

A fully Bayesian sparse polynomial chaos expansion approach with joint priors on the coefficients and global selection of terms



Paul-Christian Bürkner^{a,*}, Ilja Kröker^{a,b}, Sergey Oladyskhin^{a,b},
Wolfgang Nowak^{a,b}

^a Stuttgart Center for Simulation Science, University of Stuttgart, Germany

^b Department of Stochastic Simulation and Safety Research for Hydrosystems, Institute for Modelling Hydraulic and Environmental Systems, University of Stuttgart, Germany

ARTICLE INFO

Article history:

Received 13 April 2022

Received in revised form 13 January 2023

Accepted 1 May 2023

Available online 15 May 2023

Keywords:

Polynomial chaos expansion

Surrogate modeling

Bayesian inference

Uncertainty quantification

Shrinkage priors

Variable selection

ABSTRACT

Polynomial chaos expansion (PCE) is a versatile tool widely used in uncertainty quantification and machine learning, but its successful application depends strongly on the accuracy and reliability of the resulting PCE-based response surface. High accuracy typically requires high polynomial degrees, demanding many training points especially in high-dimensional problems through the curse of dimensionality. So-called sparse PCE concepts work with a much smaller selection of basis polynomials compared to conventional PCE approaches and can overcome the curse of dimensionality very efficiently, but have to pay specific attention to their strategies of choosing training points. Furthermore, the approximation error resembles an uncertainty that most existing PCE-based methods do not estimate. In this study, we develop and evaluate a fully Bayesian approach to establish the PCE representation via joint shrinkage priors and Markov chain Monte Carlo. The suggested Bayesian PCE model directly aims to solve the two challenges named above: achieving a sparse PCE representation and estimating uncertainty of the PCE itself. The embedded Bayesian regularizing via the joint shrinkage prior allows using higher polynomial degrees for given training points due to its ability to handle underdetermined situations, where the number of considered PCE coefficients could be much larger than the number of available training points. We also explore multiple variable selection methods to construct sparse PCE expansions based on the established Bayesian representations, while globally selecting the most meaningful orthonormal polynomials given the available training data. We demonstrate the advantages of our Bayesian PCE and the corresponding sparsity-inducing methods on several benchmarks.

© 2023 Elsevier Inc. All rights reserved.

1. Introduction

The polynomial chaos expansion (PCE) was originally introduced through the theory of homogeneous chaos by Norbert Wiener [77] in 1938. Convergence in the L^2 -sense was shown by Cameron and Martin [68]. Since Ghanem and Spanos

* Corresponding author at: Universitätsstr. 32, 70569 Stuttgart, Germany.

E-mail address: paul.buerkner@gmail.com (P.-C. Bürkner).

[24] popularized PCE-based stochastic Galerkin techniques in applications, many authors have investigated the construction [44,81] and limitations of PCE [18,45] and applied PCE techniques to many application problems [31,36,53,67,79,82]. Most applications use the PCE as a response surface technique to provide surrogates for computationally expensive models, which can then be used for uncertainty quantification and related tasks. Beyond fundamental properties like fast evaluation and regularity of PCE surrogates, the PCE can provide stochastic moments of the approximated model response (like mean and variance) and global sensitivity indices [67] in analytical form, i.e. without further evaluations of the original model.

Despite these advantages, the feasibility and efficiency of PCE approximations tend to suffer because the number of expansion terms and the corresponding required training runs of the original model grow fast with increasing input dimension (curse of dimensionality). To extend the practical applicability of PCE to higher-dimensional problems, several later techniques seek sparse PCE approximations, where only selected expansion terms have non-zero coefficients. Available ideas include sparse quadrature and stochastic collocation [16,41,84] or stratified sampling methods [25,30]. In particular, there are sparse approaches based on classical regression [2,8] and Bayesian sparse PCE approaches [15,37,47,59]. To the best of our knowledge, the existing Bayesian sparse PCE approaches see the PCE's expansion coefficients as Gaussian random variables that offers the pathway for Bayesian sparse learning as for relevance vector machine [71] to obtain the posterior distribution of the PCE coefficients. Some real-world applications require regularization strategies on the PCE to obtain reliable surrogate models. Therefore, the PCE Ansatz was combined with several regularization strategies, beginning regularization via Moore–Penrose inverse [38,50] or ℓ_2 -norm regularized least-squares classical regularization in [19], up to Elastic Net [69] and Entropy-based [85] regularization concepts. Among these, the Elastic net idea demonstrates that regularization and sparsity can be addressed together, as Elastic net includes a sparsity-inducing ℓ^1 -term in its regularization.

In this paper, we focus on Bayesian PCE approaches as they can quantify the epistemic uncertainty of the PCE (remaining after training, not the uncertainty of model responses due to potentially uncertain parameters of the true model) through the implied posterior distribution of the PCE coefficients. However, to compute the posterior analytically, the corresponding prior distributions of the PCE coefficients have to be Gaussian, in which case the posterior is Gaussian as well [23] similarly to the existing Bayesian sparse PCEs [15,37,47,59]. This severely limits the flexibility of Bayesian PCEs. In particular, it is prone to overfitting (i.e., fitting too closely on the training data and thus generalizing poorly) in cases with few training points and correspondingly high number of considered PCE terms, as the implied ridge-regularization (purely ℓ^2 -based) is insufficient in such cases [54].

In the wider field of statistics, much more strongly regularizing (shrinkage) priors have been demonstrated to provide good approximation performance even in highly under-determined problems [7,52,73,87], but so far have not found their way into the PCE literature. Due to complex geometry of the implied posteriors, sampling techniques (in particular Markov chain Monte-Carlo (MCMC)) are required to achieve trustworthy approximations of the posterior [23,26]. Gaussian approximations are insufficient in these cases, as neither the marginal posterior distributions would be sufficiently Gaussian, nor would the joint distribution necessarily follow a Gaussian dependence structure [52,83,87].

In this study, we develop and evaluate a new Bayesian sparse PCE approach. Our key idea is to use a joint shrinkage prior on the PCE coefficients, specifically the so-called R2D2 prior [87] as summarized in Section 2.4. The proposed R2D2-based PCE model aims to achieve both a sparse PCE representation and accurate uncertainty estimation of the PCE itself. The embedded Bayesian regularization via the R2D2 prior allows using higher polynomial degrees for given training points. It can even handle under-determined situations, where the number of considered PCE coefficients is much larger than the number of available training points. Moreover, our fully Bayesian approach estimates uncertainty of the PCE itself via MCMC estimation of the posterior. And, by doing so, it avoids Gaussian assumptions on both the posterior of the coefficients and the implied response surface. On top of that, we also explore multiple variable selection methods to construct sparse PCE expansions with our Bayesian approach, globally selecting the most meaningful polynomials given the available training data (see Section 2.5).

We demonstrate the advantages of the R2D2-based Bayesian PCE and of the corresponding sparsity-inducing methods on several benchmarks in Section 3. In these benchmarks, we employ the arbitrary polynomial chaos (aPC) generalization of the conventional PCE. The aPC satisfies orthonormality in data-driven situations and so assures flexibility for various applications (see Section 2.1). The test cases range from widespread low- to high-dimensional benchmark functions up to a simulation-based benchmark in the area of carbon dioxide storage in the subsurface.

2. Methods

2.1. Polynomial chaos expansion

The Polynomial Chaos Expansion (PCE) [77] expresses the dependence of a model response (output) on a set of input parameters using a projection onto an orthonormal polynomial basis [24,68]. Formally, we will consider a univariate response $y = y(x, t)$ generated from some true – but potentially unknown – model $\mathcal{M}(x, t, \omega)$ that depends on space x , time t , and on a vector of uncertain input parameters $\omega = \{\omega_1, \dots, \omega_N\}$ from the input space Ω with a given probability measure Γ_ω . In practice, we often do not know the exact analytical form of $\mathcal{M}(x, t, \omega)$ or cannot evaluate it due to computational challenges. Hence, we seek to find a surrogate model that approximates the true model as closely as possible while having some favorable practical or computational properties. When using PCE as a surrogate [11,77], we approximate the true model as

$$\mathcal{M}(x, t, \boldsymbol{\omega}) \approx \sum_{i=0}^M c_i(x, t) \Psi_i(\boldsymbol{\omega}), \tag{1}$$

where $\Psi_i(\boldsymbol{\omega})$ are basis functions of a multivariate polynomial basis $\{\Psi_0(\boldsymbol{\omega}), \dots, \Psi_M(\boldsymbol{\omega})\}$ that is orthonormal in (Ω, Γ_ω) , and $c_i = c_i(x, t)$ are corresponding coefficients that determine the form of the expansion in Equation (1). If the model \mathcal{M} contains space or time coordinates (x, t) , then we consider the coefficients $c_i(x, t)$ as defined independently for each point in space and time. For easier notation, we will omit these coordinates from now on unless explicitly required.

In the widespread total-degree truncation, the total number M of expansion terms in Equation (1) depends combinatorically on the number of input parameters N and on the degree of expansion d (i.e., maximal degree of considered polynomials) as

$$M = \frac{(N + d)!}{N! d!}. \tag{2}$$

Assuming statistical independence of the input parameters, the corresponding multivariate polynomial basis $\{\Psi_0(\boldsymbol{\omega}), \dots, \Psi_M(\boldsymbol{\omega})\}$ is the tensor product of univariate orthonormal polynomials $\{\phi_j^{(0)}, \dots, \phi_j^{(d)}\}$ up to degree d for each input parameter ω_j :

$$\Psi_\alpha(\boldsymbol{\omega}) = \prod_{j=1}^N \phi_j^{(\alpha_j)}(\omega_j), \quad \sum_{j=1}^N \alpha_j \leq d, \tag{3}$$

where α is a multivariate index that contains the combinatoric information how to enumerate all possible products of individual univariate orthonormal polynomials $\{\phi_j^{(0)}, \dots, \phi_j^{(d)}\}$ for each input parameter ω_j .

This set of univariate polynomials forms an orthonormal basis of polynomial degree d for each input parameter ω_j ($j = 1, \dots, N$):

$$\int_{\Omega} \phi_j^{(p)}(\omega_j) \phi_j^{(q)}(\omega_j) d\Gamma(\omega_j) = \delta_{p,q}, \quad \forall p, q = 0, \dots, d \tag{4}$$

where $\delta_{p,q}$ represents the Kronecker delta.

There are several ways to construct the required univariate polynomials. The original PCE theory [77] was developed for normally distributed input parameters ω_j , leading to Hermite polynomials as basis. Later works ([81] and [82]) suggested the Askey scheme [4] of polynomials that extend the original PCE theory to several other common parametric distributions (Gamma, Beta, Uniform, etc.). Like Hermite polynomials are optimal for the normal distribution, Legendre polynomials are optimal for the uniform distribution of inputs parameter ω_j . Going beyond these distributions, univariate orthogonal polynomials can be also obtained using three-term recurrences [20], Gram–Schmidt orthogonalization [79], the Stieltjes procedure [66] or by a data-driven approach [44] know as arbitrary polynomial chaos (aPC). For a comprehensive discussion of the requirements for existence and completeness of orthonormal polynomial bases, we refer to [18,22,68].

Any of the above procedures can be used to construct the required univariate orthonormal polynomials. However, in many applications, exact knowledge of the involved probability distribution [68] is not available, or would require additional assumptions [55]. To assure highly parsimonious and, if necessary, purely data-driven representation, we will pay especial attention to the aPC generalization. The aPC technique adapts to arbitrary probability distribution shapes of input parameters. Its univariate orthonormal basis can be inferred from limited input data through the $2d$ raw stochastic moments [44] as follows:

$$\phi_j^{(\alpha_j)}(\omega_j) = \frac{1}{\sqrt{\kappa_{\alpha_j}}} \sum_{i=0}^{\alpha_j} m_i^{(\alpha_j)} \omega_j^i, \quad \alpha_j = 0, \dots, d, \tag{5}$$

where $\kappa_{\alpha_j} = m_{\alpha_j}^{(\alpha_j)}$ is a constant representing the norm of the univariate polynomial. The corresponding monomial coefficients $m_i^{(\alpha_j)}$ could be defined according to the raw moments of the inputs ω_j as follows [44,45]:

$$\begin{bmatrix} \mu_0(\omega_j) & \mu_1(\omega_j) & \dots & \mu_{\alpha_j}(\omega_j) \\ \vdots & \vdots & \vdots & \vdots \\ \mu_{\alpha_j-1}(\omega_j) & \mu_{\alpha_j}(\omega_j) & \dots & \mu_{2\alpha_j-1}(\omega_j) \\ 0 & 0 & \dots & 1 \end{bmatrix} \begin{bmatrix} m_0^{(\alpha_j)} \\ \vdots \\ m_{\alpha_j-1}^{(\alpha_j)} \\ m_{\alpha_j}^{(\alpha_j)} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}, \tag{6}$$

where $\mu_k(\omega_j)$ denotes the k -th raw stochastic moment of the input ω_j . The matrix on the left-hand side of Equation (6) is known as the Hankel matrix of moments [29], and its properties have been analyzed in [61]. The polynomials are real if, and

only if, the Hankel matrix of moments is positive definite, see also the related Hamburger moment problem, e.g. [3,18,60,68]. The aPC formulation offers a certain flexibility in modeling, as the required statistical moments in the Hankel matrix can be evaluated directly from an available data set of observed input parameter values, but they could also be determined analytically from suitable, parametric distributions. For example, Equation (5) reconstructs the Legendre polynomials when inserting the raw moments $\mu_k(\omega_j)$ of the uniform distribution.

Regardless of the specific PCE technique, each expansion coefficient c_i in Equation (1) indicates how much variance one or another expansion term brings into the overall composition. Therefore, the PCE is closely related to global sensitivity analysis [67,86]. Namely, due to the orthonormal representation, Sobol [63] sensitivity indices S_i for each individual expansion term is directly given by [43]:

$$S_i = \frac{c_i^2}{\sum_{i'=1}^M c_{i'}^2}, \tag{7}$$

where the Sobol index S_i reflects the relative partial contribution of each single PCE term to the total variance of the response y .

2.2. Conventional deterministic computation of expansion coefficients

Methods to compute the expansion coefficients c_i in Equation (1) can be sub-divided into intrusive and non-intrusive approaches. The intrusive approach (e.g. Galerkin projection) requires manipulation of the governing equations for stochastic analysis in application cases where (partial) differential equations are analyzed [24,31]. Due to the necessity of symbolic manipulations, this procedure may become analytically cumbersome and transfer from one application to another requires additional efforts. For that reason, non-intrusive approaches like sparse quadrature [30], the probabilistic collocation method [35,41,75] or numerical regression [5,58] have received increasing attention for applied tasks, as they do not demand any modification of original model/simulation codes. For the sake of flexibility, we will choose non-intrusive approaches in our current work.

The non-intrusive approaches employ so-called collocation points that could be seen as training points corresponding to sets of parameter values ω for which the original model is run. Let us denote the training input parameters by $\Omega_T = \{\omega_1, \dots, \omega_T\}$ and the corresponding training responses by $\mathbf{y}_T = \{y_1, \dots, y_T\}$, where T is the number of the training points. It is well known [44] that the selection of training points strongly influences the performance of the overall PCE representation in Equation (1). According to Gaussian integration theory [75], the optimal training points are the roots of the orthonormal polynomial one degree higher than the overall expansion degree d . The Gaussian integration strategy avoids additional oscillations known as the Runge phenomenon [57]. However, our current work is not limited to any particular selection strategy of the training points and, hence, we will keep the formulation as general as possible.

Regardless of their design, the chosen training inputs Ω_T and the corresponding training responses \mathbf{y}_T can be inserted into Equation (1) to obtain the following linear system of equations that determines the expansion coefficients c_i :

$$\begin{bmatrix} \Psi_0(\omega_1) & \dots & \Psi_M(\omega_1) \\ \vdots & \vdots & \vdots \\ \Psi_0(\omega_T) & \dots & \Psi_M(\omega_T) \end{bmatrix} \begin{bmatrix} c_1 \\ \vdots \\ c_M \end{bmatrix} = \begin{bmatrix} y_1 \\ \vdots \\ y_T \end{bmatrix}, \tag{8}$$

or in the following matrix form:

$$\Psi \cdot \mathbf{c} = \mathbf{y}_T, \tag{9}$$

where Ψ is the co-called design matrix of size $T \times M$ and \mathbf{c} is the vector of size M containing all unknown expansion coefficients.

There are several deterministic ways to obtain the expansion coefficients c_i from the linear system (9). If the number M of expansion terms is exactly equal to the number T of training points, the system (9) is fully determined, e.g., as in the probabilistic collocation method [35,41,75]. Therefore, the expansion coefficients could be found by:

$$\mathbf{c} = \Psi^{-1} \cdot \mathbf{y}_T, \tag{10}$$

When the number T of training points is higher than the number M of expansion terms, the linear system (9) is over-determined. Then, the expansion coefficients can be obtained by least-squares minimization [25,39] via Moore-Penrose inversion [38,50] (also known as pseudoinverse) on the T regression points:

$$\mathbf{c} = (\Psi^T \cdot \Psi)^{-1} \cdot \Psi^T \cdot \mathbf{y}_T. \tag{11}$$

However, since the training points do not necessarily correspond to optimal Gaussian roots [75] [5,42], additional regularization [70] helps to mitigate the Runge phenomenon [57]. With Gaussian regularization, Equation (9) becomes

$$\mathbf{c} = (\mathbf{\Psi}^T \cdot \mathbf{\Psi} + \lambda \mathbf{I})^{-1} \cdot \mathbf{\Psi}^T \cdot \mathbf{y}_T, \quad (12)$$

where λ is the regularization parameter and \mathbf{I} is the identity matrix of appropriate dimension. Equation (12) is a generalized formulation of Equation (11), and returns to the original least-squares formulation at $\lambda = 0$. It also contains the fully-determined problem in Equation (10) for $\lambda = 0$ and $M = T$. Moreover, Equation (12) also works when the original linear system (9) is under-determined, i.e., when $T < M$. For these reasons, Equation (12) has been used in practice [46,48,58], relying on a publicly available aPC code implementation [40]. Alternative ideas to regularize the deterministic PCE coefficients could be used following the ideas in [19,85].

2.3. Bayesian PCE

PCE can also be performed in a fully probabilistic (Bayesian) framework, by treating it as a Bayesian linear regression model with the evaluated polynomials as predictors [59]. For this purpose, we assume the training responses \mathbf{y}_T to be normally distributed with mean vector μ equal to the PCE approximation and residual variance σ^2 . Together, this forms the likelihood $p(\mathbf{y}_T | \mathbf{c}, \sigma^2)$ of the data \mathbf{y}_T , given the PCE coefficients \mathbf{c} and residual variance σ^2 as model parameters. The latter parameter captures both potential noise in the observed responses and approximation errors caused by the truncation of the PCE to a finite number of M polynomials. When additionally considering prior distributions for all the parameters, the following Bayesian PCE model arises:

$$\mathbf{y}_T \sim \text{normal}(\mu, \sigma^2) \quad (13)$$

$$\mu = \sum_{i=0}^M c_i \Psi_i(\boldsymbol{\omega}) \quad (14)$$

$$\mathbf{c} \sim p(\mathbf{c}) \quad (15)$$

$$\sigma^2 \sim p(\sigma^2), \quad (16)$$

which implies the posterior distribution of the PCE parameters given the data:

$$p(\mathbf{c}, \sigma^2 | \mathbf{y}_T) = \frac{p(\mathbf{y}_T | \mathbf{c}, \sigma^2) p(\mathbf{c}) p(\sigma^2)}{\int p(\mathbf{y}_T | \mathbf{c}, \sigma^2) p(\mathbf{c}) p(\sigma^2) d\mathbf{c} d\sigma^2} \quad (17)$$

The specific properties of the posterior and hence of the Bayesian PCE solution will depend on the choice of priors $p(\mathbf{c})$ and $p(\sigma^2)$. For example, if all the priors were (improper) flat, $p(\mathbf{c}) \propto 1$ and $p(\sigma^2) \propto 1$, the posterior mode would be equal to the maximum likelihood estimate, which for the here considered PCE, is also equal to the least-squares solution (10). If multivariate normal priors are used as $p(\mathbf{c})$ and inverse-gamma priors are used as $p(\sigma^2)$, the priors are conjugate to the likelihood [23]. This means that the posterior will be of the same functional form as the prior, with multivariate normal posterior for \mathbf{c} (conditionally on σ^2) and inverse-gamma posterior for σ^2 . If, for example, the prior mean and covariance matrix for \mathbf{c} were $\mu_{\text{prior}} = 0$ and $\Sigma_{\text{prior}} = \frac{1}{\lambda} \mathbf{I}$, respectively, with regularization parameter λ and identity matrix \mathbf{I} , the posterior mean would be equal to the regularized PCE estimate (12). As such, there exist natural connections between classical and Bayesian PCE despite conceptually different approaches by which they are derived.

The posterior forms the main target of Bayesian inference but is, with the exception of conjugate cases, usually analytically intractable due to the intractability of the integral that occurs as denominator in (17). In the current work we will go beyond the Gaussian assumption behind the prior and posterior of PCE coefficients that have been used in previous versions of Bayesian PCE approaches [15,37,47,59]. For general cases, we will have to resort to posterior approximation methods that only require evaluating the likelihood and prior. Out of all such approximation methods, MCMC is perhaps the most important class [23]. MCMC enables us to draw dependent random samples from the posterior, which can subsequently be used to compute relevant posterior quantities such as posterior means or variances. Their approximation errors approach zero for increasing number of samples under very general conditions, but for small number of samples or poor convergence, the approximation errors may be substantial [23,74]. As such, in practice, convergence of the MCMC procedure and sufficient sampling efficiency is verified by both graphical and numerical diagnostics before estimates are computed [21,74].

2.4. The R2D2 prior

We are specifically interested in scenarios with sparse training data (i.e., small number T of training points), where we need to find sparse PCE solutions that include as few polynomials as possible without sacrificing too much approximation performance. Sparsity does not always imply that PCE coefficients of irrelevant polynomials are found to be exactly zero. Instead, approximately zero (*weakly sparse*) will usually be good enough and is easier to achieve, since the underlying parameter space remains continuous rather than reducing to a point mass [52]. Priors that induce weak sparsity by strong regularization of irrelevant (or little relevant) terms are known as *shrinkage priors* [52,73]. We will focus specifically on one of these priors, the R^2 -Dirichlet-Decomposition (R2D2) prior [87], whose assumptions match very well with the needs of

sparse PCE. The R2D2 prior belongs to the class of global-local shrinkage priors, which has been shown to scale well to high-dimensional linear models even for sparse training data [1,7]. In particular, such priors can achieve strong regularization of irrelevant coefficients (shrinkage towards zero, i.e. weak sparsity) while providing only mild regularization of truly non-zero coefficients, even if the corresponding predictor variables are highly correlated [7,52]. Below, we will introduce the R2D2 prior in more detail.

The coefficient of determination R^2 expresses the proportion of variance explained by a model (here: PCE) in relation to the total variance of the response σ_{total}^2 . For Gaussian linear models with residual variance σ^2 , we define

$$R^2 = 1 - \frac{\sigma^2}{\sigma_{\text{total}}^2}, \quad (18)$$

which can be rewritten as

$$R^2 = \frac{\text{var}(\mu)}{\text{var}(\mu) + \sigma^2}, \quad (19)$$

with linear predictor μ as defined in (14). Since the PCE polynomials are by construction orthonormal, we can rewrite the variance of μ as

$$\text{var}(\mu) = \sum_{i=1}^M c_i^2. \quad (20)$$

Now, we exploit this variance decomposition to impose an intuitive joint prior on the PCE coefficients, as any prior on R^2 will imply a joint regularization of all the PCE coefficients c_i (excluding the zero-degree coefficient c_0).

To set up the R2D2 prior for PCE, we first assign normal priors to the coefficients as $c_i \sim \text{normal}(0, \lambda_i^2)$ and then decompose the prior variance λ_i^2 of the i th coefficient as $\lambda_i^2 = \phi_i \tau^2 \sigma^2$. Here, ϕ_i denotes the proportion of explained variance assigned to the i th component, with $\phi = (\phi_1, \dots, \phi_M)$ forming a simplex, that is, $\phi_i \geq 0$ and $\sum_{i=1}^M \phi_i = 1$, and $\tau^2 = \frac{R^2}{1-R^2}$ is the variance explained by the PCE. Assigning a Beta prior on R^2 with mean parameter ζ and precision parameter ν and a Dirichlet prior on ϕ with concentration parameter vector θ , we obtain the following complete prior specification [87]:

$$c_i \sim \text{normal}\left(0, \phi_i \tau^2 \sigma^2\right) \quad \forall i = 1, \dots, M \quad (21)$$

$$\tau^2 = \frac{R^2}{1-R^2} \quad (22)$$

$$R^2 \sim \text{Beta}(\zeta, \nu) \quad (23)$$

$$\phi \sim \text{Dirichlet}(\theta) \quad (24)$$

$$c_0 \sim p(c_0) \quad (25)$$

$$\sigma^2 \sim p(\sigma^2) \quad (26)$$

Different choices of ζ and ν imply different a priori assumptions on R^2 and subsequently on τ (which is BetaPrime distributed) as illustrated in Fig. 1. Choosing a uniform prior on R^2 via $(\zeta, \nu) = (0.5, 2)$ usually implies sufficient global regularization. However, if the signal to noise ratio is very large (e.g., because training data was generated without noise), larger prior R^2 values (as realized via larger ζ) may be favorable as illustrated in Section 3.4.

The Dirichlet concentration vector θ controls the sparsity of the solution where smaller concentration values imply stronger sparsity [1,87]. Here, we set $\theta = (1, \dots, 1)$ implying a uniform prior over all simplexes of the corresponding dimension. Smaller values (e.g., $\theta = (1/2, \dots, 1/2)$) can be used as well if stronger sparsity is desired, but since we enforce exact sparsity anyway in a second step (see below), we did not further explore this option. Because the prior on R^2 controls the global regularization while the prior on θ controls the local regularization (sparsity), both priors can be chosen independently of each other. For completeness, densities and a few mathematical properties of the Beta and Dirichlet distributions are provided in Appendix A.

The priors on the zero-degree coefficient c_0 and on the residual variance σ^2 can be assigned separately. In our case studies, we will specify $p(c_0)$ as a normal prior with mean and standard deviation equal to the empirical mean and standard deviation of the training responses, respectively. This represents a weakly-informative prior also chosen as default in the R package brms [9]. The prior $p(\sigma^2)$ can be of any common distribution for continuous, non-negative random variables (e.g., an inverse-gamma, half-normal, or half-Cauchy prior). In our case studies, we use a half-normal prior on σ , with scale parameter appropriately chosen to take both the scale of observed responses and the prior on R^2 into account, since $\sigma^2 = \sigma_{\text{total}}^2 (1 - R^2)$.

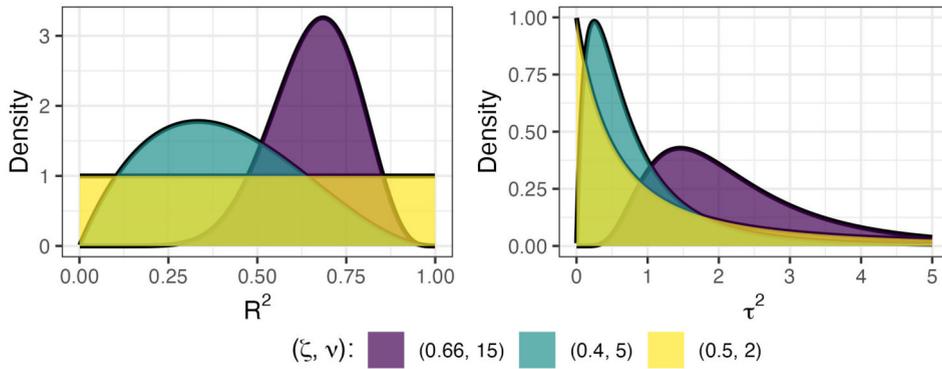


Fig. 1. Exemplary densities of the Beta prior for R^2 (left) and corresponding BetaPrime prior for $\tau^2 = R^2/(1 - R^2)$ (right) with varying mean and precision parameters (ζ, ν) . (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

2.5. Statistical methods for variable selection

Imposing the R2D2 prior on the PCE coefficients implies a PCE solution that is only weakly sparse in the sense that coefficients identified as irrelevant by the Bayesian PCE model will only be shrunken to values very close to zero, not exactly to zero. Thus, the R2D2 prior alone (or other continuous shrinkage priors) do not result in exactly sparse solutions. Instead, to achieve the latter, variable selection methods are required [51,52]. In general, exact sparsity may be favorable for two primary reasons: (1) If obtaining all the predicting variables comes with large costs of some kind, for example, monetary, ethical, computational, or time-wise; and (2) if the full reference model is not identified or yields bad predictions due to overfitting [51]. As evaluation of the PCE polynomials for given input parameters is straightforward, the first motivation does not really apply to PCE – unless when space or time coordinates at very fine resolutions result in a parallel processing of many coordinate-wise PCE surrogates. What is more, if the prior-induced shrinkage of Bayesian PCE was just right, that is, strong shrinkage was applied to all irrelevant coefficients and comparably little shrinkage to all relevant coefficients, the second motivation would not apply either. However, as the complexity of the problem grows, either by decreasing the number T of training points or by increasing the number M of considered polynomials, the prior-implied shrinkage may become far from optimal, thus either shrinking too much or too little [52,87]. In this case, applying a variable selection method in a second step after initially obtaining the full, weakly sparse PCE solution (up to a given degree d) may be beneficial to further improve approximation performance of the PCE.

Given any PCE solution, a natural variable selection method is to select the M_{sel} out of all M polynomials with the highest expected (posterior averaged) Sobol indices (7). This is a greedy selection approach in the sense that it does not account for potential dependencies between the polynomials. By construction of the PCE polynomials, they are mutually independent in the limiting case of infinite training data (see Section 2.1). However, in the practical case of only finite (potentially very small) training data, the polynomials may not be empirically independent for two reasons. First, small dependencies may occur because the chosen finite training points do not perfectly resemble their limiting distribution. Second, and perhaps more severe, for T training points only a maximum $T - 1$ polynomials can ever be linearly independent, and so any number of $M \geq T$ polynomials will be necessarily dependent. As we are primarily interested in sparse data scenarios where $M \gg T$, this can be a serious limitation for any greedy variable selection method, as we will also demonstrate in our case studies.

As a fully Bayesian non-greedy variable selection method, we consider the projective prediction (projpred) approach [51], which has been shown to provide superior results compared to alternative procedures [13,14,49,51]. In short, projpred aims to find a minimal subset of variables with comparable or better predictive performance than the full reference model (i.e., the PCE solution based on all considered M polynomials) while taking into account the epistemic uncertainty encoded in the posterior distribution of the full model. It does so by aiming to minimize the KL-divergence between the reference model's posterior predictive distribution and the predictive distribution of the sparse sub-model; a problem whose solution is usually intractable. However, for exponential-family models (including PCE expressed as a linear regression model), the KL-minimizing solution can be approximated via maximum-likelihood estimation. This estimation uses the samples from the posterior predictive distribution of the reference model as responses, instead of the actually observed responses [51]. Within the projpred framework, a lot of choices can be made by the user, for example, about how to cluster posterior predictive samples of the reference model, how to define the search space of relevant sub-models to investigate, and how to decide when the minimally sufficient subset of variables has been reached [51].

Here, we consider only a single-cluster solution (all posterior samples projected at once) in combination with a Lasso-type ℓ^1 -regularized model to construct the variable search path [72]. Compared to other alternatives in projpred, this is a computationally less demanding option that still yields robust variable selection, although the number of selected variables may be slightly higher than minimally sufficient [51]. For comparison with greedy Sobol index search and to further reduce computational cost, we fix the number M_{sel} of selected polynomials in the sub-models a priori instead of applying more principled but also more expensive decision rules, for example, based on cross-validation. Since we can again regularize

the chosen sparse sub-models with an R2D2 prior, selecting a few more terms than minimally sufficient does not do much harm as long as we also select all the relevant polynomials.

2.6. Fully Bayesian sparse aPC with R2D2 priors

Combining the methods introduced above, we propose a fully Bayesian sparse PCE. It uses the aPC for constructing orthonormal polynomials (see Section 2.2) and the R2D2 prior to provide joint shrinkage to the PCE coefficients (see Section 2.4). Bayesian sparse PCE models are specified with the R package `brms` [9,10] that internally writes highly optimized code in the probabilistic programming language Stan [12,65]. Bayesian estimation of the posterior distribution is conducted by MCMC (see Section 2.3), more precisely, by the adaptive, recursive Hamiltonian-Monte-Carlo (HMC) algorithm implemented in Stan [65]. This flavor of MCMC exploits the fact that the (Bayesian) PCE and its posterior distribution in Equation (17) is differentiable with respect to all PCE parameters. Therefore, it scales well to high-dimensional posteriors [6,26].

In addition to the full PCE models that include all the PCE coefficients up to a given maximal degree, we obtain sparse sub-models either using greedy selection via ordering of the Sobol coefficients or using the `projpred` method (see Section 2.5).

3. Experiments and results

In the following subsections, we test and discuss our sparse Bayesian PCE model with R2D2 prior and variable selection on a total of four test cases. These are the Ishigami function, the Sobol function, the signum function and a CO₂ benchmark case. All these test cases have been used for testing purposes of PCE concepts and other surrogate approximation techniques. The specific reasoning for the selection of each test case (why it is hard and/or useful, and why it serves specifically to test a sparse form of PCE) is provided in the respective subsections. All code and data required to reproduce the case studies can be found on GitHub (<https://github.com/paul-buerkner/Bayesian-sparse-PCE>).

Throughout the case studies, we consider the following three error metrics to evaluate the PCE approximation's performance:

- Error of the PCE *mean estimate* $\mu(\hat{y}) - \mu(y)$, where $\mu(y)$ is the true mean of the target model \mathcal{M} under the given true input measure Γ_ω and $\mu(\hat{y})$ is the corresponding PCE approximation. In some cases, we consider variations of this metric, for example, only showing the absolute error $|\mu(\hat{y}) - \mu(y)|$.
- Error of the PCE *standard deviation (SD) estimate* $\sigma(\hat{y}) - \sigma(y)$, where $\sigma(y)$ is the true SD of the target model \mathcal{M} under the given true input measure Γ_ω and $\sigma(\hat{y})$ is the corresponding PCE approximation. Like for the mean, we sometimes consider variations of this metric.
- The *out-of-sample RMSE* across a sufficiently large number T_{test} of randomly sampled test points (new ω sampled from the true input measure Γ_ω) not previously seen by the PCE:

$$\text{RMSE} = \frac{1}{S} \sum_{s=1}^S \sqrt{\frac{1}{T_{\text{test}}} \sum_{k=1}^{T_{\text{test}}} (y_k - \hat{y}_k^{(s)})^2}, \quad (27)$$

where y_k is the k -th response in the test set and $\hat{y}_k^{(s)}$ is the s th posterior sample of the corresponding PCE approximation out of a total of S posterior samples.

- Additionally, as MCMC sampling can be time-intensive, we also report the Bayesian PCE estimation times (in minutes).

3.1. Ishigami function

As first benchmark, we consider the Ishigami function [27]:

$$y = \sin(\omega_1) + a \sin(\omega_2)^2 + b \omega_3^4 \sin(\omega_1), \quad (28)$$

with $N = 3$ input variables $\omega_1, \omega_2, \omega_3$ being independently uniformly distributed in $[-\pi, \pi]$, and with control parameters a, b being set to $a = 7$ and $b = 0.1$. In this configuration, the Ishigami function has already been applied to benchmark other PCE approaches [67]. It serves as a good initial benchmark because its input dimensionality allows for illustrative visualization and it can be nicely approximated by PCE due to the function's smoothness properties. Also, it is a sparse function in the sense that ω_2 has no cross-terms with ω_1 or ω_3 , which produces a testable sparsity pattern for our method. What is more, statistics such as mean and variance are analytically known [27].

For training, we consider training data sets with $T = (10, 25, 50, 100, 200, 286, 400, 800)$ training points. The training points are constructed as the first T values of the 3-dimensional Sobol Sequence. We choose the Sobol sequence as it provides samples more evenly distributed in multidimensional hypercubes than random uniform samples [62].

On these training points, we evaluate the Ishigami function as target function for approximation. For simplicity, we linearly scale the prior, the Ishigami function, and the training points to the standard scaling of the Legendre polynomials

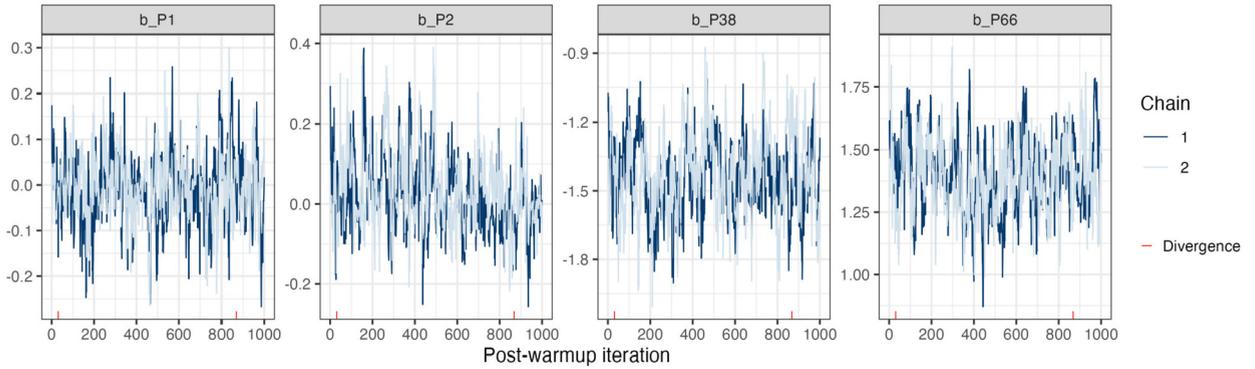


Fig. 2. Trace plot of four selected regression coefficients (two irrelevant and two relevant) for the full reference model on the Ishigami function. This shows good convergence of the reference model. The y-axis depicts the corresponding coefficient values.

(i.e., $[-1, 1]$). As PCE basis, we consider all 3-dimensional Legendre polynomials up to a degree of $d = 10$, which implies $M = 286$ polynomials in total. These values of T and M were chosen to create a varying range of inference complexity, ranging from strongly under-determined scenarios with $M \gg T$, exactly determined scenarios with $M = T$, to relatively simple over-determined scenarios with $M < T$. We also perform additional experiments with noisy training data, which can be found in Appendix B.

As full reference PCE representation for comparison with the later-computed sparse PCEs, we used all M polynomials as PCE basis terms. We equip the corresponding PCE coefficients with an R2D2 prior on the $M - 1$ non-constant polynomials as described in Section 2.6. Specifically, we used an R2D2 prior with $R^2 \sim \text{Beta}(\zeta = 0.5, \nu = 2)$ that implies a-priori uniformity of R^2 in $[0, 1]$ and provides some overall regularization while still allowing high R^2 to be possible.

Based on this full PCE Ansatz, we construct one PCE representation for each training set of size T . Then, for each full PCE representation, we fitted two sparse sub-PCEs each that merely differ in sparse selection method. The first uses the ordered Sobol indices [67] and the second uses projective prediction (projpred) for variable selection [51]. For each selection method, we extracted the $M_{\text{sel}} = 25$ most important polynomial terms out of all non-constant ones plus the constant, 0 th-degree polynomial. Only the selected polynomials were included in the respective sparse PCEs. For these selected polynomials and their corresponding coefficients, we again used a uniform R2D2 prior with $R^2 \sim \text{Beta}(\zeta = 0.5, \nu = 2)$. This still provides some regularization on the PCE coefficients, even for the already sparse models. Even though this is the same prior on R^2 as for the reference models, the implied regularization on each of the sparse models' coefficients is comparably milder because the explained variance is distributed among much fewer coefficients.

To fit these PCE models to the training data sets, we produced posterior samples using MCMC as explained in the previous section. Two independent Markov Chains were run each for 2000 iterations. From these, the first 1000 were discarded as warmup, leading to a total of 2000 post-warmups samples used for inference. All models converged well as indicated by standard algorithm-specific [6] and algorithm-agnostic [74] convergence diagnostics. Trace plots of the PCE coefficients illustrate excellent mixing of the MCMC chains also graphically (see Fig. 2 for trace plots of four selected PCE coefficients of the reference model based on $T = 100$ training points).

Posterior histograms of the reference model's PCE coefficients demonstrate the characteristic shrinkage properties of the R2D2 prior as the posteriors of irrelevant coefficients are strongly centered around zero with heavier-than-Gaussian tails (see the first two diagonal elements of Fig. 3). What is more, the bivariate pairs plots in Fig. 3 show approximate pairwise independence of the PCE coefficients' posteriors. This is caused by the independence of the corresponding Legendre polynomials.

From Fig. 4, we see that the full reference model is mostly competitive with the sparse sub-models in terms of mean and SD bias. However, it has uniformly worse out-of-sample RMSE especially for small T . This demonstrates that the R2D2-implied regularization alone (i.e., before selection of sparse subsets) may be insufficient to achieve satisfactory function approximation in $M \gg T$ scenarios, at least under the conditions considered here. In contrast, both sparse sub-models perform much better. Specifically the sparse projpred model achieves very good out-of-sample RMSE, while only slightly underestimating the true Ishigami SD already for $T = 50$ training points (again see Fig. 4).

For the sparse projpred model based on $T = 100$ training points, Sobol and total Sobol indices are shown in Fig. 5. As can be seen, only 10 non-constant polynomials (plus the constant polynomial) seem to be sufficient to achieve an almost perfect approximation. Table 1 shows more details on the corresponding polynomial degrees), and one can see that, as ideally expected by knowing the true underlying Ishigami function, no cross-terms with ω_2 occur. The impression of a good approximation is confirmed in Fig. 6, showing conditional predictions of the sparse projpred model against the corresponding true Ishigami function values. As T increases further, the Ishigami mean and SD estimates do not improve anymore compared to the $T = 100$ sparse projpred model, presumably because the remaining approximation error is due to polynomials of degree higher than $d = 10$, which were not included in the PCE approximation here.

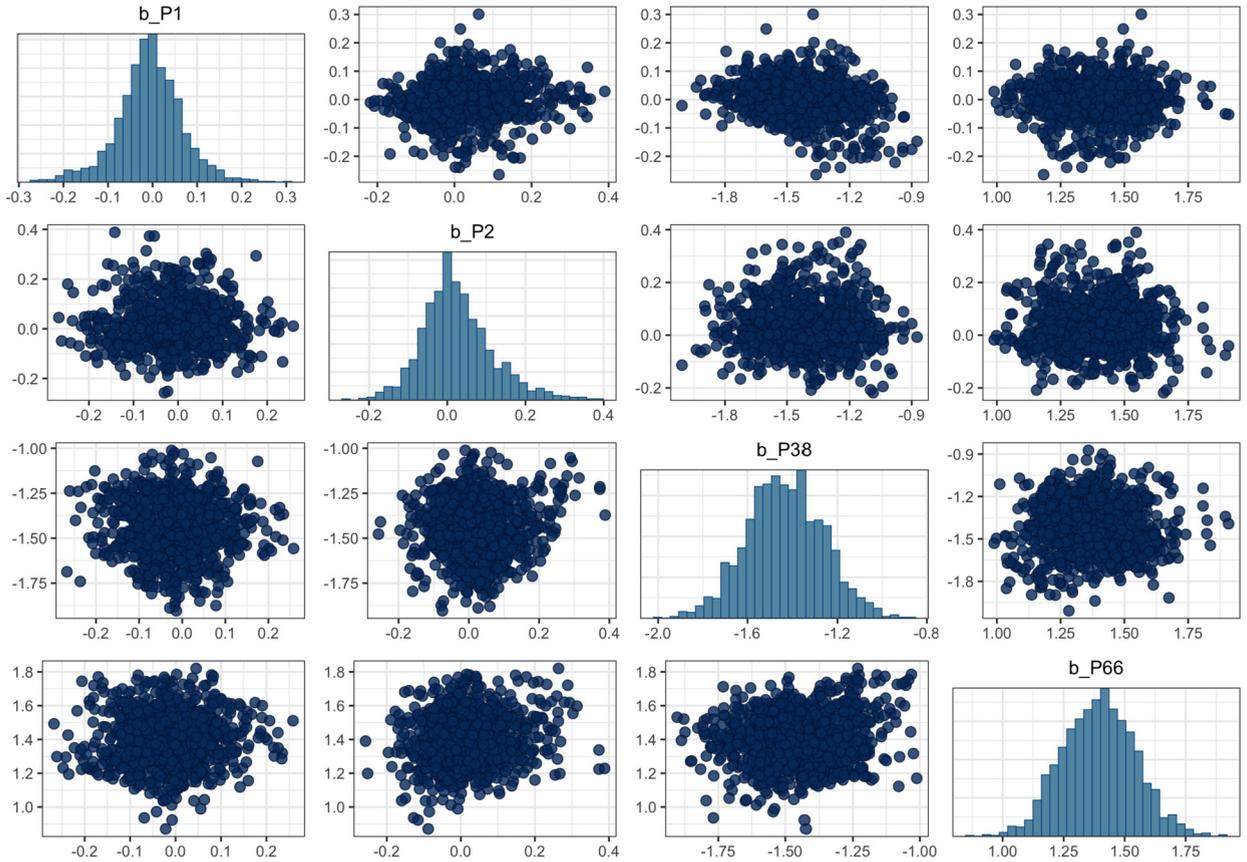


Fig. 3. Pairs plot of four selected regression coefficients (two irrelevant and two relevant) for the full reference model on the Ishigami function based on $T = 100$ training points. This illustrates (a) approximate pairwise independence of the coefficients even for $M > T$ and (b) strong shrinkage for irrelevant coefficients. In the diagonal plots, the x-axis depicts the coefficient values and the y-axis their binned probabilities. In the off-diagonal plots, the axes depict coefficients values of the corresponding column and row, respectively.

Table 1

Posterior mean Sobol indices (Mean), 95% posterior credible intervals (95%-CI), and polynomial degrees $d(\omega_m)$ by input variable of the 10 most important non-constant polynomials as selected by the sparse projpred model for the Ishigami function based on $T = 100$ training points.

Mean	95%-CI	$d(\omega_1)$	$d(\omega_2)$	$d(\omega_3)$
3.73	[3.73, 3.73]	0	4	0
2.59	[2.58, 2.59]	1	0	0
1.85	[1.84, 1.85]	1	0	2
1.81	[1.81, 1.81]	0	6	0
1.63	[1.63, 1.64]	3	0	0
1.16	[1.16, 1.17]	3	0	2
0.35	[0.34, 0.35]	0	2	0
0.16	[0.16, 0.16]	1	0	4
0.11	[0.11, 0.11]	0	8	0
0.10	[0.10, 0.10]	3	0	4

Together, these results demonstrate that, with our proposed methods, performing global polynomial search for sparse selection of PCE terms (up to a fixed maximal degree) is possible and can lead to good approximations even if $M > 6T$. What is more, nearly perfect approximations can already be achieved for $T = 100$, which corresponds to $M \approx 3T$ in the scenario considered here.

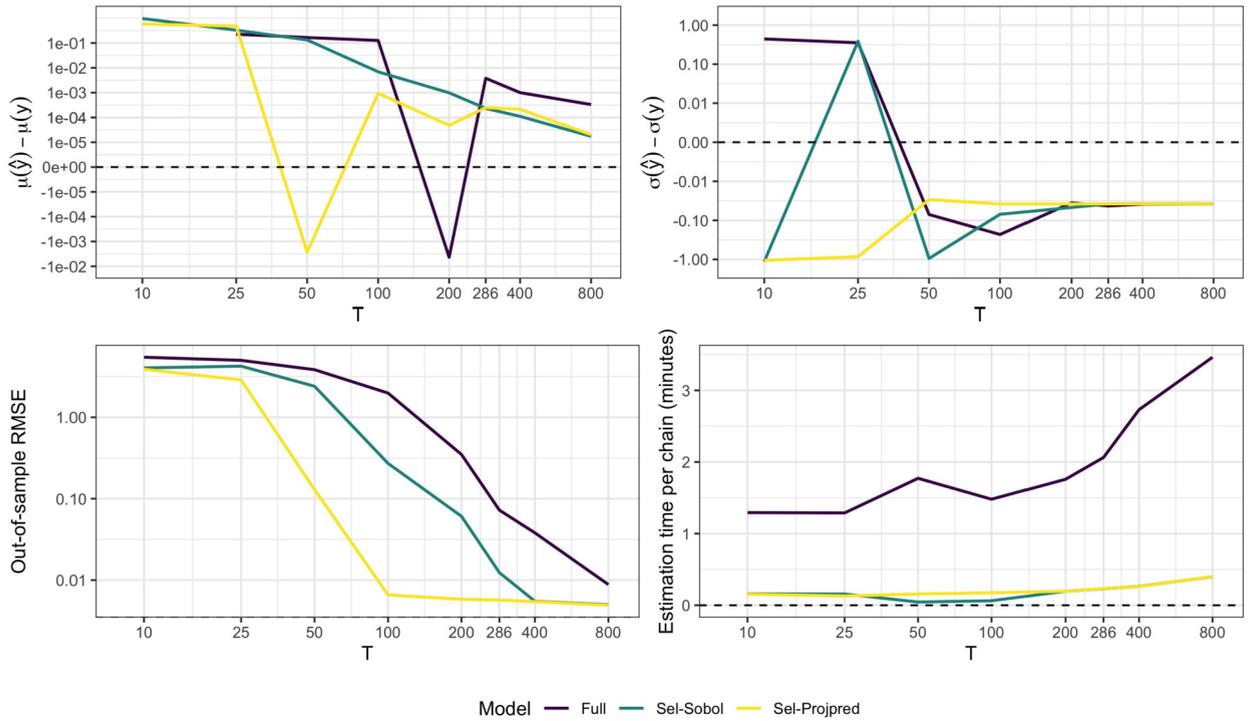


Fig. 4. Summarized results for the Ishigami function ($a = 7, b = 0.1$) by the number of training points T (log-scaled; evaluated sizes are equal to the displayed axis ticks), and model type (colors). Top-left: bias of the mean μ ; top-right: bias of the standard deviation σ ; bottom-left: average out-of-sample predictive RMSE; bottom-right: average estimation times per fitted model.

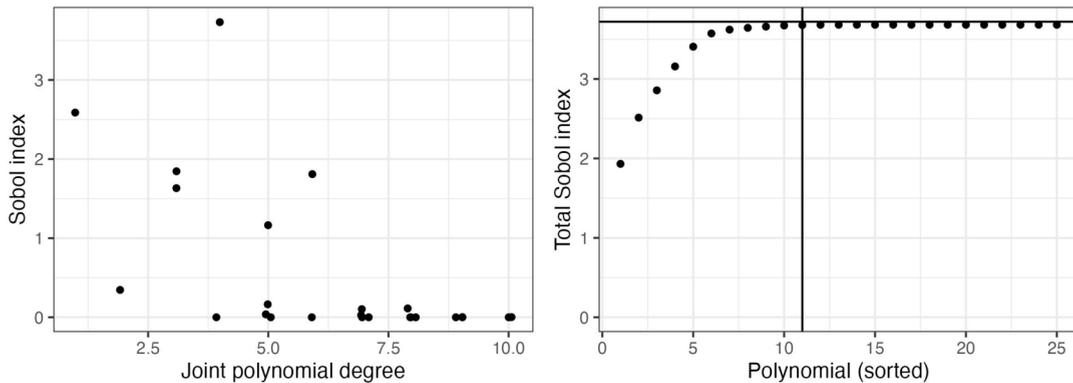


Fig. 5. Posterior mean Sobol indices and total Sobol indices for the sparse projpred model on the Ishigami function based on $T = 100$ training points and the $M_S = 25$ most important non-constant polynomials. The vertical line indicates the first index with only 0.5% or less increase in the total Sobol index. Error bars indicate 99% posterior credible intervals but they are too small to be clearly visible.

3.2. Sobol function

As second benchmark, we consider the Sobol function [62,64]:

$$y = \prod_{j=1}^N \frac{|4\omega_j - 2| + a_j}{1 + a_j}, \tag{29}$$

with $N = 8$ input variables $\omega_1, \dots, \omega_8$ being independently uniformly distributed in $[0, 1]$, and with its control parameter vector a being set to $a = (1, 2, 5, 10, 20, 50, 100, 500)$. In this configuration, the Sobol function has already been applied to benchmark other PCE approaches in previous studies [e.g., 67]. As can be readily seen from Equation (29), this function is symmetric around $\omega_j = 0.5$. Thus, only even-degree polynomials should contribute to its approximation. Accordingly, any sparse PCE method performing iterative local search (i.e., stopping their search as soon as they cannot find a polynomial of a given degree) would fail immediately: because there are no odd-degree polynomials contributing to the approximation,

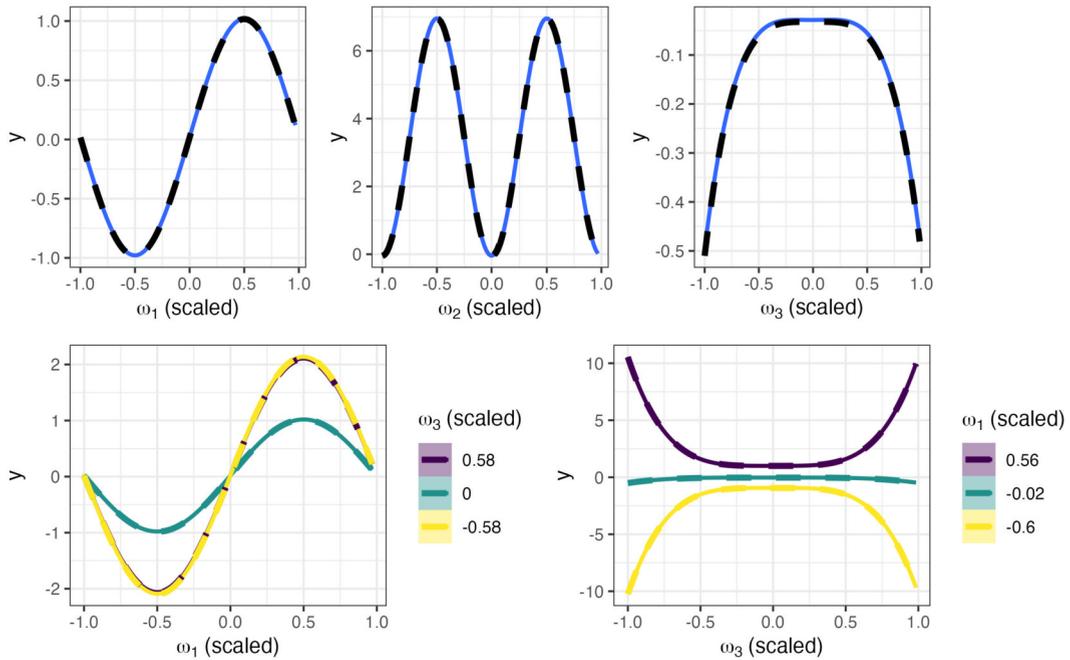


Fig. 6. Conditional prediction for the sparse projpred model of the Ishigami function based on $T = 100$ training points and the $M_S = 25$ most important polynomials. In each plot, all non-displayed input variables are fixed to their mean value. Dashed lines indicate the corresponding true Ishigami function values. Shaded areas indicate 95% posterior credible intervals but they are too small to be clearly visible. The conditional predictions match the true values almost perfectly.

such a search would stop at $0th, 2nd, 4th, \dots$ degree, depending on the initialization. Additionally, the Sobol function is an interesting benchmark because it has a relatively high-dimensional input space ($N = 8$ in our case) and is not continuously differentiable in $\omega_j = 0.5$, thus making it harder to approximate this function via PCE.

Similar to the Ishigami test case, we provide $T = (100, 300, 900, 1001, 2700, 3003, 8100)$ training points given by the first N values of the 8-dimensional Sobol Sequence [62] for training the upcoming PCE approximations. On these training points, we evaluated the Sobol function to provide actual training data.

As a first non-sparse PCE basis, we consider all 8-dimensional Legendre polynomials (scaled to fit the input range) up to a total degree of $d = 6$. This leads to $M = 3003$ possible polynomials in total. Additionally, we use a second set of polynomials as PCE basis, this time only considering the first 4 input variables in the PCE and corresponding 4-dimensional Sobol Sequence training points. In contrast to the first version, we now consider all 4-dimensional Legendre polynomials up to a total degree of $d = 10$, which implies $M = 1001$ polynomials in total. The underlying true Sobol function remains 8-dimensional, i.e. the training outcomes stay the same. This means that the last 4 input parameters are still random variables, but that the PCE assumes them to be irrelevant/constant. This is a valid choice, because the last 4 input variables are comparably less important in the Sobol function, and instead the basis approximates the relevant variables up to degree $d = 10$ instead of only $d = 6$. These two PCE bases explore the trade-off between the number of considered input variables and the maximal degree of polynomials included in the approximation.

With both PCE bases and the different sizes N of training data, we again create a varying range of inference complexity, as we already did for the Ishigami test case. Arguably, our second set of PCE models ($M = 1001$) constitutes an overall simpler scenario than the first set ($M = 3003$), because it has only about 1/3 of the number of polynomials in total.

Prior choices, sparsifying procedures, and MCMC sampling methods were the same as for the Ishigami case study (see Section 3.1). Again, all models converged well as indicated by standard algorithm-specific and algorithm-agnostic convergence diagnostics.

The approximation performance of the resulting PCEs was evaluated by comparing their approximations of the Sobol function’s mean and SD to the corresponding true values [62,64]. Additionally, we computed the out-of-sample RMSE across $T_{\text{test}} = 500$ randomly sampled test points not previously seen by the models as per Equation (27). From Fig. 7 we see that, similar to the Ishigami case, the full reference model is mostly competitive with the sparse sub-models in terms of mean and SD bias. But, again, it has worse out-of-sample RMSE, especially for small T . This demonstrates again that the R2D2-implied regularization alone (i.e., before sparse selection) may be insufficient to achieve satisfactory function approximation in $M \gg T$ scenarios. In contrast, both sparse sub-models perform much better, with the projpred models having a slight edge for smaller N (again see Fig. 4). As compared to the models considering all 8 input parameters, the models using only the first 4 variables show better mean approximation overall, better SD approximation and RMSE, as well as shorter estimation times for the full reference model.

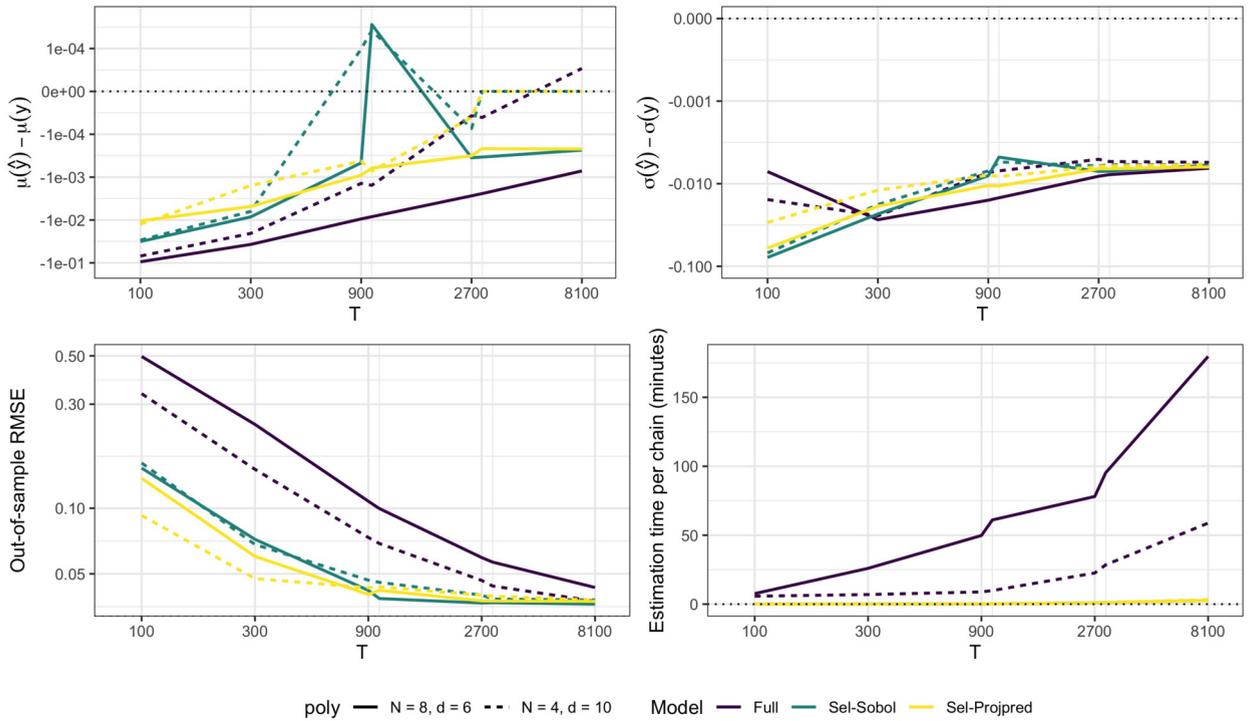


Fig. 7. Summarized results for the Sobol function ($a = (1, 2, 5, 10, 20, 50, 100, 500)$) by the number of training points T (log-scaled; evaluated sizes are equal to the displayed axis ticks), model type (colors), and considered polynomials (line type). Top-left: bias of the mean μ ; top-right: bias of the standard deviation σ ; bottom-left: average out-of-sample predictive RMSE; bottom-right: estimation times per fitted model.

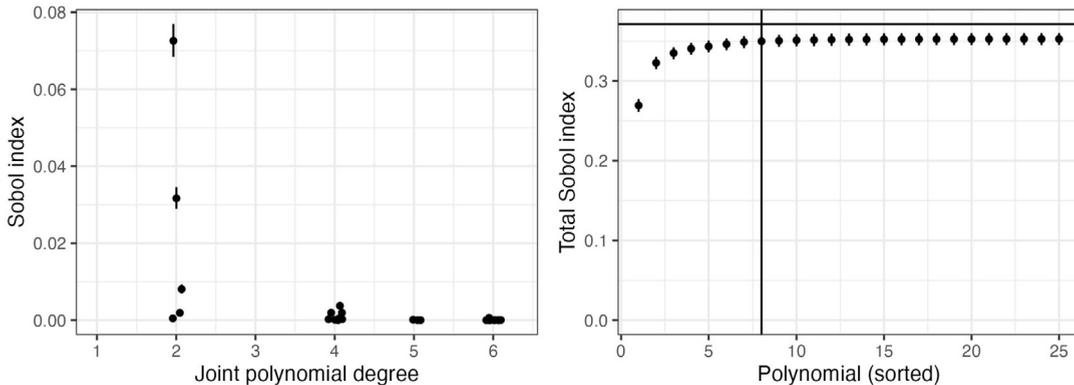


Fig. 8. Posterior mean Sobol indices and total Sobol indices for the sparse projpred model ($T = 300$) on the Sobol function ($N = 8$ variables) using the $M_S = 25$ most important non-constant polynomials constructed from all 8 variables. The vertical line indicates the first index with only 0.5% or less increase in the total Sobol indices. Error bars indicate 99% CIs.

For the sparse projpred models based on $T = 300$ training points, Sobol and cumulative Sobol indices are shown in Figs. 8 and 9 using all 8 or only the first 4 input parameters, respectively. As can be seen from the Figures, only 8 polynomials (including the constant polynomial) seem to be sufficient to achieve a good approximation of the Sobol function; and more polynomials (within the considered degrees) do not further improve this approximation. All of the relevant polynomials are even, as one would expect because of the symmetry of the Sobol function in terms of total Sobol indices (see Tables 2 and 3 for details on the polynomial degrees).

Within the total of $M_{sel} = 25$ selected polynomials, only very few are odd. These few odd terms are included simply because M_{sel} is larger than the number of actually relevant polynomials. As can be seen in Figs. 8 and 9, those odd polynomials are correctly shrunk to basically zero with almost no remaining uncertainty. Going back to Fig. 4, as T increases, the PCE estimates for the mean and SD of the Sobol function continue to improve slowly as compared to the sparse $T = 300$ models. However, the SD bias flattens out at around $\text{bias}(\sigma) = -0.1$. This asymptotic performance is about the same for both sets of models. So considering all 8 input parameters, but ignoring polynomials above degree $d = 6$ leads to about the same asymptotic SD bias as considering only the first 4 variables but ignoring all polynomials above degree $d = 10$. As can

Table 2

Posterior mean Sobol indices (Mean), 95% posterior credible intervals (95%-CI), and polynomial degree $d(\omega_j)$ by input variable of the 12 most important non-constant polynomials constructed from all 8 input variables as selected by the sparse projpred model for the Sobol function ($N = 8$ variables) based on $T = 300$ training points. All of these polynomials are even, but beyond the 7 most important of them, none of the further polynomials contributes to the PCE approximation in relevant manner.

Estimate	95%-CI	$d(\omega_1)$	$d(\omega_2)$	$d(\omega_3)$	$d(\omega_4)$	$d(\omega_5)$	$d(\omega_6)$	$d(\omega_7)$	$d(\omega_8)$
0.073	[0.069, 0.076]	2	0	0	0	0	0	0	0
0.032	[0.029, 0.034]	0	2	0	0	0	0	0	0
0.008	[0.007, 0.009]	0	0	2	0	0	0	0	0
0.004	[0.003, 0.004]	4	0	0	0	0	0	0	0
0.002	[0.001, 0.002]	2	2	0	0	0	0	0	0
0.002	[0.001, 0.002]	0	4	0	0	0	0	0	0
0.002	[0.001, 0.002]	0	0	0	2	0	0	0	0
0.001	[0.000, 0.001]	6	0	0	0	0	0	0	0
0.000	[0.000, 0.001]	0	0	0	0	2	0	0	0
0.000	[0.000, 0.001]	2	0	2	0	0	0	0	0
0.000	[0.000, 0.000]	0	0	4	0	0	0	0	0
0.000	[0.000, 0.000]	0	2	2	0	0	0	0	0

Table 3

Posterior mean Sobol indices (Mean), 95% posterior credible intervals (95%-CI), and polynomial degree $d(\omega_j)$ by input variable of the 12 most important non-constant polynomials constructed from the first 4 input variables (ω_1 to ω_4) as selected by the sparse projpred model for the Sobol function ($M = 8$ variables) based on $T = 300$ training points. All of these polynomials are even, but beyond the 7 most important of them, none of the further polynomials contributes to the PCE approximation in relevant manner.

Estimate	95%-CI	$d(\omega_1)$	$d(\omega_2)$	$d(\omega_3)$	$d(\omega_4)$
0.075	[0.072, 0.078]	2	0	0	0
0.033	[0.031, 0.035]	0	2	0	0
0.008	[0.007, 0.009]	0	0	2	0
0.004	[0.003, 0.004]	4	0	0	0
0.003	[0.002, 0.003]	2	2	0	0
0.002	[0.001, 0.002]	0	0	0	2
0.002	[0.001, 0.002]	0	4	0	0
0.001	[0.000, 0.001]	6	0	0	0
0.001	[0.000, 0.001]	2	0	2	0
0.000	[0.000, 0.001]	0	0	4	0
0.000	[0.000, 0.001]	8	0	0	0
0.000	[0.000, 0.000]	0	6	0	0

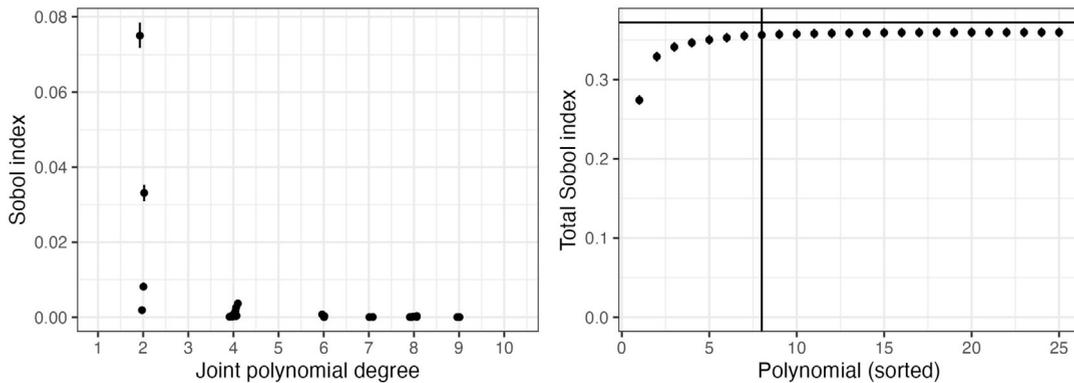


Fig. 9. Posterior mean Sobol indices and total Sobol indices for the sparse projpred model ($T = 300$) on the Sobol function ($N = 8$ variables) using the $M_S = 25$ most important non-constant polynomials constructed from the first 4 variables only. The vertical line indicates the first index with only 0.5% or less increase in the cumulative Sobol indices. Error bars indicate 99% CIs but they are too small to be clearly visible in most cases.

be seen from Tables 2 and 3, this is because both of these models identify the same polynomials as relevant, all of which are constructed from the first 4 input variables and have degree 6 or less.

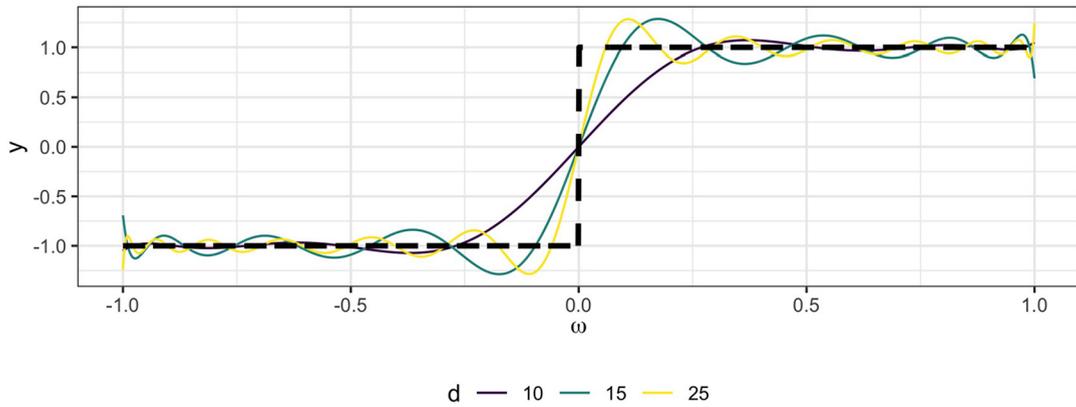


Fig. 10. Standard PCE expansion of the Signum function for $d = (10, 15, 25)$ and $T = d + 1$ training points given by Gaussian integration. The dashed line indicates the Signum function.

Together, these results demonstrate that with our proposed methods, performing global polynomial search (up to a fixed maximal degree) is possible and can lead to good approximations even if $M > 30T$. In specific, we showed that an “even-degree only” function is correctly identified – an achievement impossible for local-type greedy search scheme in sparse PCE methods. What is more, very close approximations can already be achieved for $T = 300$, which corresponds to $M \approx 10T$ in the models considering all $N = 8$ input variables. From the comparison between $N = 8$ at moderate polynomial degree ($d = 6$) and $N = 4$ at higher polynomial degree ($d = 10$), we learn that some form of manual trial-and-error for an appropriate search space can be helpful. The reason is that the global search is highly beneficial, but the corresponding search space still grows combinatorially fast with N and d .

3.3. Signum function

Polynomial-based surrogate models such as PCE are highly popular approximators of complex target functions. If the target functions are continuous, PCE ensures uniform convergence on compact sets and at least L^2 -convergence for functions in L^2 according to the Weierstrass theorem [68]. Nevertheless, these theoretical requirements are not always given in many practical applications. In particular, two effects are responsible for many drawbacks in the accuracy of the polynomial-based expansions: It is known that insufficient smoothness of the target function as well as selection of sub-optimal evaluation (training) points can lead to oscillations of the response surface, the so called Gibbs’ and Runge phenomena [68]. More precisely, the Gibbs’ phenomenon describes oscillations related to the discontinuities, and the Runge phenomenon describes oscillations caused by the selection of the evaluation points. One of the classical examples of the Gibbs’ phenomenon is the approximation of the Signum (i.e., sign) function as displayed in Fig. 10. Here, we can see that the increasing polynomial degree increases the number of oscillation and does not necessarily reduce their amplitude. As the signum function is one-dimensional, we can fully visualize the problem. Also, by construction, it is a purely odd function (unlike the Sobol function is purely even), and again provides an easily testable sparsity pattern.

In this case study, we consider PCEs of polynomial degree $d = (1, \dots, 10)$ with both $T = M = d + 1$ training points and polynomials. There are given either by Gaussian quadrature points to investigate the effect on the Gibbs’ phenomenon around the discontinuity at 0. Or they are given by the Sobol sequence that causes the Runge phenomenon on the right-hand side of the input space (i.e., close to 1). We deliberately investigate these $T = M$ cases as those are the minimal-data scenarios where standard PCE expansion is directly solvable (see Section 2.2). In addition to Standard PCE, we estimated Bayesian sparse PCE with the R2D2 prior (*Bayesian-R2D2*) and Bayesian PCE with improper flat priors on the coefficients (*Bayesian-flat*). The latter differs from Bayesian-R2D2 only in this choice of prior and can also be understood as a sampling-based equivalent of Standard PCE. Thus, the Bayesian-flat approach can be used to disentangle the influence of the R2D2 prior from the noise introduced by the MCMC algorithm. The settings of the MCMC algorithm were the same as in the Ishigami case study (see Section 3.1). As the purpose of this case study was only to study the regularizing effect of the R2D2 prior in case of a non-smooth target function, no further variable selection was performed.

In Fig. 11, we see exemplary for the case of $d = 10$ and $T = 11$ that the surrogate-models provided by classical PCE, Bayesian-Flat and Bayesian-R2D2 models trained on Gaussian quadrature points show very similar behavior, at least visually. In contrast, on the training points given by the Sobol sequence, the Bayesian-R2D2 models provide a significantly more reliable surrogate model than the other two here-considered PCE models in areas with sparse training points.

A closer investigation of the approximation errors for varying degree d (all with $T = M = d + 1$) is provided in Fig. 12. For the Sobol training points, we see almost uniformly higher accuracy of the Bayesian-R2D2 model as compared to the other two PCE approaches. The superiority of Bayesian-R2D2 is particularly striking for the RMSE, but clearly visible for all evaluated error metrics. For the training points given by Gaussian quadrature, all models perform very well overall as expected due to those training points’ provable optimality. However, due to the sampling-based estimation procedure, both

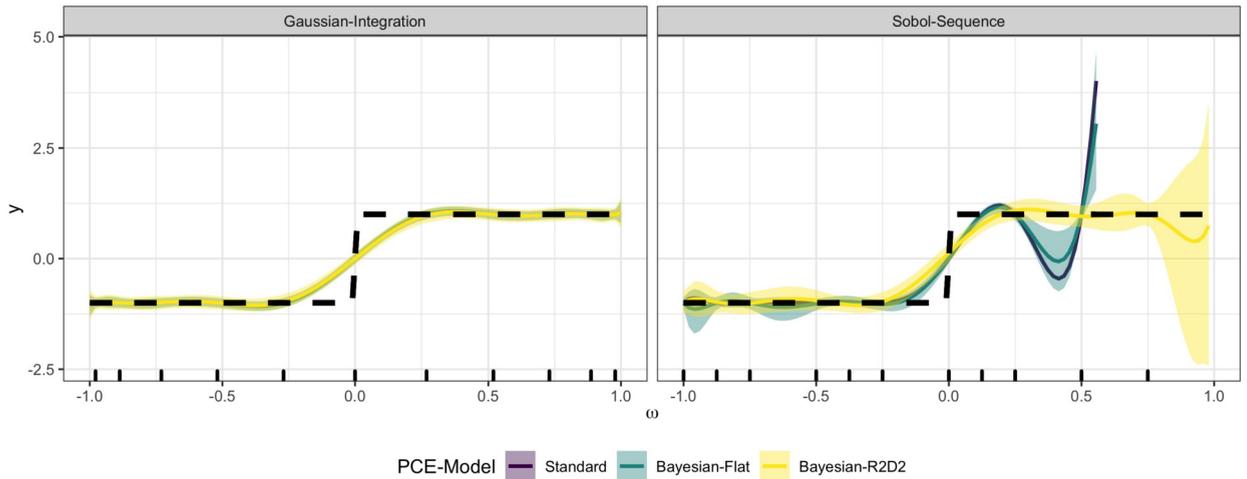


Fig. 11. Conditional predictions for different PCE models of the Signum function based on $M = 11$ polynomials and $T = 11$ training points, either given by Gaussian integration (left) or the Sobol Sequence (right). The y-axis was truncated at 5 from above and at -5 from below to keep the plot readable. No predictions are shown after the first exceedance of these thresholds for a given PCE model. The dashed line indicates the Signum function. The small vertical dashes on the x-axis indicate the training point locations.

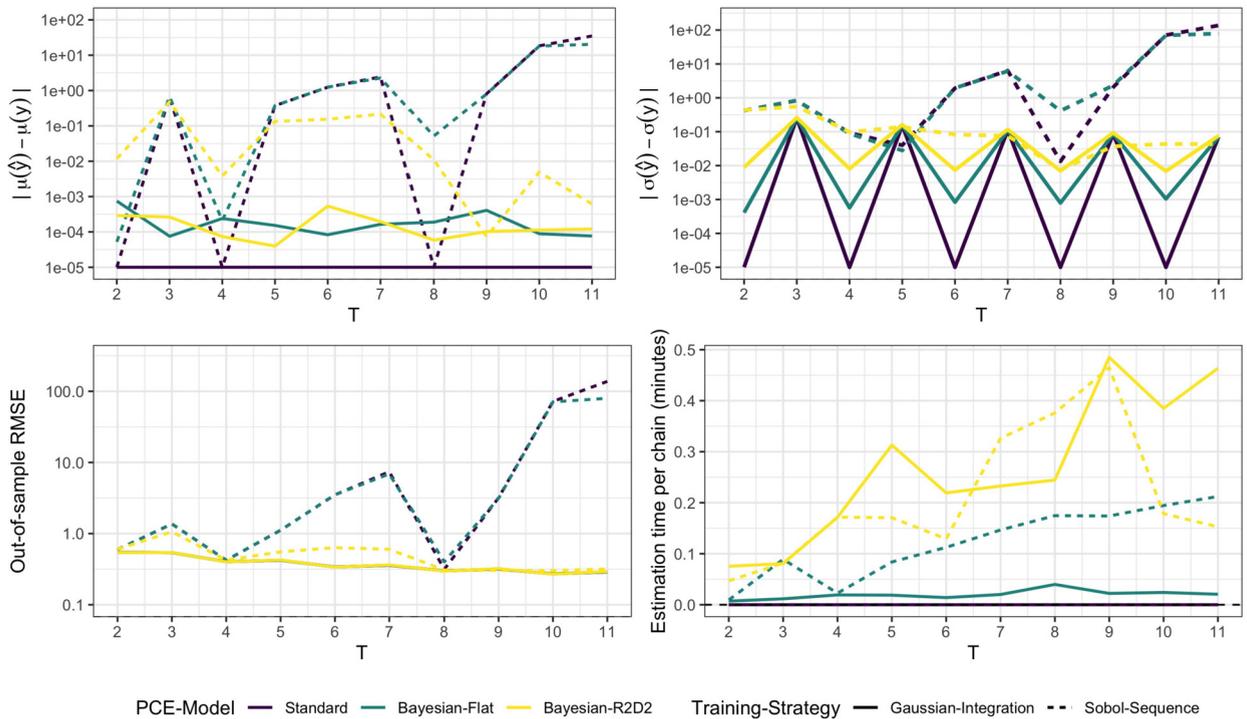


Fig. 12. Summarized results for the Signum function by the number of training points T (x-axis; log-scaled), type of PCE model of degree $d = T - 1$ (colors), and strategy for choosing training points (line type). Top-left: absolute error of the mean μ ; top-right: absolute error of the standard deviation σ ; bottom-left: average out-of-sample predictive RMSE; bottom-right: estimation times per fitted model. Smaller values are better. Absolute errors of mean and standard deviation are truncated from below at $1e - 5$ to improve readability of the figure but actually go as far as $1e - 14$.

Bayesian approaches do not achieve the extremely high accuracy of Standard PCE in certain scenarios (absolute errors of mean and standard deviation below $1e - 14$ in some cases). Increasing the number of MCMC samples S would reduce the sampling error of the Bayesian approaches arbitrarily in theory, but scaling only by a factor of roughly \sqrt{S} , which renders errors below $1e - 4$ impractical to achieve (assuming unit scale of the target function). Whether the error is $1e - 4$ or $1e - 14$ is unlikely to matter in practice, but may still be of theoretical interest.

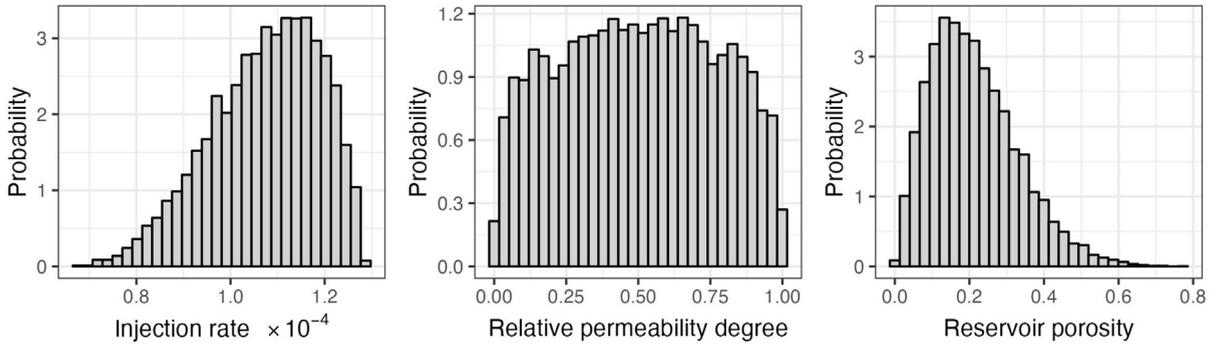


Fig. 13. Marginal distributions of the input variables in the CO₂ storage benchmark.

3.4. Carbon dioxide storage benchmark

Finally, we consider a carbon dioxide storage benchmark [31,32] for demonstration on a more realistic problem. This benchmark case describes the injection of carbon dioxide gas (CO₂) into a deep geological formation to mitigate greenhouse gas effects. Upon injection, the CO₂ spreads within the formation, which is initially filled with brine (i.e., with very salty, deep groundwater).

This test case allows to analyze the joint effects of several sources of uncertainty in the modeling of hydrosystems. In particular, these are uncertain boundary conditions (here: injection rate), uncertain material properties (here: porosity of the formation) and uncertainties in constitutive relations of multi-phase fluid flow (here: the degree to which the formation's permeability occurs in the so-called fractional flux function). It is assumed, that the density and viscosity of the fluids (CO₂, brine) are constant and that all processes are isothermal, such that CO₂ and brine can be described as a two-phase flow. When additionally assuming that the fluids' pressure can be determined in analytic form and upon transformation into radial coordinates, one can obtain the following governing equation:

$$\phi \frac{\partial S_g}{\partial t} - \frac{1}{r} \frac{\partial}{\partial r} (q_{CO_2} C_p f_g) - q_{CO_2} = 0, \quad \text{in } D \times (0, T_{\max}), \quad (30)$$

which resembles a non-linear, hyperbolic partial differential equation for the (scalar) saturation of CO₂, denoted by S_g (with subscript g for gas). In this expression, the uncertain porosity of the formation is denoted by ϕ , the time coordinate is denoted by $t \in (0, T_{\max})$, the radial space location (coordinate) by $r \in D = [0, 250]$, the unknown saturation of the gas phase by S_g , and the uncertain injection rate by q_{CO_2} . Recall that the non-linear fractional flux function, here denoted by f_g , depends on the unknown CO₂ saturation S_g and on the formation's permeability (symbols not shown here) to an uncertain degree. The constant C_p represents the pressure decay.

The numerical solution of the non-linear equation (30) is obtained via a finite-volume method with central-upwind flux [34]. The complete input and output data are available in [32].

Following the original benchmark definition [31], we analyzed the CO₂ saturation data after $T_{\max} = 100$ days at $L = 250$ discretized spatial locations. We trained independent PCE models per coordinate on $T = (27, 64, 125, 216, 1331)$ training points given by 3-dimensional grids of equidistant points across the $N = 3$ input dimensions. The training data are available in the original benchmark definition [31].

Due to strong discontinuities of the underlying physical model at certain combinations of parameter values and space/time coordinates, the approximation performance for PCE-based models is limited by design [31]. Hence, we cannot necessarily expect our proposed methods to perform competitively with methods specifically designed for such discontinuous target functions [31]. Rather, the goal is to show the performance achievable despite this limitation and compare the results to those obtained by other methods with the same limitation. A nice property of this benchmarking case is that it has $L = 250$ spatial locations, so that it actually represents 250 test cases in one single scenario.

To obtain the polynomial basis for our PCE, we used the arbitrary polynomial chaos method (aPC) [44] on the marginal distributions of input parameters as depicted in Fig. 13. As aPC basis, we considered 3-dimensional aPC polynomials up to a total degree of $d = 10$, which implies $M = 286$ polynomials in total. As for the other case studies, these values were chosen to create a varying range of inference complexity.

As full reference models (two for each value of T), we used all M considered polynomials as basis, equipped with two different R2D2 priors on the $M - 1$ non-constant polynomials as described in Section 2.6. The first R2D2 prior uses $R^2 \sim \text{Beta}(\zeta = 0.5, \nu = 2)$, while the second R2D2 prior uses $R^2 \sim \text{Beta}(\zeta = 0.9, \nu = 10)$ for reasons explained later. Subsequently, with our two variable selection methods for comparison, we extracted the $M_{\text{sel}} = (25, 50)$ most important out of all non-constant polynomials plus the constant polynomial. These selected polynomials were used as PCE bases for corresponding sparse aPC models. We investigated two different values for M_{sel} , because the local discontinuities of the underlying true function were likely to require many polynomials to be approximated acceptably well. We used the same R2D2 priors for

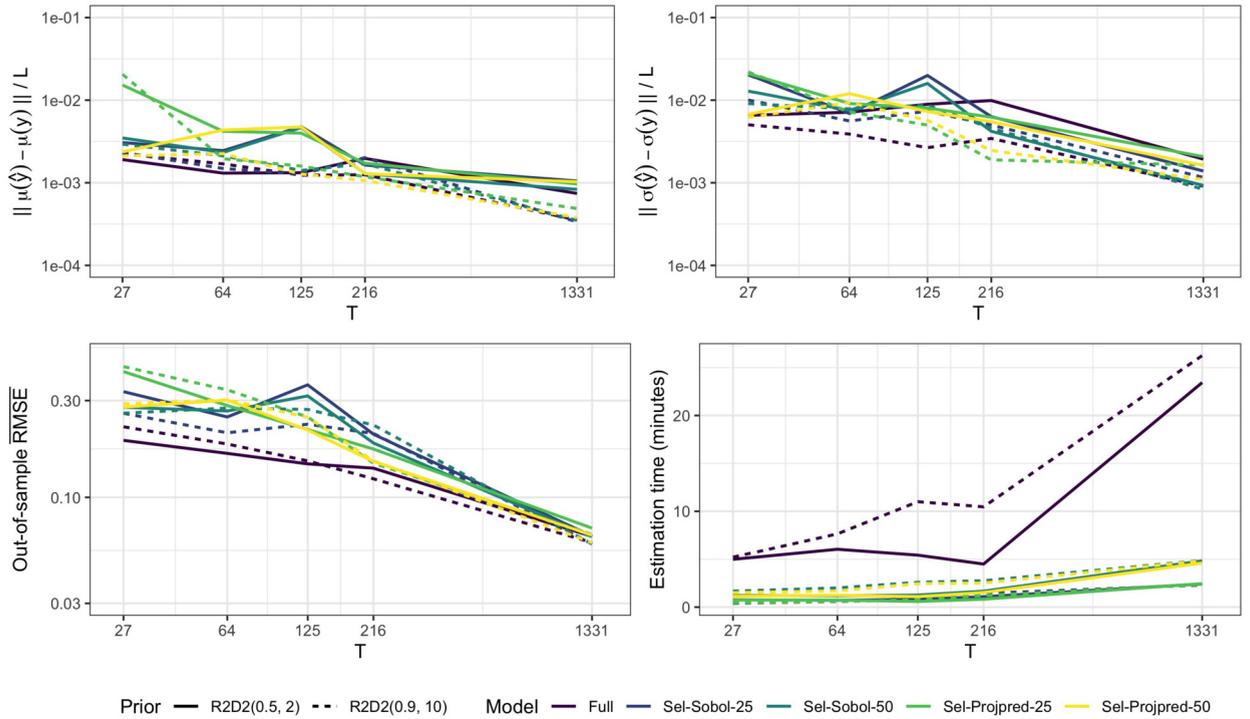


Fig. 14. Summarized results for the CO2 data by the number of training points (log-scaled; evaluated at $T = (27, 64, 125, 216, 1331)$), model type (colors), and prior on R^2 (line type). Top-left: normalized absolute error of the mean μ ; top-right: normalized absolute error of the standard deviation σ ; bottom-left: average out-of-sample predictive RMSE; bottom-right: average estimation times per fitted model.

the sparse sub-models as for the corresponding full reference models. MCMC sampling methods were the same as for the other case studies. However, now we use only a single Markov chain per model to simplify parallelization across models on a computing cluster. Each chain was run for 3000 iterations, out of which the first 1000 were discarded as warmup. This provides a total of 2000 post-warmup draws used for inference for each model version. Once again, all models converged well as indicated by standard algorithm-specific and algorithm-agnostic convergence diagnostics.

As outcome measures, we considered the normalized ℓ^2 norm of the mean and SD biases over the $L = 250$ locations

$$\frac{\|\mu(\hat{y}) - \mu(y)\|}{L} = \frac{\sqrt{\sum_{l=1}^L (\mu(\hat{y}_l) - \mu(y_l))^2}}{L} \tag{31}$$

$$\frac{\|\sigma(\hat{y}) - \sigma(y)\|}{L} = \frac{\sqrt{\sum_{l=1}^L (\sigma(\hat{y}_l) - \sigma(y_l))^2}}{L}. \tag{32}$$

The variables $\mu(\hat{y})$ and $\sigma(\hat{y})$ are the length- L vectors of model-implied mean and SD estimates, respectively. The vectors $\mu(y)$ and $\sigma(y)$ are the corresponding vectors of “true” reference values obtained by [31] via Monte Carlo simulations using 10^4 samples. These measures were chosen for comparability with the results obtained in [31]. Additionally, we considered the mean out-of-sample RMSE across locations

$$\overline{\text{RMSE}} = \frac{1}{L} \sum_{l=1}^L \text{RMSE}_l, \tag{33}$$

where RMSE_l is given by Equation (27) applied to the l -th location.

The summarized results across coordinates are depicted in Fig. 14. Especially for small to medium T , our PCE approaches are competitive with the state-of-the-art, but fall off a bit for larger T , where kernel greedy interpolation methods [17,76,78] have an edge in this benchmark [compare our results with those displayed in 31, Figure 7]. While not uniformly reaching the state-of-the-art on this benchmark due to the discussed discontinuities, our methods perform equally well or better than other PCE-based approaches [compare with the blue lines displayed in 31, Figure 7] and provide better scaling potential to higher-dimensional problems and higher polynomial degrees.

As can be readily seen, the full reference models show equal or better approximation performance than their respective sparse sub-models in terms of mean and SD bias as well as RMSE almost uniformly across T . Further inspection of the approximation performance for the individual locations, as displayed in Figs. 15, 16, and 17, reveal that this is primarily due

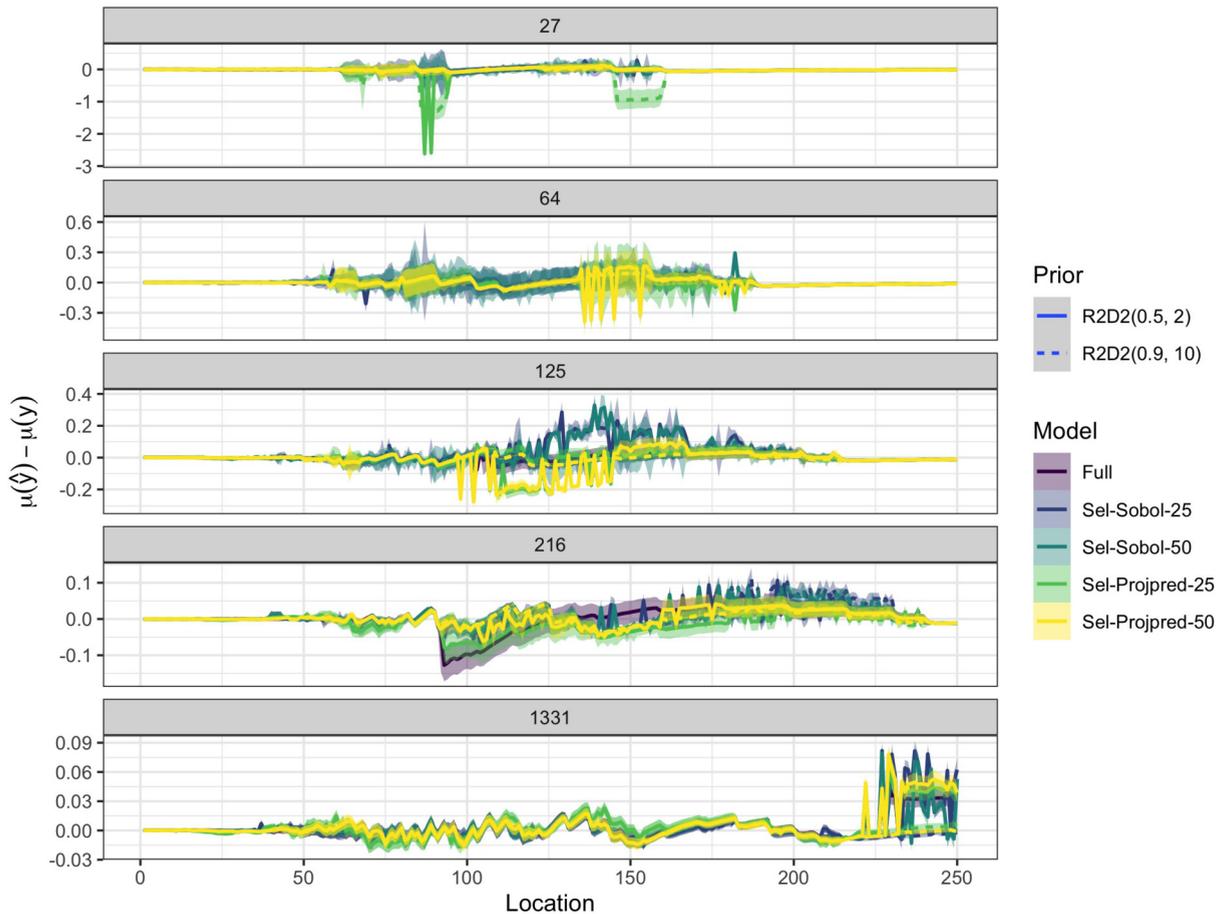


Fig. 15. Difference between model-implied and true mean for the CO₂ data by location (x-axis), number of training points T (facets), and model type (colors). Values closer to zero are better.

to approximation outliers for some coordinates $l \in [50, 200]$. These are the places where strong discontinuities and shocks are happening in the underlying true model. These outliers seem to occur less frequently for the sparse $M_{sel} = 50$ models compared to the sparse $M_{sel} = 25$ models, especially clearly so for projpred selection and small T . This suggests that many polynomials are required to provide an acceptable approximation of the true model for these locations, which could explain why the reference models perform better. Despite most of the $M = 286$ polynomials being likely redundant, still including them in the model appears to do little harm in the considered scenarios. This points to the success of the R2D2 prior in shrinkage those redundant coefficient sufficiently to prevent overfitting.

Figs. 15 and 16 reveal another problematic pattern that is present only for models with an $R^2 \sim \text{Beta}(\zeta = 0.5, \nu = 2)$ prior: For larger T and specific locations ($l > 90$ for $T = 218$ and $l > 220$ for $T = 1331$) the mean and SD estimates of the reference model, and as a result also of the corresponding sparse sub-models, have substantial bias. This is because all PCE coefficients had been shrunk to almost zero, thus preventing both sensible prediction and coefficient selection. Similar problems have been observed for other PCE-based approaches applied to this CO₂ data set [see Figure 3 in 31]. Further inspection revealed that this behavior occurred in our models because of a bimodal posterior of the PCE coefficients, with one mode at zero for all coefficients and the other “sensible” mode away from zero for at least some of the truly relevant coefficients.

In additional experiments not shown here, we were able to find this second mode at least for $T = 1331 > M$ by setting the initial values of the MCMC chains to the maximum-likelihood estimate (MLE). In this case, the resulting mean and SD estimates of the Bayesian models were almost unbiased again, closely resembling the corresponding values of the models with an $R^2 \sim \text{Beta}(\zeta = 0.9, \nu = 10)$, which implies a weaker R2D2 prior in the sense that its shrinkage effect is reduced (expected R^2 is $\zeta = 0.9$ instead of $\zeta = 0.5$). However, since the MLE does not exist for $T < M$, it is non-trivial to find such suitable initial values in general. Instead, one either has to apply a weaker R2D2 prior, as done here, to practically eliminate the all-zero mode; or one has to run multiple MCMC chains with diffuse initial values in the hope that at least one finds the desired mode.

In summary, the results demonstrate that our sparse Bayesian PCE can perform well even when approximating discontinuous target functions that are not well approximated by polynomials in general. However, care must be taken that the

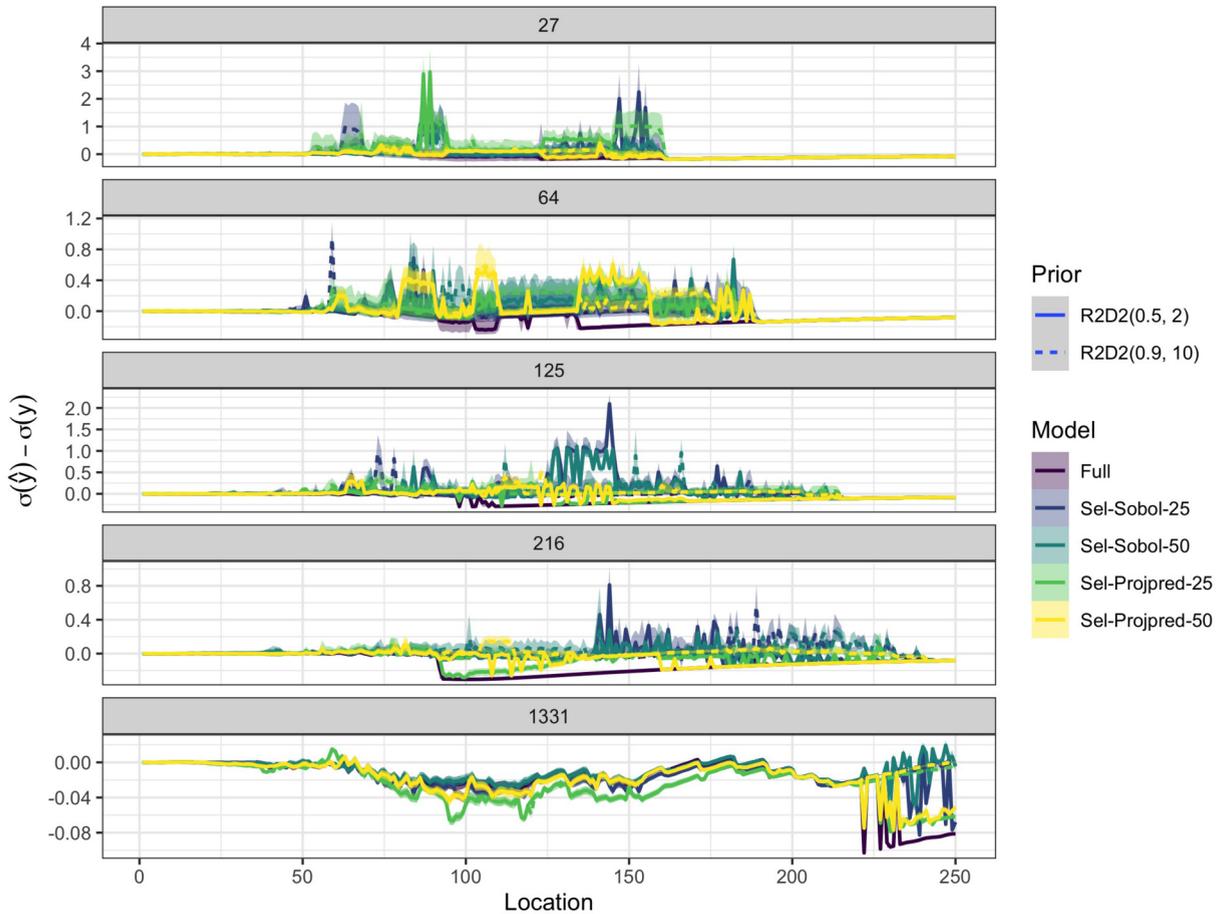


Fig. 16. Difference between model-implied and true standard deviation for the CO2 data by location (x-axis), number of training points T (facets), and model type (colors). Values closer to zero are better.

applied shrinkage priors do not “over-shrink” all PCE coefficients to zero (thus preventing any signal to be encoded in the PCE) and that enough polynomials are still included in the sparse representations.

4. Discussion

In this paper, we presented a fully Bayesian approach to sparse polynomial chaos expansion (PCE). In particular, we define the PCE coefficients as random variables to be inferred probabilistically based on training data (e.g., simulations performed by a computationally expensive model for which one wants to build a surrogate). The desired sparsity is induced by a specific choice of prior over the PCE coefficients, known as the R2D2 prior [87]. Below we recap and summarize our methods and obtained results.

The R2D2 prior belongs to the class of priors known as shrinkage priors in statistics. It consists of two parts: (1) It regularizes the proportion of variance in the true responses (e.g., of the results of the expensive model) that is captured by the PCE, that is, it regularizes the coefficient of determination R^2 as known from statistical regression. (2) It controls how the captured variance is distributed over the expansion coefficients. If desired, the latter could be informed by assumptions on how the PCE spectrum decays with increasing polynomial degree. As the variance captured by individual PCE terms is given by the square of the corresponding coefficient (and the overall captured variance is the sum of squares, without the zeroth-degree term), both parts of the R2D2 prior link naturally to PCE theory.

The R2D2 induces only weak sparsity in the sense that, upon training, many coefficients are *close* to zero, but not necessarily exactly zero. Thus, we post-process the weakly sparse PCE with a procedure of statistical variable selection to obtain truly sparse expansions. Finally, we re-train these truly sparse expansions with the same training data, and again with the same R2D2 prior; thus, we obtain the final, sparse expansion. For variable selection, we compare a well-known Sobol-based, greedy variable selection and a more recent, promising variable selection method called projective prediction [51]. The latter resembles a more global view on the sparse selection of PCE terms, and is linked to Bayesian cross-validation.

In combination, the R2D2 prior and the variable selection step allow for a global search for sparse PCEs. The global search can consider all candidate terms defined by a high-degree PCE, based on rigorously statistical tools. This goes far be-

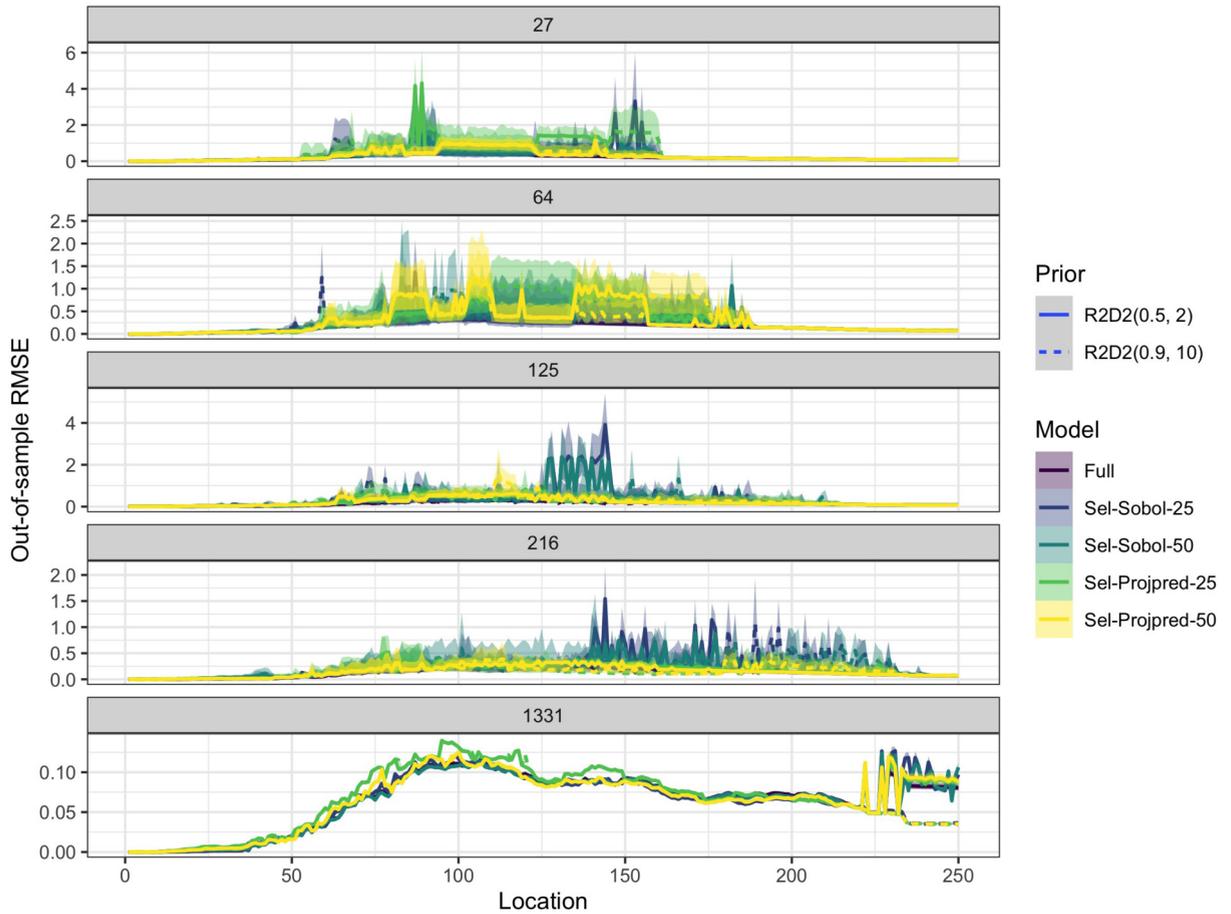


Fig. 17. Out-of-sample predictive RMSE for the CO₂ data by location (x-axis), number of training points T (facets), and model type (colors). Values closer to zero are better.

yond current, rather heuristic, search strategies for relevant PCE terms, which mostly search only in the vicinity of already identified terms. Also, our approach allows to decouple the computational design of training points (i.e., via collocation or integration rules) from the search for sparse PCEs: it can be applied even on given sets of training data, and can search through a space of possible PCE expansions that even includes strongly under-determined cases. The R2D2 prior not only enables sparse selection afterwards, but also directly regularizes against overfitting and oscillations, such that sparse selection may not even be required necessarily.

We demonstrated these favorable properties in four different test cases. The test cases included the Signum function to assess how our approach handles the Gibbs and Runge phenomena of oscillation in polynomial approximations. We also used the Ishigami function and the Sobol function; these are known, hard cases for benchmarking PCE techniques. These three functions all contain known patterns of sparsity (e.g., only even or only odd degree, independent variables without cross terms, or sets of close-to-insignificant variables). In all these cases, our sparse Bayesian PCE with R2D2 prior performed very well. For the Signum function, it was equally good as the standard PCE when optimal training points are used, and it outperformed the standard PCE for suboptimal points, as the R2D2 reduces oscillations.

For exhaustive comparison to many other surrogate methods (including different versions of the PCE), the fourth test case was a benchmark on CO₂ storage in geological formations. On this test case, we compared the results and performance of our sparse Bayesian PCE to an extensive list of benchmarked surrogate techniques published in a comparison paper [31]. Here, we could observe that our method outperforms existing polynomial-type surrogate methods. In all our test cases, the projective prediction method for sparse selection outperformed the Sobol-based selection.

Care must be taken in the specific hyperparameter choice of the R2D2 prior. In our study, we explored different Beta priors for R^2 , and found that the choice of mean and variance in the Beta distribution may require some expert knowledge or tuning (see also [1,87]). With the R2D2 prior, the PCE coefficients are no longer found by analytically solving a linear system, but by sampling via Markov Chain Monte-Carlo methods. This is much more expensive computationally, but also provides information about the uncertainty of the inferred PCE without making any distributional assumptions of the posterior.

Detailed analyses of the computational costs revealed that they are indeed mainly driven by MCMC sampling of the full model that contains all PCE polynomials up to a given degree. This took around 1-3 minutes in simple cases to around 3 hours in the most complicated case we evaluated (i.e., Sobol function with 8100 training points and 3003 polynomials). Running MCMC again for the sparse sub-models took no more than some seconds in all cases due to the small number of remaining model parameters. In comparison, constructing the PCE polynomials themselves took only a fraction of a second in all cases and constructing the search path via projpred took less than a second in simple cases up to around 30 seconds in the most complex case.

Whether the effort of MCMC sampling is worth the effort depends primarily on the computational complexity of the original simulator that produces the training data. If simulations are fast (i.e., seconds or minutes per simulation), then we do not have to bother with any sparse PCE approach at all and can just use standard (unregularized) PCE or any other surrogate model of our liking. However, when each simulation takes many hours or even days (which is not uncommon in physics applications; e.g., [33]), then we would rather invest a few minutes or hours for our fully Bayesian sparse PCE approach than to wait many days for sufficient additional simulations to finish before standard PCE (or comparable approaches) would achieve similar accuracy.

In terms of scaling, our proposed method behaves like other basis-function approximations, whether those are PCEs, wavelets, splines, or basis-function approximations to Gaussian processes [28,56,80]. That is, their computational costs and memory requirements scale well (linearly) with the number of training points, but badly (exponentially) with the number of input dimensions due to the use of tensor product terms (see also Section 2.1). Because of the latter, we would not recommend using full PCE – or any other basis function approximation – for more than about 4 input dimensions (see also [56]). In these cases, other surrogate modeling approaches such as exact Gaussian processes or deep neural networks might be favorable. Yet, enforcing sparsity in problems that are rather sparse by nature can postpone the effect of the curse of dimensionality towards higher dimensions.

The idea of sparse construction of basis function approximations by means of Bayesian linear models equipped with the R2D2 prior and subsequent variable selection is quite general. We applied it to the arbitrary polynomial chaos expansion, which is not bound to specific cases for probability distributions of model input parameters. In future research, our approach could be applied to adaptive multi-element PCE versions, could be equipped with schemes for active learning, or could be combined with Gaussian processes for handling the residuals of the found sparse PCEs.

CRediT authorship contribution statement

Paul-Christian Bürkner: Conceptualization, Data curation, Formal analysis, Methodology, Project administration, Software, Visualization, Writing – original draft, Writing – review & editing. **Ilja Kröker:** Data curation, Formal analysis, Software, Visualization, Writing – original draft, Writing – review & editing. **Sergey Oladyshkin:** Data curation, Writing – original draft, Writing – review & editing. **Wolfgang Nowak:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing.

Declaration of competing interest

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

Data availability

Data sources are referenced in the manuscript

Acknowledgements

Partially funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 - 390740016 and DFG Project 432343452.

Appendix A. Beta and Dirichlet distributions

Below, we show densities and the first two moments of the Beta and Dirichlet distributions. The density of the Beta distribution for scalar $x \in [0, 1]$ with positive shape parameters a_1 and a_2 is given by

$$p_{\text{Beta}}(x | a_1, a_2) = \frac{x^{a_1-1} (1-x)^{a_2-1}}{B(a_1, a_2)}, \quad (34)$$

where $B(a_1, a_2)$ is the two-dimensional Beta function. In this study, we are using the parameterization of the Beta distribution that arises when we set $a_1 = \zeta \nu$ and $a_2 = (1 - \zeta) \nu$. In this so-called “mean-precision” parameterization, the mean and variance of the Beta distribution are given by

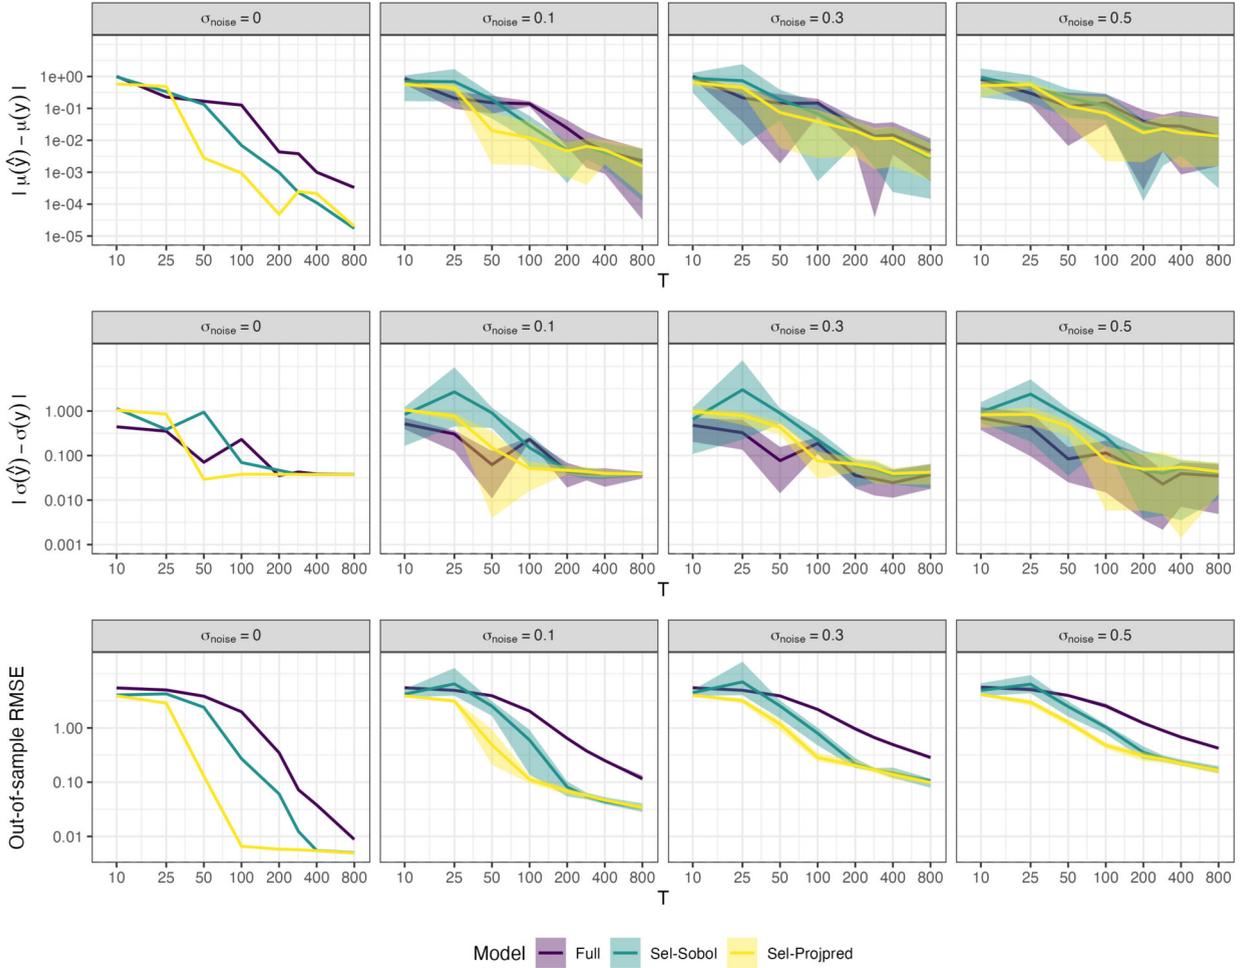


Fig. 18. Summarized results for the Ishigami function ($a = 7, b = 0.1$) by the number of training points T (log-scaled; evaluated sizes are equal to the displayed axis ticks), model type (colors), and standard deviation of noise (facets). Top: absolute error of the mean μ ; middle: absolute error of the standard deviation σ ; bottom: average out-of-sample predictive RMSE.

$$\mathbb{E}(x) = \zeta \quad \text{and} \quad \text{Var}(x) = \frac{\zeta(1 - \zeta)}{1 + \nu}. \tag{35}$$

The density of the Dirichlet distribution for a simplex $x \in [0, 1]^M$ and $\sum_{i=1}^M x_i = 1$ with concentration parameter vector $\theta \in \mathbb{R}_+^M$ is given by

$$p_{\text{Dirichlet}}(x | \theta) = \frac{1}{B(\theta)} \prod_{i=1}^M x_i^{\theta_i - 1}, \tag{36}$$

where $B(\theta)$ is the M -variate Beta function. When we set $\theta_0 = \sum_{i=1}^M \theta_i$, the mean and variance of the Dirichlet distribution for the i th component x_i are given by

$$\mathbb{E}(x) = \frac{\theta_i}{\theta_0} \quad \text{and} \quad \text{Var}(x) = \frac{\theta_i / \theta_0 \times (1 - \theta_i / \theta_0)}{1 + \theta_0}. \tag{37}$$

Appendix B. Robustness to noisy training data

In order to test the robustness of our method against noise in the training data, we re-ran our experiments of the Ishigami function from Section 3.1 with varying amounts of noise. Specifically, we added noise from a normal($0, \sigma_{\text{noise}}$) distribution where the standard deviation took on values of $\sigma_{\text{noise}} = (0.1, 0.3, 0.5)$. To reduce dependency on individual noisy simulations, we created 10 simulated training data sets per condition. The results depicted in Fig. 18 demonstrate the relative robustness of both the full reference model and the sparse projpred model against noise. This is in line with previous

research that demonstrated that strongly regularized reference models act as a noise filter that reduces overfitting and that projpred variable selection is an especially powerful reference model approach [49]. While the results show roughly the same overall patterns and orderings of the methods as the noiseless experiments, having added noise of first-digit order, accuracy in the third digit or lower becomes hard to achieve. Interestingly, the sparse sub-models resulting from greedy selection of the largest Sobol coefficients show unstable approximation behavior for $T = 25$ training points for all degrees of added noise. That is, for some of the noisy data sets, the error metrics were multiple times worse than for the other methods. This adds more evidence that greedy selection of terms may be a sub-optimal and potentially unstable approach even if selection was done based on a well-performing reference model.

References

- [1] J.E. Aguilar, P.-C. Bürkner, Intuitive Joint Priors for Bayesian Linear Multilevel Models: The R2D2M2 prior, arXiv preprint 2022.
- [2] R. Ahlfeld, B. Belkouchi, F. Montomoli, SAMBA: sparse approximation of moment-based arbitrary polynomial chaos, *J. Comput. Phys.* 320 (2016) 1–16.
- [3] N. Akhiezer, *The Classical Moment Problem*, Hafner Publication Company, New York, 1965.
- [4] R. Askey, J. Wilson, *Some Basic Hypergeometric Polynomials That Generalize Jacobi Polynomials*, *Memoirs of the American Mathematical Society*, AMS, Providence, 1985.
- [5] F. Beckers, A. Heredia, M. Noack, W. Nowak, S. Wierprecht, S. Oladyshkin, Bayesian calibration and validation of a large-scale and time-demanding sediment transport model, *Water Resour. Res.* 56 (2020) 7.
- [6] M. Betancourt, A conceptual introduction to Hamiltonian Monte Carlo, arXiv preprint 2017.
- [7] A. Bhadra, J. Datta, N.G. Polson, B. Willard, Default Bayesian analysis with global-local shrinkage priors, *Biometrika* 103 (4) (2016) 955–969.
- [8] G. Blatman, B. Sudret, Sparse polynomial chaos expansions and adaptive stochastic finite elements using a regression approach, *C. R., Méc.* 336 (6) (2008) 518–523.
- [9] P.-C. Bürkner, brms: an R package for Bayesian multilevel models using Stan, *J. Stat. Softw.* 80 (1) (2017) 1–28.
- [10] P.-C. Bürkner, Advanced Bayesian multilevel modeling with the R package brms, *R J.* 10 (1) (2018) 395–411.
- [11] R.H. Cameron, W.T. Martin, The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals, *Ann. Math.* 48 (1947) 385–392.
- [12] B. Carpenter, A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M.A. Brubaker, J. Guo, P. Li, A. Ridell, Stan: a probabilistic programming language, *J. Stat. Softw.* 76 (1) (2017) 1–32.
- [13] A. Catalina, P. Bürkner, A. Vehtari, Latent space projection predictive inference, arXiv preprint 2021.
- [14] A. Catalina, P.-C. Bürkner, A. Vehtari, Projection predictive inference for generalized linear and additive multilevel models, in: *Artificial Intelligence and Statistics (AISTATS) Conference Proceedings*, 2022.
- [15] K. Cheng, Z. Lu, S. Xiao, S. Oladyshkin, W. Nowak, Sparse Gaussian process model with mixed covariance function for uncertainty quantification, *Int. J. Uncertain. Quantificat.* (2021).
- [16] P.R. Conrad, Y.M. Marzouk, Adaptive Smolyak pseudospectral approximations, *SIAM J. Sci. Comput.* 35 (6) (2013) A2643–A2670.
- [17] S. De Marchi, R. Schaback, H. Wendland, Near-optimal data-independent point locations for radial basis function interpolation, *Adv. Comput. Math.* 23 (3) (2005) 317–330.
- [18] O.G. Ernst, A. Mugler, H.-J. Starkloff, E. Ullmann, On the convergence of generalized polynomial chaos expansions, *ESAIM: Math. Model. Numer. Anal.* 46 (2) (2012) 317–339.
- [19] L. Fagiano, M. Khammash, Nonlinear stochastic model predictive control via regularized polynomial chaos expansions, in: *IEEE Conference on Decision and Control (CDC)*, vol. 51, 2012, pp. 142–147.
- [20] J. Favard, *Sur les Polynomes de Tchebicheff*, CR Academic Science, Paris 200, 1935.
- [21] J. Gabry, D. Simpson, A. Vehtari, M. Betancourt, A. Gelman, Visualization in Bayesian workflow, *J. R. Stat. Soc., Ser. A, Stat. Soc.* 182 (2) (2019) 389–402.
- [22] W. Gautschi, *Orthogonal Polynomials: Computation and Approximation*, Numerical Mathematics and Scientific Computation, Oxford University Press, New York, 2004.
- [23] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, D.B. Rubin, *Bayesian Data Analysis*, 3rd edition, Chapman and Hall/CRC, London, 2013.
- [24] R.G. Ghanem, P.D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer, New York, 1991.
- [25] M. Hadigol, A. Doostan, Least squares polynomial chaos expansion: a review of sampling strategies, *Comput. Methods Appl. Mech. Eng.* 332 (2018) 382–407.
- [26] M.D. Hoffman, A. Gelman, et al., The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* 15 (1) (2014) 1593–1623.
- [27] T. Ishigami, T. Homma, An importance quantification technique in uncertainty analysis for computer models, in: *Proceedings. First International Symposium on Uncertainty Modeling and Analysis*, 1990, pp. 398–403.
- [28] G. James, D. Witten, T. Hastie, R. Tibshirani, *An Introduction to Statistical Learning*, vol. 112, Springer, 2013.
- [29] S. Karlin, *Total Positivity*, vol. 1, Stanford University Press, 1968.
- [30] A. Keese, H.G. Matthies, Sparse quadrature as an alternative to Monte Carlo for stochastic finite element techniques, *Proc. Appl. Math. Mech.* 3 (2003) 493–494.
- [31] M. Köppel, F. Franzelin, I. Kröker, S. Oladyshkin, G. Santin, D. Wittwar, A. Barth, B. Haasdonk, W. Nowak, D. Pflüger, C. Rohde, Comparison of data-driven uncertainty quantification methods for a carbon dioxide storage benchmark scenario, *Comput. Geosci.* (2019).
- [32] M. Köppel, F. Franzelin, I. Kröker, S. Oladyshkin, D. Wittwar, G. Santin, A. Barth, B. Haasdonk, W. Nowak, D. Pflüger, C. Rohde, Datasets and executables of data-driven uncertainty quantification benchmark in carbon dioxide storage, in: *Zenodo*, 2017.
- [33] M. Köppel, I. Kröker, C. Rohde, Intrusive uncertainty quantification for hyperbolic-elliptic systems governing two-phase flow in heterogeneous porous media, *Comput. Geosci.* 21 (4) (2017) 807–832.
- [34] A. Kurganov, G. Petrova, Central-upwind schemes on triangular grids for hyperbolic systems of conservation laws, *Numer. Methods Partial Differ. Equ.* 21 (3) (2005) 536–552.
- [35] H. Li, D. Zhang, Probabilistic collocation method for flow in porous media: comparisons with other stochastic methods, *Water Resour. Res.* 43 (9) (2007) 1–13.
- [36] H.G. Matthies, A. Keese, Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations, *Comput. Methods Appl. Mech. Eng.* 194 (12–16) (2005) 1295–1331.
- [37] F. Mohammadi, E. Eggenweiler, B. Flemisch, S. Oladyshkin, I. Rybak, M. Schneider, K. Weishaupt, Uncertainty-aware validation benchmarks for coupling free flow and porous-medium flow, arXiv preprint 2022.
- [38] E.H. Moore, On the reciprocal of the general algebraic matrix, *Bull. Am. Math. Soc.* 26 (1920) 394–395.
- [39] R. Mura, T. Ghisu, S. Shahpar, Least squares approximation-based polynomial chaos expansion for uncertainty quantification and robust optimization in aeronautics, in: *AIAA Aviation Forum*, 2020.

- [40] S. Oladyskhin, aPC Matlab toolbox: data-driven arbitrary polynomial chaos, <https://www.mathworks.com/matlabcentral/fileexchange/72014-apc-matlab-toolbox-data-driven-arbitrary-polynomial-chaos>, 2022.
- [41] S. Oladyskhin, H. Class, R. Helmig, W. Nowak, An integrative approach to robust design and probabilistic risk assessment for CO₂ storage in geological formations, *Comput. Geosci.* 15 (3) (2011) 565–577.
- [42] S. Oladyskhin, H. Class, W. Nowak, Bayesian updating via Bootstrap filtering combined with data-driven polynomial chaos expansions: methodology and application to history matching for carbon dioxide storage in geological formations, *Comput. Geosci.* 17 (4) (2013) 671–687.
- [43] S. Oladyskhin, F. De Barros, W. Nowak, Global sensitivity analysis: a flexible and efficient framework with an example from stochastic hydrogeology, *Adv. Water Resour.* 37 (2012) 10–22.
- [44] S. Oladyskhin, W. Nowak, Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion, *Reliab. Eng. Syst. Saf.* 106 (2012) 179–190.
- [45] S. Oladyskhin, W. Nowak, Incomplete statistical information limits the utility of high-order polynomial chaos expansions, *Reliab. Eng. Syst. Saf.* 169 (2018) 137–148.
- [46] S. Oladyskhin, P. Schroeder, H. Class, W. Nowak, Chaos expansion based Bootstrap filter to calibrate CO₂ injection models, *Energy Proc.* 40 (2013) 398–407.
- [47] Q. Pan, X. Qu, L. Liu, D. Dias, A sequential sparse polynomial chaos expansion using Bayesian regression for geotechnical reliability estimations, *Int. J. Numer. Anal. Methods Geomech.* 44 (6) (2020) 874–889.
- [48] F. Papi, F. Balduzzi, G. Ferrara, A. Bianchini, Uncertainty quantification on the effects of rain-induced erosion on annual energy production and performance of a multi-mw wind turbine, *Renew. Energy* 165 (2021) 701–715.
- [49] F. Pavone, J. Piironen, P.-C. Bürkner, A. Vehtari, Using reference models in variable selection, *arXiv preprint* 2020.
- [50] R. Penrose, On Best Approximate Solutions of Linear Matrix Equations, *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 52, Cambridge University Press, 1956, pp. 17–19.
- [51] J. Piironen, M. Paasiniemi, A. Vehtari, Projective inference in high-dimensional problems: prediction and feature selection, *Electron. J. Stat.* 14 (1) (2020) 2155–2197, Publisher: Institute of Mathematical Statistics and Bernoulli Society.
- [52] J. Piironen, A. Vehtari, Sparsity information and regularization in the horseshoe and other shrinkage priors, *Electron. J. Stat.* 11 (2) (2017) 5018–5051.
- [53] G. Poëtte, B. Després, D. Lucor, Uncertainty quantification for systems of conservation laws, *J. Comput. Phys.* 228 (7) (2009) 2443–2467.
- [54] N.G. Polson, J.G. Scott, Shrink globally, act locally: sparse Bayesian regularization and prediction, in: J.M. Bernardo, M.J. Bayarri, J.O. Berger, A.P. Dawid, D. Heckerman, A.F.M. Smith, M. West (Eds.), *Bayesian Statistics 9*, Oxford University Press, 2011, pp. 501–538.
- [55] J. Red-Horse, A. Benjamin, A probabilistic approach to uncertainty quantification with limited information, *Reliab. Eng. Syst. Saf.* 85 (1) (2004) 183–190.
- [56] G. Riutort-Mayol, P.-C. Bürkner, M.R. Andersen, A. Solin, A. Vehtari, Practical Hilbert space approximate Bayesian Gaussian processes for probabilistic programming, *Stat. Comput.* 33 (1) (Dec. 2022) 17.
- [57] C. Runge, Über empirische Funktionen und die Interpolation zwischen äquidistanten Ordinaten, *Z. Angew. Math. Phys.* 46 (224–243) (1901) 20.
- [58] S. Scheurer, A. Schäfer Rodrigues Silva, F. Mohammadi, J. Hommel, S. Oladyskhin, B. Flemisch, W. Nowak, Surrogate-based Bayesian comparison of computationally expensive models: application to microbially induced calcite precipitation, *Comput. Geosci.* 25 (6) (2021) 1899–1917.
- [59] Q. Shao, A. Younes, M. Fahs, T.A. Mara, Bayesian sparse polynomial chaos expansion for global sensitivity analysis, *Comput. Methods Appl. Mech. Eng.* 318 (2017) 474–496.
- [60] J. Shohat, J. Tamarkin, *The Problem of Moments*, *Mathematical Surveys No. 1*, American Mathematical Society, New York, 1943, p. 1950.
- [61] W.M. Siebert, On the determinants of moment matrices, *Ann. Stat.* 17 (2) (1989) 711–721.
- [62] I.M. Sobol', On the distribution of points in a cube and the approximate evaluation of integrals, *Ž. Vychisl. Mat. Mat. Fiz.* 7 (4) (1967) 784–802.
- [63] I.M. Sobol', On sensitivity estimation for nonlinear mathematical models, *Mat. Model.* 2 1 (1990) 112–118.
- [64] I.M. Sobol', *Theorems and Examples on High Dimensional Model Representation*, *Reliab. Eng. Syst. Saf.* 79 (2) (2003) 187–193, Publisher: Elsevier.
- [65] Stan Development Team, *Stan Modeling Language: User's Guide and Reference Manual*, 2022.
- [66] T.J. Stieltjes, Quelques recherches sur la théorie des quadratures dites mécaniques, *Ann. Sci. Éc. Norm. Supér.* 1 (1884) 409–426.
- [67] B. Sudret, Global sensitivity analysis using polynomial chaos expansions, *Reliab. Eng. Syst. Saf.* 93 (7) (2008) 964–979.
- [68] T.J. Sullivan, *Introduction to Uncertainty Quantification*, *Texts in Applied Mathematics*, vol. 63, Springer, Cham, 2015.
- [69] A. Tarakanov, A.H. Elsheikh, Regression-based sparse polynomial chaos for uncertainty quantification of subsurface flow models, *J. Comput. Phys.* 399 (2019).
- [70] A.N. Tikhonov, V.Y. Arsenin, *Solutions of Ill-Posed Problems*, *Scripta Series in Mathematics*, V. H. Winston & Sons, Washington, D.C., 1977, John Wiley & Sons, New York-Toronto, Ont.-London, translated from the Russian, Preface by translation editor Fritz John.
- [71] M.E. Tipping, Sparse Bayesian learning and the relevance vector machine, *J. Mach. Learn. Res.* 1 (Jun 2001) 211–244.
- [72] M.-N. Tran, D.J. Nott, C. Leng, The predictive Lasso, *Stat. Comput.* 22 (5) (2012) 1069–1084.
- [73] S. Van Erp, D.L. Oberski, J. Mulder, Shrinkage priors for Bayesian penalized regression, *J. Math. Psychol.* 89 (2019) 31–50.
- [74] A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, P.-C. Bürkner, Rank-normalization, folding, and localization: an improved \hat{R} for assessing convergence of MCMC (with discussion), *Bayesian Anal.* 16 (2) (2021) 667–718.
- [75] J. Villadsen, M.L. Michelsen, *Solution of Differential Equation Models by Polynomial Approximation*, Prentice-Hall, 1978.
- [76] H. Wendland, *Scattered Data Approximation*, vol. 17, Cambridge University Press, 2004.
- [77] N. Wiener, The homogeneous chaos, *Am. J. Math.* 60 (4) (1938) 897–936.
- [78] D. Wirtz, B. Haasdonk, A vectorial kernel orthogonal greedy algorithm, *Dolom. Res. Notes Approx.* 6 (2013).
- [79] J.A. Witteveen, H. Bijl, Modeling arbitrary uncertainties using Gram-Schmidt polynomial chaos, in: *44th AIAA Aerospace Sciences Meeting and Exhibit*, 2006.
- [80] S.N. Wood, *Generalized Additive Models: an Introduction with R*, Chapman and Hall/CRC, London, 2017.
- [81] D. Xiu, G.E. Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, *SIAM J. Sci. Comput.* 24 (2) (2002) 619–644.
- [82] D. Xiu, G.E. Karniadakis, Modeling uncertainty in flow simulations via generalized polynomial chaos, *J. Comput. Phys.* 187 (2003) 137–167.
- [83] Y. Yao, A. Vehtari, D. Simpson, A. Gelman, Yes, but did it work?: evaluating variational inference, in: *International Conference on Machine Learning*, 2018, pp. 5581–5590.
- [84] L. Zeng, L. Shi, D. Zhang, L. Wu, A sparse grid based Bayesian method for contaminant source identification, *Adv. Water Resour.* 37 (2012) 1–9.
- [85] S. Zeng, X. Duan, J. Chen, L. Yan, Optimized sparse polynomial chaos expansion with entropy regularization, *Adv. Aerodyn.* 4 (2022) 1.
- [86] Y. Zhang, Y. Liu, G. Pau, S. Oladyskhin, S. Finsterle, Evaluation of multiple reduced-order models to enhance confidence in global sensitivity analyses, *Int. J. Greenh. Gas Control* 49 (2016) 217–226.
- [87] Y. Zhang, B.P. Naughton, H.D. Bondell, B.J. Reich, Bayesian regression using a prior on the model fit: the R2-D2 shrinkage prior, *J. Am. Stat. Assoc.* (2020).