Polynomial regression with derivative information in nuclear reactor uncertainty quantification*

M. Anitescu[a†], O. Roderick[a], P. Fischer[a], W.S. Yang[b]

[a] Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, USA

[b] Nuclear Engineering Division, Argonne National Laboratory, Argonne, IL, USA

Abstract. We introduce a novel technique of uncertainty quantification using polynomial regression with derivative information and apply it to analyze the performance of a model of a sodium-cooled fast reactor. We construct a surrogate model as a goal-oriented projection onto an incomplete space of polynomials, find coordinates of projection by collocation, and use derivative information to reduce the number of sample points required by the collocation procedure. This surrogate model can be used to estimate range, sensitivities and the statistical distribution of the output. Numerical experiments show that the suggested approach is significantly more computationally efficient than random sampling, or approaches that do not use derivative information, and that it has greater precision than linear models.

1. Introduction.

Uncertainty quantification (i.e., the task of relating the available information on the uncertainties in the input parameters to the resulting variation in the arbitrary chosen outputs of the model) plays an important role in the currently expanding field of nuclear engineering, where sophisticated simulation codes are used to implement mechanisms of observation and control, increase the efficiency in the use of resources, and ensure safety[1][2][3].

In any complex simulation code, the inputs and the intermediate parameters include experimental errors and design simplifications, resulting in uncertainties in the outputs. The usual difficulties in modeling the work of the nuclear reactor models include the large size of the solved systems of equations, the nonlinearity, and the implicit dependence of the equations on input parameters. As a result, one can normally afford to run the computational model only for a small number of scenarios involving the values of the physical parameters. In addition, although the information on the behavior of parameters is available in formats convenient for experimental physics and engineering purposes, such formats are not necessarily appropriate for uncertainty analysis.

Traditionally, the influence of the uncertainties in the inputs on the outputs is described either by linear approximations (linear sensitivity) using first-order derivative information and disregarding the non-linear effects[4], or by random sampling methods (pure Monte-Carlo, or using other sampling strategies[5]) requiring many runs of the model[6]. We introduce a novel technique of uncertainty quantification (UQ) using polynomial regression with derivative information (PRD) and use it to predict the effect of uncertainties in multiple physical parameters on the performance of a model of sodium-cooled fast reactor.

The experiments show that the suggested method is significantly more computationally efficient than random sampling, and it performs with greater precision than linear approximations. The constructed polynomial approximation correctly reproduces the range of the output values, and the most

* This work was supported by Contract No. DE-AC02-06CH11357 of the U.S. Department of Energy.
† Current address: MCS division, Argonne National Laboratory, 9700 S. Cass Ave, Argonne, IL 60439.

significant sensitivites of the output with respect to parameters. This indicates the possibility of using the polynomial approximation for such tasks as estimating the statistical distribution of the output, determining the confidence and tolerance levels for the quantities of interest, making design decisions, performing verification, validation and safety analysis for reactor models.

We also note that the method is very flexible, and can be applied to a wide class of reactor models, and for arbitrary inputs and outputs. The associated development cost consists only of introducing a convenient uncertainty structure, and obtaining first-order derivative information.

2. Uncertainty quantification by surrogate models

2.1 Definition of the problem

We first present the mathematical model in the most general form, as a system of algebraic-differential equations:

$$\begin{aligned} &F(T,R)=0 \\ &R=R(T)\cdot(1+\Delta R(T,\alpha)) \\ &J=J(T) \end{aligned} \tag{1}$$

Here, the variables $T=(T_1,T_2,...,T_n)$ characterize the model state (for example, the temperature field of the reactor core). The expressions for the dependence of intermediate physical parameters $R=(R_1,R_2,...,R_N)$ (such as heat capacity, thermal conductivity, density) on the state of the model include experimental errors $\Delta R=(\Delta R_1,\Delta R_2,...,\Delta R_N)$. An observation of interest on the state of the model is expressed by the merit function $J(T)$. The dependency of the experimental errors on the state of the model is described by a set of stochastic uncertainty quantifiers $\alpha=(\alpha_1,\alpha_2,...,\alpha_m)$. For a given merit function, we redefine the output as a function of uncertainty quantifiers, $\Im(\alpha):=J(T)$, and use a polynomial approximation $\hat{\Im}(\alpha)\approx\Im(\alpha)$ as a computationally cheap alternative to evaluation of the full model.

The first task is to construct a convenient parametric representation of uncertainty. A typical example is based on reviewing the literature on material properties. In the available sources[7][8][9], the dependencies of thermodynamical properties on temperature are estimated from the experimental data, and provided in the Laurent series form

$$R=\sum_i r^{(i)}T^i\left(1+\frac{\Delta R}{R_0}\right) \tag{2}$$

Here, $R_0=\sum_i r^{(i)}T^i$ is the reference value, and $\frac{\Delta R}{R_0}$ is the relative error uncertainty term, the value of which also depends on temperature, and is recorded only for a few values of T, in the form

$$\xi_{\min}\le\frac{\Delta R(T)}{R_0(T)}\le\xi_{\max} \tag{3}$$

There is no information about the correlations between the values of uncertainty at different temperatures. To represent the relative error term, we choose a well-conditioned set of basis functions $\{C(T)\}$ (for example, a set of Chebyshev polynomials), and expand:

$$\frac{\Delta R}{R_0}=\sum_{j=0}\alpha_j C_j(T) \tag{4}$$

This expansion can be stopped at any term, depending on how much nonlinearity in (2) we would like to preserve in the representation. In the end, a few uncertainty quantifiers α are assigned to each physical parameter.

The region of admissible values of the uncertainty quantifiers, $\Omega \subset R^m$, is estimated by large-scale sampling inside of a sufficiently large (multidimensional, cubic) region, and rejecting the value combinations that contradict (3), with pessimistic extrapolations of (3) used for the values of T for which the relative error has not been observed.

The introduced uncertainty structure is consistent with the geometry and physics of the model: a determenistic, smooth dependence of the error on the model state (4) prevents unrealistic differences in the values of parameters used in the description of the neighboring geometric elements, or of the interdependent physical processes. The structure is very flexible: the algebraic structure of the expansion (3), or the shape of the validity region can be adjusted to fit additional experimental data.

Given that specific statistical information on the distribution of parameters is unavailable, we assume uniform distribution of uncertainty quantifiers inside the validity region. This assumption is conservative, but not as pessimistic as the worst-case assumption, according to which only the values of uncertainty quantifiers leading to largest relative error values are realized.

2.2 Polynomial regression with derivative information

Polynomial regression, or a polynomial chaos expansion[10] is a standard approach for constructing flexible, nonlinear approximations with convenient algebraic structure. To approximate a model output by polynomial regression, we choose a set $\{\Psi_q\}$ of multivariate polynomials and set

$$\hat{\Im}(\alpha) = \sum_q x_q \Psi_q(\alpha) \tag{5}$$

We use Hermite multivariable polynomial basis given by

$$\Psi_q(\alpha_1, \alpha_2, \alpha_3, ...) = \prod_i H^{k_i}(\alpha_i)$$
$$H^{(0)}(\alpha_i) = 1 \qquad H^{(1)}(\alpha_i) = 2\alpha_i \qquad H^{(2)}(\alpha_i) = 4\alpha_i^{\,2} - 2 \tag{6}$$
$$H^{(3)}(\alpha_i) = 8\alpha_i^{\,3} - 12\alpha_i \qquad H^{(4)}(\alpha_i) = 16\alpha_i^{\,4} - 48\alpha_i^{\,2} + 12 \qquad ...$$

By default, all polynomials up to a fixed total degree are used. The size of the basis grows with an increase of this maximal degree r, $O(n^r)$ asymptotically. For example, there are $\frac{1}{6}n^3 + n^2 + \frac{11}{6}n + 1$ polynomials of degree up to 3 on n variables.

In a setup without derivative information, the coefficients x_q are found by enforcing the polynomial fitting conditions at the nodes $S_1, S_2, ..., S_m$, resulting in a system of linear equations

$$\begin{pmatrix} \Psi(S_1) \\ \Psi(S_2) \\ \vdots \\ \Psi(S_m) \end{pmatrix} \cdot \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{pmatrix} = \begin{pmatrix} \Im(S_1) \\ \Im(S_2) \\ \vdots \\ \Im(S_m) \end{pmatrix} \tag{7}$$

where $S_i = (\alpha_{i1}, \alpha_{i2}, ..., \alpha_{iN})$ is an uncertainty state of the system, i.e. a point in the validity region,

and $\Im(S_i) = \Im(\alpha_{i1}, \alpha_{i2}, ..., \alpha_{iN})$ is the corresponding exact value of the output. The rows of the linear system (7) are constructed by evaluating all the basis functions at different uncertainty states:

$$\Psi(S_i) = \left(\Psi_1(\alpha_{i1}, \alpha_{i2}, ..., \alpha_{iN}), \Psi_2(\alpha_{i1}, \alpha_{i2}, ..., \alpha_{iN}), ..., \Psi_m(\alpha_{i1}, \alpha_{i2}, ..., \alpha_{iN})\right) \tag{8}$$

The computational cost of the collocation consists essentially of the evaluation of the right side, requiring m runs of the full model; the other operations have comparatively negligible computational costs. We show that it is possible to use each full model evaluation more efficiently, by also using first-order sensitivities of the output function to find the collocation coefficients. We call this novel approach Polynomial Regression with Derivative information (PRD).

We compute derivatives of the output function, and add derivative information to this polynomial fitting matrix. Then, each right-side entry $\Im(S_i)$ will generate a subcolumn of n entries $\frac{d\Im(S_i)}{d\alpha_j}$, $j = 1,2,...,n$, providing right-side information for several rows at once; resulting in an augmented system of $m(n+1)$ equations:

$$\begin{pmatrix} \Psi_1(S_1) & \Psi_2(S_1) & \cdots \\ \frac{\partial\Psi_1(S_1)}{\partial\alpha_1} & \frac{\partial\Psi_2(S_1)}{\partial\alpha_1} & \cdots \\ \frac{\partial\Psi_1(S_1)}{\partial\alpha_2} & \frac{\partial\Psi_2(S_1)}{\partial\alpha_2} & \cdots \\ \vdots & & \\ \frac{\partial\Psi_1(S_1)}{\partial\alpha_n} & \frac{\partial\Psi_2(S_1)}{\partial\alpha_n} & \cdots \\ \Psi_1(S_2) & \Psi_1(S_2) & \cdots \\ \frac{\partial\Psi_1(S_2)}{\partial\alpha_1} & \frac{\partial\Psi_2(S_2)}{\partial\alpha_1} & \cdots \\ \vdots & & \\ \Psi_1(S_m) & \Psi_2(S_m) & \cdots \\ \vdots & & \\ \frac{\partial\Psi_1(S_m)}{\partial\alpha_1} & \frac{\partial\Psi_2(S_m)}{\partial\alpha_1} & \end{pmatrix} \cdot x = \begin{pmatrix} \Im(S_1) \\ \frac{\partial\Im(S_1)}{\partial\alpha_1} \\ \frac{\partial\Im(S_1)}{\partial\alpha_2} \\ \vdots \\ \frac{\partial\Im(S_1)}{\partial\alpha_n} \\ \Im(S_2) \\ \frac{\partial\Im(S_2)}{\partial\alpha_1} \\ \vdots \\ \Im(S_m) \\ \\ \frac{\partial\Im(S_2)}{\partial\alpha_m} \end{pmatrix} \tag{9}$$

The system (9) is then solved using a generalized pseudo-inverse approach[11] to account for either type of ill-posedness (overdetermined or underdetermined).

In comparison with an approach without derivative information, the minimal required number of sample points drops by a factor of $1/(1+n)$. An additional computation advantage of the augmented system is its sparse structure (many partial derivatives of basis polynomials are identically equal to zero). In practice, it turns out that finding all first-order partial derivatives for a sample point is computationally more expensive by only a small factor compared to adding another point. This bound on computational overhead is also theoretically confirmed[12]. With an increase in the number of sources of uncertainty, the computational advantage achieved by using the derivative information overtakes the computational cost of obtaining the derivatives.

We note that the use of computationally cheap derivative information to construct surrogate models was previously described in [13], but without relating the idea to high-order polynomial interpolation, and for a very different application. Polynomial regression for parametric analysis and uncertainty quantification was also used previously[14][15], but without the use of our uncertainty structure, or derivative information.

2.3 Basis truncation

Determination of the size of the basis set is a tradeoff decision: the nonlinear dependence of the output on parameters implies the use of high-degree polynomials, but this also results in an unrealistically large size of the required training set. The use of derivative information only partially resolves the contradiction. To further decrease the size of the required training set, we also use a smaller basis, which includes high-order polynomials only in some important variables (and excludes high-order polynomials that depend on variables of low importance only).

We use the following simple, first-order test of importance: if we observe

$$\left|\frac{\partial \Im(S)}{\partial \alpha_i}\sigma_i^2\right| > \left|\frac{\partial \Im(S)}{\partial \alpha_j}\sigma_j^2\right| \tag{10}$$

at a few representative uncertainty states S, we conlcude that the uncertainty quantifier α_i is relatively more important for correct representation of the output. Here, σ_i^2 is the variance of the parameter, if available; otherwise estimated assuming uniform distribution in the validity region.

There is a number of possible ways to compute the derivative information, either directly, or through the use of the adjoint variable[12]. The factors taken into account when choosing the approach are the computational efficiency and flexibility, i.e. the ease at which the procedure can be adapted for a class of examples. Given that the subject model is likely to be a complex numerical code, designed without the goal to make differentiation convenient, the development cost may be high. In the test cases, where every part of the model was well-documented, we used a simple chain-rule approach based on augmentation of the code with partial derivatives of every mathematical procedure. In effect, together with evaluation of the mathematical model (1), we also resolved the equations

$$\begin{aligned}&\left(\frac{\partial F}{\partial T}+\frac{\partial F}{\partial R}\cdot\frac{\partial R}{\partial T}\right)\cdot\frac{dT}{d\alpha}+\frac{\partial F}{\partial R}\cdot\frac{\partial R}{\partial \alpha}=0\\ &\frac{d\Im(\alpha)}{d\alpha}=\nabla_\alpha J(T)=\frac{dJ(T)}{dT}\cdot\frac{dT}{d\alpha}\end{aligned} \tag{11}$$

For the cases where the underlying model equations (1) are very complex, or not available explicitly, we are developing an approach based on the existing automatic (or algorithmic) differentiation (AD) tools [16][17]. AD tools augment the model code by adding derivatives to each elementary function that involves the variables of interest; as the processed code runs, the derivative is assembled by chain rule. In an ideal situation, human involvement is limited to identification of variables of interest, and debugging the code at the complilation stage of the augmented model. In practice, applying the existing AD tools to nuclear reactor models still takes considerable development effort, and has not been fully validated[18]. More details will be made available in additional publications.

2.4 Global measure of sensitivity

A computationally cheap, local estimate of importance (10) is insufficient to address a question of whether an advanced uncertainty quantificaiton method can correctly assess which variables are the most important globally, i.e. over the entire validity region. Instead, we introduce a more relevant index of sensitivity[19]:

$$I_i = \frac{\mathrm{var}[E[\Im_i]]}{\mathrm{var}[E[\Im(\alpha_1,\alpha_2,...,\alpha_n)]]} \tag{12}$$

based on a marginal expected value defined as

$$E[\Im_i] = \frac{1}{|\Theta|}\int_{\Theta} \Im(\alpha_1, \alpha_2, ..., \alpha_n) d\alpha_1 d\alpha_2 ... d\alpha_{i-1} d\alpha_{i+1} ... d\alpha_n \quad (13)$$

where Θ is the validity region, as described in Section 2.1. The expression is normalized, $\sum_i I_i = 1$.

Due to high dimension of uncertainty space, numerical evaluation of (13) requires many runs of the model: very computationally expensive for the full model, not so for the surrogate.

3. Subject model

3.1 Model setup

We apply our approach to a test model, developed with the goal to exhibit the typical behavior of more complex systems, but at the same time to be as basic as possible, so as to avoid model-specific complexities of nuclear reactor analysis.

We use a simplified steady-state, 3-dimensional finite-volume model of a reactor core, with uniform fuel elements and no control mechanisms. In Figure 1, we visualize a single horizonal layer of the model, a representative point for each volume element is shown as a dot in its center. A basic unit of the core is a cylindrical fuel pin, surrounded by flowing coolant. We used a setup with 7 fuel pins.

There are two sources of heat: nuclear fisssion inside the fuel pin, and thermal energy carried by the coolant. The model takes into account heat transport (including uniform convection, and diffusion), and fission, described by multi-group neutronic diffusion equations. Once the distribution of temperature over the core is obtained by solving a coupled system of heat balance and neutronic diffusion equations, the temperatures in the central fuel element are resolved to a finer mesh. The operational parameters of the model were chosen to correspond to sodium-cooled fast reactor with realistic temperature excursions.

The intermediate parameters R are divided into two groups: thermodynamical and neutronic. Thermodynamical parameters are: heat conductivity K in fuel and coolant, specific coolant heat c_p and the heat transfer coefficient h. Neutronic parameters are (per energy group g, or per group pair g, g'): fission, scattering and absorbtion-removal cross-sections ($\Sigma_{fg}, \Sigma_{g\to g'}, \Sigma_{tg}$ correspondingly), energy spectrum parameter χ_g, yield of electrons per fission ν_g and neutron diffusion coefficient D_g. We used 4 energy groups.

The underlying mathematical model consists of a coupled system:

$$\begin{aligned}
&\int_{\partial\Omega} \left(K\nabla T + \rho c_p T\vec{u}\right)\cdot \vec{n} dS = \int_{\Omega} q''' dV \\
&-\nabla\cdot D_g \nabla\phi_g + \Sigma_{tg}\phi_g - \sum_{g\neq g'} \Sigma_{g\to g'}\phi_{g'} = -\frac{1}{\kappa}\sum_g \chi_g \sum_{g'} \nu_{g'}\Sigma_{fg}\phi_{g'}: \quad g, g' = 1,2,3,4 \\
&\sum_{\partial\Omega}\sum_g \phi_g = \int_{\Omega} q''' dV
\end{aligned} \quad (14)$$

where ϕ_g is the neutronic flux corresponding to energy group g, q''' is the total nuclear heat source; Ω is a finite volume element with boundary $\partial\Omega$, $\vec{n}$ is an outward normal vector depending on the shape of the cell; ∇T is approximated by a linear expression for heat exchange rate between two volume elements (depends on h if the exchange is between fuel and coolant); density ρ is assumed to be constant.

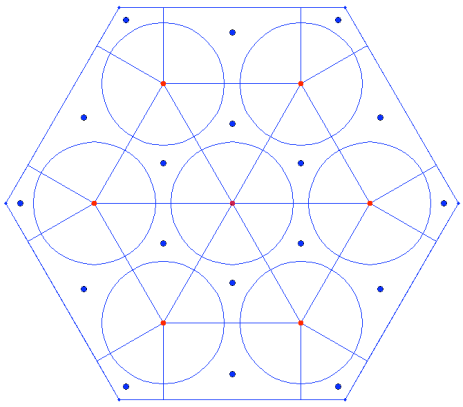

FIG. 1. Finite-volume model of the reactor core

The information on the model's uncertainties was taken from the available material properties reports[7][8], and then interpreted as shown in (4). Due to relatively smaller observed experimental error, neutronic parameters were characterzied by only 1 uncertainty quantifier each (constant error), thermodynamical parameters received a second-order characterization by 3 uncertainty quantifiers. The total is 38 uncertainty quantifiers, some tests were performed with just the uncertainty in thermodynamical properties, with 12 uncertainty quantifiers.

The coupled system (14) is solved using a fixed-point iteration procedure starting from a uniform distribution of $T.R$ and a prescribed spatial distribution $q'''(z) = C_s \sin(\pi z / H)$ along the pin length $0 \le z \le H$, $C_s = const$:

$$T := T(R, \phi) \quad , \quad R := R(T) \quad , \quad \phi := \phi(T, R) \tag{15}$$

After convergence of (15), we obtain the distribution of heat inside the central pin by solving an additional heat equation in cylindrical coordinates (r, z) (with boundary conditions taken from the computed heat distribution over the core):

$$\begin{aligned} -\nabla \cdot K \nabla T_{pin} &= q''' \\ T_{centerline}(z) &= T_{pin}(r = 0, z) \end{aligned} \tag{16}$$

We use the maximal fuel centerline temperature as a merit function; non-differentiability is resolved by using a differentiable approximation by a vector norm:

$$\max(T_{centerline}) \approx J(T) = \left\| \vec{T}_{centerline} \right\|_{p=100} \tag{17}$$

We compute the complete gradient of the output by applying the augmented equations (11) to the last step of the iteration (15), and to the solution of (16). Note that for the chosen differentiation approach, this represents the minimal intrusion into the model code: all computed partial derivatives of elementary functions are required to evaluate the answer.

3.2 Performance of the model

The developed model demonstrates the temperature excursion of approximately 200 K in the coolant (770 K at the inflow); the centerline maximal temperature is in the 2200 – 2500 K range. According to trial runs of the model with different unceratinty states, the most important source of uncertainty for the chosen output is the fuel heat conductivity. On the other side of the scale, the effect of uncertainties in neutronic parameters is almost negligible (due to very small maximal uncertainty values, and absense of fuel depletion effects that would result in non-uniform heat generation).

One full run of the model takes 1 to 2 minutes of computational time (higher estimate corresponds to an average desktop computer). The calculation of the full gradient takes between 150% and 200% of the time of the function evaluation (compare with theoretical bound of 500%[12]).

4. Surrogate model

4.1 Surrogate model construction

The surrogate model was created and validated using the following sequence of actions:

- Define a training set: choose a list of uncertainty quantifiers, $S = (\alpha_1, \alpha_2, ..., \alpha_m)$ as explained in comments to (4); generate a set $\{S_i\}$ of admissible values according to conditions on maximal uncertainty (3).
- Initialize and run the physics model: for each value S_i, evaluate $J(T(S_i))$ by solving (14), (15), (16). Store the required intermediate components, and compute the derivatives $\frac{\partial \Im}{\partial \alpha_i}$.
- Evaluate the importance of indvidiual variables according to (10), at a neutral state of the system with no uncertainty: $S = (0,0,...,0)$.
- Basis construction: assemble a polynomial basis (6) of degree up to 3 (2 in the tests that use all 38 parameters). The next step, basis truncation, is optional.
- *Basis truncation*: using the magnitude of the parametric sensitivity, rank components of uncertainty state S by importance. We used three groups: high importance (top 3 values), intermediate importance (the following 3 values), low importance (the rest of the variables). Our truncated basis has maximal degree 3: each polynomial of third degree must include variables of high importance, each polynomial of second degree must include variables of high or intermediate importance, the rest of the polynomials are linear. The specifics of the approach are arbitrary: the groups of importance, and the degrees used can be adjusted to fit an available computational budget. In the setup used in our experiments, a PRD model on 12 parameters required 12 full runs of the model; a PRD model on 38 parameters required 8 full runs.
- Create polynomial fitting equations (9). The minimal number of matrix rows is equal to the size of the basis: we use a taller matrix (slight over-sampling).
- *Additional analysis:* use derivative information to also create a linear approximation:

$$\hat{\Im}_{linear} = J_0 + \sum_i \frac{\partial \Im}{\partial \alpha_i} \cdot \alpha_i \qquad (18)$$

- *Validation:* create a validating set $\{\hat{S}_i\}$ by the same rules as the training set, and evaluate the full model at every point, to compare with the results of PRD (trained on $\{S_i\}$), and of the linear approximation.
- *Global sensitivity analysis:* evaluate the sensitivity (12) of the full model, and of the surrogate model with respect to individual variables. For the full model, the estimation was performed using 8550 runs of the full model (at the computational cost of several days). To achieve the same precision, we used 15200 evaluations of the surrogate model (less than 1 minute of computational time).

4.2 Performance of the PRD model

In our numerical experiments, we seek to compare the performance of the introduced PRD method in reproducing the output (17) with two basic approaches to uncertainty quantificaiton: random sampling, and linear approximation. To simulate a sampling approach, we observe the outputs of the model on a validating set of 100 points. The linear approximation is constructed as shown in (18). We present the results for PRD constructed using the full and the truncated basis. We record the observed range and variability of the output, the average magnitude and variability of the error. We also record the number

of full model evaluations required for each approach. For random sampling, the corresponding computational cost of each run is slightly lower, because derivative information does not need to be computed.

The results of one such test, for a model with 12 uncertainty quantifiers (i.e. only thermodynamical uncertainties enabled) are presented in Table 1. The temperatures are measured in degrees Kelvin. The tests for a model with 38 uncertainty quantifiers produced virtually the same performance, due to low magnitude of uncertainty in neutronic parameters, and overall low effect of this uncertainty. A visualization of the performance of surrogate models is also given in Figure 2: we show the magnitudes of error for the linear approximation, and for the PRD approaches, for 40 randomly selected points from the validation set.

In another experiment, we measured how well the PRD model preserves the global sensitivities (12) of the full model. We used a full model on 38 uncertainty quantifiers, and a PRD model with full basis of order 2. See Table 2 and Figure 3 for the numerical data, and a visualization (logarithmic plot).

In the measurement of global sensitivites, the uncertainty parameters corresponding to fuel heat conductivity are responsible for 99.99% of the variation in the PRD model, and for 98% in the full model. While we would also like to reproduce the small effects well, the large effects have greater practical significance.

Table 1. Performance of the 12-parameter PRD model.

	Random sampling	Linear approximation	PRD, full basis	PRD, truncated basis
Sample size	100	1	72	12
Output range	2237.82 – 2460.54	2227.43 – 2450.09	2237.82 – 2460.55	2237.51 – 2459.63
Model standard deviation	59.05	59.12	59.05	58.96
Error range		-0.01 – +10.38	-0.02 – + 0.02	- 0.43 – + 0.90
Error standard deviation		2.99	0.01	0.29

Table 2. Global sensitivities of the model

Material property (first 3 uncertainty quantifiers, 1 in case of macroscopic cross-sections)	Sensitivity of full model	Sensitivity predicted by PRD
Coolant heat capacity	$0.2187 \cdot 10^{-3}$ $0.1239 \cdot 10^{-5}$ $0.1203 \cdot 10^{-3}$	$0.0001 \cdot 10^{-3}$ $0.0017 \cdot 10^{-5}$ $0.0001 \cdot 10^{-3}$
Fuel heat conductivity	0.3876 0.0005 0.6097	0.4090 0.0017 0.5893
Coolant heat capacity	$0.4963 \cdot 10^{-4}$ $0.8723 \cdot 10^{-6}$ $0.1416 \cdot 10^{-4}$	$0.1232 \cdot 10^{-4}$ $0.0207 \cdot 10^{-6}$ $0.1349 \cdot 10^{-4}$
Heat transfer coefficient	$0.3635 \cdot 10^{-4}$ $0.2890 \cdot 10^{-5}$ $0.2941 \cdot 10^{-6}$	$0.5180 \cdot 10^{-9}$ $0.2593 \cdot 10^{-12}$ $0.6660 \cdot 10^{-9}$
Macroscopic cross-sections	On the order of: 10^{-6} to 10^{-3}	On the order of: 10^{-10} to 10^{-7}

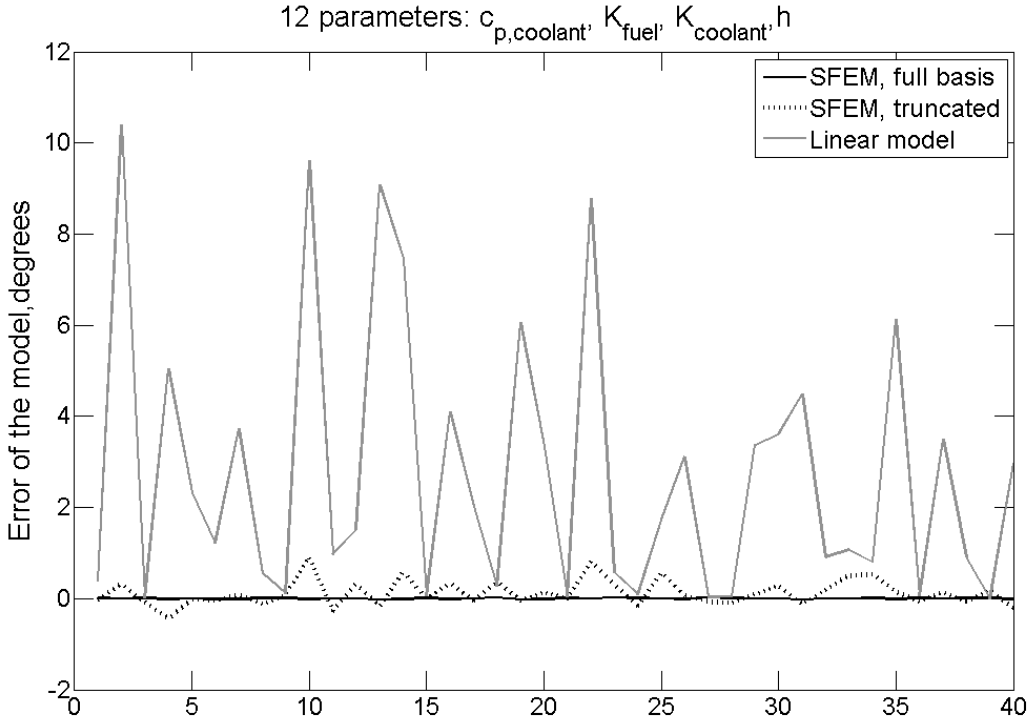

FIG. 2. Performance of the 12-parameter PRD model.

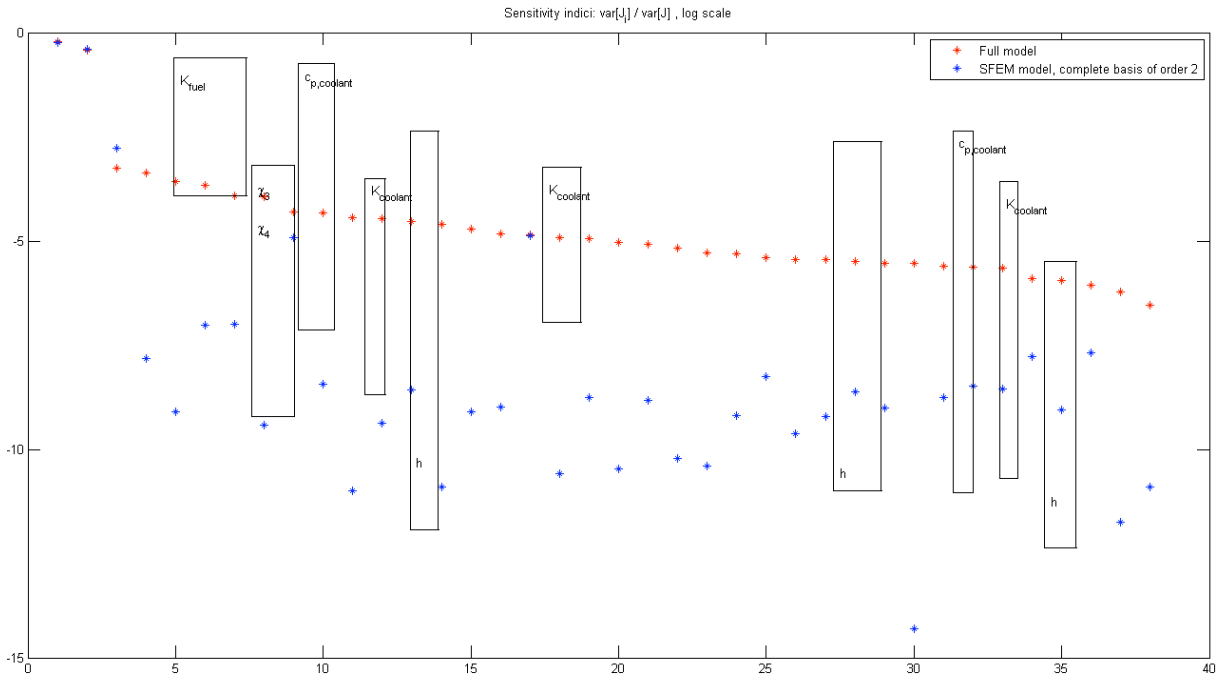

FIG. 3. Global sensitivities, log. scale

5. Discussion

Our numerical experiments show (at least, for a simplified model) that the truncated PRD approach is an attractive choice for uncertainty quantification, both in absolute terms (low approximation error), and in relative terms (error per amount of effort). The truncated basis still captures a substantial portion of nonlinearity, without needing the large computational expense associated with the full basis. As the available computational resources increase, the PRD with full basis will eventually outperform both the PRD with truncated basis, and polynomial regression with no derivative information. We point out, however, that for complex multiphysics models, it is unlikely that we will be allowed a larger computational budget, i.e. the most likely allowed training set is on the order of 100 points.

To characterize the performance of each method, we note the following. The range size is approximately 223 K: when predicting the range, the first-order sensitivity method produces an error of 4.5%, the PRD with truncated basis produces an error of 0.5%, and the PRD with full basis is almost error-free. If the goal is to predict correctly the range to within 10% of the error, the linear model is adequate. If the required threshhold if 1%, only PRD models perform well, at at no additional development cost (the derivative information needs to be computed in any case), though it may need about 10 times more computational time. We conclude that our hybrid (sampling-with-derivatives) approach is the most convenient way to extend the capabilities of multiphysics codes.

Our conclusions are largely unaffected by the assumptions on the probability distribution of the uncertainty. As the number of parameters increases, and their relative importance becomes more balanced (for example, as in the model discussed in [18]) , we expect the performance specifications described above to change significantly, and in the favor of PRD approaches of the type presented here.

6. Conclusions and research plans

In our work, we found that effects of a moderate number of uncertain parameters on a complex nuclear reactor simulation model can be efficiently approximated using a hybrid method, combining polynomial regression with the use of derivative information and polynomial basis truncation. We have successfully applied this method for a coupled model of neutronics and thermohydraulics of a nuclear reactor, where the uncertainty originated in physical parameters (material properties) of the system. From the available information we have created a probabilistic model, consistent with the available experimental information. We observed the substantial advantage of PRD methods computed using derivative information over classifcal methods of uncertainty quantification, such as pure random sampling, linear approximation, and polynomial regression that does not use derivative information. In addition, we observed that our basis truncation heutristic efficiently produces a far better approximation of the output function than the linear model, while using far less computational resources than the PRD model with full basis.

Overall, we have shown that PRD approaches have good potential for industrial applications. Further research is needed to determine whether the trends observed in this work hold for other applications, and for larger instances of unceratinty assessment in nuclear reactors. In our current efforts, we are extending the method to models of the nuclear reactor core that take into account additional uncertainty coming from the description of the non-uniform flow of the coolant, structural deformations of the reactor elements, and fuel depletion.

The introduced approach to uncertainty quantification leads to multiple mathematical and engineering questions. The issues of interest include creation of optimal uncertainty models from available data about the physical properties of the reactor; optimal selection of collocations nodes (given limited computational budget for analysis of the model's properties at each such point); efficient estimation of importance of individual variables leading to optimal basis truncation; choice of the best basis functions; correct processing of incomplete or unreliable derivative and state information. We also hope to provide additional theoretical justification for the suggested uncertainty quantification techniques, and derive error estimates for the performance of the surrogate model.

ACKNOWLEDGEMENTS

We are grateful to A. Siegel, T. Fanning, and our other SHARP team colleagues for valuable comments. We are grateful to J. Utke, and the rest of the OpenAD development team for their help in applying automatic differentiation tools to legacy codes.

REFERENCES

[1] FREPOLI C., An overview of Westinghouse realistic large break LOCA evaluation model, Science and Technology of Nuclear Installations, Hidawi Publishing Corporation (2008).
[2] DOWNING D.J., GARDNER R.H. and HOFFMAN F.O., An examination of response-surface methodologies for uncertainty analysis in assessment models, Technometrics, 27, 151 (1985).
[3] BOYACK B.E., DUFFEY R.B., GRIFFITH P., KATSMA K.R., LELLOUCHE G.S., LEVY S., ROHATGI U.S., WILSON G.E., WULFF W., and ZUBER N., "Quantifying reactor safety margins: part 1: an overview of the code scaling applicability, and uncertainty evaluation methodology", Los Alamos, (1988).
[4] ELDRED M.S., GIUNTA A.A., VAN BLOEMEN WAANDERS B.G., WOJTKIEWICZ S.F., HART W.E. and ALLEVA M.P., Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis, SAND2001-3796, Sandia (2002).
[5] LIU J.S., Monte Carlo strategies in scientific computing, Springer, Berlin (2001).
[6] CACUCI D.G., Sensitivity and uncertainty analysis, Chapman & Hall/CRC, Boca Raton, Florida (2003).
[7] FINK J.K., Thermophysical properties of Uranium dioxide., J. Nuclear Mater., 279, 1, 1 (2000).
[8] FINK J.K. and LEIBOWITZ L., Thermodynamic and transport properties of sodium liquid and vapor., ANL/RE-95/2, Argonne (1995).
[9] TANG Y.S., COFFIELD R.D. Jr. and MARKLEY R.A. Jr., Thermal analysis of liquid-metal fast breeder reactors, ANS, LaGrange Park, Illinois (1978).
[10] GHANEM R. and SPANOS P.D., The stochastic finite element method: a spectral approach, Springer, New York (1991).
[11] GOLUB G.H. and VAN LOAN C.F., Matrix computations, John Hopkins University Press (1996).
[12] GRIEWANK A., Evaluating derivatives: principles and techniques of algorithmic differentiation, 19, Frontiers in applied mathematics, SIAM, Pensilvania (2000).
[13] ISUKAPALLI S.S., ROY A. and GEORGOPOULOS P.G., Efficient sensitivity/uncertainty analysis using the combined stochastic response surface method and automated differentiation: application to environmental and biological systems, Risk Analysis, 20, 5, 591 (2000).
[14] ANITESCU M., Constrained optimization formulation of stochastic finite element methods for minimum eigenvalue problems, Argonne (2006).
[15] XIU D. and HESTHAVEN J.S., High-order collocation methods for differential equations with random inputs, SIAM J. Sci. Comput., 27, 1118 (2005).
[16] BISCHOF C.H., HOVLAND P.D. and NORRIS B., On the implementation of automatic differentiation tools, High-order symbolic computation, 21, 3, 311 (2008).
[17] BISCHOF C.H., BUCKER H.M., HOVLAND P.D. and NAUMANN J., Advanced in automatic differentiatio, Springer Verlag (2008).
[18] ALEXE M., RODERICK O., UTKE J., ANITESCU M., FANNING T. and HOVLAND P.D. Automatic differentiation of the nuclear reactor code MATWS, Argonne (2009).
[19] RODERICK O. and ANITESCU M., A note on sensitivity analysis of the multiphysics model, Tech. Rep, Argonne (2008).