publications, 2003.
- [23] Loïc Giraldi, Olivier P. Le Maître, Kyle T. Mandli, Clint N. Dawson, Ibrahim Hoteit, and Omar M. Knio. Bayesian inference of earthquake parameters from buoy data using a polynomial chaos-based surrogate. *Computational Geosciences*, 21(4):683–699, Aug 2017.
- [24] Heikki Haario, Eero Saksman, and Johanna Tamminen. An adaptive metropolis algorithm. *Bernoulli*, 7(2):223–242, 04 2001.
- [25] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

- [26] Y. Ho and R. Lee. A bayesian approach to problems in stochastic estimation and control. *IEEE Transactions on Automatic Control*, 9(4):333–339, Oct 1964.
- [27] Ibrahim Hoteit, Dinh-Tuan Pham, and Jacques Blum. A simplified reduced order kalman filtering and application to altimetric data assimilation in tropical pacific. *Journal of Marine Systems*, 36(1–2):101 – 127, 2002.
- [28] G Korres I Hoteit, G Triantafyllou. Using low-rank ensemble kalman filters for data assimilation with high dimensional imperfect models. *JNAIAM*, 2(1-2):67–78, 2007.
- [29] Chester P Jelesnianski. Numerical computations of storm surges without bottom stress. *Monthly Weather Review*, 94(6):379–394, 1966.
- [30] Jari P Kaipio and Erkki Somersalo. Statistical inversion theory. *Statistical and Computational Inverse Problems*, pages 49–114, 2005.
- [31] K. Karhunen. *Ueber lineare Methoden in der Wahrscheinlichkeitsrechnung*. Annales Academiae scientiarum Fennicae. Series A. 1, Mathematica-physica. 1947.
- [32] A. B. Kennedy, U. Gravois, B. C. Zachry, J. J. Westerink, M. E. Hope, J. C. Dietrich, M. D. Powell, A. T. Cox, R. A. Luettich, and R. G. Dean. Origin of the hurricane ike forerunner surge. *Geophysical Research Letters*, 38(8), 2011.
- [33] M. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *J. of the Royal Stat. Soc. Series B (Stat. Method.)* 63 (2001), 63:425–464, 2011.
- [34] Ingemar P.E. Kinnmark and William G. Gray. The $2\delta x$ -test: A tool for analyzing spurious oscillations. *Advances in Water Resources*, 8(3):129 – 135, 1985.
- [35] O. Knio and O. Maitre. Uncertainty propagation in cfd using polynomial chaos decomposition. *Physica D*, 38:616–640, 2006.
- [36] P. Lax. *Linear Algebra*. Wiley-Interscience, 1996.
- [37] O. Le Maitre and O. Knio. *Spectral Methods for Uncertainty Quantification with Applications to Computational Fluid Dynamics*. Springer, Berlin, 2010.
- [38] O. Le Maitre, O. Knio, H Najm, and R. Ghanem. A stochastic projection method for fluid flow. i. basic formulation. *Journal of Computational Physics*, 173:481–511, 2001.
- [39] O. Le Maitre, H. Najm, P. Pébay, R. Ghanem, and O. Knio. Multi-resolution-analysis scheme for uncertainty quantification in chemical systems. *SIAM Journal Scientific Computation*, 29(2):864–889, 2007.
- [40] M Loeève. *Fonctions aleatoires de second order*. Processus Stochastiques et Mouvement Brownien. Hermann, Paris, 1947.
- [41] R. Luettich and J. Westerink. *Formulation and numerical implementation of the 2D/3D ADCIRC finite element model version 44*. XX. R. Luettich, 2004.
- [42] RA Luettich Jr, JJ Westerink, and Norman W Scheffner. Adcirc: An advanced three-dimensional circulation model for shelves, coasts, and estuaries. report 1. theory and methodology of adcirc-2ddi and adcirc-3dl. Technical report, COASTAL ENGINEERING RESEARCH CENTER VICKSBURG MS, 1992.
- [43] Daniel R. Lynch and William G. Gray. A wave equation model for finite element tidal computations. *Computers & Fluids*, 7(3):207 – 228, 1979.
- [44] Youssef M. Marzouk and Habib N. Najm. Dimensionality reduction and polynomial chaos acceleration of bayesian inference in inverse problems. *Journal of Computational Physics*, 228(6):1862 – 1902, 2009.
- [45] Youssef M. Marzouk, Habib N. Najm, and Larry A. Rahn. Stochastic spectral methods for efficient bayesian solution of inverse problems. *Journal of Computational Physics*, 224(2):560 – 586, 2007.
- [46] Talea Mayo, Troy Butler, Clint Dawson, and Ibrahim Hoteit. Data assimilation within the advanced circulation (adcirc) modeling framework for the estimation of manning’s friction coefficient. *Ocean Modelling*, 76:43–58, April 2014.
- [47] H. Najm, B. Debusschere, Y. Marzouk, S. Widmer, and O Le Maitre. Multi-resolution-analysis scheme for uncertainty quantification in chemical systems. *International Journal of Numerical Methods in Engineering*, 80(6):789–814, 2009.
- [48] B. Phenix, J. Dinaro, M. Tatang, J. Tester, J. Howard, and G. McRae. Incorporation of parametric uncertainty into complex kinetic mechanisms: application to hydrogen oxidation in supercritical water. *Combustion and Flame*, 112:132–146, 1998.
- [49] Derek J Posselt and Craig H Bishop. Nonlinear parameter estimation: Comparison of an ensemble kalman smoother with a markov chain monte carlo algorithm. *Monthly Weather Review*, 140(6):1957–1974, 2012.
- [50] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- [51] G. Robert, P. Casella. Monte carlo statistical methods. *Applied Mathematical Sciences*, 160(ISBN 0-387-22073-9):344, 2004.
- [52] Gareth O. Roberts and Jeffrey S. Rosenthal. Examples of adaptive mcmc. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.
- [53] Maher Salloum, Alen Alexanderian, Olivier P. Le Maitre, Habib N. Najm, and Omar M. Knio. Simplified csp analysis of a stiff stochastic ode system. *Computer Methods in Applied Mechanics and Engineering*, 217-220:121 – 138, 2012.
- [54] Huiyan Sang and Jianhua Z. Huang. A full scale approximation of covariance functions for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(1):111–132, 2012.
- [55] G. Y. H. El Serafy and A. E. Mynett. Improving the operational forecasting system of the stratified flow in osaka bay using an ensemble kalman filter-based steady state kalman filter. *Water Resource Research*, 44:W06416, doi:10.1029/2006WR005412, 2008.
- [56] B. W. Silverman. On the estimation of a probability density function by the maximum penalized likelihood method. *Ann. Statist.*, 10(3):795–810, 09 1982.
- [57] Adil Siripatana, Talea Mayo, Omar Knio, Clint Dawson, Olivier Le Maitre, and Ibrahim Hoteit. Ensemble kalman filter inference of spatially-varying manning’s n coefficients in the coastal ocean. *Journal of Hydrology*, 562:664 – 684, 2018.

- [58] Adil Siripatana, Talea Mayo, Ihab Sraj, Omar Knio, Clint Dawson, Olivier Le Maitre, and Ibrahim Hoteit. Assessing an ensemble kalman filter inference of manning’s n coefficient of an idealized tidal inlet against a polynomial chaos-based mcmc. *Ocean Dynamics*, pages 1–28, 2017.
- [59] J. V. T. Sorensen and H. Madsen. Parameter sensitivity of three kalman filter schemes for assimilation of water levels in shelf sea models. *Ocean Modelling*, 11:441–463, 2006.
- [60] Ihab Sraj, Mohamed Iskandarani, Ashwanth Srinivasan, W Carlisle Thacker, Justin Winokur, Alen Alexanderian, Chia-Ying Lee, Shuyi S Chen, and Omar M Knio. Bayesian inference of drag parameters using axbt data from typhoon fanapi. *Monthly Weather Review*, 141(7):2347–2367, 2013.
- [61] Ihab Sraj, OlivierP Le Maître, OmarM Knio, and Ibrahim Hoteit. Coordinate transformation and polynomial chaos for the bayesian inference of a gaussian process with parametrized prior covariance function. *Computer Methods in Applied Mechanics and Engineering*, 298:205–228, 2016.
- [62] Ihab Sraj, KyleT Mandli, OmarM Knio, Clint Dawson, and Ibrahim Hoteit. Uncertainty quantification and inference of manning’s friction coefficients using dart buoy data during the tohoku tsunami. *Ocean Modelling*, 83:82–97, 2014.
- [63] Ihab Sraj, SarahE. Zedler, OmarM. Knio, CharlesS. Jackson, and Ibrahim Hoteit. Polynomial chaos-based bayesian inference of k-profile parameterization in a general circulation model of the tropical pacific. *Monthly Weather Review*, 0(0):null, 2016.
- [64] M.L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics. Springer New York, 1999.
- [65] D. Stoyan. Matérn, b.: Spatial variation. 2nd ed., springer-verlag, berlin, heidelberg, new york, london, paris, tokyo 1986, 151 s., dm 33,—. *Biometrical Journal*, 30(5):594–594, 1988.
- [66] Jacob Viborg Tornfeldt Sørensen and Henrik Madsen. Efficient kalman filter techniques for the assimilation of tide gauge data in three-dimensional modeling of the north sea and baltic sea system. *Journal of Geophysical Research: Oceans (1978-2012)*, 109(C3), 2005.
- [67] Piyush M. Tagade and Han-Lim Choi. A generalized polynomial chaos-based method for efficient bayesian calibration of uncertain computational models. *Inverse Problems in Science and Engineering*, 22(4):602–624, 2014.
- [68] Piyush. M. Tagade and Han-Lim Choia. A generalized polynomial chaos-based method for efficient bayesian calibration of uncertain computational models. *Inverse Problems in Science and Engineering*, 22(4):602–624, 2014.
- [69] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. siam, 2005.
- [70] M Verlaan and AW Heemink. Tidal flow forecasting using reduced rank square root filters. *Stochastic Hydrology and Hydraulics*, 11(5):349–368, 1997.
- [71] J. J. Westerink, R. A. Luettich, J. C. Feyen, J. H. Atkinson, C. Dawson, H. J. Roberts, M. D. Powell, J. P. Dunion, E. J. Kubatko, and H. Pourtaheri. A basin-to channel-scale unstructured grid hurricane storm surge model applied to southern louisiana. *Monthly Weather Review*, 136(3):833–864, 2008.
- [72] N. Wiener. The homogeneous chaos. *American Journal of Mathematics*, 60(4):897–936, 1938.
- [73] D. Xiu and G. Karniadakis. The weiner-askey polynomial chaos for stochastic differential equations. *SIAM J. of Sci. Comp.*, 24:619–644, 2002.
- [74] Dongbin Xiu. *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton University Press, Princeton, NJ, USA, 2010.
- [75] T. Yanagi. *Coastal oceanography, Vol. 1*. Springer, 1999.

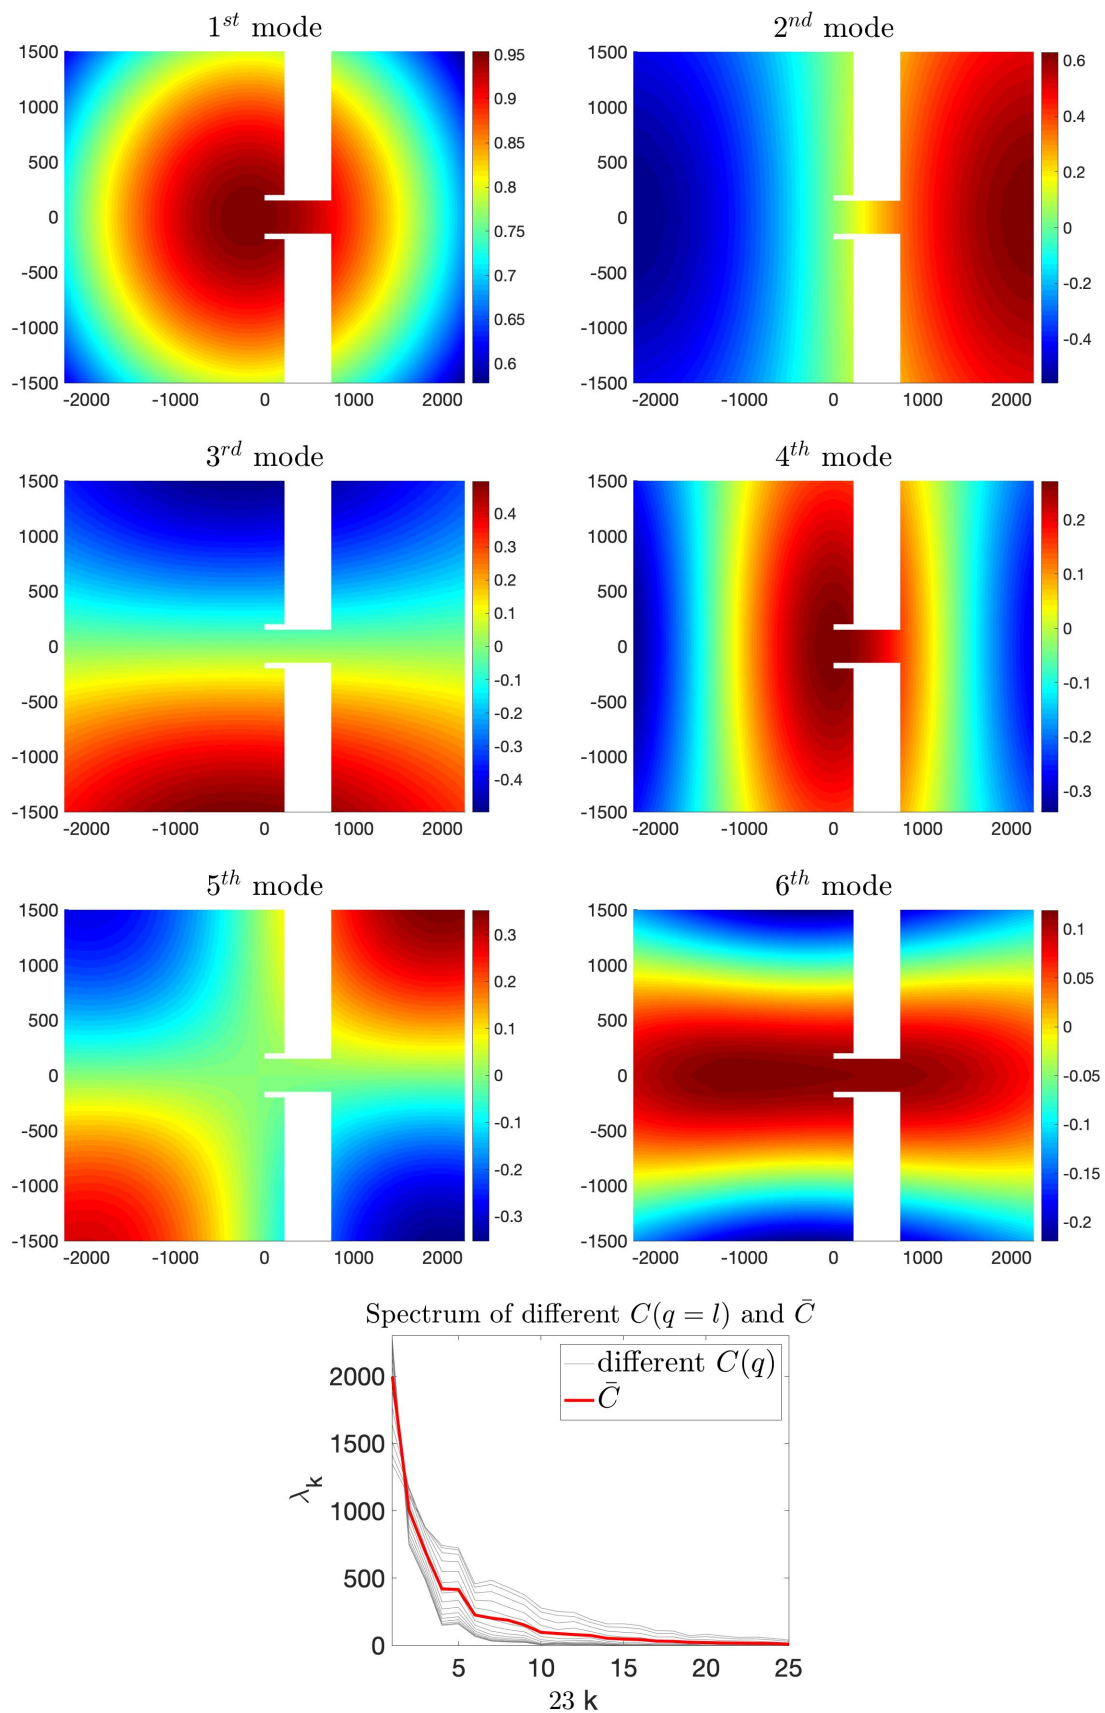


Figure 3: Resulting basis from KL decomposition of \bar{C} for the first test case, M^{ran} . First six panels: the first six eigenmodes (from left to right, top to bottom), bottom: the spectrum of \bar{C} .

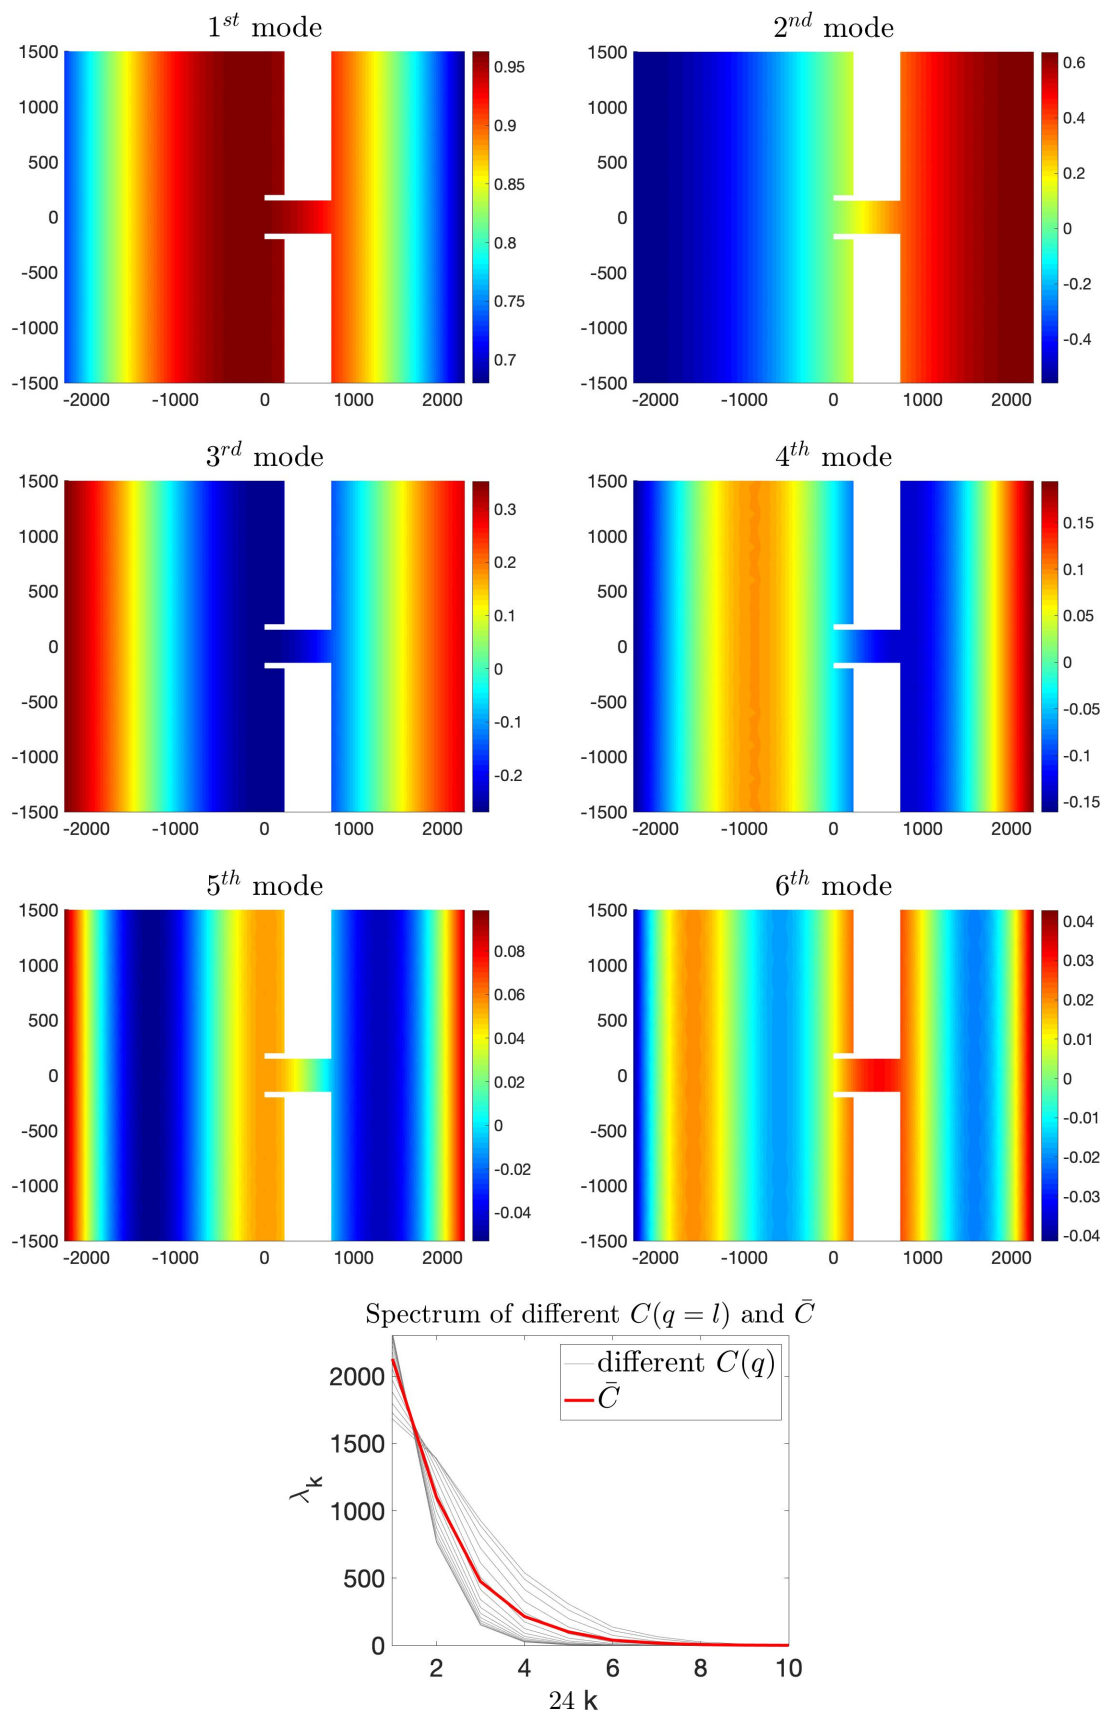


Figure 4: Resulting basis from KL decomposition of \bar{C} for the second test case, M^{zone} . First six panels: the first six eigenmodes (from left to right, top to bottom), bottom: the spectrum of \bar{C} .

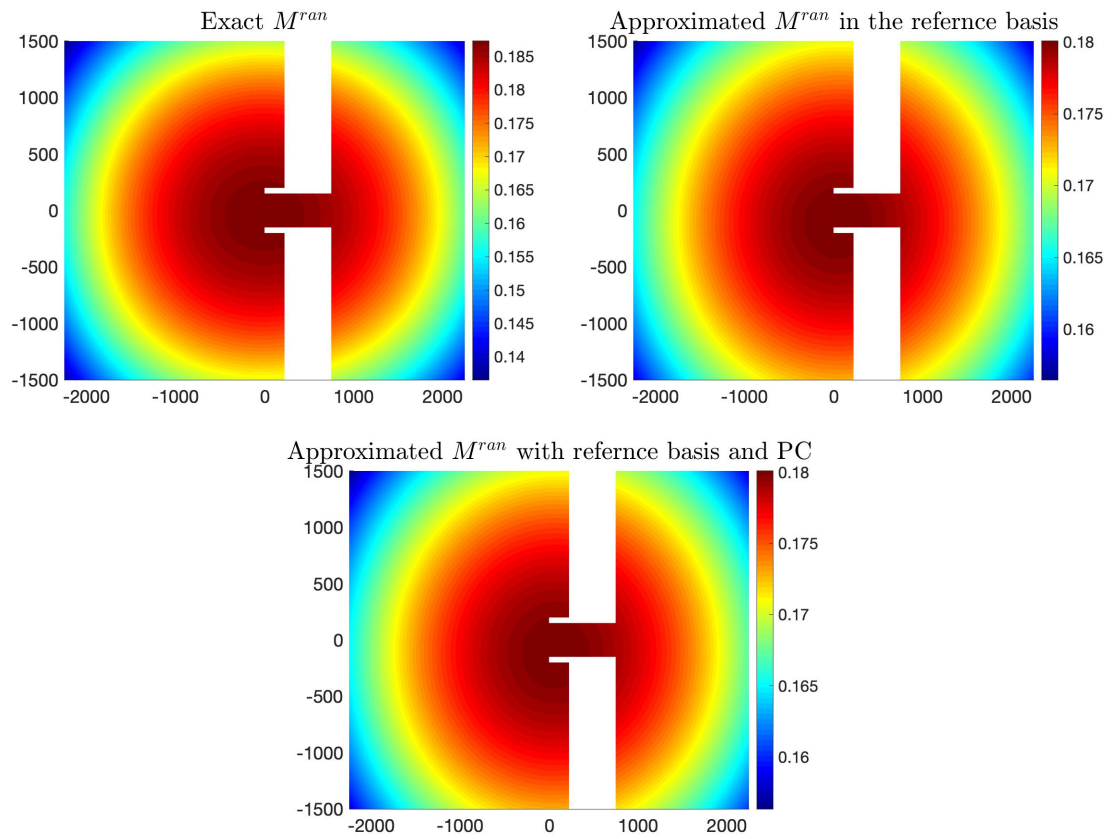


Figure 5: Approximation of M^{ran} , using the reference basis and change of coordinates. Top-left: M_K , the true profile, top-right: \widehat{M}_K , the approximation of M_K in the reference basis, and bottom: \widehat{M}_K^{PC} , the approximation of M_K in the reference basis with the PC approximation of the change of coordinates, $\widetilde{\mathbf{B}}(\mathbf{q})$.

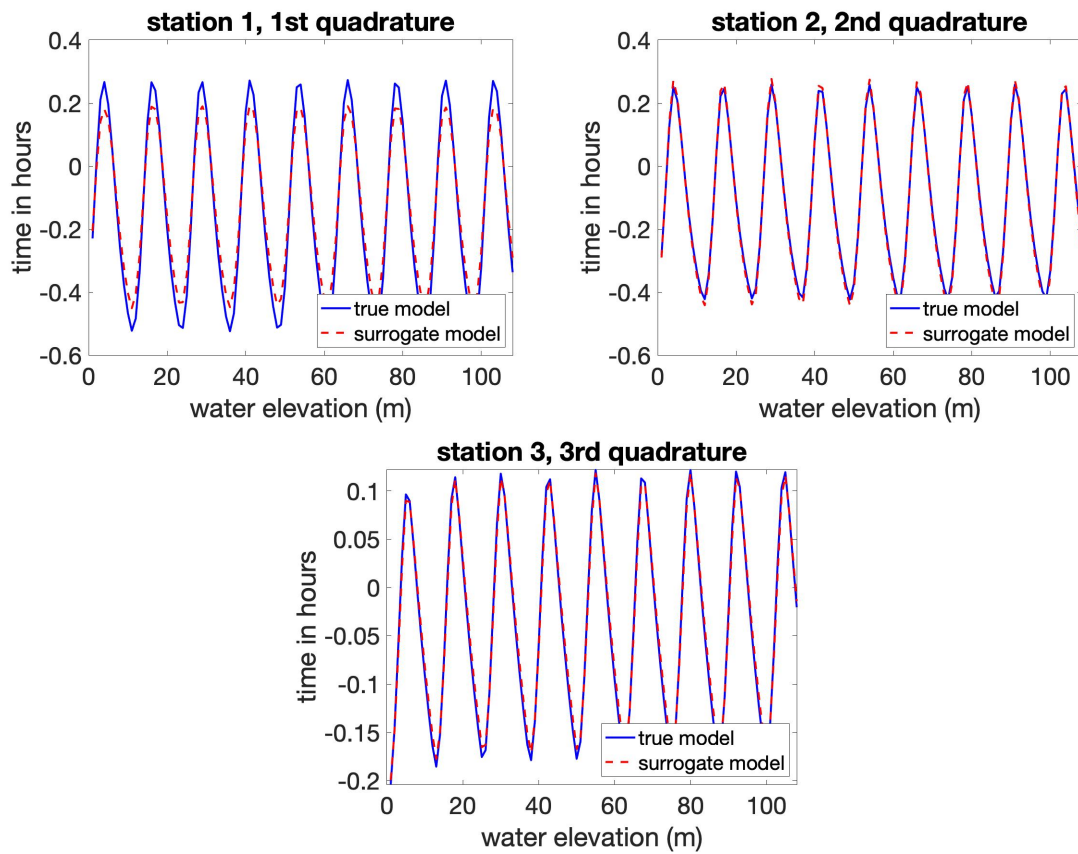


Figure 6: Time-evolution of the water elevation as produced by the full model (blue) and the surrogate model (red). Each plot shows the water elevation at different observation stations and different quadrature points (i.e., different realization of ξ), which characterizes the Manning's n field.

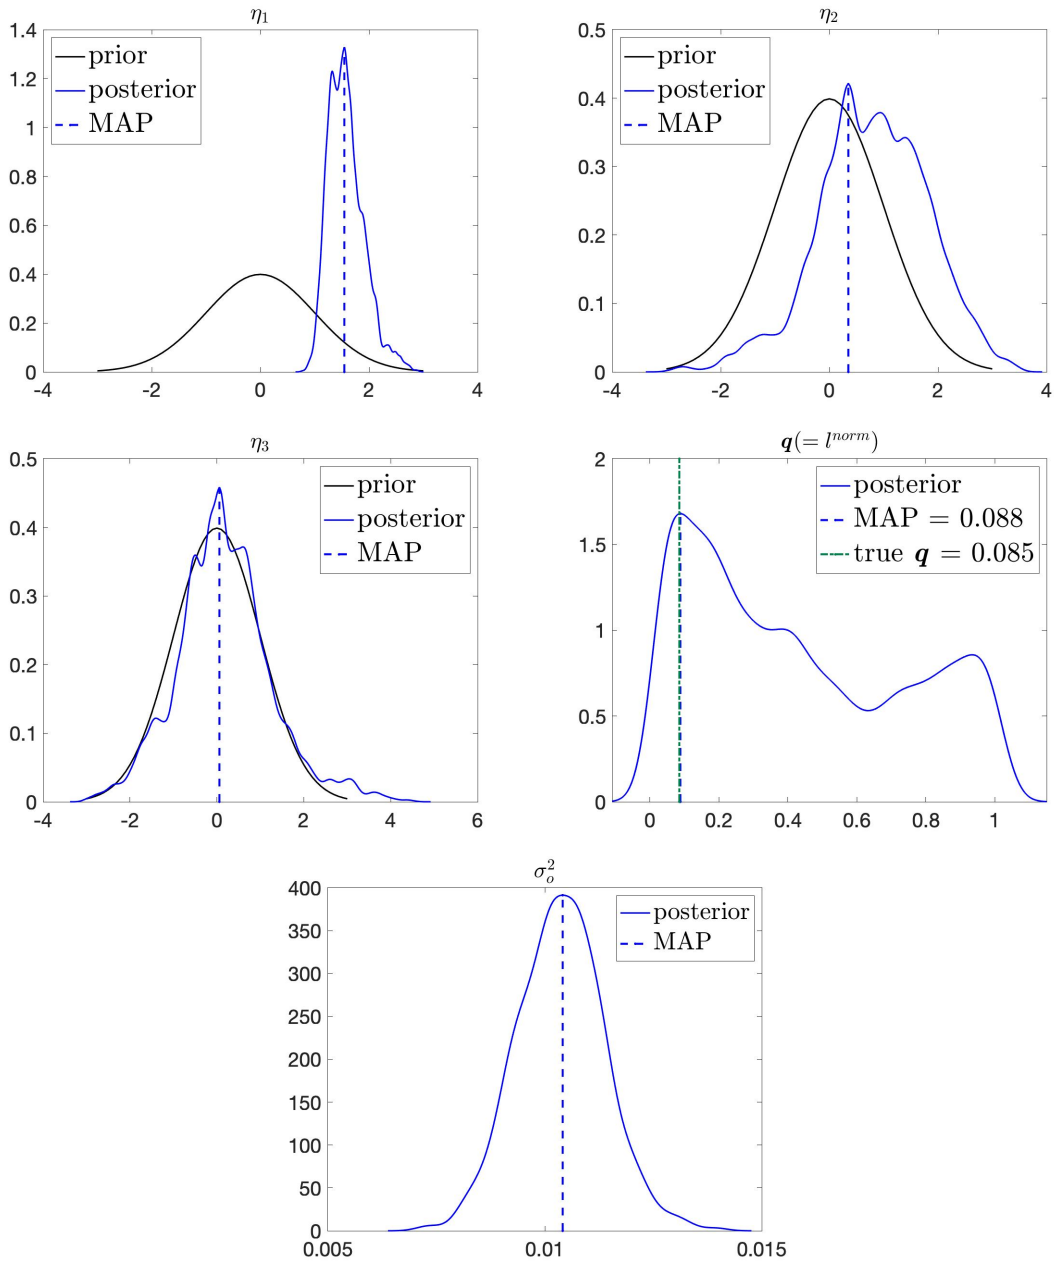


Figure 7: Comparison between the priors and marginal posteriors of the 3 KL coordinates η_k , the length-scale hyper-parameter l^{norm} , and the noise variance hyper-parameter σ_σ^2 (only the posteriors are shown for l^{norm} and σ_σ^2) for the inference of M^{ran} using the COC method. The maximum a posteriori probability (MAP) estimate for each parameter are also presented using dashed blue line. The true l^{norm} used to generate the true field is plotted in dashed green.

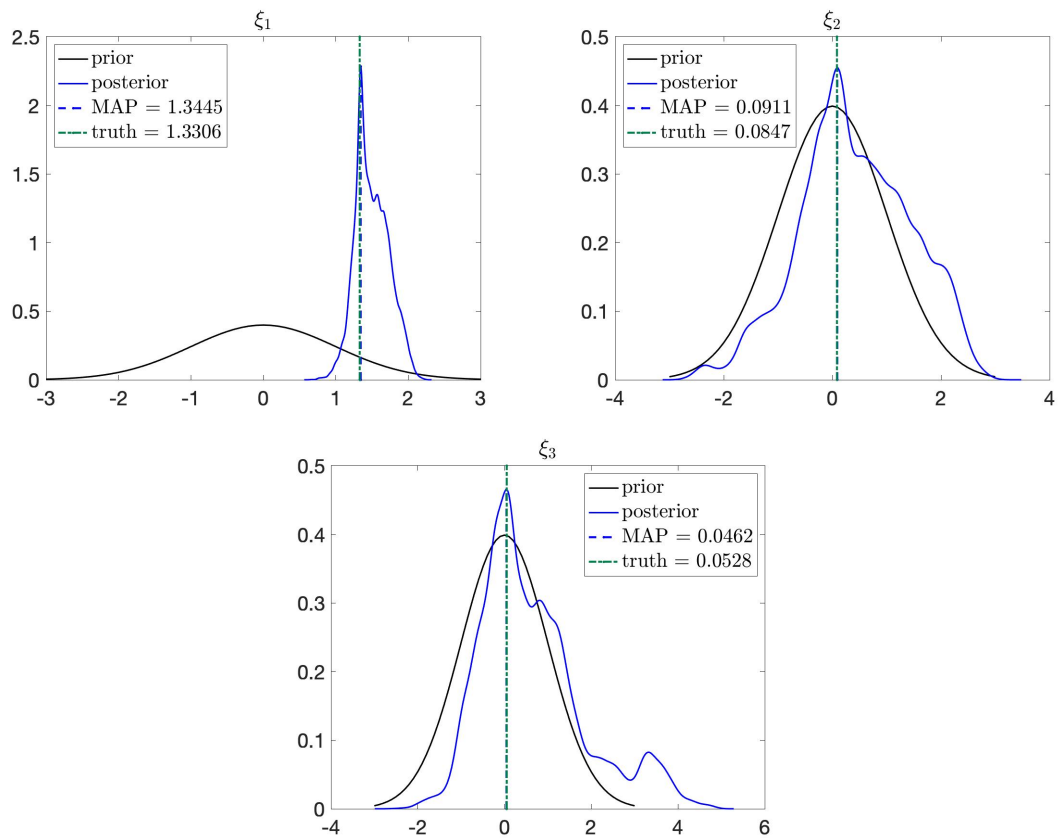


Figure 8: Comparison between the priors, marginal posteriors of the 3 KL coordinates η_k and the true KL coordinates of M^{ran} in the \mathbf{q} -averaged reference subspace (ξ). In each plot, the prior is the black curve, the posterior is in blue, the MAP estimate is in dashed blue, and the truth is in dashed green.

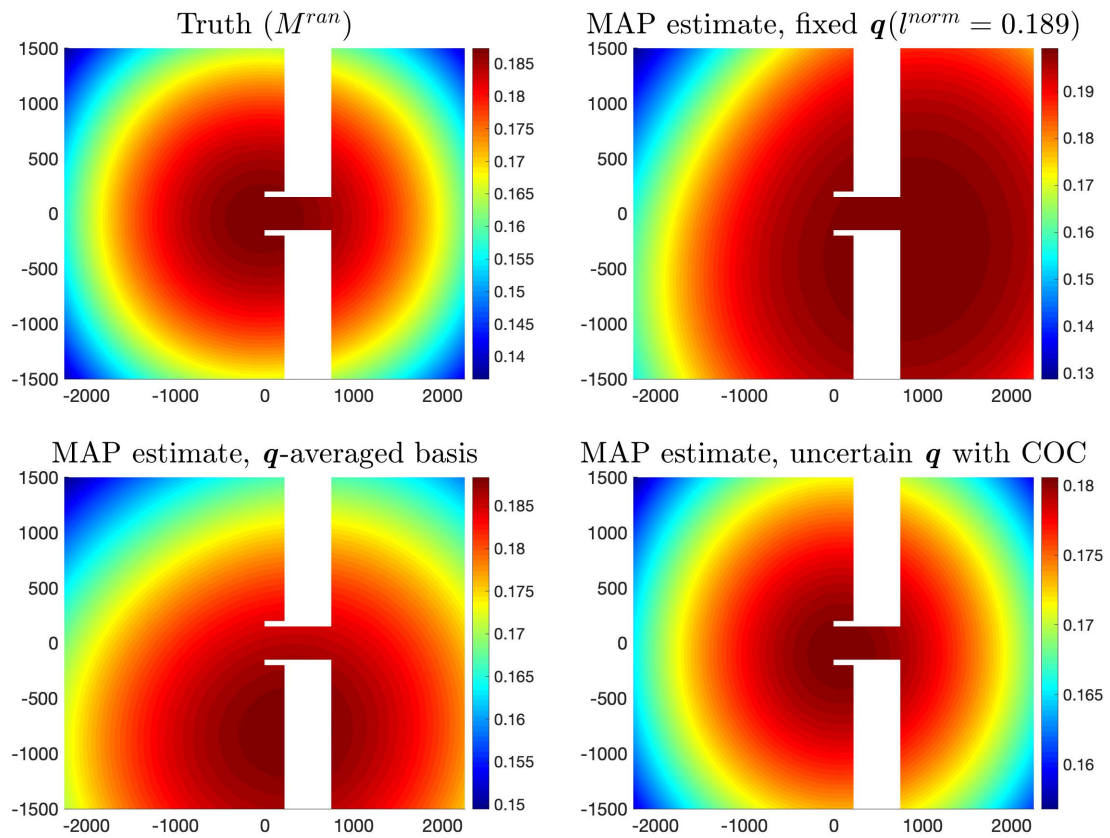


Figure 9: Spatial maps of several inferred Manning's n fields compared to the truth: (top-left) the true Manning's n field, (top-right) the inferred Manning's n field using a fixed correlation length scale, (bottom-left) the inferred Manning's n field based on the basis derived from $\bar{\mathbf{C}}$ without using COC, and (bottom-right) the inferred Manning's n field from using PC-MCMC with COC.

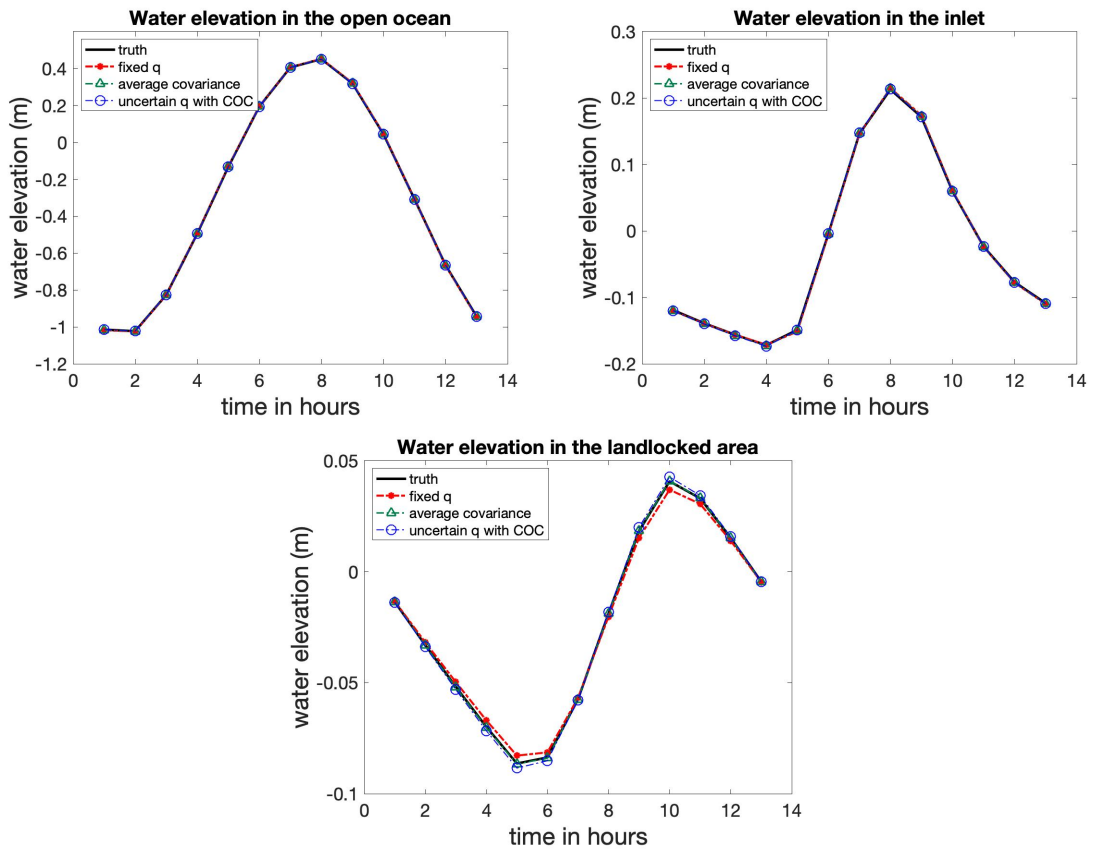


Figure 10: Time-evolution of the water elevation in the final 12 hours of the ADCIRC simulations subjected to different Manning's n field.

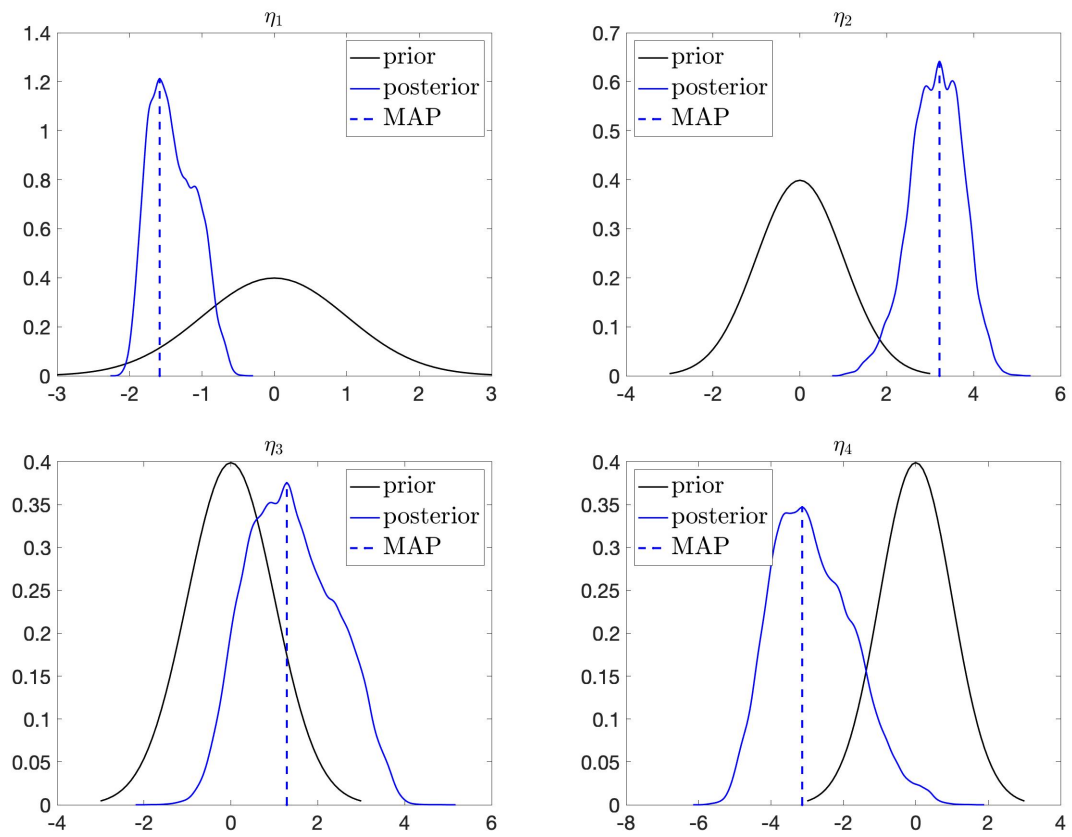


Figure 11: Comparison between the priors and marginal posteriors of the 4 KL coordinates for the inference of M^{ran} using \mathbf{q} -averaged covariance basis. The maximum a posteriori probability (MAP) estimate for each parameters are also presented using dashed blue line.

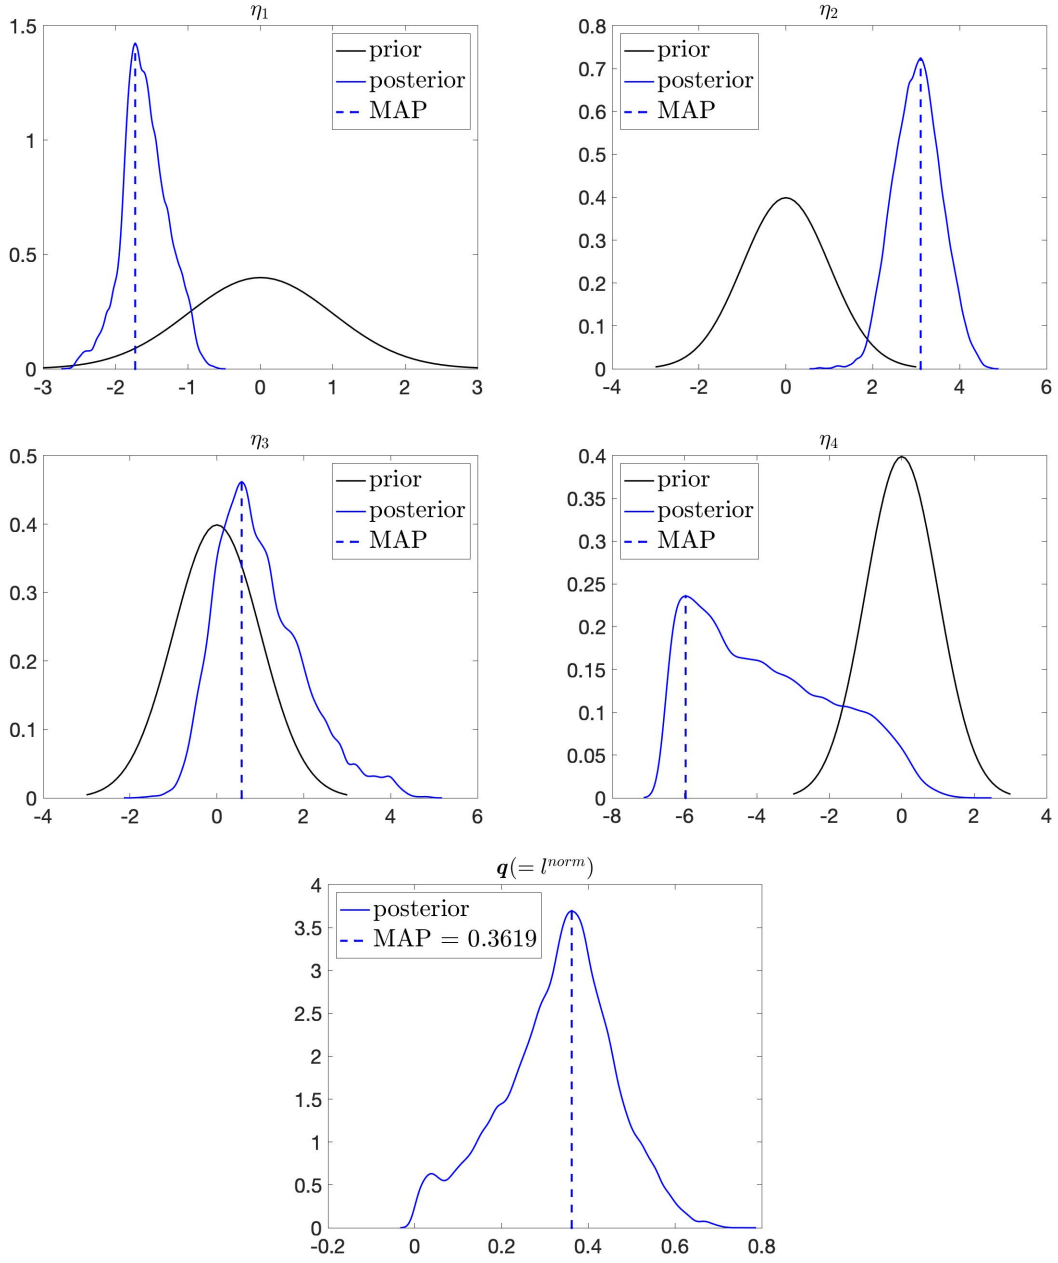


Figure 12: Comparison between the priors and marginal posteriors of the 4 KL coordinates and the length-scale hyper-parameter l (only the posterior is shown) for the inference of M^{ran} using the COC method. The maximum a posteriori probability (MAP) estimate for each parameters are also presented using dashed blue line.

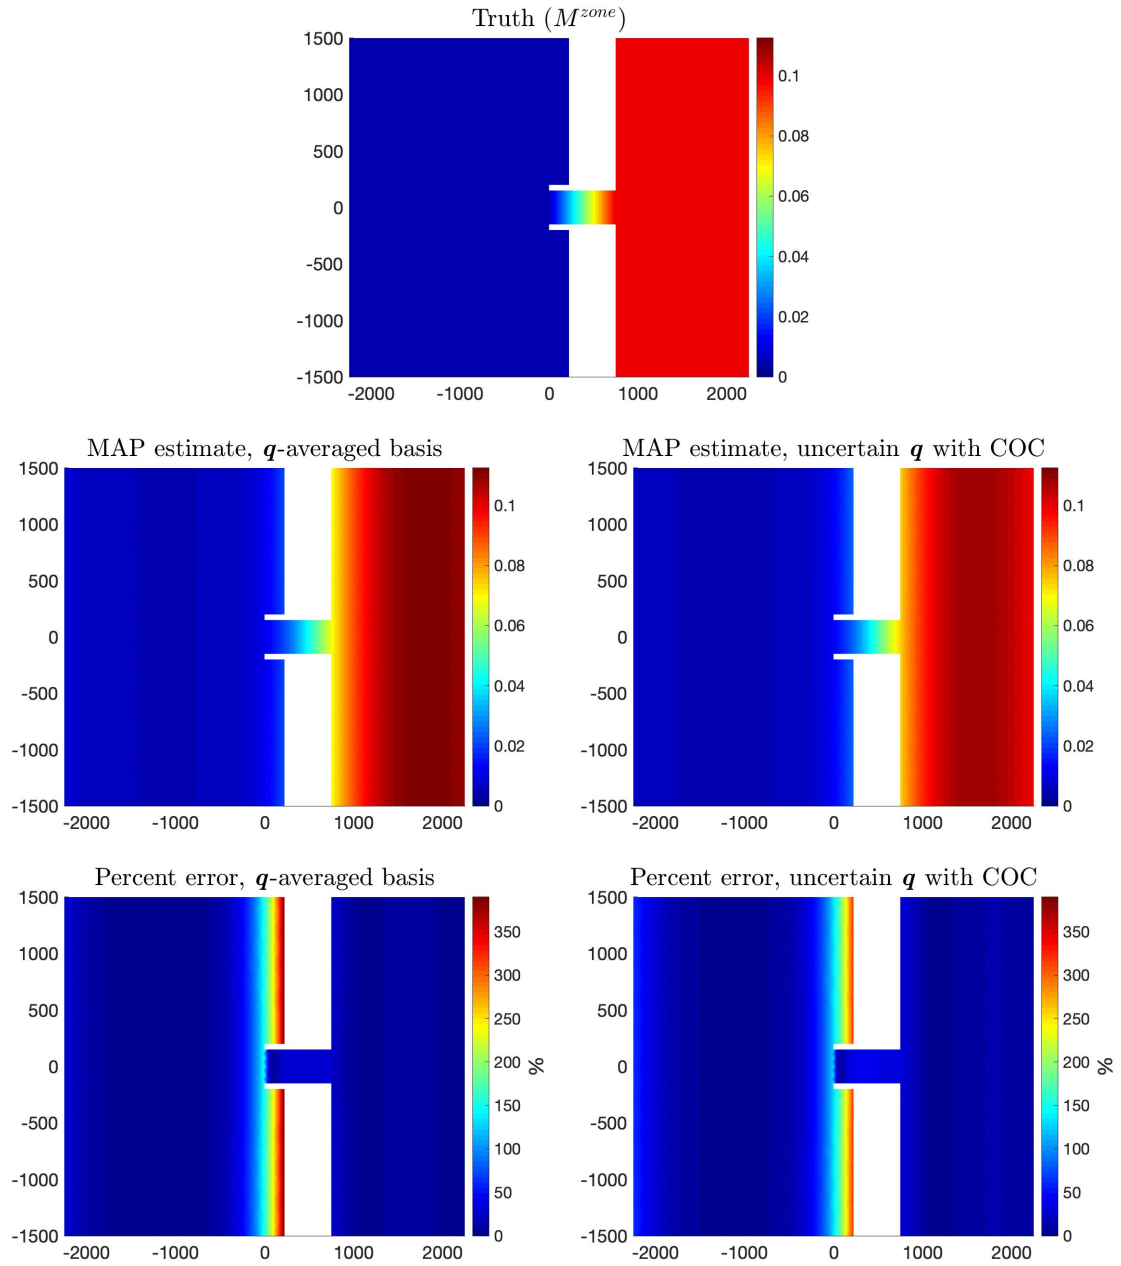


Figure 13: Spatial maps of several inferred Manning's n fields compared to the truth: (top) the true Manning's n field, (middle-left) the inferred Manning's n field based on the basis derived from $\bar{\mathcal{C}}$ without using COC, and (middle-right) the inferred Manning's n field from using PC-MCMC with COC. Bottom panel, spatial maps of the percent error with respect to the truth: (bottom-left) the inferred Manning's n field based on the basis derived from $\bar{\mathcal{C}}$ without using COC, and (bottom-right) the inferred Manning's n field from using PC-MCMC with COC.

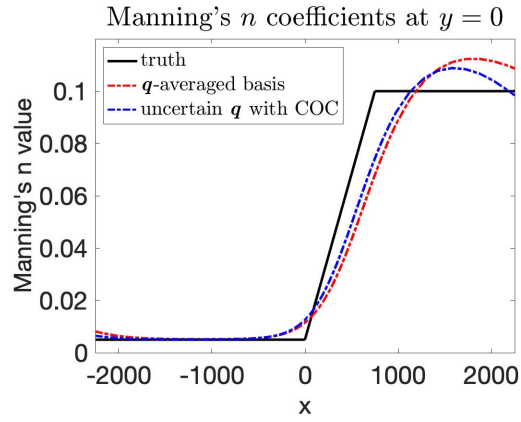


Figure 14: Cross section plot at $y = 0$ (midway between north and south) of the inferred Manning's n fields compared to the true field M^{zone} (black). The inferred field based on the basis derived from \bar{C} without using COC and with COC are plotted in red and blue, respectively.

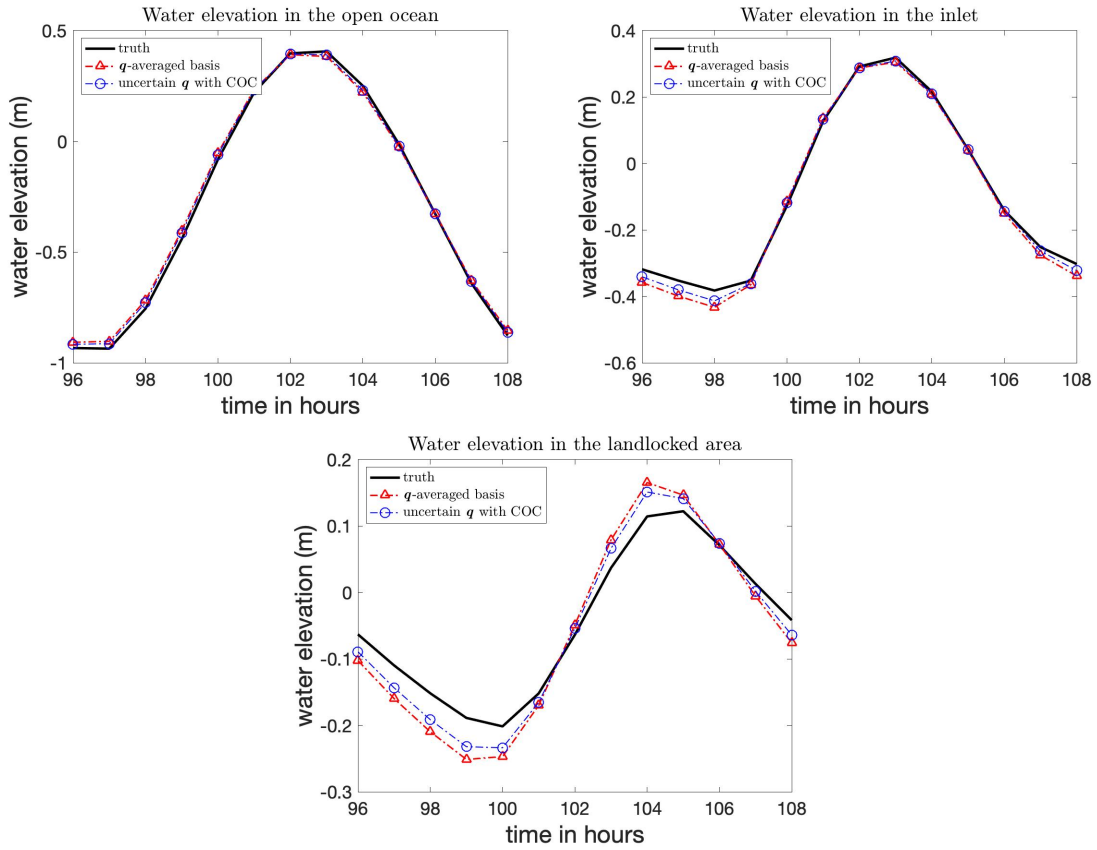


Figure 15: Time-evolution of the water elevation in the final 12 hours of the ADCIRC simulations subjected to different Manning's n field.