

Homogeneous chaos basis adaptation for design optimization under uncertainty: Application to the oil well placement problem

CHARANRAJ THIMMISSETTY,¹ PANAGIOTIS TSILIFIS,² AND ROGER GHANEM²

¹Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, California, USA

²Sonny Astani Department of Civil and Environmental Engineering, University of Southern California, Los Angeles, California, USA

(RECEIVED December 19, 2016; ACCEPTED February 19, 2017)

Abstract

A new method is proposed for efficient optimization under uncertainty that addresses the curse of dimensionality as it pertains to the evaluation of probabilistic objectives and constraints. A basis adaptation strategy previously introduced by the authors is integrated into a design optimization framework that construes the optimization cost function as the quantity of interest and computes stochastic adapted bases as functions of design space parameters. With these adapted bases, the stochastic integrations at each design point are evaluated as low-dimensional integrals (mostly one dimensional). The proposed approach is demonstrated on a well-placement problem where the uncertainty is in the form of a stochastic process describing the permeability of the subsurface. An analysis of the method is carried out to better understand the effect of design parameters on the smoothness of the adaptation isometry.

Keywords: Adapted Bases; Karhunen–Loeve Expansion; Optimization Under Uncertainty; Optimal Well Placement; Polynomial Chaos Expansions

1. INTRODUCTION

Design optimization deals with exploring multiple instances of a design with the aim of identifying the best among them. In the presence of uncertainty, each of these instances entails the exploration of several nominally equivalent realizations intended to cover the range of anticipated scatter. In cases of interest to the present study, this uncertainty and the associated scatter are described as probabilistic measures on parameters or environmental conditions associated with each design instance. In the context of model-based optimization, each instance of the design involves the evaluation of a model, typically within a computational setting. Under conditions of uncertainty, the model is probabilistic, and each evaluation involves the exploration of the set of possible events, thus considerably increasing the number of function evaluations required to identify the optimal design. Clearly, the characterization of the probability measure on parameter space is a critical step in proceeding with any optimization under uncertainty (OUU) problem. That step itself can be

cast as an iteration over a model, requiring the repetitive evaluation of an objective function or a likelihood function. This calibration or update step notwithstanding, the added computational burden is a key challenge when introducing uncertainty to a computational problem. For large-scale computational models, this challenge is on the critical path of credible model-supported decision making and design.

Typically, the objective function in an OUU problem is the expectation of a quantity of interest (QoI), Q , with respect to the probability measure of the random variables that describe the uncertainty in the problem. This expectation could, for instance, represent average performance, or the probability of some particular event, often associated with failure. Monte Carlo sampling techniques are standard tools for evaluating these expectations, and several optimization procedures have been integrated as sampling strategies resulting in algorithms such as stochastic gradient descent (Bottou, 2010), sample average approximation (Kleywegt, 2002), Robbins–Monro (Robbins & Monro, 1951), the simultaneous perturbation stochastic approximation (Spall, 1952), and random search (Heyman & Sobel, 2003). Although adapted to specific problems, and generally more efficient than standard sampling strategies, these algorithms can still be prohibitive

Reprint requests to: Roger Ghanem, Sonny Astani Department of Civil and Environmental Engineering, University of Southern California, Los Angeles, CA 90089, USA. E-mail: ghanem@usc.edu

for problems with large uncertainties or expensive function evaluations. In addition to uncertainty in the objective functions, uncertainty can also be present in the constraints. In such *chance-constrained problems*, the task of verifying the constraints incurs a significant computational burden, and tolerances on satisfying probabilistic constraints must be introduced. Robust optimization problems have been introduced to alleviate the computational burden by replacing the stochastic problem with a deterministic equivalent (Ben-Tal et al., 2009). This equivalence often pertains to low-order statistics and may fail to capture extreme behavior, relevant in many design situations.

Reservoir modeling and subsurface hydrology are two of the areas where design optimization problems are of particular interest and where the computational model involved is usually very expensive to solve as a result of the complex multiphase/multicomponent physics of flow and transport in porous media that give rise to coupled systems of equations and the fine scale discretization of the domain that is imposed on them. As such, the need for efficient methods in related design optimization problems is crucial, and an extended literature has been devoted to it. The objectives are typically to infer subsurface properties such as permeability and porosity or to optimize decision-making strategies such as oil production and soil remediation. The design problem that often arises in this case is that of determining the optimal locations to place the wells in an oil reservoir or the boreholes in an aquifer such that some oil production-related function or some information-based criterion is maximized (Tsilifis et al., 2017). Variations on the oil well placement problem using optimization algorithms (Rian & Hage, 1994; Beckner & Song, 1995), including uncertainty (Güyagüler & Horne, 2001; Güyagüler, 2002; Yeten et al., 2002; Bangerth et al., 2006), have explored well placement scenarios, details of the flow model, or specifics of the cost function.

One way to cope with the computational challenges described above is to replace the forward model with a cheap-to-evaluate surrogate such as polynomial chaos (PCE; Eldred, 2009; Keshavarzzadeh et al., 2016) or Gaussian process (Bilonis & Zabarar, 2013). Once such a surrogate is available, the computational cost required to perform an optimization algorithm is reduced dramatically; however, there is no free lunch. Constructing the surrogate will often be an expensive procedure and can become even prohibitive for high-dimensional inputs. In this work, we develop an alternative framework based on polynomial chaos expansions that reduces further the cost of constructing the surrogate and therefore can be applicable in higher dimensional settings where standard nonintrusive methods would fail. The idea is based on constructing adapted polynomial chaos expansions that was first introduced in Tipireddy and Ghanem (2014), which relies on rotating the Gaussian input of the chaos expansion and constructing a reduced expansion dependent on only a few Gaussian inputs, thus reducing the input dimension and requiring a significantly smaller number of forward evaluations for the estimation of its coefficients. The novel contribution of the

present work is in considering the optimization objective function as the QoI used in adaptation. In this manner, an adaptation is constructed for each iteration in the design space. This procedure permits the efficient evaluation of probabilistic operations (expectations or probabilities of failure) at each design point. We note that the proposed formalism is very different from sparse PCE procedures where a subset of the initial random variables is used in representing the solution. Substantial additional reduction is achieved in the present work by seeking direction in the parameter space along which the solution is concentrated. Typically, all of the original components are active along these directions and will thus contribute to the final solution.

The paper is structured as follows: Subsection 2.1 describes the design optimization objectives. In Subsection 2.2 we present the polynomial chaos properties and the basis adaptation concept along with the computational algorithm for its implementation. In Subsection 2.3 we describe how the polynomial chaos basis adaptation framework can be applied in a design optimization problem and a computational algorithm is also provided. In Section 3 we apply our methodology in a standard oil well placement problem that consists of a two-dimensional oil reservoir with random permeability where one production well and one injection well need to be placed such that threshold above which the cumulative production rate will be with a high confidence level, is maximized. Conclusions are presented in Section 4.

2. METHODOLOGY

2.1. Design optimization objectives

We present the general objective of the design optimization problem under uncertainty that we are investigating. We consider a scalar QoI $Q(\mathbf{p}, \boldsymbol{\theta}) \in \mathbb{R}$ that can be thought of as the outcome of a physical model or a computer code that depends on design parameters $\mathbf{p} \in \mathbb{R}^n$ and some random parameters $\boldsymbol{\theta} \in \mathbb{R}^m$ that are defined on some probability space (Ω, \mathcal{F}, P) and follow some probability distribution that is assumed to be known. Let $F_Q(q) = P(Q(\mathbf{p}, \boldsymbol{\theta}) \leq q)$ be the cumulative probability distribution of the quantity of interest evaluated at \mathbf{p} . It is a common task to seek the estimation of the value q_α that provides a one-sided confidence at level $\alpha \in (0, 1)$ (typically small). Thus, denoting the cumulative distribution function with F_Q and the quantile with q_α ,

$$F_Q(q) = P(Q(\mathbf{p}, \boldsymbol{\theta}) \leq q), \quad q_\alpha = F_Q^{-1}(\alpha), \quad (1)$$

and the objective is to find q_α such that $Q(\mathbf{p}, \boldsymbol{\theta})$ is above q_α with probability $1 - \alpha$. Clearly, q_α depends on \mathbf{p} . Motivated by this, we are interested in solving the optimization problem of the form

$$\mathbf{p}^* = \arg \max_{\mathbf{p} \in \mathcal{X}} J(\mathbf{p}), \quad J(\mathbf{p}) = q_\alpha, \quad (2)$$

where $\mathbf{X} \subset \mathbb{R}^n$ is the domain of the design parameters (typically bounded). The dependence of q_α on \mathbf{p} comes through the distribution F_Q . For evaluation of the objective function subject to minimization, knowledge of this distribution is required. In all but the simplest cases, it will not be possible to compute F_Q analytically and only an estimate of it will be feasible. This could typically be estimated based on empirical procedures, for instance, a histogram obtained from Monte Carlo (MC) sampling or similar sampling-based techniques such as kernel density estimation (Scott, 2015). The main drawback in these approaches is that the number of MC estimates required for such estimations can be prohibitive in terms of the computational resources necessary to carry out such tasks. Among other characteristics that affect the required number of samples are the complexity of the computational model involved and the dimensionality and the distribution of $\boldsymbol{\theta}$. These both affect the structure of the probability distribution that we seek to estimate. In the case where the computational model involved is replaced by a surrogate, the computational burden of MC sampling does not necessarily disappear as the cost of actually constructing the surrogate can be equally high, if not higher. Our ultimate goal in this work is to provide an alternative methodology that attempts to reduce the cost of estimating the distribution of F_Q that is applicable in the case where the model uncertainties can be written as functions of Gaussian noise and the QoI admits a homogeneous chaos expansion.

2.2. Polynomial chaos basis adaptation

2.2.1 Homogeneous chaos

In the previous paragraph, we introduced uncertainty in our QoI through the random vector $\boldsymbol{\theta}$ whose distribution is assumed known. In order to develop our methodology, it is further necessary to assume that $\boldsymbol{\theta}$ can be written as a function of d uncorrelated Gaussian variables $\boldsymbol{\xi} \in \mathbb{R}^d$ and write $\boldsymbol{\theta} = \mathbf{g}(\boldsymbol{\xi})$. Note that d can be considered to be infinite, but here we will only consider the finite case. Clearly Q will then be also written as a function of the same Gaussian variables as $Q(\mathbf{p}, \mathbf{g}(\boldsymbol{\xi}))$. For the sake of simplicity we will denote $Q(\mathbf{p}, \boldsymbol{\xi})$, ignoring the dependence through \mathbf{g} when there is no confusion. One necessary condition for the above assumption to hold is for $\boldsymbol{\theta}$ to have finite second moment and be measurable with respect to the σ -algebra $\mathcal{F}(\boldsymbol{\xi})$, generated by $\boldsymbol{\xi}$. Then, the Cameron–Martin theorem (Cameron & Martin, 1947) ensures that $\boldsymbol{\theta}$ admits a representation of the form

$$\boldsymbol{\theta} = \sum_{|\alpha|=0}^p \boldsymbol{\theta}_\alpha \Psi_\alpha(\boldsymbol{\xi}), \quad j = 1, \dots, m, \quad \boldsymbol{\theta}_\alpha \in \mathbb{R}^m, \quad (3)$$

that is, mean-square convergent as $p \rightarrow \infty$, $\alpha = (\alpha_1, \dots, \alpha_d)$, $|\alpha| = \sum_{i=1}^d \alpha_i$, and $\Psi_\alpha(\boldsymbol{\xi})$ are multidimensional Hermite

polynomials of order α defined as

$$\Psi(\boldsymbol{\xi}) = \prod_{i=1}^d \psi_{\alpha_i}(\xi_i). \quad (4)$$

Here, $\psi_{\alpha_i}(\xi_i)$ are the normalized one-dimensional Hermite polynomials of order α_i . Assuming that q has also finite second moment, it is clear that a similar representation can be obtained:

$$Q(\mathbf{p}, \boldsymbol{\xi}) = \sum_{|\alpha|=0}^p Q_\alpha(\mathbf{p}) \Psi_\alpha(\boldsymbol{\xi}). \quad (5)$$

The total number P of terms in the expansion is given by

$$P + 1 = \frac{(d + p)!}{d!p!}. \quad (6)$$

That the functions $\{\Psi_\alpha\}_\alpha$ form an orthonormal basis on the space of square integrable random variables $L^2(\Omega, \mathcal{F}(\boldsymbol{\xi}), P)$ gives a convenient expression for the coefficients

$$Q_\alpha(\mathbf{p}) = \langle Q(\mathbf{p}, \cdot), \Psi_\alpha \rangle, \quad (7)$$

which can be estimated using several techniques, the most popular of which are Galerkin projections (Ghanem & Spanos, 1991; Le Maître, 2002) and pseudo-spectral projections (Reagan et al., 2003) that are based on numerical integration using quadrature rules, known also as intrusive and nonintrusive approaches, respectively. In the application included in this work, we will focus on the second approach, which allows treating the computational model as a black box. Once the coefficients $Q_\alpha(\mathbf{p})$ are estimated, one can reconstruct the density function of Q through samples of $\boldsymbol{\xi}$, as evaluation of the polynomial series is easy and fast and one can proceed in solving optimization problems of the form in Equation (2). However, the number of evaluations required for estimation of the coefficients grows exponentially fast as a function of the dimensionality of $\boldsymbol{\xi}$ and quickly becomes prohibitive or at least not cheaper, for the purpose of solving Equation (2), than direct sampling using the forward model. One way to address this issue, as we demonstrate in the next section, is to adapt the Gaussian basis of the expansion in a fashion that reduces the dimensionality and consequently requires fewer evaluations.

2.2.2 Basis adaptation

Let $\mathbf{A} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be an isometry on \mathbb{R}^d , or equivalently, an orthogonal matrix ($\mathbf{A}\mathbf{A}^T = \mathbf{I}_d$) and define $\boldsymbol{\eta} = \mathbf{A}\boldsymbol{\xi}$. It follows (Janson, 1999) that the components of $\boldsymbol{\eta}$ are also uncorrelated standard Gaussian variables that span the same Gaussian Hilbert space with $\boldsymbol{\xi}$ and generate the same σ -algebra $\mathcal{F}(\boldsymbol{\xi})$; therefore, a polynomial chaos expansion of any element $Q \in L^2(\Omega, \mathcal{F}(\boldsymbol{\xi}), P)$ can be taken with respect to the new

germ $\boldsymbol{\eta}$ as

$$Q^A(\mathbf{p}, \boldsymbol{\eta}) = \sum_{\beta=0}^p Q_{\beta}^A(\mathbf{p})\psi_{\beta}^A(\boldsymbol{\xi}), \tag{8}$$

where we use the superscript on Q^A and Q_{β}^A to indicate that the new basis has been obtained after applying rotation \mathbf{A} on $\boldsymbol{\xi}$ and $\psi_{\beta}^A(\boldsymbol{\xi}) = \psi_{\beta}(\mathbf{A}\boldsymbol{\xi})$. Then, we have that $Q^A(\mathbf{p}, \boldsymbol{\eta}) = Q(\mathbf{p}, \boldsymbol{\xi})$ almost surely and the coefficients of the new expansion can be written in terms of $\{Q_{\alpha}(\mathbf{p})\}_{\alpha=0}^p$ as

$$Q_{\beta}^A(\mathbf{p}) = \sum_{|\alpha|=0}^p Q_{\alpha}(\mathbf{p})\langle \psi_{\alpha}, \psi_{\beta}^A \rangle. \tag{9}$$

We have $\langle \psi_{\alpha}, \psi_{\beta}^A \rangle = 0$ for $|\alpha| \neq |\beta|$, while for $|\alpha| = |\beta|$ the inner products can be computed explicitly (Tsilifis & Ghansem, 2016). The motivation behind applying a change of basis as described above is that, by choosing an appropriate matrix \mathbf{A} , the new expansion can concentrate the dominant effects of the original expansion to fewer coefficients of Q^A and thus provide sparse representations depending on only a few components of $\boldsymbol{\eta} = (\eta_1, \dots, \eta_d)$, say, $\boldsymbol{\eta}_{d_0} = (\eta_1, \dots, \eta_{d_0})$ with $d_0 < d$, and therefore, we can discard terms of the adapted expansion via projection (see Tipieddy & Ghanem, 2014; Tsilifis & Ghanem, 2016, for details as well as several popular choices of \mathbf{A}). If we let $\mathbf{J}_{p,d} = \{|\alpha| = (\alpha_1, \dots, \alpha_d) : |\alpha| \leq p\}$ be the set of d -dimensional multi-indices of order up to p and take any \mathbf{J}_{p_0,d_0} for some $p_0 < p$, $d_0 < d$, which implies that $\mathbf{J}_{p_0,d_0} \subset \mathbf{J}_{p,d}$, then for given \mathbf{A} , the reduced decomposition

$$Q^A(\mathbf{p}, \boldsymbol{\eta}_{d_0}) = \sum_{\beta \in \mathbf{J}_{p_0,d_0}} Q_{\beta}^A(\mathbf{p})\psi_{\beta}(\boldsymbol{\eta}_{d_0}) \tag{10}$$

approaches

$$Q^A(\mathbf{p}, \boldsymbol{\eta}) \stackrel{[0pt] \text{ almost surely}}{=} Q(\mathbf{p}, \boldsymbol{\xi})$$

as $p_0 \rightarrow p$ and $d_0 \rightarrow d$. One then can carefully choose p_0, d_0 such the reduced decomposition in Eq. (10) can be sufficiently close to $Q(\mathbf{p}, \boldsymbol{\xi})$.

Obtaining this expression requires the estimation of $Q_{\beta}^A(\mathbf{p})$, $\beta \in \mathbf{J}_{p_0,d_0}$. In practice, using Eq. (9) does not deliver us from the computational cost of obtaining a full polynomial chaos expansion because the coefficients $Q_{\alpha}(\mathbf{p})$ are required. Instead, once an isometry \mathbf{A} is chosen, one can compute the coefficients $Q_{\beta}^A(\mathbf{p})$ of the adapted lower dimensional expansion by following a nonintrusive implementation as described in Tipieddy and Ghanem (2014). The computational benefits of doing so is that for $d_0 = d$, the number of quadrature points required for a nonintrusive estimation of the coefficients $Q_{\beta}^A(\mathbf{p})$ is much smaller than that of the original expansion while the adapted expansion serves as an accurate estimate of our QoI. The details of the procedure are summarized in Algorithm 1.

Algorithm 1: Nonintrusive implementation

Input : Matrix \mathbf{A} , adapted expansion dimension $d_0 < d$, quadrature coefficients $\{\boldsymbol{\eta}_{d_0}^{(r)}\}_{r=1}^s$ and their weights $\{w_r\}_{r=1}^s$ corresponding to d_0 -dimensional, s -point quadrature rule.

Step 1 : Complete the vectors $\boldsymbol{\eta}_{d_0}^{(r)}$ to d -dimensional $\boldsymbol{\eta}^{(r)}$:

$$\boldsymbol{\eta}^{(r)} = \begin{pmatrix} \boldsymbol{\eta}_{d_0}^{(r)} \\ \mathbf{0} \end{pmatrix},$$

where $\mathbf{0}$ is $(d - d_0)$ -dimensional vector.

Step 2 : Compute $\boldsymbol{\xi}^{(r)} = \mathbf{A}^T \boldsymbol{\eta}^{(r)}$ for $r = 1, \dots, s$.

Step 3 : Evaluate QoI at $\{\boldsymbol{\xi}^{(r)}\}_{r=1}^s$ to obtain $\{Q(\mathbf{p}, \boldsymbol{\xi}^{(r)})\}_r^s = 1$.

Step 4 : Estimate the coefficients $Q_{\beta}(\mathbf{p})$, $\beta \in \mathbf{J}_{p_0,d_0}$ as

$$Q_{\beta}(\mathbf{p}) \approx \sum_{r=1}^s Q(\mathbf{p}, \boldsymbol{\xi}^{(r)})\psi_{\beta}(\boldsymbol{\eta}_{d_0}^{(r)})w_r.$$

Throughout this paper we restrict ourselves with obtaining \mathbf{A} by using the Gaussian adaptation. This consists of first estimating the coefficients $\{Q_{\boldsymbol{\varepsilon}_i}\}_{i=1}^d$ of the first-order polynomials $\{\psi_{\boldsymbol{\varepsilon}_i}(\boldsymbol{\xi}_i)\}_{i=1}^d$, where $\boldsymbol{\varepsilon}_i = (0, \dots, 1, \dots, 0)$, that is, the multi-index with 1 at location i and 0 everywhere else. This is achieved using a low-level quadrature rule with very few points. Next, we define the vector $\mathbf{Q}_1 = [Q_{\boldsymbol{\varepsilon}_1}, \dots, Q_{\boldsymbol{\varepsilon}_d}]$, and we set the first row of \mathbf{A} , denoted with $\mathbf{A}_{1,\cdot}$, to be

$$\mathbf{A}_{1,\cdot} = \frac{\mathbf{Q}_1}{\|\mathbf{Q}_1\|}. \tag{11}$$

The matrix \mathbf{A} can then be completed by performing a Gram-Schmidt procedure.

2.3. Design optimization using adapted polynomial chaos

The adaptation framework described in the previous paragraph can be implemented in order to obtain cheap surrogates of the QoI at different locations \mathbf{p} , which can then be used for estimation of F_Q and its inverse. The optimization problem described in Section 2.1 can then be approached as follows: first, one needs to employ an optimization algorithm to be used for maximization of $J(\mathbf{p})$. At each iteration k , the adapted chaos expansion of the QoI needs to be computed at design parameters \mathbf{p}_k , for a dimension value d_0 that achieves a satisfactory level of accuracy. Then F_Q can be estimated with MC sampling, and $F_Q^{-1}(\cdot)$ can be evaluated at α , thus giving the values of the objective function $J(\mathbf{p})$. The choice of d_0 is not known *a priori*, and it needs to be identified. One way to do this is to simply start by constructing a one-dimensional adapted expansion and then increase the dimension until some relative error of the consecutive adaptations falls below some threshold. Because

the number of evaluations required increases exponentially as a function of the dimension, the evaluations required to compute all adaptations up to dimension d_0 will be still much smaller than those required for the d -dimensional one as long as d_0 remains significantly smaller than d . Algorithm 2 displays the pseudocode with all the steps required for one evaluation of $J(\mathbf{p})$.

Algorithm 2: Evaluation of $J(\mathbf{p})$.

Input : Design parameters \mathbf{p} , quadrature coefficients $\{\xi^{(r)}\}_{r=1}^s$ and their weights $\{w_r\}_{r=1}^s$ corresponding to d -dimensional, s -point quadrature rule.

Step 1 : Evaluate QoI at $\{\xi^{(r)}\}_{r=1}^s$ to obtain $\{Q(\mathbf{p}, \xi^{(r)})\}_{r=1}^s$ and estimate coefficients $Q_{\epsilon i}(\mathbf{p})_{i=1}^d$

$$Q_{\epsilon i}(\mathbf{p}) \approx \sum_{r=1}^s Q(\mathbf{p}, \xi^{(r)}) \xi_i^{(r)} w_r$$

Step 2 : Construct **A**.

Initialize: Set $d_0 = 1$.

repeat

Obtain quadrature coefficients $\{\eta_{d_0}^{(r)}\}_{r=1}^s$ and their weights $\{w_r\}_{r=1}^s$

Perform Algorithm 1

Estimate the density function of $Q(\mathbf{p}, \eta_{d_0})$

Generate enough realizations of $Q(\mathbf{p}, \xi^{(r)})$ to estimate $q_{0.05}$

$d_0 = d_0 + 1$

until relative change in $q_{0.05}$ is less than ϵ

The outside loop, where d_0 is being incremented, may seem to present an additional workload beyond standard OOU procedures. It is noted, however, that for values of d_0 well below the value of d , the computational burden associated with evaluations for dimensions up to and including d_0 is still a fraction of the cost for an evaluation at dimension d . While the numerical value of the tolerance, ϵ , can be selected to become smaller as d_0 is approached, in the present paper we fix ϵ to a predetermined value of 0.05.

3. APPLICATION: OIL WELL PLACEMENT PROBLEM

3.1. The oil reservoir model

To show the computational efficiency of the proposed method, we consider a multiphase flow in a petroleum reservoir. The governing equations of the corresponding physics describing black oil flow through a heterogeneous porous me-

dium are the following (Chen et al., 2008):

$$\begin{aligned} \frac{\partial}{\partial t} \left[\frac{\phi \rho_{W_s}}{B_w} S_w \right] + \nabla \cdot \left(\frac{\rho_{W_s}}{B_w} \vec{u}_w \right) - q_w &= 0, \\ \frac{\partial}{\partial t} \left[\frac{\phi \rho_{O_s}}{B_o} S_o \right] + \nabla \cdot \left(\frac{\rho_{O_s}}{B_o} \vec{u}_o \right) - q_o &= 0, \\ \frac{\partial}{\partial t} \left[\phi \left(\frac{R_{so} \rho_{G_s}}{B_o} S_o + \frac{\rho_{G_s}}{B_g} S_g \right) \right] + \nabla \cdot \left(\frac{R_{so} \rho_{G_s}}{B_o} u_o + \frac{\rho_{G_s}}{B_g} u_g \right) - q_G &= 0. \end{aligned} \quad (12)$$

where ϕ denotes the porosity; $S_\alpha, \mu_\alpha, p_\alpha, B_\alpha, \rho_{\alpha_s}$, and q_α are the saturation, viscosity, pressure, formation volume factor, density, and volumetric rate at each phase α (water, gas, and oil); R_{so} is the gas solubility; u_w, u_g , and u_o are the respective Darcy velocities obtained by solving Darcy’s law for each phase

$$\begin{aligned} u_\alpha &= - \frac{k_{r\alpha} k(\mathbf{x}, \omega)}{\mu_\alpha} (\nabla p_\alpha - \rho_\alpha \gamma \nabla z) \\ \alpha &= o, w, g, \quad \mathbf{x} \in \Gamma, \quad \omega \in \Omega, \end{aligned} \quad (13)$$

where ρ_α and $k_{r\alpha}$ are the respective mass density and relative permeabilities of each α phase, γ is the magnitude of the gravitational acceleration, z is the depth, and $k(\mathbf{x})$ is the intrinsic permeability. The relative permeabilities $k_{r\alpha}$ are functions of saturation, and their values generally lie between 0 and 1. They are obtained either experimentally or by using analytical expressions such as the Corey (1994) and Stone (1970, 1973) models. In the above equations, fluid injection wells and production wells are represented by point sources and sinks whenever they are present (Peaceman, 2000). Uncertainty is introduced in our model by considering the absolute permeabilities $k(\mathbf{x})$ to be log-normally distributed, that is, $\log k(\mathbf{x}) : \mathbf{N}(\mathbf{m}(\mathbf{x}), \mathbf{C}(\mathbf{x}, \mathbf{x}))$ for all \mathbf{x} . Here we use the SPE-10 data set (Christie & Blunt, 2001), which is discretized in a $60 \times 220 \times 85$ regular Cartesian grid with real range being $600 \times 2200 \times 170$ (ft.)³. We treat the 85 layers as different realizations of our two-dimensional domain to construct our sample mean and sample covariance matrix to be used as our $\mathbf{m}(\mathbf{x})$ and $\mathbf{C}(\mathbf{x}, \mathbf{x}')$, respectively. Next we parameterize our random field by considering a Karhunen–Loeve expansion of $k(\mathbf{x})$ as

$$k(\mathbf{x}) = \mathbf{m}(\mathbf{x}) + \sum_{i=1}^{20} \sqrt{\lambda_i} \phi_i(\mathbf{x}) \xi_i, \quad (14)$$

truncated up to 20 terms. The mean and variance of the generated log-permeability field are shown in Figure 1, and two sample realizations drawn from the distribution are shown in Figure 2. The permeability values used in our simulations are measured in mDarcy. We should note that the specific nature of the distribution is not relevant for the applicability of the proposed method. Under general conditions on their moments, most data-driven distributions can be represented as polynomials of Gaussian random variables.




Fig. 1. Mean (left) and variance (right) of the log-permeability field.

3.2. Oil production as a QoI

The quantity of interest $Q(\mathbf{p}, \boldsymbol{\theta})$ considered in this example is taken to be the cumulative oil production over $T = 2000$ days, where the design parameters $\mathbf{p} = (x_1^i, x_2^i, x_1^p, x_2^p)$ are the locations (x_1^i, x_2^i) and (x_1^p, x_2^p) of one injection and one production well, respectively, and $\boldsymbol{\theta}$ is taken here to be the random permeability field. The cumulative production is defined as

$$Q(\mathbf{p}, \boldsymbol{\theta}) = \int_E \int_T u_o(\mathbf{p}, \boldsymbol{\theta}) \times d\mathbf{E} dt, \tag{15}$$

where $u_o(\mathbf{p}, \boldsymbol{\theta})$ is the Darcy velocity of the oil phase, E is the cross-sectional area of the wellbore, and $\mathbf{E} = E\hat{n}$ is the vector

area, where \hat{n} is the unit vector normal to the area. The objective function that we are interested in maximizing is

$$J(\mathbf{p}) = q_\alpha = F_Q^{-1}(\alpha). \tag{16}$$

3.3. Optimization results

Because the objective function considered here is complex in nature and its theoretical properties are not known, it can exhibit nonconvex behavior. In order to avoid exploring only an area around the local minimum, we start our convex optimization algorithm with an initial guess obtained from the brute force grid parameter study. For this study, the parameter space \mathbf{p} is divided into four evenly spaced zones $(x_{1g}^i = \{0, 200, 400, 600\})$,




Fig. 2. Sample realizations of the log-permeability field.

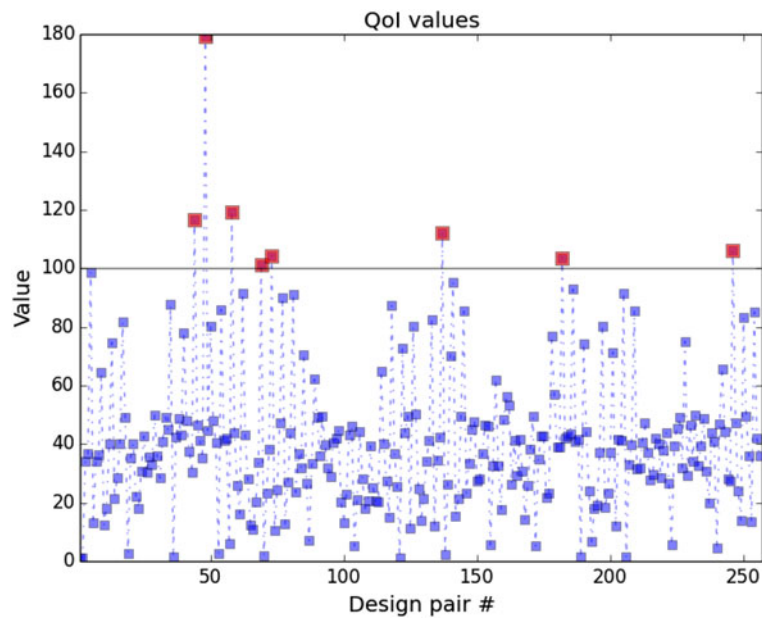


Fig. 3. Values of q_α at the 256 design points. The values with red “W” are above 100.

$x_{2g}^i = \{0, 733.3, 1466.7, 2200\}$, $x_{1g}^p = \{0, 200, 400, 600\}$, $x_{2g}^p = \{0, 733.3, 1466.7, 2200\}$, and the objective function is evaluated at the corresponding grid design points $\mathbf{p}_g = (x_{1g}^i \times x_{2g}^i \times x_{1g}^p \times x_{2g}^p)$. The values of the objective function evaluated at the 256 grid design points \mathbf{p}_g are shown in Figure 3,

highlighting with red, the ones that give oil production above 100 with probability 0.93. The possible locations of the wells where $J(\mathbf{x})$ are evaluated and the location of the well combinations resulting in maximum q_α are shown in Figure 4 (left and middle). The largest value of the objective function corre-

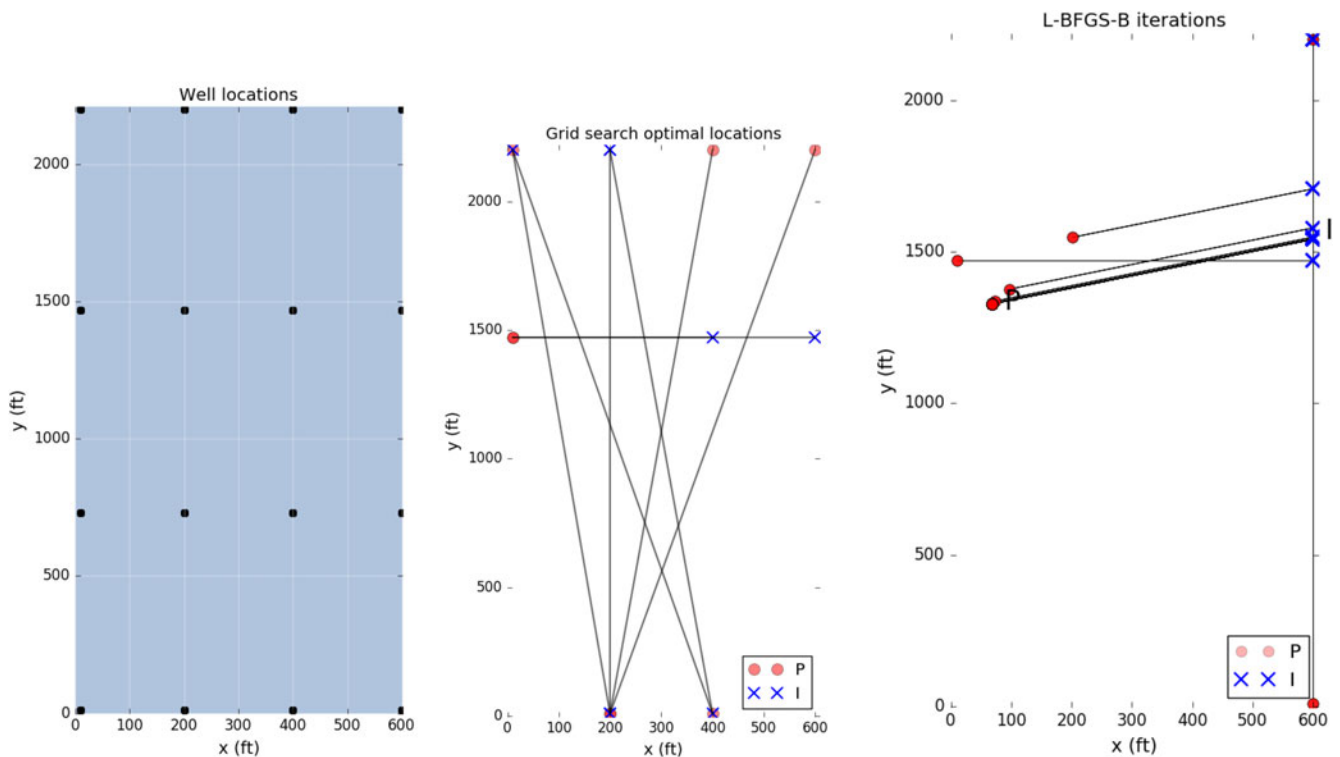


Fig. 4. Left: Injection and production well locations generating 256 possible combinations. Middle: 8 pairs of well locations resulting in $q_\alpha > 100$. Right: Pairs of well locations produced by running limited memory Broyden–Fletcher–Goldfarb–Shanno optimization.

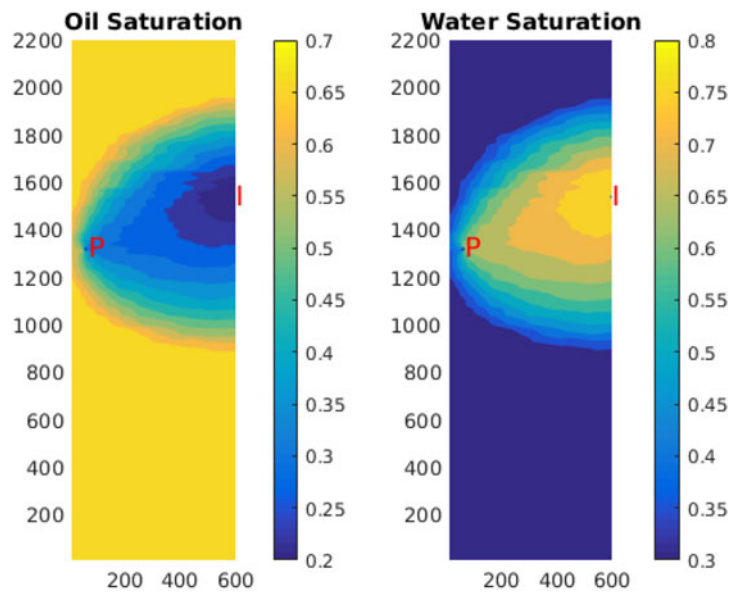


Fig. 5. Realization of the oil saturation and water saturation on the 2000th day at the optimal design point using mean permeability.

sponds to the design point $(0, 146.67, 60, 146.67)$, which we use as the initial guess to run the limited memory Broyden–Fletcher–Goldfarb–Shanno algorithm (Byrd et al., 1995) that can perform bound constrained optimization, in order to find the optimal design point. The well locations generated from iterations of the algorithm and the resulting optimal design location of the injection and production wells (I and P labels, respectively) are shown in the Figure 4 (right). Figure 5 shows the water and oil saturation on the 2000th day at the optimal design point corresponding to the mean permeability field, and Figure 6 shows a sample realization of the water and oil saturation on the same day and at the same design point.

Finally, in order to perform a comparison of the computational needs of our approach with possible alternatives, it is necessary to look at the dimensionality of the adapted expansions used in the evaluations of the objective function. Figure 7 shows the histogram with the dimensionality of the adapted expansions used in the estimation of $F_Q(\cdot)$ at all design points. We see that the dimensionality of the η components necessary for accurate estimation of $F_Q(\cdot)$ does not exceed 7, and no more than 30 design points require dimension larger than 3. Furthermore, Figure 8 depicts the contribution of each ξ_i to η_1, \dots, η_7 at one of the design configuration where the dimension of the adapted dimension is 7. These contributions

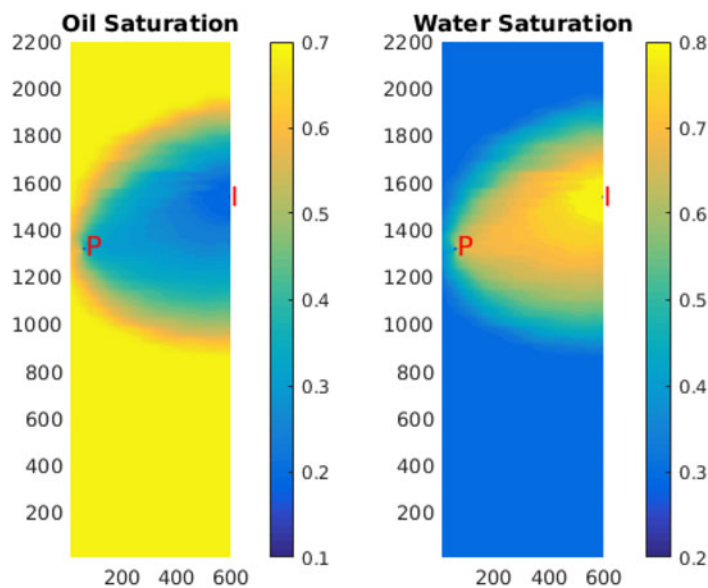


Fig. 6. Sample realization of the oil and water saturation on the 2000th day at the optimal design point.

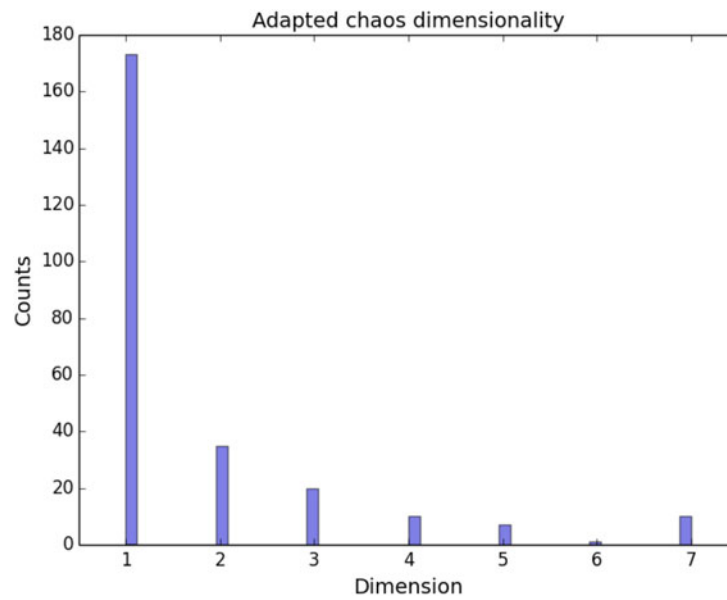


Fig. 7. Dimension of the adapted chaos expansions used in all 256 design points.




Fig. 8. Contribution of each ξ_i to η_1, \dots, η_7 .

Table 1. Number of forward evaluations required for estimation of d -dimensional adapted chaos expansions for $d = 1, \dots, 7$, with polynomial order = 3 and level 2 Smolyak quadrature

Dimension	No. Quadr. Points
1	9
2	45
3	117
4	233
5	401
6	629
7	925

are obtained as the respective rows of the adaptation matrix \mathbf{A} . This figure reveals that all 20 ξ s are contributing to the adapted dimensions, showing that the quantity of interest depends on all 20 dimensions in the ξ space. A polynomial order of 3 in the reduced dimension d_0 is pursued in all subsequent calculations. The number of quadrature points, and consequently the number of forward evaluations required for the estimation of the coefficients of the adapted expansions of order 3 for different dimensionality of $\boldsymbol{\eta}_{d_0}$, is shown in Table 1. Smolyak sparse quadrature rule at Level 2 is used for estimation of multidimensional integrals.

3.4. Variability of the isometry \mathbf{A}

In this subsection, we investigate the dependence of the isometry \mathbf{A} on the design \mathbf{p} , and write $\mathbf{A}(\mathbf{p})$. This dependence

is clearly a result of the dependence of the coefficients $q_\alpha(\mathbf{p})$ of $q(\mathbf{p}, \boldsymbol{\xi})$ on \mathbf{p} and particularly the ones corresponding to the first-order Hermite polynomials that are used in the first row of \mathbf{A} . We perform a simple test that aims to provide some further insight on how \mathbf{A} varies as a function of \mathbf{p} . The motivation for doing so is that if this relation is known *a priori* to the experimenter, then the computation cost involved in the algorithms described above, which is related to the evaluation of \mathbf{A} , could be omitted, thus resulting in even more efficient implementation. The idea is simply that $\mathbf{A}(\mathbf{p})$, or at least its first row \mathbf{A}_1 , which is defined in Eq. (11), is often a smooth function of \mathbf{p} , and therefore, we should be able to estimate it; for instance, we could evaluate \mathbf{A}_1 at several points and then apply some interpolation scheme. A brief theoretical discussion

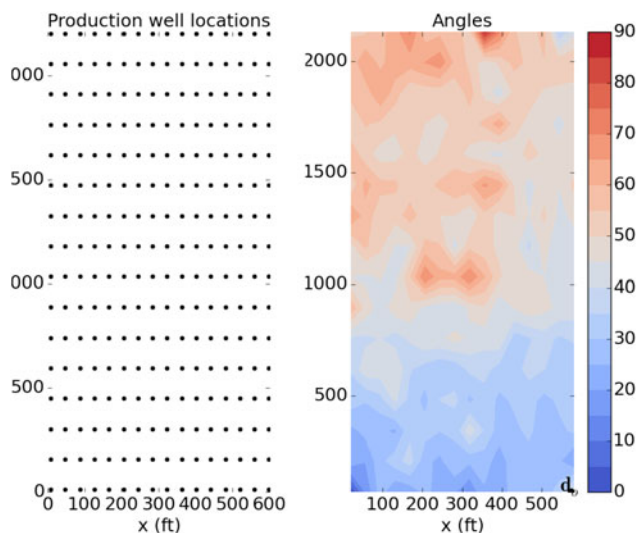


Fig. 9. Spatial distribution of rotation angle for a fixed injection well, as a function of the production well.

on the smoothness of the entries of \mathbf{A} can also be found in (Tsifiis & Ghanem, 2016), where it is shown that if $Q_\alpha(\mathbf{p})$ are square integrable functions of \mathbf{p} , then the covariance kernels of η_i s define Hilbert–Schmidt operators. Here, we only demonstrate that in our example this claim is true and the entries of \mathbf{A}_1 or some properties of them could be estimated. Further development of interpolation or estimating procedures is beyond the scope of the present paper.

To simplify the dependence on \mathbf{p} , we fix the location of the injection well in our example to be at the point (1, 2200) and thus $\mathbf{p} = (x_1^p, x_2^p)$ becomes two dimensional. We compute the Gaussian part $q_\alpha(\mathbf{p})$, $|\alpha| = 1$ at 256 locations as shown in

Figure 8 (right), and from that we find the angles

$$\phi = \arccos(\mathbf{A}_1(\mathbf{p}_o) \times \mathbf{A}_1(\mathbf{p})), \tag{17}$$

that is, the angles between the row $\mathbf{A}_1(\mathbf{p})$ and $\mathbf{A}_1(\mathbf{p}_o)$ for $\mathbf{p}_o = (600, 1)$ and all other \mathbf{p} . There are shown in Figure 9 (left). We observe that the angle between the first rows of the rotations $\mathbf{A}(\mathbf{p})$ and $\mathbf{A}(\mathbf{p}_o)$ is increasing as the production well moves to the upper area of the domain, that is, it approaches the injection well.

3.5. Effect of the confidence level on the optimal design

In this section we are interested in the effect of the confidence level $(1 - \alpha)$ on the design. For this study, similarly to Section 3.3 the parameter space \mathbf{p} is divided into four evenly spaced zones $(x_{1g}^i = \{0, 200, 400, 600\}, x_{2g}^i = \{0, 733.3, 1466.7, 2200\}, x_{1g}^p = \{0, 200, 400, 600\}, x_{2g}^p = \{0, 733.3, 1466.7, 2200\})$, and the objective function is evaluated at the corresponding grid design points $\mathbf{p}_g = (x_{1g}^i \times x_{2g}^i \times x_{1g}^p \times x_{2g}^p)$. Each of the grid points designated as design number 1, 2, . . . , 256. The values of the objective function evaluated at the 256 grid points for different confidence levels 30, 35, 40, . . . , 95. For each confidence level, the optimal design is picked where the objective function is maximum. Figure 10 (left) shows how the optimal design (here its index number) changes for different value of the confidence level. We observe that there are essentially only three possible optimal designs along the whole confidence level spectrum. The production and injection wells $(\{P_i, I_i\}_{i=1,2,3})$ for these three optimal designs are shown in Figure 10 (right). Here, the optimal design one, two, and three corresponds to confidence levels 0.2 to 0.4, 0.45 to 0.65 and 0.70 to 0.95, respectively. Figure 11 (left) shows the quantile values q_α obtained at the




Fig. 10. Confidence level versus the optimal design number (left) and the locations of the optimal designs on the domain (right).




Fig. 11. Dependence of the optimal production level on confidence level (1- α) (left) and probability density function Q_α of the three optimal designs (right).

optimal design point as a function of the confidence level, and Figure 11 (right) shows the probability density functions of q_α for the three optimal designs. It is worth mentioning here that the support and the shape of the probability density functions indicates that although a low confidence level will result most likely in a low minimum production q_α , there is still a positive probability that a large value (up to around 4000) of oil production might be achieved. As the confidence level increases, the support of the probability density function shrinks, indicating a lower upper bound for q_α that, however, can be achieved with higher probability, and thus the minimum oil production is steadily increasing while the associated risk is decreasing.

4. CONCLUSIONS

We presented a polynomial chaos based approach to design optimization that makes use of the basis adaptation capabilities of the PCE in order to obtain highly efficient reduced expansions that offer an alternative to design optimization problems for a special case of objective functions that consist of failure probabilities and thus require full knowledge of the density function of the QoI. We applied our methodology to the oil well placement problem with an expensive-to-solve reservoir model, and our numerical results indicate that our approach can provide solutions that dramatically decrease the computational cost that would be required by computing a polynomial chaos expansion on the initial basis. Much is left to be explored and understood in this novel direction in design optimization under uncertainty based on stochastic expansions, including the behavior of the rotation matrix as a function of the design parameters as well possible constructions of the rotation matrix other than the Gaussian adaptation that is used here.

ACKNOWLEDGMENTS

This research was supported by the ScramJet-UQ project funded under DARPA EQUIPS Program.

REFERENCES

- Bangerth, B., Klie, H., Wheeler, M.F., Stoffa, P.L., & Sen, M.K. (2006). On optimization algorithms for the reservoir oil well placement problem. *Computational Geosciences* 10, 303–319.
- Beckner, B.L., & Song, X. (1995). Field development planning using simulated annealing-optimal economic well scheduling and placement. *Proc. SPE Annual Technical Conf. Exhibition*. Richardson, TX: Society of Petroleum Engineers.
- Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust Optimization*. Princeton, NJ: Princeton University Press.
- Bilionis, I., & Zabaras, N. (2013). Solution of inverse problems with limited forward solver evaluations: a bayesian perspective. *Inverse Problems* 30, 015004.
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. *Proc. COMPSTAT'2010*. Heidelberg, Germany: Physica-Verlag.
- Byrd, R.H., Lu, P., Nocedal, J., & Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 1190–1208.
- Cameron, R., & Martin, W. (1947). The orthogonal development of nonlinear functionals in series of Fourier-hermite functionals. *Annals of Mathematics* 48, 385–392.
- Chen, Z., Huan, G., & Ma, Y. (2006). *Computational Methods for Multiphase Flows in Porous Media*, Vol. 2. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Christie, M.A., & Blunt, M.J. (2001). Tenth SPE comparative solution project: a comparison of upscaling techniques. *Proc. SPE Reservoir Evaluation & Engineering*, Vol. 4. Richardson, TX: Society of Petroleum Engineers.
- Corey, A.T. (1994). *Mechanics of Immiscible Fluids in Porous Media*. Highlands Ranch, CO: Water Resources Publications.
- Eldred, M.S. (2009). Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design. *Proc. 50th AIAA, ASME, ASCE, AHS, ASC Structures, Structural Dynamics, and Materials and Co-Located Conf.*, Palm Beach, CA.
- Ghanem, R., & Spanos, P. (1991). *Stochastic Finite Elements: A Spectral Approach*. New York: Springer-Verlag.
- Güyagüler, B. (2002). *Optimization of well placement and assessment of uncertainty*. PhD thesis. Stanford University.
- Güyagüler, B., & Home, R.N. (2001). Uncertainty assessment of well placement optimization. *Proc. SPE Annual Technical Conf. Exhibition*. Richardson, TX: Society of Petroleum Engineers.
- Heyman, D.P., & Sobel, M.J. (2003). *Stochastic Models in Operations Research: Stochastic Optimization*, Vol. 2. North Chelmsford, MA: Courier Corporation.
- Janson, S. (1999). *Gaussian Hilbert spaces*. Cambridge: Cambridge University Press.
- Keshavarzadeh, V., Meidani, H., & Tortorelli, D.A. (2016). Gradient based design optimization under uncertainty via stochastic expansion methods. *Computer Methods in Applied Mechanics and Engineering* 306, 47–76.
- Kleywegt, A.J., Shapiro, A., & Homem-de Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12, 479–502.

- Le Maître, O.P., Reagan, M.T., Najm, H.N., Ghanem, R.G., & Knio, O.M. (2002). A stochastic projection method for fluid flow: II. Random process. *Journal of Computational Physics* 181, 9–44.
- Peaceman, D.W. (2000). *Fundamentals of Numerical Reservoir Simulation*, Vol. 6. New York: Elsevier.
- Reagan, M.T., Najm, H.N., Ghanem, R.G., & Knio, O.M. (2003). Uncertainty quantification in reacting-flow simulations through non-intrusive spectral projection. *Combustion and Flame* 132, 545–555.
- Rian, D.T., & Hage, A. (1994). Automatic optimization of well locations in a north sea fractured chalk reservoir using a front tracking reservoir simulator. *Proc. Int. Petroleum Conf. Exhibition of Mexico*. Richardson, TX: Society of Petroleum Engineers.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics* 22(3), 400–407.
- Scott, D.W. (2015). *Multivariate Density Estimation: Theory, Practice, and Visualization*. Hoboken, NJ: Wiley.
- Spall, J.C. (1992). Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control* 37, 332–341.
- Stone, H.L. (1970). Probability model for estimating three-phase relative permeability. *Journal of Petroleum Technology* 22(2), 214–218.
- Stone, H.L. (1973). Estimation of three-phase relative permeability and residual oil data. *Journal of Petroleum Technology* 12(4).
- Tipireddy, R., & Ghanem, R.G. (2014). Basis adaptation in homogeneous chaos spaces. *Journal of Computational Physics* 259, 304–317.
- Tsilifis, P., & Ghanem, R. (2016). Reduced Wiener chaos representation of random fields via basis adaptation and projection. *Journal of Computational Physics*. Advance online publication.
- Tsilifis, P., Ghanem, R., & Hajali, P. (2017). Efficient Bayesian experimentation using an expected information gain lower bound. *SIAM/ASA Journal on Uncertainty Quantification* 5(1), 30–62.
- Yeten, B., Durlofsky, L.J., & Aziz, K. (2002). Optimization of nonconventional well type, location and trajectory. *Proc. SPE Annual Technical Conf. Exhibition*. Richardson, TX: Society of Petroleum Engineers.

Charanraj Thimmisetty is a Postdoctoral Associate at the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory. Dr. Thimmisetty has a PhD in civil and environmental engineering from the University of Southern California (USC).

Panagiotis Tsilifis is a Postdoctoral Associate in the Department of Civil Engineering at USC, where he attained his PhD in mathematics.

Roger Ghanem is the Gordon S. Marshall Professor of Engineering Technology at USC. He obtained his PhD from Rice University. Dr. Ghanem's current research deals with probabilistic uncertainty quantification and data-based and model-based predictability.