# A GRADIENT-BASED SAMPLING APPROACH FOR DIMENSION REDUCTION OF PARTIAL DIFFERENTIAL EQUATIONS WITH STOCHASTIC COEFFICIENTS 

M. Stoyanov \& C. G. Webster*<br>Applied Mathematics Group, Computer Science and Mathematics Division, Oak Ridge National Laboratory, 1 Bethel Valley Road, P.O. Box 2008, Oak Ridge TN 37831-6164

Original Manuscript Submitted: 02/20/2014; Final Draft Received: 07/17/2014


#### Abstract

We develop a projection-based dimension reduction approach for partial differential equations with high-dimensional stochastic coefficients. This technique uses samples of the gradient of the quantity of interest (QoI) to partition the uncertainty domain into "active" and "passive" subspaces. The passive subspace is characterized by near-constant behavior of the quantity of interest, while the active subspace contains the most important dynamics of the stochastic system. We also present a procedure to project the model onto the low-dimensional active subspace that enables the resulting approximation to be solved using conventional techniques. Unlike the classical Karhunen-Loève expansion, the advantage of this approach is that it is applicable to fully nonlinear problems and does not require any assumptions on the correlation between the random inputs. This work also provides a rigorous convergence analysis of the quantity of interest and demonstrates: at least linear convergence with respect to the number of samples. It also shows that the convergence rate is independent of the number of input random variables. Thus, applied to a reducible problem, our approach can approximate the statistics of the QoI to within desired error tolerance at a cost that is orders of magnitude lower than standard Monte Carlo. Finally, several numerical examples demonstrate the feasibility of our approach and are used to illustrate the theoretical results. In particular, we validate our convergence estimates through the application of this approach to a reactor criticality problem with a large number of random cross-section parameters.


KEY WORDS: representation of uncertainty, stochastic model reduction method, stochastic sensitivity analysis, high-dimensional approximation, stochastic partial differential equations, Karhunen-Loève expansion, Monte Carlo

## 1. INTRODUCTION

A large number of phenomena in science and engineering are modeled by a map from a set of input data, i.e., model coefficients, forcing functions, boundary and initial conditions, geometry, etc., to an output quantity of interest (QoI). This mapping is typically achieved by virtue of one or more differential and/or partial differential equations (PDEs). However, in practice, the deterministic exact values of the input data are seldom known as they are affected by uncertainty. Such uncertainties can be included in the mathematical model by adopting a probabilistic setting, provided enough information is available for a complete statistical characterization of the physical system. Once the probability distribution, either given or through a calibration procedure, of the input random data is known, the goal of the computational simulation becomes the prediction of statistics (mean, variance, covariance, etc.) of a QoI, or the probability of some given responses of the system. In this setting, stochastic formulations are utilized to account for this random behavior, enabling uncertainty quantification (UQ) in practical applications. In particular, the input data are modeled

[^0]as random variables, and the underlying dynamics originally described by a set of PDEs, are naturally transformed into stochastic parameterized PDEs (SPDEs).

In this work we focus on QoIs coming from the solution of SPDEs whose coefficients and/or forcing terms are described by a finite dimensional random vector; either because the problem itself can be described by a finite number of random variables or because the input coefficients are modeled as truncated random fields. We especially address the situation where the input data are assumed to depend on a large number of random variables.

One particular problem of interest comes from the physics of a nuclear reactor. The behavior of a nuclear reactor depends on the flow of neutrons through the reactor core and the flow of neutrons itself depends on a large number of cross-section parameters that describe the way various types of materials (i.e., nuclear fuel, control rods, coolant, etc.) interact with the neutron field. The number of the parameters is oftentimes in the thousands or even tenths of thousands and recent advances in simulation techniques allow us to resolve the neutron flow for problems with realistic complexity [1]. However, transport solvers still assume that the values of the cross sections are known exactly, while in practice, those are measured experimentally and hence they come with a potentially wide range of uncertainty. The neutron transport problem is an excellent candidate to motivate our gradient-based reduction approach, as it offers a particular set of challenges: including a significantly large number of (potentially uncorrelated) input parameters with a wide range of uncertainty, and significant computational cost associated with each realization.

Sensitivity analysis (SA) is technique for estimating the first two moments of the QoI [2]. The derivative of the QoI with respect to the random input data is used to form a local linearization of the PDE, and then the distributions of the input parameters are propagated through the simplified model. In essence, the expectation is approximated by evaluating the QoI at the nominal (mean) values of the random inputs and the variance is approximated via a "sandwich rule" using the gradient of the QoI and the covariance matrix. Computationally, SA is relatively cheap, however, for a general nonlinear QoI, the accuracy of the linearization away from the nominal point deteriorates quickly, and therefore is only feasible for problems where the size of the noise is relatively small. Yet, many engineering and science applications are affected by a relatively large amount of uncertainty in the input data, and an accurate approximation of the QoI over the full range of uncertain inputs is a desired result.

The Monte Carlo (MC) family of methods (see, e.g., [3]) for random sampling is the classical and most popular approach for estimating statistics of QoIs that depend on the solution to a SPDE with a large number of random inputs. When coupled with a discretization in the physical domain, e.g., finite elements, finite difference, finite volume, spectral, or even $h-p$, the MC approximation is based on independent identically distributed (iid) realizations of the input parameters; approximations of the expectation or other moments of QoIs are obtained by simply averaging over the corresponding realizations of that quantity. As such, the MC method requires a deterministic solution of the PDE for each realization of the input parameters, making it simple to implement, allows for maximal code reusability, and is straightforward to parallelize. The resulting numerical error is proportional to $(1 / \sqrt{k})$, where $k$ is the number of realizations. The advantage of using the Monte Carlo sampling approach is that the convergence rate does not deteriorate with respect to the number of random variables in the problem, making the method very attractive for problems with a large number of random inputs. When solving large-scale applications for which numerical solutions of the PDE are expensive to obtain, the exponent $1 / 2$ in the rate of convergence generates a tremendous amount of computational work in order to achieve accurate solutions. Other ensemble-based methods such as quasi-Monte Carlo (QMC), Latin hypercube sampling, lattice rules, orthogonal arrays, etc. (see, e.g., [4, 5] and the references cited therein), have been devised to produce "faster" convergence rates, e.g., proportional to $\left(\log (k)^{r(N)} / k\right)$, where $r(N)>0$ grows with the number $N$ of the random input variables, and thus, deteriorating convergence as $N$ increases.

More recently, alternative approaches for approximating SPDEs, that utilize standard approximations in the physical space, and stochastic polynomial approximations, using either Galerkin projections or Lagrange interpolation, in the probabilistic domain, have gained considerable attention. These methods are typically known as stochastic Galerkin (SG) and stochastic collocation (SC), respectively, and both techniques exploit the regularity of the solution to acquire faster convergence rates. Moreover, to combat the explosion in computational effort, caused by the curse of dimensionality, approximations are employed in sparse [6-16] or piecewise polynomial spaces [17-23]. However, the intrusive nature of the SG approach requires solving a system of equations that couples all degrees of freedom in the approximation to the stochastic solution. As such, as the number of random inputs grows, then the corresponding number of degrees of freedom is prohibitively large. SC offers a nonintrusive, ensemble-based approach, similar to

Monte Carlo, for constructing a fully discrete approximation over the entire range of probabilistic inputs. Moreover, as shown in [11, 12], the sparse grid SC approximation of SPDEs in which the input random variables come, e.g., from Karhunen-Loève -type truncations of "smooth" random fields, the convergence rate is at least sub-exponential, and in some particular examples is independent of dimension. However, in general, when the parameter space is truly high-dimensional and/or the solution exhibits steep gradients, sharp transitions/bifurcations, or jump discontinuities, all SG and SC methods converge very slowly or even fail to converge. As such, to effectively exploit the fast convergence of both the SC and SG approaches, it becomes necessary to reduce the parameter dimensions to a moderate size of the most important random variables.

For a second-order stochastic process, the Karhunen-Loéve (KL) [24, 25] expansion is the most common dimension reduction technique associated with random input data of a SPDE. KL creates a lower dimensional representation to the inputs that preserves the mean and yields an increasing accurately approximation of the variance. However, KL requires prior knowledge of the correlation between the inputs and the existence of a suitable low-dimensional representation is contingent upon the inputs being strongly correlated. Therefore, this approach is not feasible if the inputs are uncorrelated and, even if a low-dimensional representation of the uncertainty exists, the relation between the error in the KL projection and the corresponding discrepancy in the statistics of the QoI is not rigorously defined for nonlinear problems.

A new reduction approach for approximation SPDEs that depend on high-dimensional parameter spaces by combining the advantages of MC sampling with SA [26-30]. Monte Carlo sampling is used to compute the sensitivities (i.e., derivatives) of the QoI in order to construct a subspace that approximates the span of the gradient of the QoI. Consecutively, the problem is projected onto the resulting low-dimensional subspace, thus reducing the number of inputs and allowing for the application of QMC, SC, and SG techniques. This method is similar to an approach proposed in [31,32], where the dominant singular values of a matrix are inferred from the action of the matrix onto a set of random Gaussian vectors, i.e., every component of the vector is sampled independently from a Gaussian distribution. However, in general, the gradient is not a linear function of the inputs and if it is represented as a product of a matrix and a nonlinear vector function (e.g. [28]), the samples from the nonlinear function do not follow the Gaussian distribution. Therefore, the error estimates in [31,32] are not applicable. Furthermore, even if the gradient depends linearly on the input parameters, these error bounds relate to the discrepancy between the computed and actual dominant singular values of the matrix, while in context of SPDEs, we are interested in the error in the projection of the QoI (which is associated with the neglected small singular values). The gradient sampling method has been successfully applied to several problems, however, it suffers from the lack of rigorous error bounds, relating the approximation of the gradient samples to the error in the statistics of the projected QoI.

The main contribution of this work is to develop a rigorous approach for gradient-based dimension reduction of SPDEs with high-dimensional random inputs. In particular, we propose an approach to identify the passive subspace, i.e., where the QoI is constant or can be accurately (within small error tolerance) approximated by a constant, and then project the problem onto the (active) orthogonal complement, subspace. Moreover, we also derive rigorous bounds relating the error in the statistics of the projected QoI to the error in the approximation of the gradient samples. In the case that the QoI is the variance of a vector that depends linearly on the inputs, our approach is equivalent to Karhunen-Loéve expansion. However, our results extend for higher statistical moments, fully nonlinear problems, and does not require any assumptions on the correlation between the random inputs. Our analysis reveals, that in the worst case, the convergence rate is proportional to $(1 / k)$ and independent from the dimensionality. As such, for problems with low-dimensional active subspace, the method has definite advantage over the classical Monte Carlo sampling. Moreover, as our numerical examples reveal, this linear rate of convergence is sometimes an overestimate and a suitable projection space can be identified with even fewer samples.

The rest of this paper is organized as follows, in Section 2 we define the abstract problem setting and introduce the mathematical problem and the main notation used throughout. In Section 3 we present the analysis that relates the error in the projection of the QoI to the error in the approximation of the gradient. In Section 4 we present a numerical sampling-based algorithm for approximation of the gradient of the QoI and provide the rigorous error analysis of our approach. In Section 5 we present three numerical examples where we apply our method to: a KL expansion involving random matrices (see Section 5.1); a highly reducible random parameter problem (see Section 5.2); and a neutron transport problem with uncertain cross sections (see Section 5.3).

## 2. PROBLEM SETTING

We begin by following the notation in $[11,12,33]$ and let $D \subset \mathbb{R}^{d}, d=1,2,3$, be a convex bounded polygonal domain in $\mathbb{R}^{d}, d=1,2,3$, and $(\Omega, \mathcal{F}, P)$ a complete probability space. We let $\mathcal{L}$ be a differential operator, linear or nonlinear, defined on a domain $D$, which depends on some coefficient(s) $\gamma(\omega, x)$ with $x \in D, \omega \in \Omega$. Here $\Omega$ is the set of outcomes, $\mathcal{F} \subset 2^{\Omega}$ is the $\sigma$ algebra of events, and $P: \mathcal{F} \rightarrow[0,1]$ is a probability measure. We are interested in the following stochastic boundary value problem: find $u: \Omega \times \bar{D} \rightarrow \mathbb{R}^{m}$ such that $P$-almost everywhere in $\Omega$

$$
\begin{equation*}
\mathcal{L}(\gamma)(u)=0 \quad \text { in } D \tag{1}
\end{equation*}
$$

equipped with suitable boundary conditions. Typically $\mathcal{L}$ defines a physical system by virtue of ordinary or partial differential equations, however, the analysis presented in this work is agnostic with respect to any specific structure in the model. We denote by $W(D)$ a Banach space of functions $v: D \rightarrow \mathbb{R}$ and define, for $q \in[1, \infty]$, the stochastic Banach spaces

$$
\begin{equation*}
L_{P}^{q}(\Omega) \otimes W(D):=\left\{u: \Omega \times \bar{D} \rightarrow \mathbb{R}^{m} \mid \int_{\Omega}\|u\|_{W(D)}^{q} d P(\omega)<+\infty\right\} \tag{2}
\end{equation*}
$$

We are particularly interested in the case when $q=2$ as we assume the underlying stochastic input data are chosen so that the corresponding stochastic partial differential equation (1) is well-posed so that it has an unique solution $u(\omega, x) \in L_{P}^{2}(\Omega) \otimes W(D)$, consisting of Banach-space valued functions that have finite second moments. Finally, we note that in this setting the solution $u$ can either be a scalar or vector-valued function depending on the system of interest.

We also assume that the stochastic coefficients $\gamma(\omega, x)$ depend on a finite-dimensional real-valued vector of independent random variables $\boldsymbol{y}=\left[y_{1}(\boldsymbol{\omega}), \ldots, y_{N}(\boldsymbol{\omega})\right]: \Omega \rightarrow \mathbb{R}^{N}$ with $N \in \mathbb{N}_{+}$. Then, the solution $u$ of (1) depends on the realization $\omega \in \Omega$ through the value taken by the random vector $\boldsymbol{y}$, i.e., $u=u(\boldsymbol{\omega}, x)=u(\boldsymbol{y}(\boldsymbol{\omega}), x)$. Below we give an example of the typical finite-dimensional noise decomposition. However, we note that this is not an assumption of our sampling approach or the accompanying convergence analysis.

## Example 1. (Stochastic input data)

In many applications, the stochastic input data may have a simple piecewise random representation whereas, in other applications, the coefficients a in (1) may have spatial variation that can be modeled as a correlated random field, making them amenable to description by a Karhunen-Loève (KL) expansion [24, 25]. In practice, one has to truncate such expansions so that they are of the form

$$
\begin{equation*}
\gamma(\omega, x)=\gamma(\boldsymbol{y}(\omega), x)=\gamma_{0}(\omega, x)+\sum_{n=1}^{N} y_{n}(\omega) b_{n}(x) \tag{3}
\end{equation*}
$$

where the number $N$ terms retained depends on the regularity of the given covariance function and the desired accuracy of the expansion. Please see [11, Section 2.1] for detailed descriptions of both types of noise.

In what follows, we denote by $\Gamma_{n} \equiv y_{n}(\Omega) \subset \mathbb{R}$ the image of the random variable $y_{n}$, then set $\Gamma \equiv \prod_{n=1}^{N} \Gamma_{n} \equiv$ $y(\Omega)$, and assume that the components of the real-valued random vector $\boldsymbol{y}=\left[y_{1}(\boldsymbol{\omega}), \ldots, y_{N}(\boldsymbol{\omega})\right]: \Omega \rightarrow \mathbb{R}^{\bar{N}}$ have a joint probability density function (PDF)

$$
\rho: \Gamma \rightarrow \mathbb{R}_{+} \quad \text { with } \quad \rho(\boldsymbol{y}) \in L^{\infty}(\Gamma)
$$

where $\rho(\boldsymbol{y})=\prod_{n=1}^{N} \rho_{n}\left(y_{n}\right)$ if the random variables are independent. Therefore, the probability space $(\Omega, \mathcal{F}, P)$ is equivalent to $(\Gamma, \mathcal{B}(\Gamma), \rho(\boldsymbol{y}) d \boldsymbol{y})$, where $\mathcal{B}(\Gamma)$ is the Borel $\sigma$-algebra on $\Gamma$ and $\rho(\boldsymbol{y}) d \boldsymbol{y}$ is the finite measure of the random vector $\boldsymbol{y}$. In this setting the stochastic Banach space $L_{P}^{q}(\Omega)$ is equivalent to $L_{\rho}^{q}(\Gamma)$, consisting of functions on $\Gamma$ with respect to the measure $\rho(\boldsymbol{y}) d \boldsymbol{y}$.

Therefore, the solution of the state equation (1) is an unique square integrable function $u=u(\boldsymbol{y}, x) \in L_{\rho}^{2}(\Gamma) \otimes$ $W(D)$, for any $\boldsymbol{y} \in \Gamma$ and $x \in \bar{D}$. The goal of this effort is to construct statistical information related to an output QoI, as a function of the random vector $\boldsymbol{y}=\left[y_{1}, \ldots, y_{N}\right]$, by evaluating the map

$$
\begin{equation*}
Q=Q(u(\boldsymbol{y}, x))=Q(\boldsymbol{y}): \Gamma \times \bar{D} \rightarrow \mathbb{R} \tag{4}
\end{equation*}
$$

that we assume depends on a high-dimensional subspace $\Gamma \subset \mathbb{R}^{N}$. Here, by high-dimensional we really mean that $N=\operatorname{dim}(\Gamma) \sim \mathcal{O}(100)$.

## Remark 1. (Support of the joint PDF)

Even though the support of $\Gamma \subset \mathbb{R}^{N}$ may be bounded, we assume that $\rho(\boldsymbol{y})$ is defined over all of $\mathbb{R}^{N}$. If $\Gamma$ is a bounded domain then we set $\rho(\boldsymbol{y})=0$ outside the region $\Gamma$. We note that the is strictly an artifact of the projection techniques described in Section 2.1 and the analysis presented in Section 4, and has no effect in the main convergence rates described in Theorems 2 and 3.

Moreover, in this effort we focus on complex stochastic problems defined by (1), where, given a sample $\boldsymbol{y}(\omega) \in \Gamma$ evaluating the $\mathrm{QoI}(4)$ is computationally expensive. However, we assume that we can compute the gradient of $Q(\boldsymbol{y})$ at a specific random vector $\boldsymbol{y}$, denoted $\nabla Q(\boldsymbol{y})$, with comparable cost to the computation of the value of $Q(\boldsymbol{y})$. We remark that in this effort the gradient operator $\nabla \equiv \partial / \partial \boldsymbol{y}$ denotes the gradient with respect to $\boldsymbol{y}$ only. Finally, we focus on constructing a numerical approximation of the expected value of the QoI, namely

$$
\begin{equation*}
\mathbb{E}[Q]=\int_{\Gamma} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y} \tag{5}
\end{equation*}
$$

however, higher order statistics of the QoI and the gradient can also be approximated by replacing $Q(\boldsymbol{y})$ and $\nabla Q(\boldsymbol{y})$ with

$$
Q_{k}(\boldsymbol{y})=(Q(\boldsymbol{y})-\mathbb{E}[Q])^{k}, \text { and } \nabla Q_{k}(\boldsymbol{y})=k(Q(\boldsymbol{y})-\mathbb{E}[Q])^{k-1} \nabla Q(\boldsymbol{y}), k \in \mathbb{N}_{+}
$$

respectively. Note from Remark 1 that since $\rho(\boldsymbol{y})=0$ when $\boldsymbol{y} \notin \Gamma$, we can extend $Q(\boldsymbol{y})$ arbitrarily outside $\Gamma$ [i.e., define $Q(\boldsymbol{y})=0$ for $\boldsymbol{y} \notin \Gamma$ ] and thus the integral (5) can be defined over all of $\mathbb{R}^{N}$.

Next we give an example problem posed in this setting:

## Example 2. (Neutron transport with stochastic cross sections)

In one spatial dimension, i.e. $D=[0,1]$, the $k$-eigenvalue transport problem in strong form with finite uncertainty in the capture, scatter, and fission cross sections, denoted $\sigma_{c}(\boldsymbol{y}, x), \sigma_{s}(\boldsymbol{y}, x)$ and $\sigma_{f}(\boldsymbol{y}, x)$ respectively, is given by [34]

$$
\begin{equation*}
\mu \frac{\partial \psi}{\partial x}(\boldsymbol{y}, x, \mu)+\sigma_{T}(\boldsymbol{y}, x) \psi(\boldsymbol{y}, x, \mu)=\sigma_{s}(\boldsymbol{y}, x) \phi(\boldsymbol{y}, x)+\frac{v}{k(\boldsymbol{y})} \sigma_{f}(\boldsymbol{y}, x) \phi(\boldsymbol{y}, x) \quad \text { for a. e. } x \in D, \theta \in[0, \pi] \tag{6}
\end{equation*}
$$

where $\mu=\cos (\theta), \sigma_{T}(\boldsymbol{y}, x)=\sigma_{c}(\boldsymbol{y}, x)+\sigma_{s}(\boldsymbol{y}, x)+\sigma_{f}(\boldsymbol{y}, x)$ measures the uncertainty in the total cross-section, $\psi(\boldsymbol{y}, x, \mu)$ is the random angular flux measuring the uncertainty in the density of neutrons at location $x \in D$ in the direction $\theta, \phi(\boldsymbol{y}, x)=1 / 2 \int_{-1}^{1} \psi(\boldsymbol{y}, x, \mu) d \mu$ is the uncertainty in the total number of neutrons at location $x$, and $v$ is the average number of neutrons emitted after a fission reaction.

If we let the cross sections have a finite stochastic representation similar to (3) that is uniformly bounded and coercive, i.e., for the total cross section $\sigma_{T}(\boldsymbol{y}, \cdot)$ there exists $\sigma_{T_{\min }}>-\infty$ and $\sigma_{T_{\max }}<+\infty$ such that

$$
\begin{equation*}
P\left[\omega \in \Omega: \sigma_{T_{\min }} \leq \sigma_{T}(\boldsymbol{y}(\boldsymbol{\omega}), x) \leq \sigma_{T_{\max }} \quad \forall x \in \bar{D}\right]=1 \tag{7}
\end{equation*}
$$

and similarly for the fission, capture, and scatter cross sections, then (6) satisfies all the above assumptions with $W(D)=L^{2}\left(H^{1}(D) ; 0, \pi\right)$. Of course, in this setting, the operator $\mathcal{L}$ from $(1)$ corresponds to the eigenvalue problem (6), the coefficients correspond to the cross sections, and the QoI described by (4) is the stochastic eigenvalue $k(\boldsymbol{y})$. Similar to (5) we are interested in computing the expected value of $k$-effective, whose value determines whether a reactor is sub-critical, $\mathbb{E}[k]<1$, super-critical, $\mathbb{E}[k]>1$, or critical, $\mathbb{E}[k]=1$.

When the amount of uncertainty is large, that is $N=\operatorname{dim}(\Gamma)$ is high-dimensional and the range of each parameter is large, and if calculating the QoI (4) is costly (e.g., see Example 2), then approximating (5) becomes computationally infeasible. Any deterministic quadrature approach, e.g., tensor products, sparse grids, quasi-Monte Carlo, latin hypercube sampling, etc. will suffer from the curse of dimensionality since the rates of convergence depend on the dimension $N$. Of course, we could directly apply a random sampling approach, such as Monte Carlo, however, the convergence rate is quite slow and a high level of accuracy is achieved only with a substantial amount of function evaluations. As such, the goal of this effort is to reduce the amount of uncertainty by accurately quantifying the $\widetilde{N} \ll N$ most active dimensions having the largest influence on statistics of the QoI. Then, one can apply any stochastic polynomial approximation technique to approximate $Q(\boldsymbol{y})$, e.g., spectral-Galerkin, stochastic collocation, piecewise locally adaptive, etc., or a deterministic sampling technique to compute (5) directly. To accomplish such stochastic dimension reduction we utilize a random sampling procedure, however, instead of sampling $Q(\boldsymbol{y})$ directly we instead sample information from the gradient $\nabla Q(\boldsymbol{y})$. In doing so, we analytically show that this approach converges at least linearly in the number of samples when approximating the expected value of the QoI given by (5).

### 2.1 Active and Passive Subspaces

In general, the dimension $N$ of the random domain $\Gamma$ may be very large a priori, however, the random parameters $y_{1} \ldots, y_{N}$ typically do not all have equal influence on the desired QoI. For example, in many practical applications, the QoI is close to invariant under perturbations of $\boldsymbol{y}$ with arbitrary size for most directions. The largest possible subspace over which the QoI exhibits constant or near-constant behavior is what we define as the passive subspace, denoted $\Lambda^{p} \subset \mathbb{R}^{N}$. Moreover, we also define the active subspace $\Lambda^{a} \subset \mathbb{R}^{N}$ as the orthogonal complement of $\Lambda^{p}$, i.e. $\Lambda^{a} \perp \Lambda^{p}$ such that every $\boldsymbol{y} \in \Lambda^{a} \oplus \Lambda^{p}=\mathbb{R}^{N}$ can be decomposed as $\boldsymbol{y}=\boldsymbol{y}^{a}+\boldsymbol{y}^{p}$ with $\boldsymbol{y}_{a} \in \Lambda^{a}, \boldsymbol{y}^{p} \in \Lambda^{p}$ and $\left\langle\boldsymbol{y}^{a}, \boldsymbol{y}^{p}\right\rangle=0$. Finally, we define the orthogonal projection operators $F_{\Lambda^{a}}$ and $F_{\Lambda^{p}}$ so that $\boldsymbol{y}^{a}=F_{\Lambda^{a}} \boldsymbol{y}$ and $\boldsymbol{y}^{p}=F_{\Lambda^{p}} \boldsymbol{y}$.

From our assumption in Section 2 that $Q(\boldsymbol{y})$ is approximately invariant under perturbations of $\boldsymbol{y}$ in the direction of $\boldsymbol{y}^{p} \in \Lambda^{p}$, we can define the projection of the QoI (4) onto the subspace $\Lambda^{a}$ as $\widehat{Q}: F_{\Lambda^{a}} \Gamma \rightarrow \mathbb{R}$. Special consideration has to be given to the case when $F_{\Lambda^{a}} \Gamma \not \subset \Gamma$, since $\widehat{Q}$ has to be defined for values $\boldsymbol{y}^{a} \in F_{\Lambda^{a}} \Gamma$ with $\boldsymbol{y}^{a} \notin \Gamma$, in which case $Q\left(\boldsymbol{y}^{a}\right)$ may not have a trivial extension, e.g., the PDE may be ill-posed. A visual representation of the projection problem can be observed in Fig. 1. Since $\boldsymbol{y}^{a} \in F_{\Lambda^{a}} \Gamma$ implies the existence of $\boldsymbol{y} \in \Gamma$ so that $F_{\Lambda^{a}} \boldsymbol{y}=\boldsymbol{y}^{a}$, and since $Q$ is assumed to be approximately invariant in directions $\boldsymbol{y}^{p}=\boldsymbol{y}-\boldsymbol{y}^{a}=F_{\Lambda^{p}} \boldsymbol{y} \in \Lambda^{p}$, we define $\widehat{Q}\left(\boldsymbol{y}^{a}\right)$ as


FIG. 1: A two-dimensional simple illustration of a bounded support $\Gamma$ of the probability density function as well as the projection on the low dimensional active subspace $\Lambda^{a}$.

$$
\widehat{Q}\left(\boldsymbol{y}^{a}\right)= \begin{cases}Q\left(\boldsymbol{y}^{a}\right), & \boldsymbol{y}^{a} \in \Gamma,  \tag{8}\\ 0, & \boldsymbol{y}^{a}+\boldsymbol{y}^{p} \notin \Gamma \quad \forall \boldsymbol{y}^{p} \in \Lambda^{p}, \\ Q\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right), & \text { for any } \boldsymbol{y}^{p} \in \Lambda^{p} \text { such that } \boldsymbol{y}^{a}+\boldsymbol{y}^{p} \in \Gamma\end{cases}
$$

By this definition, $Q\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) \approx \widehat{Q}\left(\boldsymbol{y}^{a}\right)$ for all $\boldsymbol{y}^{p} \in \Lambda^{p}$ and we can project (5) onto $\Lambda^{a}$ by virtue of

$$
\begin{align*}
& \int_{\mathbb{R}^{N}} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}=\int_{\Lambda^{a}} \int_{\Lambda^{p}} Q\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) \rho\left(\boldsymbol{y}^{p}+\boldsymbol{y}^{a}\right) d \boldsymbol{y}^{p} d \boldsymbol{y}^{a} \\
\approx & \int_{\Lambda^{a}} \widehat{Q}\left(\boldsymbol{y}^{a}\right) \int_{\Lambda^{p}} \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) d \boldsymbol{y}^{p} d \boldsymbol{y}^{a}=\int_{\Lambda^{a}} \widehat{Q}\left(\boldsymbol{y}^{a}\right) \widehat{\rho}\left(\boldsymbol{y}^{a}\right) d \boldsymbol{y}^{a} \tag{9}
\end{align*}
$$

where

$$
\begin{equation*}
\widehat{\rho}\left(\boldsymbol{y}^{a}\right)=\int_{\Lambda^{p}} \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) d \boldsymbol{y}^{p} \tag{10}
\end{equation*}
$$

is the projected probability density function defined over $\Lambda^{a}$. Therefore, our goal is to construct suitable $\Lambda^{p}$ and, using (8)-(10), efficiently project the high-dimensional integral (5) onto the low dimensional subspace $\Lambda^{a}$, with $\operatorname{dim}\left(\Lambda^{a}\right) \ll$ $N$, such that

$$
\begin{equation*}
\left|\int_{\mathbb{R}^{N}} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}-\int_{\Lambda^{a}} \widehat{Q}\left(\boldsymbol{y}^{a}\right) \widehat{\rho}\left(\boldsymbol{y}^{a}\right) d \boldsymbol{y}^{a}\right| \leq \epsilon \tag{11}
\end{equation*}
$$

with $\epsilon$ a predefined error tolerance. The advantage of working with the lower dimensional projected integral is that we can compute (9) using various collocation or polynomial approximation techniques. However, assuming we are given $\Lambda^{a}$ and $\Lambda^{p}$, the next remark describes specifically how we project the PDF multivariate Gaussian distribution.

## Remark 2. (Projecting a Probability Density Function)

For the specific case that the PDF is a multivariate Gaussian distribution, i.e., $\rho(\boldsymbol{y})=\left[e^{-1 / 2\|\boldsymbol{y}\|^{2}}\right] /\left[(2 \pi)^{N / 2}\right]$, we want to project the PDF on the low-dimensional space $\Lambda^{a}$. We can exploit the orthogonality of $\Lambda^{a}$ and $\Lambda^{p}$ and obtain

$$
\widehat{\rho}\left(\boldsymbol{y}^{a}\right)=\int_{\Lambda^{p}} \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) d \boldsymbol{y}^{p}=\frac{e^{-\frac{1}{2}\left\|\boldsymbol{y}^{a}\right\|^{2}}}{(2 \pi)^{\operatorname{dim}\left(\Lambda^{a}\right) / 2}}
$$

Similar result holds for other distributions such as uniform distribution on an $l^{2}$ ball centered at the origin and truncated Gaussian distribution (so long as the truncation is done at the boundary of an $l^{2}$ ball). A more general distribution is not trivial to project and one may be unable to derive a density function in closed form. In many practical applications, even if computing realizations of the QoI, given by (4), are computationally expensive, computing samples of the random vector $\boldsymbol{y}$ may be cheap. Hence, we can use these samples to approximate $\widehat{\rho}$ by an integral over a small region $V_{\boldsymbol{y}^{a}}$ around $\boldsymbol{y}^{a}$, i.e., for any $\tilde{\boldsymbol{y}}^{a} \in \Lambda^{a} \cap V_{\boldsymbol{y}^{a}}$

$$
\begin{equation*}
\widehat{\rho}\left(\boldsymbol{y}^{a}\right) \approx \frac{1}{\left|V_{\boldsymbol{y}^{a}}\right|} \int_{\Lambda^{a} \cap V_{\boldsymbol{y}^{a}}} \int_{\Lambda^{p}} \rho\left(\tilde{\boldsymbol{y}}^{a}+\boldsymbol{y}^{p}\right) d \boldsymbol{y}^{p} d \tilde{\boldsymbol{y}}^{a}=\frac{1}{\left|V_{\boldsymbol{y}^{a}}\right|} \int_{\mathbb{R}^{N}} I_{V_{\boldsymbol{y}^{a}}}(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y} \tag{12}
\end{equation*}
$$

where $\left|V_{\boldsymbol{y}^{a}}\right|$ indicates the multidimensional volume of $V_{\boldsymbol{y}^{a}} \subset \Lambda^{a}$ and

$$
I_{V_{\boldsymbol{y}^{a}}}(\boldsymbol{y})= \begin{cases}1, & F_{\Lambda^{a}} \boldsymbol{y} \in V_{\boldsymbol{y}^{a}} \\ 0, & \text { o.w }\end{cases}
$$

Without considering the computational cost, assume we can generate a large number of samples $\left\{\boldsymbol{y}_{i}\right\}_{i=1}^{k}$, then (12) can be approximated by a Monte Carlo method. Furthermore, this sampling technique can also be useful to find the value of $\widehat{Q}\left(\boldsymbol{y}^{a}\right)$, when $\boldsymbol{y}^{a} \notin \Gamma$. That is, if $I_{V}\left(\boldsymbol{y}_{i}\right)=1$ and if $\boldsymbol{y}^{a}+F_{\Lambda^{p}} \boldsymbol{y}_{i} \in \Gamma$, then $\widehat{Q}\left(\boldsymbol{y}^{a}\right)=Q\left(\boldsymbol{y}^{a}+F_{\Lambda^{p}} \boldsymbol{y}_{i}\right)$. We note that the details of such a sampling approach are described in Section 4.

## 3. THE GRADIENT AND THE ERROR

To accomplish our objective of constructing $\Lambda^{p}$ and $\Lambda^{a}$ we utilize random samples of the gradient of the QoI, denoted $\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}_{i=1}^{k}$, described in Section 4. First, we need a result regarding the relationship between the error $\epsilon$ in our approximation (11) and the gradient of the QoI. In what follows $\langle\cdot, \cdot\rangle$ denotes the standard inner product in $\mathbb{R}^{N}$.

## Theorem 1. (Gradient bound)

Let subspaces $\Lambda^{a}, \Lambda^{p} \subset \mathbb{R}^{N}$ be such that $\Lambda^{a} \perp \Lambda^{p}$ and $\Lambda^{a} \oplus \Lambda^{p}=\mathbb{R}^{N}$. Furthermore, suppose the probabilistic domain $\Gamma \subset \mathbb{R}^{N}$ is convex, $Q(\boldsymbol{y})$ is continuously differentiable over $\Gamma$, and define for every $\boldsymbol{y} \in \Gamma$ the set $\Lambda^{p}(\boldsymbol{y})=$ $\left\{\boldsymbol{y}^{p} \in \Lambda^{p}: \boldsymbol{y}+\boldsymbol{y}^{p} \in \Gamma\right\}$, Then, if there exists a probability density function $g: \mathbb{R}^{N} \rightarrow \mathbb{R}_{+}$, and $\in \geq 0$, that satisfy either

$$
\begin{equation*}
|Q(\boldsymbol{v}+\boldsymbol{w})-Q(\boldsymbol{v})| \rho(\boldsymbol{v}+\boldsymbol{w}) \leq \epsilon g(\boldsymbol{v}+\boldsymbol{w}), \quad \forall \boldsymbol{v} \in \Gamma, \forall \boldsymbol{w} \in \Lambda^{p}(\boldsymbol{v}) \tag{13}
\end{equation*}
$$

or

$$
\begin{equation*}
\left|\int_{0}^{1}\langle\nabla Q(\boldsymbol{v}+s \boldsymbol{w}), \boldsymbol{w}\rangle d s\right| \rho(\boldsymbol{v}+\boldsymbol{w}) \leq \epsilon g(\boldsymbol{v}+\boldsymbol{w}), \quad \forall \boldsymbol{v} \in \Gamma, \forall \boldsymbol{w} \in \Lambda^{p}(\boldsymbol{v}) \tag{14}
\end{equation*}
$$

or

$$
\begin{equation*}
|\langle\nabla Q(\boldsymbol{v}+s \boldsymbol{w}), \boldsymbol{w}\rangle| \rho(\boldsymbol{v}+\boldsymbol{w}) \leq \epsilon g(\boldsymbol{v}+\boldsymbol{w}), \quad \forall \boldsymbol{v} \in \Gamma, \forall \boldsymbol{w} \in \Lambda^{p}(\boldsymbol{v}), \forall s \in[0,1], \tag{15}
\end{equation*}
$$

then we have that

$$
\begin{equation*}
\left\|Q(\boldsymbol{y})-\widehat{Q}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)\right\|_{L_{\rho}^{1}\left(\mathbb{R}^{N}\right)}=\int_{\mathbb{R}^{N}}\left|Q(\boldsymbol{y})-\widehat{Q}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)\right| \rho(\boldsymbol{y}) d \boldsymbol{y} \leq \epsilon \tag{16}
\end{equation*}
$$

as well as

$$
\begin{equation*}
\left|\int_{\mathbb{R}^{N}} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}-\int_{\Lambda^{a}} \widehat{Q}\left(\boldsymbol{y}^{a}\right) \widehat{\rho}\left(\boldsymbol{y}^{a}\right) d \boldsymbol{y}^{a}\right| \leq \epsilon \tag{17}
\end{equation*}
$$

where the operator with $\widehat{Q}\left(\boldsymbol{y}^{a}\right)$ and $\widehat{\rho}\left(\boldsymbol{y}^{a}\right)$ given by (8) and (10) respectively.
Proof. We begin by noting that

$$
\left|\int_{\mathbb{R}^{N}} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}-\int_{\Lambda^{a}} \widehat{Q}\left(\boldsymbol{y}^{a}\right) \widehat{\rho}\left(\boldsymbol{y}^{a}\right) d \boldsymbol{y}^{a}\right| \leq \int_{\mathbb{R}^{N}}\left|Q(\boldsymbol{y})-\widehat{Q}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)\right| \rho(\boldsymbol{y}) d \boldsymbol{y}
$$

and therefore (16) implies (17). Using the convexity of $\Gamma$ and the differentiability of $Q(\cdot)$ we apply the Fundamental Theorem of Calculus to the left-hand-side of (13) which yields

$$
|Q(\boldsymbol{v}+\boldsymbol{w})-Q(\boldsymbol{v})| \rho(\boldsymbol{v}+\boldsymbol{w})=\left|\int_{0}^{1}\langle\nabla Q(\boldsymbol{v}+s \boldsymbol{w}), \boldsymbol{w}\rangle d s\right| \rho(\boldsymbol{v}+\boldsymbol{w}) \leq \int_{0}^{1}|\langle\nabla Q(\boldsymbol{v}+s \boldsymbol{w}), \boldsymbol{w}\rangle| \rho(\boldsymbol{v}+\boldsymbol{w}) d s
$$

Therefore, (15) implies (14), which in turn implies (13) and thus all we need to show is that (13) implies (16). To get this, we assume (13), we define $D(\boldsymbol{y})=\left|Q(\boldsymbol{y})-\widehat{Q}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)\right| \rho(\boldsymbol{y})$ as the integrand of (16) and we consider the four possible scenarios for $\boldsymbol{y}$ and the corresponding projection $F_{\Lambda^{a}} \boldsymbol{y}$ (see Fig. 1 for a visual of these cases):

Case 1: if $\boldsymbol{y} \notin \Gamma$, then $\rho(\boldsymbol{y})=0$ and thus $D(\boldsymbol{y})=0$;
Case 2: if $\boldsymbol{y} \in \Gamma \cap \Lambda^{a}$, then $\widehat{Q}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)=Q\left(F_{\Lambda^{a}} \boldsymbol{y}\right)=Q(\boldsymbol{y})$ and thus, $D(\boldsymbol{y})=0$;
Case 3: if $\boldsymbol{y} \in \Gamma, \boldsymbol{y} \notin \Lambda^{a}$ and $\boldsymbol{y}^{a}=F_{\Lambda^{a}} \boldsymbol{y} \in \Gamma$, then $\boldsymbol{y}-\boldsymbol{y}^{a} \in \Lambda^{p}\left(\boldsymbol{y}^{a}\right)$, and we can apply condition (13) with $\boldsymbol{v}=\boldsymbol{y}^{a}$ and $w=\boldsymbol{y}-\boldsymbol{y}^{a}$, to get that $D(\boldsymbol{y}) \leq € g(\boldsymbol{y})$; and
Case 4: if $\boldsymbol{y} \in \Gamma, \boldsymbol{y} \notin \Lambda^{a}$, and $\boldsymbol{y}^{a}=F_{\Lambda^{a}} \boldsymbol{y} \notin \Gamma$, then by definition (8), there is $\boldsymbol{y}^{p} \in \Lambda^{p}$ so that $\boldsymbol{y}^{a}+\boldsymbol{y}^{p} \in \Gamma$ and $\widehat{Q}\left(\boldsymbol{y}^{a}\right)=Q\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right)$. Since $\boldsymbol{y}-\boldsymbol{y}^{a}-\boldsymbol{y}^{p} \in \Lambda^{p}\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right)$, we can apply condition (13) with $\boldsymbol{v}=\boldsymbol{y}^{a}+\boldsymbol{y}^{p}$ and $\boldsymbol{w}=\boldsymbol{y}-\boldsymbol{y}^{a}-\boldsymbol{y}^{p}$, to also get that $D(\boldsymbol{y}) \leq \epsilon g(\boldsymbol{y})$.

As such, in all cases for $\boldsymbol{y}$ we have that $D(\boldsymbol{y}) \leq € g(\boldsymbol{y})$, and therefore,

$$
\int_{\mathbb{R}^{N}}\left|Q(\boldsymbol{y})-\widehat{Q}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)\right| \rho(\boldsymbol{y}) \leq \int_{\mathbb{R}^{N}} \epsilon g(\boldsymbol{y}) d \boldsymbol{y}=\epsilon
$$

## Remark 3. (Global support)

If $\Gamma=\mathbb{R}^{N}$, then $\boldsymbol{y}^{a}$ will always be in $\Gamma$ and we do not need to consider the fourth case of Theorem 1. Therefore, conditions (13)-(15) can be weakened, and need only hold for $v \in \Lambda^{a}$ (as opposed to $v \in \Gamma$ ).

## Remark 4. (Alternative condition)

Condition (15) is equivalent to

$$
\begin{equation*}
\left|\left\langle\nabla Q\left(\boldsymbol{y}+\boldsymbol{y}^{p}\right), s \boldsymbol{y}^{p}\right\rangle\right| \rho\left(\boldsymbol{y}+s \boldsymbol{y}^{p}\right) \leq \epsilon g\left(\boldsymbol{y}+s \boldsymbol{y}^{p}\right), \quad \forall \boldsymbol{y} \in \Gamma, \forall \boldsymbol{y}^{p} \in \Lambda^{p}(\boldsymbol{y}), \forall s>1 \tag{18}
\end{equation*}
$$

which in turn implies (16) and (17). We will utilize (18) in the sampling algorithm described in Section 4, since for point $\boldsymbol{y}$ and corresponding value of $\nabla Q(\boldsymbol{y})$, we can examine all values of s by considering $\rho(\cdot)$ and $g(\cdot)$ only. This is unlike conditions (14) and (15) that require knowledge of $\nabla Q(\boldsymbol{y})$ for a range of the uncertainties $\boldsymbol{y}^{a}+s \boldsymbol{y}^{p}$.

## Corollary 1. (True dimension)

Let $\Lambda_{n u l l}^{p}=\left\{\boldsymbol{v} \in \mathbb{R}^{N}:\langle\boldsymbol{v}, \nabla Q(\boldsymbol{y})\rangle=0, \forall \boldsymbol{y} \in \Gamma\right\}$, then conditions (14) and (15) are satisfied with $\epsilon=0$ for any arbitrary function $g: \mathbb{R}^{N} \rightarrow \mathbb{R}_{+}$. Therefore, we can project $Q: \Gamma \times \bar{D} \rightarrow \mathbb{R}$ without loss of information and the actual dimension of the QoI is $N-\operatorname{dim}\left(\Lambda_{\text {null }}^{p}\right)$.

Corollary 2. (Low-dimensional surrogate)
Suppose there exists a function $\mathcal{I}\left(\boldsymbol{y}^{a}\right)$ that approximates $\widehat{Q}\left(\boldsymbol{y}^{a}\right)$ in the low-dimensional space $\Lambda^{a}$. Then $\mathcal{I}\left(F_{\Lambda^{a}} \boldsymbol{y}\right)$ approximates $Q(\boldsymbol{y})$ and satisfies the following estimate:

$$
\left\|\mathcal{I} \circ F_{\Lambda^{a}}-Q\right\|_{L_{\rho}^{1}\left(\mathbb{R}^{N}\right)} \leq\|\mathcal{I}-\widehat{Q}\|_{L_{\hat{\rho}}^{1}\left(\Lambda^{a}\right)}+\left\|\widehat{Q} \circ F_{\Lambda^{a}}-Q\right\|_{L_{\rho}^{1}\left(\mathbb{R}^{N}\right)}
$$

Our main focus remains the approximation of the integral (5), however, Corollary 2 allows us to utilize the lowdimensional space $\Lambda^{a}$ to create a surrogate model for $Q(\boldsymbol{y})$. Since the dimension of $\Lambda^{a}$ is small relative to $\Gamma$, we can apply various stochastic polynomial methods, such as spectral Galerkin or collocation approximations. In the next section we explain how our our gradient-based approach, for stochastic dimension reduction, can be viewed as a generalization to the classic finite-dimensional Karhunen-Loéve expansion [24, 25].

### 3.1 Relationship to the Finite-Dimensional Karhunen-Loéve Expansion

Consider the case where the components of $\boldsymbol{y} \in \mathbb{R}^{N}$ are identically and independently distributed as normalized Gaussian variables, i.e., $y_{1} \in N(0,1), \ldots, y_{N} \in N(0,1)$ with $\rho(\boldsymbol{y})=e^{-(1 / 2)\|\boldsymbol{y}\|^{2}} /(2 \pi)^{N / 2}$. Next, we define

$$
\begin{equation*}
u(\boldsymbol{y})=L^{T} \boldsymbol{y}, \quad \text { and } \quad Q(u(\boldsymbol{y}))=\|u\|^{2}=\boldsymbol{y}^{T} L L^{T} \boldsymbol{y}=\boldsymbol{y}^{T} C \boldsymbol{y} \tag{19}
\end{equation*}
$$

where $C=L L^{T}$ is the covariance matrix of $u$ and the integral of the QoI $Q$ is the variance of $u$. Here we assume $C$ is nonsingular, that is, the actual dimension of the problem is indeed $N$. Given the simple expression for QoI allows us to easily find the gradient of $Q(\boldsymbol{y})$ as

$$
\begin{equation*}
\nabla Q(\boldsymbol{y})=2 C \boldsymbol{y} \tag{20}
\end{equation*}
$$

From Theorem 1 and Remark 3, we seek $\Lambda^{a}$ and $\Lambda^{p}$ that minimize $\epsilon \geq 0$. First, we consider the simplest case where $\operatorname{dim}\left(\Lambda^{p}\right)=1$. In this case, we want to find $\boldsymbol{y}^{p}$ that minimizes (14), that is,

$$
\begin{equation*}
\min _{\Lambda^{p}}\left|\int_{0}^{1}\left\langle\nabla Q\left(\boldsymbol{y}^{a}+s \boldsymbol{y}^{p}\right), \boldsymbol{y}^{p}\right\rangle d s\right| \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right), \quad \forall \boldsymbol{y}^{p} \in \Lambda^{p} \text { and } \boldsymbol{y}^{a} \perp \boldsymbol{y}^{p} \tag{21}
\end{equation*}
$$

Volume 5, Number 1, 2015

Substituting (21) into (20) yields

$$
\begin{align*}
\min _{\Lambda^{p}}\left|\int_{0}^{1}\left\langle\nabla Q\left(\boldsymbol{y}^{a}+s \boldsymbol{y}^{p}\right), \boldsymbol{y}^{p}\right\rangle d s\right| \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) & =\min _{\Lambda^{p}}\left|\int_{0}^{1} 2\left(\boldsymbol{y}^{p}\right)^{T} C \boldsymbol{y}^{a}+s 2\left(\boldsymbol{y}^{p}\right)^{T} C \boldsymbol{y}^{p} d s\right| \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right)  \tag{22}\\
& =\min _{\Lambda^{p}}\left|2\left(\boldsymbol{y}^{p}\right)^{T} C \boldsymbol{y}^{a}+\frac{1}{2} 2\left(\boldsymbol{y}^{p}\right)^{T} C \boldsymbol{y}^{p} d s\right| \rho\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right)
\end{align*}
$$

The minimum of (22) is achieved when $\Lambda^{p}$ is the eigenspace that is associated with the smallest eigenvalue of $C$, i.e., $C \boldsymbol{y}^{p}=\lambda_{\min } \boldsymbol{y}^{p}$. Since $C$ is symmetric positive definite, $\boldsymbol{y}^{a} \perp C \boldsymbol{y}^{p}$, and therefore, we get that

$$
\begin{align*}
\left|2\left(\boldsymbol{y}^{a}\right)^{T} C \boldsymbol{y}^{p}+\left(\boldsymbol{y}^{p}\right)^{T} C \boldsymbol{y}^{p}\right| \frac{e^{-\frac{1}{2}\left\|\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right\|^{2}}}{(2 \pi)^{N / 2}} & =\left(\boldsymbol{y}^{p}\right)^{T} C \boldsymbol{y}^{p} \frac{e^{-\frac{1}{2}\left\|\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right\|^{2}}}{(2 \pi)^{N / 2}} \\
& =\lambda_{\min }\left\|\boldsymbol{y}^{p}\right\|^{2} \frac{e^{-\frac{1}{2}\left\|\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right\|^{2}}}{(2 \pi)^{N / 2}} \equiv \lambda_{\min } g(\boldsymbol{y}) . \tag{23}
\end{align*}
$$

Since $\int g(\boldsymbol{y}) d \boldsymbol{y}=1$, if we apply Theorem 1 with condition (14) and $\epsilon=\lambda_{\min }$, the final error associated with reducing the uncertainty domain by one dimension is bounded by $\epsilon$. By extrapolating this result, we can reduce the dimensions recursively and deduce that the passive subspace is the span of the eigenvectors of $C$ associated with the smallest eigenvalues. Moreover, the theoretical error bound is the sum of the neglected eigenvalues. This result is equivalent to the classical finite-dimensional Karhunen-Loéve expansion, however, our projection approach extends to problems with far more complex structure than linear functions $u$, quadratic functionals $Q(u)$, and Gaussian random variables $\boldsymbol{y}$ [24].

## 4. THE SAMPLING APPROACH AND ERROR ANALYSIS

In practice, we seldom have analytical form of the gradient of the QoI (4), and therefore, we have to create an approximation to both $\Lambda^{p}$ and $\Lambda^{a}$. More importantly, the possible choices for $\Lambda^{p}$ and $\Lambda^{a}$ may not be unique and so we define $\Lambda^{a}$ and $\Lambda^{p}$ to be any two subspaces that satisfy any of the condition in Theorem 1 . Without loss of generality, we will focus our attention on the condition (18), defined in Theorem 1, however, our approach and analysis extends to conditions (14) and (15) as well. Furthermore, we want to project the QoI on a subspace with smallest possible dimension. As such, we attempt to discover subspaces such that $\Lambda^{a}$ has the smallest possible dimension, or alternatively, $\Lambda^{p}$ has the largest possible dimension.

To accomplish these goals we propose a Monte-Carlo-based random sampling approach. That is, given a desired tolerance $\epsilon>0$, we take $k$ random samples, i.e., $\left\{\boldsymbol{y}_{i}\right\}_{i=1}^{k} \in \mathbb{R}^{N}$, where each $\boldsymbol{y}_{i}$ is independently sampled from distribution with probability density $\rho$. For each sample, we compute $\nabla Q\left(\boldsymbol{y}_{i}\right)$ and use the gradients to find a decomposition that approximates $\Lambda^{p}$ and $\Lambda^{a}$. Since we cannot analytically verify condition (18) over the entire domain, we weaken the requirement so that it holds only with respect to the computed samples $\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}_{i=1}^{k}$. Hence we need a procedure to form $\Lambda^{a}$ and $\Lambda^{p}$ from an already computed set of gradient samples.

Our specific approach for forming this approximation is problem dependent. We require the definition of a map $J: \mathbb{R}^{N} \rightarrow \mathbb{R}^{N}$ that associates a finite subsets of $\mathbb{R}^{N}$ with a subspace of $\mathbb{R}^{N}$. That is, let $T$ be a finite set of vectors in $\mathbb{R}^{N}$ and $\Lambda^{a}=J(T)$, a subspace of $\mathbb{R}^{N}$, so that if $\Lambda^{p}$ is the orthogonal complement of $\Lambda^{a}$. Then (18) is satisfied for all $\nabla Q\left(\boldsymbol{y}_{i}\right) \in T$. Of course, the structure of the map $J$ depends on $\epsilon$ and $g(\cdot)$. One possible choice is to take

$$
\begin{equation*}
J(T)=\operatorname{span}\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}_{\boldsymbol{y}_{i} \in T} \tag{24}
\end{equation*}
$$

which guarantees (18) remains valid for all $\epsilon>0$. However, it is more desirable to choose the map $J$ that returns a subspace with smallest possible dimension. Hence, a more practical approach would be to use a procedure that weights $\nabla Q\left(\boldsymbol{y}_{i}\right)$ and returns the subspace spanned by only some of the samples. A common approach to this type of problem is to look at different eigenvalue problems. Below, we give two examples of specific choices for the map $J(T)$. Here we suppose we are given a set of vectors $T$ and tolerance $\epsilon$, and we wish to construct the map $J(T)$ that decomposes $\mathbb{R}^{N}$ so that condition (14) or (18) is satisfied for all $\boldsymbol{y}_{i} \in T$.

Example 3. (Compact support and the choice of $\boldsymbol{J}(\boldsymbol{T})$ )
Suppose the PDF $\rho$ has compact support in $\mathbb{R}^{N}$, i.e., there exists a constant $r$ such that $\rho(\boldsymbol{y})=0, \forall\|\boldsymbol{y}\|>r$. In this case, if we let $g(\boldsymbol{y})=\rho(\boldsymbol{y})$ the (18) is satisfied for all $\|\boldsymbol{y}\|>r$, regardless of $\epsilon$, and therefore, we only need to consider $\left\|\boldsymbol{y}^{p}\right\| \leq r$ or $s \leq r /\left\|\boldsymbol{y}^{p}\right\|$. As such, we arrange the vectors $\nabla Q\left(\boldsymbol{y}_{i}\right)$ into the columns of a matrix $H$. We define $\Lambda_{a}$ to be the space spanned by the $l$ dominant singular vectors of $H$, where $l$ is chosen so that the $(l+1)$ st dominant singular value $\lambda_{l+1}$ satisfies

$$
r \lambda_{l+1} \leq \epsilon
$$

Then, if we take $\boldsymbol{y}^{p} \perp \Lambda^{a}$ such that $\left\|\boldsymbol{y}^{p}\right\| \leq r$ and consider $\left\langle\nabla Q\left(\boldsymbol{y}_{i}\right)\right.$, s $\left.\boldsymbol{y}^{p}\right\rangle$, which is largest when $s=r /\left\|\boldsymbol{y}^{p}\right\|$, we get that

$$
\left\langle\nabla Q\left(\boldsymbol{y}_{i}\right), s \boldsymbol{y}^{p}\right\rangle \leq\left\langle\nabla Q\left(\boldsymbol{y}_{i}\right), \frac{r}{\left\|\boldsymbol{y}^{p}\right\|} \boldsymbol{y}^{p}\right\rangle \leq r \lambda_{l+1} \leq \epsilon
$$

Therefore, condition (18) is satisfied with $\Lambda^{a}=J(T)$ for all $\boldsymbol{y}_{i} \in T$.
Example 4. (Finite-dimensional Karhunen-Loéve expansion and the choice of $\boldsymbol{J}(\boldsymbol{T})$ )
Similar to Section 3.1 we let $Q(\boldsymbol{y})=\boldsymbol{y}^{T} C \boldsymbol{y}$ and $\rho(\boldsymbol{y})=e^{-(1 / 2)\|\boldsymbol{y}\|^{2}} /(2 \pi)^{N / 2}$. We want to construct a procedure that will map a set of vectors $T$ to a subspace, with smallest possible dimension, that satisfies condition (14). That is, for $\lambda \in \mathbb{R}_{+}$:

$$
\begin{equation*}
\left|<\nabla Q\left(\boldsymbol{y}_{i}\right), \boldsymbol{y}^{p}>\right| \leq 2 \lambda\left\|\boldsymbol{y}^{p}\right\|^{2}, \quad \forall \boldsymbol{y}^{p} \perp J(T), \quad \forall \boldsymbol{y}_{i} \in T . \tag{25}
\end{equation*}
$$

First, consider $\Lambda^{p}=J(T)^{\perp}$, and observe that the $\operatorname{span}\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}^{\perp}$ satisfies (25) for all $\lambda \geq 0$, and thus, $\operatorname{span}\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}^{\perp} \subset \Lambda^{p}$. Second, we need to classify the $\operatorname{span}\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}$. Let $v \in \operatorname{span}\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}$, i.e.,

$$
\begin{equation*}
v=\sum_{i=1}^{m} c_{i} \nabla Q\left(\boldsymbol{y}_{i}\right) \equiv H \boldsymbol{c} \tag{26}
\end{equation*}
$$

where $m=\operatorname{card}(T), \boldsymbol{c} \in \mathbb{R}^{m}$ and $H$ is the matrix with columns $\nabla Q\left(\boldsymbol{y}_{i}\right)$. Substituting (26) into (25) and observing that $\boldsymbol{y}^{p}=\left(<v, \boldsymbol{y}^{p}>/\|v\|^{2}\right) v$, we have that

$$
\begin{equation*}
\left|<\nabla Q\left(\boldsymbol{y}_{i}\right), H \boldsymbol{c}>|\leq 2 \lambda|<\boldsymbol{y}_{i}, H \boldsymbol{c}>\right| . \tag{27}
\end{equation*}
$$

By letting $E$ be the matrix with columns $\boldsymbol{y}_{i}(27)$ reduces to the generalized eigenvalue problem

$$
\begin{equation*}
H^{T} H \boldsymbol{c}=2 \lambda E^{T} H \boldsymbol{c} \tag{28}
\end{equation*}
$$

where the eigenvalues $\left\{\lambda_{l}\right\}_{l=1}^{m}$, placed in descending order, are real and positive. We then take $J(T)$ to be the space associated with $H c_{l}$, where $c_{l}, l=1, \ldots, N^{a} \ll N$, are the $N^{a}$ dominant eigenvectors. To select the cutoff value of $\lambda$, we use the eigenvectors associated with the dominant eigenvalues. Let $F_{J(T) \perp}$ be the operator for orthogonal projection onto $J(T)^{\perp}$ and multiply both sides of (25) by $\rho\left(\boldsymbol{y}_{i}\right)$

$$
\left|<\nabla Q\left(\boldsymbol{y}_{i}\right), \boldsymbol{y}^{p}>\right| \rho\left(\boldsymbol{y}_{i}\right) \leq 2 \lambda \int\left\|F_{J(T)^{\perp}} \xi\right\|^{2} \rho(\xi) d \xi \frac{\left\|\boldsymbol{y}^{p}\right\|^{2} \rho\left(\boldsymbol{y}_{i}\right)}{\int\left\|F_{J(T)^{\perp}} \xi\right\|^{2} \rho(\xi) d \xi}
$$

which matches condition (14) with $g(\boldsymbol{y})=\left\|F_{J(T) \perp} \boldsymbol{y}\right\|^{2} \rho(\boldsymbol{y}) / \int\left\|F_{J(T) \perp}{ }^{\perp}\right\|^{2} \rho(\xi) d \xi$ and

$$
\begin{equation*}
\epsilon=2 \lambda \int\left\|F_{J(T)^{\perp}}\right\|^{2} \rho(\xi) d \xi=2 \lambda(N-n), \quad \Rightarrow \quad \lambda=\frac{\epsilon}{2(N-n))}, \tag{29}
\end{equation*}
$$

where $n=\operatorname{dim}(J(T))$. We can use the estimate (29) to select a cutoff value for $\lambda$, however, we should note that when $\operatorname{dim}(J(T)) \ll N$ and the eigenvalues of $H$ decay very fast, (29) is a big overestimate.

Given an appropriate mapping $J(T)$, we consider the $k$ samples $\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}_{i=1}^{k}$. We split the samples into two groups. First we look for the smallest subset of the samples $\widetilde{T} \subset\left\{\nabla Q\left(\boldsymbol{y}_{i}\right)\right\}_{i=1}^{k}$ such that $J(\widetilde{T})$ is a subspace that satisfies condition (18) not only with respect to $\widetilde{T}$ but also with respect to all other samples. Thus, the samples in $\widetilde{T}$ we call essential and the remaining ones are the set of the over-samples. In nontechnical terms, the essential samples are used to identify the dynamics of the QoI and since they are randomly selected, it is possible that they fail to capture all aspects of the behavior of $Q(\boldsymbol{y})$. The over-samples had an opportunity to discover any missing dynamics and no such dynamics were found, hence the more over-samples we have, the more likely it is that we have identified the full behavior of the QoI. The essential samples are needed to identify the approximate passive and active subspaces, while the over-samples are related to the confidence that we have in the approximation. This relationship is quantified in Theorems 2 and 3.

Algorithm 1 summarizes the sampling procedure, where $\epsilon$ is the desired tolerance, and at iteration $k$ we have $\Lambda_{k}^{p}$, $\Lambda_{k}^{a}$ and $d_{k}$ as, respectively, the approximate passive and active subspaces and the number of over-samples.

```
Algorithm 1: Approximate the passive subspace.
    Set \(\Lambda_{0}^{p}=\mathbb{R}^{N}, \Lambda_{0}^{a}=\{0\}, d_{0}=0\) and the tolerance, \(\epsilon\).
    for \(k=1,2, \ldots\), do
        Sample the random vector \(\boldsymbol{y}^{k} \in \Gamma\) with \(\operatorname{PDF} \rho(\boldsymbol{y})\), and evaluate the gradient at the sample point, i.e.,
        \(\nabla Q\left(\boldsymbol{y}^{k}\right)\).
        if (18) is satisfied using \(\nabla Q\left(\boldsymbol{y}^{k}\right)\) with \(\Lambda_{k-1}^{p}, \Lambda_{k-1}^{a}\) then
            Set \(\Lambda_{k}^{p} \leftarrow \Lambda_{k-1}^{p}\) and \(\Lambda_{k}^{a} \leftarrow \Lambda_{k-1}^{a}\) and increment \(d_{k}=d_{k-1}+1\).
        else
            if (18) is not satisfied using \(\nabla Q\left(\boldsymbol{y}^{k}\right)\) with \(\Lambda_{k-1}^{p}, \Lambda_{k-1}^{a}\) then
                Define \(\mathcal{T}=\left\{T \in 2^{\left\{\boldsymbol{y}^{i}\right\}_{i=1}^{k}}: J(T)\right.\) satisfies (18) \(\left.\forall \boldsymbol{y}_{i}\right\}\) and \(\widetilde{T} \in \mathcal{T}\) with \(\operatorname{card}(\widetilde{T}) \leq \operatorname{card}(S)\),
                \(\forall S \in \mathcal{T}\).
                        Set the active subspace \(\Lambda_{k}^{a} \leftarrow J(\widetilde{T})\) with \(\Lambda_{k}^{p}\) being its orthogonal complement.
                        Set \(d_{k}=k-\operatorname{card}(\widetilde{T})\).
            endif
        endif
        if \(d_{k}\) is sufficiently large then
            Stop the iteration and use projection (9) with \(\Lambda^{a}=\Lambda_{k}^{a}\) and \(\Lambda^{p}=\Lambda_{k}^{p}\) to reduce the dimension of the
            problem.
        endif
    endfor
```

Obviously, Algorithm 1 requires a proper convergence criterion with respect to the number of over-samples $d_{k}$, More importantly, recall from (11) that what we really want to know the error we commit in approximating the expectation of our QoI, when using our projection into the active subspace $\Lambda^{a}$. However, when using the sampling approach described by Algorithm 1 we construct an approximation $\Lambda_{k}^{a}$ and thus, we are interested in the error at the $k$ th iteration, defined by

$$
\begin{equation*}
e_{k}=\left|\int_{\mathbb{R}^{N}} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}-\int_{\Lambda_{k}^{a}} \widehat{Q}_{k}\left(\boldsymbol{y}^{a}\right) \widehat{\rho}_{k}\left(\boldsymbol{y}^{a}\right) d \boldsymbol{y}^{a}\right| \tag{30}
\end{equation*}
$$

where $\widehat{\rho}_{k}\left(\boldsymbol{y}^{a}\right)=\int_{\Lambda_{k}^{p}} Q\left(\boldsymbol{y}^{a}+\boldsymbol{y}^{p}\right) d \boldsymbol{y}^{p}$ and $\widehat{Q}_{k}$ denotes the projection of the QoI (4) onto the subspace $\Lambda_{k}^{a}$, and $\widehat{Q}_{k}\left(\boldsymbol{y}^{a}\right)$ is defined by setting $\Lambda^{p}=\Lambda_{k}^{p}$ in (8). Next, using (30) we present two theoretical results that describe the probability of finding the active subspace $\Lambda_{k}^{a}$, using Algorithm 1, and the distribution of the error with respect to the number of samples. These are given in the following two theorems.

## Theorem 2. (Probability of failure)

Given a realization of Algorithm 1 with tolerance $\epsilon$ and $e_{k}$ described by (30), there is a sequence of numbers $m_{k} \in \mathbb{R}$ such that the discrete (or boolean) probability measure satisfies

$$
\begin{equation*}
\mathbb{P}\left(e_{k}>\epsilon\right) \leq\left(1-m_{k}\right)^{d_{k}} \tag{31}
\end{equation*}
$$

where $d_{k}$ is the number of over-samples at step $k$. Moreover, there exists an $n \in \mathbb{N}_{+}$that is independent from the realization of the samples (i.e., it only depends on the properties of $\nabla Q(\boldsymbol{y})$ and not $\boldsymbol{y}_{i}$ ), so that if $k<n$ then $m_{k}=0$ and if $k \geq n$ then $m_{k}>0$.

Proof. If $m_{k}=0$, then (31) is the trivial statement that a probability of an event is bounded by 1 . Consider $\Lambda^{a}$ with smallest dimension that will satisfy condition (18) and let this dimension be $n=\operatorname{dim}\left(\Lambda^{a}\right)$. The largest subspace that the mapping $J$ can return is given by (24) and hence $\operatorname{dim}\left(\Lambda_{k}^{a}\right) \leq \operatorname{card}(T) \leq k$. Therefore, if $k<n$, then $\Lambda_{k}^{a}$ necessarily fails (18) and hence Theorem 1 does not apply. Therefore, we can make only the trivial statement

$$
P\left(e_{k}>\epsilon\right) \leq 1=(1-0)^{d_{k}}=\left(1-m_{k}\right)^{d_{k}}
$$

If $k<n$, then (31) holds only for the trivial choice of $m_{k}=0$.
Suppose $k \geq n$ and define

$$
Z_{k}=\left\{\boldsymbol{y} \in \Gamma: \text { there is a pair } s>1 \text { and } \boldsymbol{y}^{p} \in \Lambda_{k}^{p} \text { such that } \boldsymbol{y}+\boldsymbol{y}^{p} \in \Gamma \text { and condition (18) fails }\right\}
$$

and let

$$
p_{k}=P\left(Z_{k}\right)=\int_{Z_{k}} \rho(\boldsymbol{y}) d \boldsymbol{y}
$$

At iteration $k$, if $p_{k}=0$, then condition (18) holds but for a set of zero probability, hence according to Theorem 1 $e_{k} \leq \epsilon$. The probability of $e_{k}>\epsilon$ given that $p_{k}=0$ is in fact zero

$$
P\left(e_{k}>\epsilon \mid p_{k}=0\right)=0
$$

and therefore, (31) will be true for any $0<m_{k} \leq 1$.
Suppose $p_{k}>0$ and then consider the number of over-samples. By construction of $\Lambda_{k}^{p}$, none of the over-samples belong to $Z_{k}$. Therefore, we have $d_{k}$ number of over-samples all of which were randomly and independently selected outside of the region $Z_{k}$. The probability of such even is $\left(1-p_{k}\right)^{d_{k}}$, therefore,

$$
P\left(e_{k}>\epsilon \mid p_{k}>0\right) \leq\left(1-p_{k}\right)^{d_{k}}
$$

Thus, we can select $m_{k}=p_{k}$.
From Theorem 2 it follows that if for large enough $k, m_{k}$ are uniformly bounded away from zero, then the sampling method has exponential convergence. This means that in the best case, Algorithm 1 could exhibit the fast convergence of the collocation methods, without the curse of dimensionality. However, in general, $m_{k}$ depends on $Z_{k}$ and thus the distribution of the samples $\left\{\boldsymbol{y}_{i}\right\}_{i=1}^{k}$ for each $k$, and the corresponding $m_{k}$ is a random variable and it is possible for the sequence of $m_{k}$ to tend to zero for very large $k$. We have to consider the probability that the error will exceed the tolerance $\epsilon$ with respect to the probability distribution of $m_{k}$.

## Theorem 3. (Distribution of the error)

Suppose there is a constant $M$ such that for any arbitrary subset $G \subset \mathbb{R}^{N}$

$$
\left|\int_{G} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}\right| \leq M \int_{G} \rho(\boldsymbol{y}) d \boldsymbol{y}
$$

Then if $k \geq n$ and $e_{k}>\epsilon$ we have the following bounds, for the expected value $\mathbb{E}\left[e_{k}-\epsilon\right] \leq M /\left(d_{k}+1\right)$, and the variance $\mathbb{V a r}\left[e_{k}-\epsilon\right] \leq M^{2} /\left(d_{k}+1\right)^{2}$, with respect to the distribution of the number of samples.

Proof. Let $d \mu_{k}$ be the probability measure associated with the distribution of $m_{k}$ at step $k$. Following our construction in Theorem 2, the measure $p_{k}$ of the set where condition (18) fails is at most $m_{k}$, therefore, the error in the approximation to the QoI is bounded by $M m_{k}$. Furthermore, according to (31) the probability $P\left(e_{k}>\epsilon\right)$ is bounded by $\left(1-m_{k}\right)^{d_{k}}$. Therefore,

$$
\mathbb{E}\left[e_{k}-\epsilon\right] \leq M \int_{0}^{1} x(1-x)^{d_{k}} d \mu_{k}
$$

The integrand is bounded and attains its maximum at $x=1 /\left(d_{k}+1\right)$ and hence we have the bound

$$
\mathbb{E}\left[e_{k}-\epsilon\right] \leq M \frac{\left(1-\frac{1}{\left(d_{k}+1\right)}\right)^{d_{k}}}{d_{k}+1} \leq \frac{M}{d_{k}+1}
$$

In an analogous way we have that

$$
\operatorname{Var}\left[e_{k}-\epsilon\right] \leq M^{2} \int_{0}^{1} x^{2}(1-x)^{2 d_{k}} d \mu_{k}-\left(\mathbb{E}\left[e_{k}-\epsilon\right]\right)^{2} \leq C^{2} \frac{\left(1-\frac{1}{\left(d_{k}+1\right)}\right)^{2 d_{k}}}{\left(d_{k}+1\right)^{2}} \leq \frac{M^{2}}{\left(d_{k}+1\right)^{2}}
$$

## Remark 5. (Comparisons to Monte Carlo sampling)

Assume that we want to estimate (5) to error tolerance $\epsilon$. The computational cost associated with classical Monte Carlo sampling is $O\left(1 / \epsilon^{2}\right)$ evaluations of $Q(\boldsymbol{y})$, and the estimate is independent from the dimension, regularity or structure of $Q(\boldsymbol{y})$.

Our approach makes a couple of assumptions. Theorem 1 assumes continuous differentiability and throughout this paper we assume the existence of a low dimensional active subspace (which may depend on on $€$, e.g., see Section 5.3), while Monte Carlo sampling cannot take advantage of such structure. However, Algorithm 1, according to Theorem 3, can identify a suitable approximation to the active subspace with $O(1 / \epsilon)$ samples of the gradient of $Q(\boldsymbol{y})$. Utilizing the adjoint method for sensitivity analysis [2], the work needed to find the gradient is comparable to the work of solving for $Q(\boldsymbol{y})$, hence, the cost of finding the active subspace is comparable to $O(2 / \epsilon)$ Monte Carlo samples.

The resulting low-dimensional projected problem can be attacked with a multitude of methods. Staying in the realm of random sampling, we could use any of the methods from the QMC family that have well-established convergence properties without the burden of additional assumptions. In low dimensions, the QMC method can find the projected integral in approximately $O(1 / \epsilon)$ evaluations of $Q(\boldsymbol{y})$, which, combined with the earlier estimate for Algorithm 1, leads a total cost of $O(3 / \epsilon)$. Thus, for this particular class of reducible differentiable problems, Algorithm 1 combined with QMC approach has total cost of $O(3 / \epsilon)$, which is much lower than the corresponding $O\left(1 / \epsilon^{2}\right)$ associated with the classical Monte Carlo. In addition, $S G$ and $S C$ methods can exploit regularity of the projected problem and hence achieve even faster convergence.

## Remark 6. (Convergence of rational functions)

Here we consider a special case where $\epsilon=0, Q(\boldsymbol{y})$ is a rational function of the components of $\boldsymbol{y}$ and the PDF $\rho(\boldsymbol{y}) \in \mathbb{L}^{1}\left(\mathbb{R}^{N}\right)$. According to Corollary 1 there is a unique pair of $\Lambda^{a}$ and $\Lambda^{p}$, namely $\Lambda^{a}=\operatorname{span}\{\nabla Q(\boldsymbol{y})\}$ and $\Lambda^{p}=\operatorname{span}\{\nabla Q(\boldsymbol{y})\}^{\perp}$. Thus for each set of sample points we pick $\Lambda_{k}^{p} \perp\left\{\nabla Q_{i}\left(\boldsymbol{y}_{i}\right\}_{i=1}^{k}\right.$, i.e. $J(T)$ is given by (24). For every $v \in \mathbb{R}^{N}$ define

$$
S_{v}=\left\{\boldsymbol{y} \in \mathbb{R}^{N}:\langle\nabla Q(\boldsymbol{y}), v\rangle \neq 0\right\}
$$

and consider $\alpha(\boldsymbol{y})=\langle\nabla Q(\boldsymbol{y}), v\rangle$, which is in itself a rational function. A rational function is zero either everywhere or on a set of measure zero, and thus if $\alpha(\boldsymbol{y}) \equiv 0$ then $S_{v}=\emptyset$, otherwise the measure of $S_{v}$ is equal to 1 . Therefore, $P\left(S_{v}\right)=\int_{S_{v}} \rho(\boldsymbol{y}) d \boldsymbol{y}$ can attain only the values of 0 and 1 , and furthermore, if $P\left(S_{v}\right)=0$, then $v \perp \nabla Q(\boldsymbol{y})$ for all $\boldsymbol{y}$ and hence $v \in \Lambda^{p}$. Conversely, if $v \notin \Lambda^{p}$, then $P\left(S_{v}\right)=1$.

Consider a realization of Algorithm 1 and suppose that at step $k \geq n=\operatorname{dim}\left(\Lambda^{a}\right)$ we have $\Lambda_{k}^{a}$, which is incomplete, i.e., there is a nonzero vector $v \in \Lambda^{a} \cap \Lambda_{k}^{p}$. By definition of $\Lambda^{a}$, we have that $v \in \operatorname{span}\{\nabla Q(\boldsymbol{y})\}$, however, since
$v \notin \Lambda_{k}^{a}$, we have that condition (18) fails for all $\boldsymbol{y} \in S_{v}$. Assuming the notation of Theorem 2, we note that $S_{v} \subset Z_{k}$ and according to the same theorem, for $k \geq n$, we have that $m_{k} \geq P\left(S_{v}\right)$ and since $v \notin \Lambda^{p}$ we have that

$$
m_{k} \geq P\left(S_{v}\right)=1
$$

Using $m_{k}=1$ in Eq. (31) we conclude that the Algorithm 1 can identify the active subspace in a finite number of steps, namely $n=\operatorname{dim}\left(\Lambda^{a}\right)$.

## 5. NUMERICAL EXAMPLES

In this section, we present three numerical examples to illustrate our numerical approach and to validate our theoretical results. The first example considers a linear problem within the classical Karhunen-Loève setting. We demonstrate that in the worst case scenario, the bound on the convergence rate predicted by Theorem 3 is indeed sharp. In the second example, we consider an output of interest with a very low dimensional active subspace. Our method consistently identifies that subspace with very few samples and hence we achieve convergence with number of over-samples $d_{k}$ that is orders of magnitude less than what is required by Theorem 3. This demonstrates that for some problems, we can find the active subspace long before we have a sample size that can give us sufficient confidence in the result. However, in all cases, our method converges significantly faster then competing sampling approaches. Finally, we apply our method to a one-dimensional physical reactor problem, described in Example 2, with a significantly large number of cross-section uncertainties. We demonstrate that for moderate error tolerance $\epsilon$, we can find a very low dimensional active subspace that preserves the dynamics of the output of interest. However, when we tighten the tolerance, the size of the active subspace grows very fast, exhibiting how our gradient-based reduction technique can only be successfully applied to the neutronics problem for moderate error tolerance.

### 5.1 Application to Classical Finite-Dimensional Karhunen-Loéve Expansion

Consider the classical finite-dimensional Karhunen-Loéve problem,

$$
Q(\boldsymbol{y})=\boldsymbol{y}^{T} S \boldsymbol{y}, \boldsymbol{y} \in \mathbb{R}^{N}, \text { and } \rho(\boldsymbol{y})=\frac{e^{-\frac{1}{2}\|\boldsymbol{y}\|^{2}}}{(2 \pi)^{N / 2}}
$$

where $N=100$ and $S \in \mathbb{R}^{100 \times 100}$ is a symmetric positive definite matrix. In order to illustrate the theoretical results given by Theorem 3, we take $S$ to be a random matrix of size 100 , specifically generated through the following procedure:

1. Generate a matrix $R \in \mathbb{R}^{100 \times 100}$, where the elements of $R$ are sampled from a standard Gaussian distribution with zero mean and unit variance;
2. Define $S=R^{T} R$;
3. Scale the eigenvalues of $S$ with use of a sequence $\left\{s_{i}\right\}_{i=1}^{N}$ to enforce rapid eigenvalue decay.

Figure 2 shows the eigenvalue decay of the two test matrices that we use for the discussion below.
In order to validate the results form Theorem 3, we need to consider the expectation and variance of the error in the projection associated with Algorithm 1 with respect to the distribution of the samples. To that end, we define

$$
e(d): \mathbb{N} \rightarrow \mathbb{R}
$$

where $e(d)$ is a possible realization of the error associated with $d$ number of over-samples. Obviously, $e(d)$ is a random variable for each $d$ and hence we are interested in the statistical expectation $\mathbb{E}[e(d)-\epsilon]$ and variance $V[e(d)]$. In order to compute the statistics, we execute the algorithm multiple times and gather samples for $e(d)$ in a manner consistent with the Monte Carlo method.

Volume 5, Number 1, 2015


FIG. 2: Log plot of the eigenvalue decay for the two random test matrices. The first 12 eigenvalues decay as $s_{i}=$ $\exp (-1.727(i-1))$, while the remaining ones are set to $10^{-14}$ (left) and $10^{-9}$ (right). All eigenvalues are normalizes so that $\lambda_{\max }=1$.

Due to the structure of the problem, the exact expectation of $Q(\boldsymbol{y})$ can be computed by summing the eigenvalues of $S$. Furthermore, suppose that we want to project the QoI onto a lower dimensional active space, then we can take any orthonormal basis for $\Lambda^{a}$ and if we arrange the basis into the columns of $V$, then the expectation of the projected QoI is the sum of the eigenvalues of $V^{T} S V$. Finding the eigenstructure of a 100 dimensional matrix is trivial, which makes it feasible to compute a large number of realizations.

We apply Algorithm 1 using $J(T)$ described in Section 4 Example 4 with cutoff $\lambda=10^{-4}$. The error bound in Theorem 3 uses $\epsilon$, which is not equal to $\lambda$, and thus we utilize the heuristic estimate

$$
\epsilon \approx \min _{d}(\mathbb{E}[e(d)])
$$

In every realization of Algorithm 1, we seek a reduced system of smallest dimension that will satisfy condition (28). However, the probabilistic nature of the algorithm results in an active sub-space of variable size. In Table 1, we show the statistics for the size of the reduced system.

According to Theorem 3, we should observe the relation

$$
\mathbb{E}[e(d)] \leq \epsilon+\frac{C}{d+1}, \quad V[e(d)] \leq \frac{C^{2}}{(d+1)^{2}}
$$

and Figs. 3 and 4 give the computed decay rate. In both cases, the error obeys the convergence bounds of Theorem 3 . Furthermore, we note that for one of our random matrices, the rate of convergence $O\left(d_{k}^{-1}\right)$ is indeed sharp. Even though Theorem 2 suggests the possibility of exponential convergence, in the general case, we cannot assume faster linear rate.

TABLE 1: Size of the reduced system

|  | Minimum | Maximum | Mean | Median |
| :---: | :---: | :---: | :---: | :---: |
| Example with faster eigenvalue decay | 4 | 6 | 4.3 | 4 |
| Example with slower eigenvalue decay | 6 | 9 | 7.2 | 7 |



FIG. 3: Log plot of $\mathbb{E}[e(d)-\epsilon]$ and $V[e(d)-\epsilon]$ as a function of $d$. The decay rate for the first 20 entries is -1.73 and -3.02 .


FIG. 4: Log plot of $\mathbb{E}[e(d)-\epsilon]$ and $V[e(d)-\epsilon]$ as a function of $d$. The decay rate for the entries 2 through 11 is -1.02 and -2.47 , respectively.

### 5.2 Highly Reducible Random Parameter Problem

We present an example of a nonlinear problem with low dimensional active subspace and we consider the performance of several common methods compared to our gradient-based reduction approach. Let $\eta(x)$ be a piecewise constant approximation to an uncorrelated noise field, given by

$$
\eta(\boldsymbol{y}, x)=\sum_{i=1}^{N} \boldsymbol{y}_{i} I_{i}(x)
$$

where each $y_{i}$ is uniformly distributed in $[-1,1]$ and $I_{i}(x)$ is the indicator function of the interval $[(i-1) / N, i / N]$, i.e.,

$$
I_{i}= \begin{cases}1, & x \in\left[\frac{(i-1)}{N}, \frac{i}{N}\right] \\ 0, & x \notin\left[\frac{(i-1)}{N}, \frac{i}{N}\right]\end{cases}
$$

Define the operator $\mathcal{C}: \mathbb{L}^{2}([0,1]) \rightarrow \mathbb{L}^{2}([0,1])$ by

$$
\mathcal{C}(f)=\sum_{k=1}^{\infty} \frac{1}{(k \pi)^{6}} \sin (k \pi x) \int_{0}^{1} \sin (k \pi s) f(s) d s
$$

and consider the $\mathbb{L}^{2}$ inner product $\langle\eta, \mathcal{C}(\eta)\rangle_{\mathbb{L}^{2}}=\int_{0}^{1} \eta \mathcal{C}(\eta) d x$ such that the QoI is given by

$$
Q(\boldsymbol{y})=e^{-\frac{1}{2}\langle\mathfrak{\eta}, \mathcal{C}(\mathfrak{\eta})\rangle_{\mathbb{L}^{2}}}
$$

Here we take $N=1000$, the random parameter domain is $\boldsymbol{y} \in \Gamma=[-1,1]^{N}$, and probability distribution is $\rho(\boldsymbol{y})=$ $1 / 2^{N}$. The gradient can be computed by formal differentiation

$$
\frac{\partial Q(\boldsymbol{y})}{\partial \boldsymbol{y}_{i}}=Q(\boldsymbol{y})\left\langle I_{i}, \mathcal{C}\left(I_{i}\right)\right\rangle_{L^{2}} \boldsymbol{y}_{i} \quad i=1,2, \cdots, N
$$

and the goal is to compute $\mathbb{E}[Q]=\int_{\Gamma} Q(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y}$.
We compare the accuracy of different approaches for computing the expectation of the QoI, against the results of a brute force random sampling with $10^{6}$ realizations. These include sensitivity analysis, Karhunen-Loéve projection, and our Algorithm 1. For the two projection schemes, if we project the QoI on a low-dimensional active subspace, we could apply the collocation method, however, this is beyond the scope of our paper. Since we are only interested in the error associated with the projection, we use Monte Carlo sampling to compute the expected value of both the full order and projected QoI.

We can compute the gradient $\nabla_{\boldsymbol{y}} Q(\boldsymbol{y})$ and hence we can utilize sensitivity analysis [2]. The nominal value for the random vector is $\mathbb{E}[\boldsymbol{y}]=\mathbf{0}$ and $Q(\mathbb{E}[\boldsymbol{y}])=Q(\mathbf{0})=1$. The nominal gradient is $\nabla_{\boldsymbol{y}} Q(\mathbf{0})=\mathbf{0}$ and therefore the method yields the approximation $\mathbb{E}[Q] \approx 1$ with variance of 0 . However, the quantity of interest in this example is nonlinear and brute force Monte Carlo sampling gives $\mathbb{E}[Q] \approx 0.8832$. Sensitivity analysis results in error of more than $13 \%$ and this single point method does not offer a strategy to improve this approximation.

The operator $\mathcal{C}$ was purposely chosen to have very fast eigenvalue decay and the structure of $\nabla_{\boldsymbol{y}} Q(\boldsymbol{y})$ suggests that the active subspace will be associated with only the dominant eigenspace of $\mathcal{C}$, and therefore it is low-dimensional. However, we cannot rigorously apply the classical KL approach because $Q(\boldsymbol{y})$ is not a quadratic functional. Therefore, the KL error bounds will not be valid. Indeed, if we let $\Lambda^{a}$ be the space spanned by the first four dominant eigenvectors of $\mathcal{C}$, then the error predicted by KL is $\approx 1.0629 \times 10^{-4}$, while the actual error is $\approx 5.2550 \times 10^{-4}$. The KL approach greatly underestimates the projection error. If we let $\epsilon=10^{-4}$ and apply Algorithm 1, then we can take $\operatorname{dim}\left(\Lambda^{a}\right)=4$ and achieve an approximation that is a full order of accuracy better than the linear KL approach (see Fig. 5).

Finally, we apply Algorithm 1 and compare the results to the error bound in Theorem 3. Figure 5 superimposes four reduction tests with different values of $\epsilon$. The reduction algorithm needs only a few samples to approximate $\Lambda^{a}$ and reach the desired tolerance. According to Theorem 3, the expected value of the error is bounded by $\epsilon+O\left(M^{-1}\right)$ and the standard deviation is less than $O\left(M^{-1}\right)$, thus our confidence in the reduced model is low. However, even though the rate of $O\left(M^{-1}\right)$ is a conservative estimate, the decay is faster than the brute force sampling approach.

### 5.3 Neutron Transport with Stochastic Cross Sections

Consider the PDE with stochastic cross sections described in Section 2 Example 2 with deterministic domain illustrated on Fig. 6. We consider two "fuel rods" and a "control rod" between them; the space between the rods is filled with "coolant." The fuel-rod regions have large fission cross-sections, the control rod region has a large capture cross


FIG. 5: Convergence of Algorithm 1 vs the confidence estimate in Theorem 3. The reduction approach needs only a few samples to find a suitable active subspace and reach the desired tolerance $\epsilon=10^{-3,-4}$, thus the algorithm can identify the active subspace long before we have a sufficient confidence in the computed projection. The dimension of the active subspace is $\operatorname{dim}\left(\Lambda^{a}\right)=3$ when $\epsilon=10^{-3}$, and $\operatorname{dim}\left(\Lambda^{a}\right)=4$ when $\epsilon=10^{-4}$.


FIG. 6: Mockup reactor problem.
section, and the coolant interacts only weakly with the neutrons. As such, we define the indicator functions for each of the three materials

$$
I_{U}(x)=\left\{\begin{array}{ll}
1, & x \in \text { fuel-rod } \\
0, & \text { o.w. }
\end{array} \quad I_{B}(x)=\left\{\begin{array}{ll}
1, & x \in \text { control-rod } \\
0, & \text { o.w. }
\end{array} \quad I_{W}(x)=\left\{\begin{array}{ll}
1, & x \in \text { coolant } \\
0, & o . w .
\end{array},\right.\right.\right.
$$

where the letters $U, B$, and $W$ are chosen to abbreviate the regions of the domain based on the common material used for nuclear fuel (Uranium), control-rod (Boron), and coolant (Water). We model the uncertainty in the cross sections as an additive, scaled, uncorrelated piecewise constant field

$$
\begin{equation*}
\sigma_{p}(x ; \omega)=\bar{\sigma}_{p}(x)+c_{p}(x) \boldsymbol{y}_{p}(x ; \omega)=\sum_{l \in\{U, B, W\}} \bar{\sigma}_{p, l} I_{l}(x)+c_{p, l} I_{l}(x) \boldsymbol{y}_{p}(x ; \omega), \quad p \in\{s, c, f\}, \tag{32}
\end{equation*}
$$

where $\bar{\sigma}_{p, l}$ are the nominal cross sections and $c_{p, l}$ are the scaling factors. The numerical values that we used are given in Table 2 . Note that the actual cross-section values used are not the physical values associated with any real

TABLE 2: Nominal values and noise scaling for the cross sections

|  | Scatter |  | Capture |  | Fission |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Nominal | Scaling | Nominal | Scaling | Nominal | Scaling |
| Fuel rod | 1.0 | $\pm 0.30$ | 0.10 | $\pm 0.05$ | 7.0 | $\pm 6.0$ |
| Control rod | 1.0 | $\pm 0.03$ | 8.15 | $\pm 4.00$ | 0.0 | $\pm 0.0$ |
| Coolant | 0.1 | $\pm 0.05$ | 0.10 | $\pm 0.05$ | 0.0 | $\pm 0.0$ |

world materials, we are using artificially selected values with very large range, in order to make the problem more challenging.

We focus on the reactor criticallity problem given by Eq. (6). The discretization is achieved by virtue of a finite difference scheme: we take $n$ points in space $\left\{x_{i}\right\}_{i=1}^{n} \subset[0,1]$ and $m$ Clenshaw-Curtis [35] points $\left\{\mu_{j}\right\}_{j=1}^{m} \subset[-1,1]$ and we approximate the neutron flux by $\psi(x, \mu) \approx \psi\left(x_{i}, \mu_{j}\right)=\psi_{i}^{j}$. The convection operator is discretized via an up-winding scheme

$$
\mu_{j} \frac{\partial \psi}{\partial x}\left(x_{i}, \mu_{j}\right) \approx\left\{\begin{array}{cc}
\mu_{j} \frac{\psi_{i+1}^{j}-\psi_{i}^{j}}{x_{i+1}-x_{i}}, & \mu_{j}<0  \tag{33}\\
\mu_{j} \frac{\psi_{i}^{j}-\psi_{i-1}^{j}}{x_{i}-x_{i-1}}, & \mu_{j}>0
\end{array}\right.
$$

where we impose zero-Dirichlet boundary condition at the inflow, i.e., the reactor is shielded from external neutron sources. We discretize $\phi$ with the quadrature rule

$$
\begin{equation*}
\phi\left(x_{i}\right)=\frac{1}{2} \int_{-1}^{1} \psi\left(x_{i}, \mu\right) d \mu \approx \frac{1}{2} \sum_{j=1}^{m} w_{j} \psi\left(x_{i}, \mu_{j}\right) \approx \frac{1}{2} \sum_{j=1}^{m} w_{j} \psi_{i}^{j} \tag{34}
\end{equation*}
$$

where $w_{j}$ are the Clenshaw-Curtis weights. At each point of the domain, we represent the noise by three random variables associated with the three types of cross section parameters. Each sample has a uniform distribution on the canonical interval $[-1,1]$. Thus we have $\boldsymbol{y} \in \mathbb{R}^{3 \times n}$ and (32) becomes

$$
\sigma_{p}\left(x_{i}\right)=\bar{\sigma}_{p}\left(x_{i}\right)+c_{p}\left(x_{i}\right) \boldsymbol{y}_{p}^{i}, \quad p \in\{s, c, f\} .
$$

The discretized version of (6) can be written as

$$
T \psi_{i}^{j}+\left(S_{s}(\boldsymbol{y})+S_{c}(\boldsymbol{y})+S_{f}(\boldsymbol{y})\right) \psi_{i}^{j}=S_{s}(\boldsymbol{y}) D \psi_{i}^{j}+\frac{\nu}{k} S_{f}(\boldsymbol{y}) D \psi_{i}^{j}
$$

where $T$ is the discrete convection operator (33), $D$ is the integral operator (34), $S_{s}(\boldsymbol{y}), S_{c}(\boldsymbol{y}), S_{f}(\boldsymbol{y})$ are diagonal matrices with the cross sections, and $v=2.4$. More generally, the eigenvalue problem can be expressed as

$$
\begin{equation*}
A \psi=\lambda B \psi \tag{35}
\end{equation*}
$$

where $A=T+\left(S_{s}(\boldsymbol{y})+S_{c}(\boldsymbol{y})+S_{f}(\boldsymbol{y})\right)-S_{s}(\boldsymbol{y}) D, B=v S_{f}(\boldsymbol{y}) D$, and $\lambda=1 / k$. Recall that our goal is to estimate the expected values $\mathbb{E}[k(\boldsymbol{y})]$

$$
\begin{equation*}
\mathbb{E}[k(\boldsymbol{y})]=\int_{\Gamma} k(\boldsymbol{y}) \rho(\boldsymbol{y}) d \boldsymbol{y} \tag{36}
\end{equation*}
$$

where $\rho(\boldsymbol{y})$ is the uniform distribution on the hyper-cube $\Gamma=[-1,1]^{3 n}$. For our examples, we use discretizations $n=1000$ and $m=14$, which means that the integral (36) is in 3000 dimensions.

In order to apply Algorithm 1, we need a way to approximate $\partial \lambda / \partial \boldsymbol{y}_{i}$. Each of the operators $S_{s}(\boldsymbol{y}), S_{c}(\boldsymbol{y}), S_{f}(\boldsymbol{y})$ depend linearly on the uncertainty $\boldsymbol{y}$ and therefore can be easily differentiated. However, $\boldsymbol{\lambda}$ has a nonlinear dependence on $\boldsymbol{y}$. Suppose $\psi$ and $\lambda$ satisfy Eq. (35) and to simplify notation define

$$
\partial A_{i}=\frac{\partial A}{\partial y_{i}}, \quad \partial B_{i}=\frac{\partial B}{\partial y_{i}}, \quad \partial \psi_{i}=\frac{\partial \psi}{\partial y_{i}}, \quad \partial \lambda_{i}=\frac{\lambda}{\partial y_{i}}
$$

then formally differentiating (35) we have

$$
\begin{equation*}
\left(\partial A_{i}-\lambda \partial B_{i}\right) \psi+(A-\lambda B) \partial \psi_{i}=\partial \lambda_{i} B \psi . \tag{37}
\end{equation*}
$$

Let $\hat{\psi}$ be the left generalized eigenvector of $A \hat{\psi}=\lambda B \hat{\psi}$ associated with $\lambda$, then we can multiply (37) by $\hat{\psi}^{T}$ and simplify

$$
\hat{\psi}^{T}\left(\partial A_{i}-\lambda \partial B_{i}\right) \psi=\partial \lambda_{i} \hat{\psi}^{T} B \psi
$$

which allows us to solve for $\partial \lambda_{i}$.
We first take 2000 samples of $\nabla_{\boldsymbol{y}} k(\boldsymbol{y})$ and arrange them into the columns of a snapshot matrix, revealing the decay of the singular values. Figure 7 shows the initial rapid decay, followed by a "plateau." Therefore, for a moderate choice of $\epsilon$, we expect to find a low-dimensional active subspace, however, if we decrease the tolerance, the dimension of $\Lambda^{a}$ should increase dramatically.

We apply the reduction algorithm with three different values of $\epsilon$, i.e., $10^{-3}, 10^{-4}$, and $10^{-4}$, and we give the results in Table 3. In each case, we stopped the iteration when the number of over-samples reached around 1000. For


FIG. 7: Decay of the singular values of a set of sample for $\nabla k(\mathbf{y})$. We observe a sharp decay until singular value 6 , followed by a "plateau" until singular value 90 . Hence, for moderate values of the tolerance $\epsilon$, we expect to identify an active subspace $\Lambda^{a}$ with low dimension; if we decrease $\epsilon$, we expect the dimension $\Lambda^{a}$ to increase dramatically.

TABLE 3: Results from three realizations of Algorithm 1 for three different values of $\epsilon$. We use discretization parameters $n=$ 1000 and $m=14$ resulting in a problem with 3000 dimensions of uncertainty. All three realizations were terminated when $d_{k}$ reached 1000

|  | $\epsilon$ | $\operatorname{dim}\left(\Lambda^{a}\right)$ | Error |
| :---: | :---: | :---: | :---: |
| Test 1 | $10^{-2}$ | 2 | $1.6 \times 10^{-3}$ |
| Test 2 | $10^{-3}$ | 4 | $3.9 \times 10^{-4}$ |
| Test 3 | $10^{-4}$ | 103 | $9.7 \times 10^{-6}$ |

the largest tolerance, we can approximate the expected value to the desired tolerance by keeping only 2 out of the 3000 dimensions. If we decrease the tolerance to $10^{-3}$, the size of the active subspace increases to 4 , however, the error in the projection decreases. As expected, when we decrease the tolerance to $10^{-4}$, the size of the active subspace increases to 103 and even though the error in the projection is considerably lower, 103 is still prohibitively large to allow for efficient application of any low dimensional integration scheme. A low-dimensional $\Lambda^{a}$ can be found only for moderate values of $\epsilon$.

A significant proportion of the dynamics of $k(\boldsymbol{y})$ is dominated by only a few modes. Therefore, we can approximate $\mathbb{E}[k(\boldsymbol{y})]$ to less than $1 \%$ by projecting the QoI onto an active subspace $\Lambda^{a}$ of no more than four dimensions. Furthermore, Algorithm 1 is a reliable method for identifying $\Lambda^{a}$ at an approximately linear computational cost. However, if higher degree of accuracy is desired, then the additional dynamics that need to be identified are associated with a much larger number of directions. For small $\epsilon$, the dimension of the active subspace increases dramatically, which renders infeasible the application of sparse grids collocation or other fast convergent low-dimensional integration techniques.

## 6. CONCLUSIONS

In this work, we presented a projection approach that utilizes the gradient, for forward uncertainty quantification of high-dimensional problems. We use Monte Carlo sampling for the sensitivity of the output of interest (i.e., the gradient at the sample point), we use this information to identify a low-dimensional active subspace and project the output of interest in a manner similar to the classical Karhunen-Loéve expansion. However, our method produces results that are valid for problems with large range of uncertainty and hence more accurate than the single-point sensitivity analysis. Moreover, unlike the classical Karhunen-Loéve expansion, our error bounds are valid for highly nonlinear problems. Finally, if the resulting projected problem is moderate-dimensional, we could apply conventional quasi-Monte Carlo or stochastic collocation sampling techniques and benefit from their fast convergence rate, which leads to total computational cost that is much lower than classical brute force Monte Carlo. The success of out method is contingent upon the existence of a low-dimensional active subspace, which in turn depends on the structure of the problem and the choice of tolerance $\epsilon$. Some problems can only be reduced for a moderate error tolerance. Furthermore, in some cases, our error bounds can be overly conservative producing low confidence in an otherwise accurate result. Nonetheless, this method can be successfully applied to PDE models with large number of uncertain parameters, such as the criticality of the nuclear reactor with a large number of uncertain cross-sections.

## ACKNOWLEDGMENTS

The preparation of the article was supported in part by the Office of Science of the US Department of Energy under grant number ERKJE45; by the US Air Force Office of Scientific Research under grant number 1854-V521-12; and by the Lab- oratory Directed Research and Development program at the Oak Ridge National Laboratory (ORNL). The ORNL is operated by UT-Battelle, LLC, for the US Department of Energy under Contract DE-AC05-00OR22725.

## REFERENCES

1. Evans, T. M., Stafford, A. S., Slaybaugh, R. N., and Clarno, K. T., Denovo: A new three-dimensional parallel discrete ordinates code in scale, Nuclear Technology, 171(2):171-200, 2010.
2. Cacuci, D. G., Sensitivity and Uncertainty Analysis, Chapman \& Hall/CRC, Boca Raton, 2003.
3. Fishman, G., Monte Carlo, Springer Series in Operations Research, Springer-Verlag, New York, 1996.
4. Niederreiter, H., Random Number Generation and Quasi-Monte Carlo Methods, Vol. 63 of CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
5. Helton, J. and Davis, F., Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems, Reliab. Eng. Syst. Safety, 81:23-69, 2003.
6. Barthelmann, V., Novak, E., and Ritter, K., High dimensional polynomial interpolation on sparse grids, Adv. Comput. Math., 12(4):273-288, 2000.
7. Smolyak, S., Quadrature and interpolation formulas for tensor products of certain classes of functions, Dokl. Akad. Nauk SSSR, 4:240-243, 1963.
8. Griebel, M., Adaptive sparse grid multilevel methods for elliptic PDEs based on finite differences, Comput., 61(2):151-179, 1998.
9. Gerstner, T. and Griebel, M., Dimension-adaptive tensor-product quadrature, Comput., 71:65-87, 2003.
10. Xiu, D. and Hesthaven, J., High-order collocation methods for differential equations with random inputs, SIAM J. Sci. Comput., 27:1118-1139, 2005.
11. Nobile, F., Tempone, R., and Webster, C. G., A sparse grid stochastic collocation method for partial differential equations with random input data, SIAM J. Numer. Anal., 46(5):2309-2345, 2008.
12. Nobile, F., Tempone, R., and Webster, C. G., An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data, SIAM J. Numer. Anal., 46(5):2411-2442, 2008.
13. Eldred, M., Webster, C. G., and Constantine, P. G., Evaluation of non-intrusive approaches for Wiener-Askey generalized polynomial chaos, Am. Inst. Aer. Astron. Pap., 1892:1-23, 2008.
14. Doostan, A. and Iaccarino, G., A least-squares approximation of partial differential equations with high-dimensional random inputs, J. Comput. Phys., 228(12):4332-4345, 2009.
15. Xiu, D. and Karniadakis, G. E., Modeling uncertainty in steady state diffusion problems via generalized polynomial chaos, Comput. Methods Appl. Mech. Eng., 191(43):4927-4948, 2002.
16. Ghanem, R. G. and Spanos, P. D., Stochastic Finite Elements: A Spectral Approach, Springer-Verlag, New York, 1991.
17. Klimke, A. and Wohlmuth, B., Algorithm 847: Spinterp: piecewise multilinear hierarchical sparse grid interpolation in matlab, ACM Trans. Math. Software (TOMS), 31(4):561-579, 2005.
18. Ma, X. and Zabaras, N., An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations, J. Comput. Phys., 228(8):3084-3113, 2009.
19. Gunzburger, M., Webster, C. G., and Zhang, G., An adaptive wavelet stochastic collocation method for irregular solutions of partial differential equations with random input data, Lect. Notes Comput. Sci. Eng., 97:137-170, 2014.
20. Babuška, I. M., Tempone, R., and Zouraris, G. E., Galerkin finite element approximations of stochastic elliptic partial differential equations, SIAM J. Numer. Anal., 42(2):800-825, 2004.
21. Webster, C. G., Zhang, G., and Gunzburger, M., An adaptive sparse-grid iterative ensemble kalman filter approach for parameter field estimation, Int. J. Comput. Math., 91(4):798-817, 2014.
22. Zhang, G., Lu, D., Ye, M., Gunzburger, M., and Webster, C., An adaptive sparse-grid high-order stochastic collocation method for bayesian inference in groundwater reactive transport modeling, Water Resour. Res., 49(10):6871-6892, 2013.
23. Burkardt, J., Gunzburger, M., Webster, C., and Zhang, G., A hyper-spherical adaptive sparse-grid method for high-dimensional discontinuity detection, Tech. Rep. ORNL/TM-2012/186, Oak Ridge National Laboratory, 2012.
24. Loève, M., Probability Theory. I, 4th ed., Springer-Verlag, New York, 1977.
25. Loève, M., Probability Theory. II, 4th ed., Springer-Verlag, New York, 1978.
26. Abdel-Khalik, H. S., On Non-linear Reduced order Modeling, International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, Rio de Jeneiro, Brazil, LAS/ANS, May 8-12, 2011.
27. Abdel-Khalik, H. S. and Hite, J. M., Reduced order modeling: Tensor-free expansion for nonlinear features identification, Trans. Am. Nucl. Soc., 104:343-346, 2011.
28. Bang, Y., Abdel-Khalik, H. S., and Hite, J. M., Hybrid reduced order modeling applied to nonlinear models, Int. J. Numer. Methods Eng., 91(9):929-949, 2012.
29. Bui-Thanh, T., Burstedde, C., Ghattas, O., Martin, J., Stadler, G., and Wilcox, L., Extreme-scale UQ for Bayesian inverse problems governed by PDEs, Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, Salt Lake City, Utah, IEEE, Nov. 10-16, 2012.
30. Constantine, P. G. and Wang, Q., Input subspace detection for dimension reduction in high dimensional approximation, eprint arXiv:1202.3508, Feb. 2012.
31. Liberty, E., Woolfe, F., Martinsson, P.-G., Rokhlin, V., and Tygert, M., Randomized algorithms for the low-rank approximation of matrices, PNAS, 104(51):20167-20172, 2007.
32. Martinsson, P.-G., Rockhlin, V., and Tygert, M., A randomized algorithm for the approximation of matrices, DTIC Document,
33. 
34. Babuška, I., Nobile, F., and Tempone, R., A stochastic collocation method for elliptic partial differential equations with random input data, SIAM J. Numer. Anal., 45(3):1005-1034, 2007.
35. Lewis, E. and Miller, W., Computational Methods of Neutron Transport, John Wiley and Sons, Inc., New York, NY, 1984.
36. Davis, P. and Rabinowitz, P., Methods of Numerical Integration, 2nd ed., Dover, New York, 2007.

[^0]:    *Correspond to C. G. Webster, E-mail: webstercg@ornl.gov, URL: http://csm.ornl.gov/~cgwebster

