A Multi-fidelity Polynomial Chaos-Greedy Kaczmarz Approach for Resource-Efficient Uncertainty Quantification on Limited Budget

Negin Alemazkoor[a,b], Arghavan Louhghalam[b], Mazdak Tootkaboni[b]

[a]*Department of Engineering Systems and Environment, University of Virginia, VA, USA*
[b]*Department of Civil and Environmental Engineering, University of Massachusetts Dartmouth, MA, USA*

Abstract

Polynomial chaos expansion (PCE) has been widely used to facilitate uncertainty quantification and stochastic computations for complex systems. Multi-fidelity approaches expedite the construction of the PCE surrogate by blending the efficiency of low-fidelity models and accuracy of high-fidelity models. In this work, we propose a novel non-intrusive multi-fidelity sampling approach that exploits the low computational cost of the low-fidelity model to select a set of high-yield sampling locations for the high-fidelity model. Particularly, the proposed approach draws upon the unique features of the Kaczmarz updating scheme to design a greedy search that explores a large pool of low-fidelity samples and iteratively removes the least contributive ones. Facilitated via a subset updating strategy, the search lands on a small subset of the initial pool which is then used to construct the PCE surrogate for the high-fidelity model. The proposed approach offers a remarkable computational performance, practically delivering accurate results with a high-fidelity sample size about the cardinality of the basis, and as such is amenable to efficient uncertainty quantification on fixed and limited budget. We provide several numerical examples that demonstrate the promise of the proposed approach in substantially reducing the number of high-fidelity samples required for accurate construction of the PCE surrogate.

1. Introduction

High-fidelity simulations of complex physical systems require extensive computing time and data storage even with advance parallelization. Additionally, only one run of the computational model is barely enough for reliable analysis of the system as the available information for accurate representation of many system parameters is incomplete and lacks confidence. Therefore, uncertainty quantification (UQ) is required to propagate the uncertainty in input variables (e.g., uncertain system parameters, and boundary or initial conditions) through the model and quantify the uncertainty in the outputs. Monte Carlo (MC) simulation, the most commonly used approach to UQ, suffers from a poor convergence rate – root mean squared error in estimates of expectations decays at the rate of $\mathcal{O}(M^{-0.5})$, with M the number of samples. As a prominent alternative to MC sampling and its variants (e.g., importance sampling [1] and quasi-Monte Carlo [2]), stochastic spectral approaches have been shown to speed-up convergence over Monte Carlo approaches.

Stochastic spectral approaches facilitate stochastic computations and UQ by constructing approximate surrogates that replace full-scale simulations [3]. One of the most widely adopted surrogates is the polynomial chaos expansion (PCE), where the stochastic output or quantity of interest (QoI) is represented through a series expansion in a generalized polynomial chaos basis [4, 5]. Different non-intrusive approaches such as spectral projection [6, 7], sparse grid interpolation [8, 9], and least squares [10, 11] have been used to construct the PCE approximation. The outstanding challenge is that these approaches suffer from the curse of dimensionality, i.e., the number of required sample points for accurate PCE surrogate construction increases dramatically with the dimensionality of random input.

Several approaches including active subspace methods [12], basis adaptation [13] optimal sampling [14, 15], adaptive sparse grid [16], and functional ANOVA decomposition [17] have been developed to alleviate the challenge associated with the curse of dimensionality and reduce the number of required samples for surrogate construction. Motivated by the fact that the PCE approximation of many high dimensional problems is often highly sparse, compressive sampling has also been extensively used to construct the PCE surrogate using a significantly smaller number of samples than the number of PCE coefficients to be estimated [18, 19, 20]. The accuracy of sparse polynomial approximation has been shown to strongly depend on incoherence properties of the measurement matrix, i.e., the Vandermonde-like matrix that is formed by evaluations of orthogonal

polynomials at sample locations [21]. Several efforts have, therefore, aimed at improving the incoherence properties of the measurement matrix and achieve better approximation accuracy through optimal sampling [22, 23], adaptive basis selection [24, 25], and preconditioning [26].

Recently, multi-fidelity approaches that exploit the low computational cost of low-fidelity simulations have attracted increasing attention in the field of uncertainty quantification. Multi-fidelity Gaussian processes [27, 28], multi-fidelity radial basis functions [29, 30], and multi-fidelity reduced basis functions [31, 32], for example, have shown great promise in reducing the number of high-fidelity samples required for uncertainty quantification. Multi-fidelity approaches have also been used to construct accurate PCE approximations of high-fidelity simulations. In [33], authors developed a multi-fidelity approach that incorporates different levels of sparse grid, where the low-level sparse grid is a subset of the higher-level sparse grid and serves as a corrector for the high-level grid. More specifically, the low- and high-level sparse grids were used to construct the PCE approximation for the low-fidelity simulation and the discrepancy between the high-fidelity model and the low-fidelity models, respectively. The two expansion were then combined into a single expansion that approximates the high-fidelity model. One of the drawbacks of proposed approach in [33] is that the number of high-fidelity samples depends on the level of the sparse grid and cannot be arbitrary. Moreover, the locations of the sample points are structured and an arbitrary set of random archived simulation results cannot be leveraged using this approach. This restricts performing UQ specially under fixed and limited computational budget.

Regression-based multi-fidelity approaches, where the sample size can be arbitrary, can be used to alleviate the challenge of limited computational budget. In [34], authors proposed using regression to construct PCE approximation for low-fidelity simulation and the discrepancy between the high-fidelity model and low-fidelity models. In [35], ℓ_1 minimization was performed on low-fidelity samples to identify the significant PCE coefficients and the number of high-fidelity samples to construct PCE approximation for the discrepancy between the high-fidelity and low-fidelity models was set at twice the cardinality of the significant coefficients. This is motivated by the fact that accurate coefficient estimation using least squares requires oversampling rate of $1.5 \sim 3$ [15].

In this work, we propose a novel sampling strategy that allows for efficient uncertainty quantification on limited budget. Specifically, we propose a multi-fidelity sampling approach that exploits the abundance of low-fidelity samples and Kaczmarz updating [36] to identify the locations of "most informative" sampling locations for high-fidelity simulation and accurate PCE approximation. Unlike [34, 35], where random sampling is used for PCE approximation of high-fidelity simulations, our proposed multi-fidelity sampling approach allows accurate construction of PCE approximation of the high-fidelity model with $\mathcal{O}(1)$ oversampling. This substantially reduces the number of required high-fidelity simulations and provides an ideal means for uncertainty quantification on limited budget. The organization of this paper is as follows. Section 2 presents general concepts in PCE approximation. In Section 3, we introduce the proposed approach. This includes the strategy for the selection of basis functions and the adaptation of Greedy Kaczmarz updating scheme to identify the sampling locations where the high-fidelity model needs to be evaluated. Section 4 includes numerical examples and discussions on the improved accuracy offered by the proposed approach and Section 5 provides the concluding remarks.

2. Problem setup

2.1. Polynomial chaos expansion

Consider the d-dimensional vector of independent random variables $\boldsymbol{\Xi} = (\Xi_1, ..., \Xi_d)$ to be the system input. Let I_{Ξ_i} and $\rho_i : I_{\Xi_i} \to \mathbb{R}^+$ be the support and probability measure of Ξ_i, respectively. Consequently, let $I_{\boldsymbol{\Xi}} = \times_{i=1}^{d} I_{\Xi_i}$ be the tensor-product domain that is the support of $\boldsymbol{\Xi}$ and let $\rho(\boldsymbol{\Xi}) = \prod_{i=1}^{d} \rho_i(\Xi_i)$ be the probability measure for variable $\boldsymbol{\Xi}$. Also, consider $u(\boldsymbol{\Xi})$ to be a square-integrable random Quantity of Interest (QoI). Then, $u(\boldsymbol{\Xi})$ can be written as an expansion in orthogonal polynomials, i.e.,

$$u(\boldsymbol{\Xi}) = \sum_{\boldsymbol{\alpha} \in \mathbb{N}_0^d} c_{\boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}(\boldsymbol{\Xi}). \tag{1}$$

with $\boldsymbol{\alpha} = (\alpha_1, ..., \alpha_d)$ a multi-index, $\{\psi_{\boldsymbol{\alpha}}\}_{\boldsymbol{\alpha} \in \mathbb{N}_0^d}$ the set of orthonormal basis functions satisfying

$$\int_{I_{\boldsymbol{\Xi}}} \psi_{\boldsymbol{m}}(\boldsymbol{\xi}) \psi_{\boldsymbol{n}}(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi} = \delta_{\boldsymbol{mn}}, \quad \boldsymbol{m}, \boldsymbol{n} \in \mathbb{N}_0^d, \tag{2}$$

and δ_{mn} the delta function.

For computation's sake, the expansion in (1) is often truncated by limiting the total order of polynomial only including basis functions with $\|\boldsymbol{\alpha}\|_1 \leqslant k$ in the expansion. The resulting truncated PCE, $u_k(\boldsymbol{\Xi})$, induces a truncation error but offers a practical approximation of $u(\boldsymbol{\Xi})$,

$$u(\boldsymbol{\Xi}) \approx u_k(\boldsymbol{\Xi}) = \sum_{\|\boldsymbol{\alpha}\|_1 \leqslant k} c_{\boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}(\boldsymbol{\Xi}). \tag{3}$$

The cardinality of basis, $\{\psi_{\boldsymbol{\alpha}} : \boldsymbol{\alpha} \in \mathbb{N}_0^d, \|\boldsymbol{\alpha}\|_1 \leq k\}$, denoted by K, will then be

$$K = \frac{(k+d)!}{k!\ d!}. \tag{4}$$

and the exact PCE coefficients in Equation (1) can be calculated by projecting $u(\boldsymbol{\Xi})$ onto the basis functions $\psi_{\boldsymbol{\alpha}}$:

$$c_{\boldsymbol{\alpha}} = \mathbb{E}_\rho \left[u(\boldsymbol{\Xi}) \psi_{\boldsymbol{\alpha}}(\boldsymbol{\Xi}) \right] = \int_{I_{\boldsymbol{\Xi}}} u(\boldsymbol{\xi}) \psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi}) \rho(\boldsymbol{\xi})\, \mathrm{d}\boldsymbol{\xi}. \tag{5}$$

where $\mathbb{E}_\rho \left[\cdot \right]$ denotes the operation of mathematical expectation with respect to probability measure $\rho(\boldsymbol{\Xi})$.

2.2. Computation of coefficients

This multidimensional integral in (5) can be numerically calculated using non-intrusive approaches such as sparse grid quadrature [5]. However, these approaches become impractical for high-dimensional problems as they require a very large number of samples [5, 37]. For high-dimensional problems, least squares regression provides a practical approach to calculate the PCE coefficients. The approach estimates the coefficients by solving an over-determined system of equations. Specifically, let us define:

$$\boldsymbol{u} := (u(\boldsymbol{\xi}^{(1)}), ..., u(\boldsymbol{\xi}^{(M)}))^T, \tag{6}$$

$$\boldsymbol{\Psi}(i,j) := \psi_{\boldsymbol{\alpha}^j}(\boldsymbol{\xi}^{(i)}), \quad i = 1, \cdots, M, \quad j = 1, \cdots, K. \tag{7}$$

where vector $\boldsymbol{u}$ is formed by evaluations of $u(\boldsymbol{\Xi})$ at random realizations of $\boldsymbol{\Xi}$ denoted by $\{\boldsymbol{\xi}^{(i)}\}$ and matrix $\boldsymbol{\Psi}$, i.e., measurement matrix, consists of evaluations of basis functions at these realizations. Least squares then solves

$$\boldsymbol{c} = \underset{\boldsymbol{c}}{\operatorname{argmin}} \|\boldsymbol{u} - \boldsymbol{\Psi}\boldsymbol{c}\|_2^2. \tag{8}$$

It has been shown that stability and optimal accuracy of least-square approximation can be achieved under the mild condition that the number of points scales linearly with the basis cardinality up to an additional logarithmic factor [38]. As a rule of thumb, it is generally accepted that the estimation of coefficients requires $1.5 \sim 3$ oversampling rate [14, 15]. However, affording that many samples may still be computationally intensive. The proposed framework in this work is premised on the fact that an abundance of low-fidelity simulations provides valuable insight not only on the choice of the dominant basis functions but also on the choice of the most informative sampling locations. It is well-known that, in many high-dimensional problems, QoI is sparse with respect to the PCE basis [18]. We, therefore, use low-fidelity simulations to examine the sparsity of PC approximation and to identify the set of basis functions with non-zero coefficients in the PCE approximation. Reducing the cardinality of basis, in the presence of strong sparsity, reduces the number of required high-fidelity simulations for PCE construction substantially. The reduced set of coefficients, however, may be still be overwhelming, particularly under scenarios where the task of uncertainty quantification is challenged with the availability of fixed or limited "computational budget". To significantly reduce the number of required high-fidelity simulations, we employ Kaczmarz updating scheme [36] to search through the low-fidelity samples and to identify the sampling locations that lead to the highest yield, that is the most accurate PCE construction, from high-fidelity simulations.

3. Multi-fidelity UQ via Kaczmarz updating

We leverage the low computational cost of low-fidelity simulations to reduce the number of required high-fidelity simulation samples for construction of PCE surrogate using a three-step approach. We first

generate a large pool of low-fidelity samples and use them to construct a polynomial approximation for the low-fidelity model and identify the reduced basis. With the cardinality of basis reduced, we then search among low-fidelity samples to select the sample locations that result in the most accurate reduced basis PCE approximation for the low-fidelity model where we leverage Kaczmarz updating scheme to guide the search. Finally, we run high-fidelity simulations at the selected sample locations and use them to construct the reduced basis PCE for the high-fidelity model. In what follows we provide a detailed explanation of these steps.

3.1. Reduction of the basis set

To select the dominant basis functions, we start with a rather large pool of sample candidates. We declare the measurement matrix $\boldsymbol{\Psi}^{\text{pool}}$ of size $M_p \times K$, where M_p is the number of candidates in the pool and K is cardinality of full basis. We evaluate the QoI at all sample candidates in the pool using the low-fidelity model, from vector $\boldsymbol{u}^{\text{pool}}$ that consists of the low-fidelity QoI at all candidate locations and solve

$$\boldsymbol{c} = \underset{\boldsymbol{c}}{\operatorname{argmin}} \left\| \boldsymbol{u}^{\text{pool}} - \boldsymbol{\Psi}^{\text{pool}} \boldsymbol{c} \right\|_2^2 . \tag{9}$$

Considering the orthonormality of basis functions, the relative contribution of the jth basis function in approximation of QoI is then calculated as

$$\Theta_j := \frac{c_j^2}{\sum_{i=2}^{K} c_i^2}, \tag{10}$$

where Θ_j denotes the relative contribution of the jth basis function [39, 40]. With the relative contribution of each basis function evaluated, the criterion to keep or discard depends on the desired approximation accuracy as well as the computational budget, given that the number of high-fidelity simulations must be larger than the cardinality of the reduced basis set. In this work, we define the reduced basis set to be $\{\psi_{\boldsymbol{\alpha}}\}^R := \{\psi_{\boldsymbol{\alpha}^j} : \Theta_j > 10^{-5}\} \subseteq \{\psi_{\boldsymbol{\alpha}}\}$, that is we discard those basis functions with relative contributions below 10^{-5}.

It must be noted that the above basis reduction is based on the premise that the low-fidelity model is able to capture the functional variation of the QoI in the parameter space, despite its low accuracy in the physical space. This implies that a basis function that corresponds to a zero coefficient for the approximation of low-fidelity simulation should have a zero coefficient for the high-fidelity simulation too. This assumption, however, may not be always true. Furthermore, while a posterior error estimates measuring the ability of the low-fidelity simulations to approximate the high-fidelity data are useful to tune the truncation tolerance for basis reduction, we leave further investigation in that direction for future research and provide empirical evidence on the impact of truncation tolerance in Section 4.3. Finally, to demonstrate the applicability and performance of the proposed strategy regardless of the adopted basis set (full or reduced) results in Section 4 are reported with and without basis reduction.

3.2. Kaczmarz updating scheme

The proposed multi-fidelity approach in this work exploits the efficiency of Kaczmarz algorithm in selection of high-fidelity samples, that is samples where evaluation of high-fidelity model provides the highest yield in increasing the accuracy of the polynomial approximation. Kaczmarz algorithm was first introduced in [36] as an iterative approach to solve overdetermined linear system of equations $\boldsymbol{\Psi c} = \boldsymbol{u}$. Because of its simplicity, Kaczmarz algorithm has gained great attention in various fields such as computer tomography and image processing [41]. In [42], Kaczmarz algorithm was successfully used for function approximation when data collection is cheap and very large volume of data is available. Precisely, in [42], cases where $M \gg K \gg 1$ were considered. In such cases, solving (8) in one step may become computationally intractable as even forming the measurement matrix $\boldsymbol{\Psi}$ may overburden the memory capacity of the system. Therefore, it was suggested to iteratively update the coefficient vector according to Kaczmarz algorithm. That is

$$\boldsymbol{c}^{(i+1)} = \boldsymbol{c}^{(i)} + \frac{u_i - \left\langle \boldsymbol{\psi}_i, \boldsymbol{c}^{(i)} \right\rangle}{\|\boldsymbol{\psi}_i\|_2^2} \boldsymbol{\psi}_i^T , \tag{11}$$

where $\boldsymbol{\psi}_i$ is the i-th row of matrix $\boldsymbol{\Psi}$ and u_i is the i-th observed QoI for the i-th drawn sample. As it can be seen in Equation (11), Kaczmarz updating only involves vector-vector multiplication, thereby demanding

very low memory and computational capacity. In the following, we explain how we exploit this updating scheme to search among a pool of low-fidelity samples and select the set of most informative high-fidelity samples.

3.3. Greedy Kaczmarz identification of high-yield sample locations

Let $\boldsymbol{\Psi}^{\text{r-pool}}$ denote the $M_p \times R$ measurement matrix associated with the evaluation of reduced basis at the M_p sample candidates. The goal is to select $\boldsymbol{\Psi}^{\text{opt}(M)}$ to be the quasi-optimal $M \times R$ row-submatrix of the "pool" matrix $\boldsymbol{\Psi}^{\text{r-pool}}$ associated with the sample locations at which the high-fidelity QoI must be evaluated. To this end, we propose the Greedy Kaczmarz Algorithm (GKA).

The proposed algorithm starts with the entire pool of sample candidates and iteratively removes the most insignificant candidate from the pool. Different approaches can be used to identify the most insignificant (the least contributive) sample. A naive approach is to try to temporarily remove each candidate, solve the least squares problem using the remaining samples, and select the most insignificant sample to be the sample which, if removed from the sample set, results in the smallest change in approximation accuracy. This approach, however, can become computationally intractable as it requires solving the least squares problem repetitively. Our proposed greedy Kaczmarz algorithm, on the other hand, selects and removes the most insignificant candidate, at each iteration, with minimal computational cost. Specifically, for each candidate, we update the coefficient vector using Kaczmarz algorithm assuming that the candidate is "newly added" to the pool. We then consider the candidate which, if added to the pool, results in the least accurate approximation to be the most insignificant candidate and permanently remove it from the pool. The removal process continues until only M "high yield" samples–with M determined by the (often fixed or limited) available computational budget–are retained in the pool. Algorithm 1 shows the pseudocode for this selection process where $\boldsymbol{\Psi}^{\text{opt}[i]}$ of size $i \times R$ denotes the sought after near-optimal submatrix at each iteration, and $\boldsymbol{\Psi}^{\text{opt}[i]}_{(j)}$ represents the jth row of the submatrix.

Algorithm 1 Greedy Kaczmarz Algorithm

Initiate $\boldsymbol{\Psi}^{\text{opt}(M_p)} = \boldsymbol{\Psi}^{\text{r-pool}}$ and $\boldsymbol{u}^{\text{opt}(M_p)} = \boldsymbol{u}^{\text{pool}}$
for $i = M_p : -1 : M+1$ **do**
  $\boldsymbol{c} = \underset{\boldsymbol{c}}{\text{argmin}} \left\| \boldsymbol{u}^{\text{opt}(i)} - \boldsymbol{\Psi}^{\text{opt}(i)} \boldsymbol{c} \right\|_2^2$
  for $j = 1 : i$ **do**
    $\boldsymbol{c}^{(j)} = \boldsymbol{c} + \frac{u^{\text{opt}(i)}_{(j)} - \left\langle \boldsymbol{\Psi}^{\text{opt}(i)}_{(j)}, \boldsymbol{c} \right\rangle}{\left\| \boldsymbol{\Psi}^{\text{opt}(i)}_{(j)} \right\|_2^2} \boldsymbol{\Psi}^{\text{opt}(i)^T}_{(j)}$ ▷ Update $\boldsymbol{c}$ assuming the jth candidate is newly added to the pool
    $e^{(j)} = \left\| \boldsymbol{u}^{\text{pool}} - \boldsymbol{\Psi}^{\text{r-pool}} \boldsymbol{c}^{(j)} \right\|_2^2$
  end for
  $j^* = \underset{j}{\text{argmax}} \{e^{(j)}\}$
  $\boldsymbol{\Psi}^{\text{opt}(i-1)} = \boldsymbol{\Psi}^{\text{opt}(i)}$ and $\boldsymbol{u}^{\text{opt}(i-1)} = \boldsymbol{u}^{\text{opt}(i)}$
  $\boldsymbol{\Psi}^{\text{opt}(i-1)}_{(j^*)} = [\,]$ and $\boldsymbol{u}^{\text{opt}(i-1)}_{(j^*)} = [\,]$; ▷ Permanently remove row j^*
end for

Further speedup via subset updating. The proposed GKA eliminates the need for solving least squares repetitively and substantially reduces the computational complexity. To further speedup the identification process, in each iteration, instead of searching among all candidates for the most insignificant one, we perform the search on a rather small subset of candidates. Later, in Section 4, we show the gain in accuracy quickly plateaus beyond a rather small number of samples in this subset. We, therefore, report results only for the case when GKA is performed in conjunction with subset updating. We will further demonstrate, through empirical evidence, the remarkable performance of GKA with subset updating as it delivers highly accurate approximations with an oversampling rate of $\mathcal{O}(1)$. Algorithm 2 shows the pseudocode for GKA with the subset search where N_{sub} is the size of search subset.

Algorithm 2 Greedy Kaczmarz Algorithm with subset updating

Initiate $\boldsymbol{\Psi}^{\text{opt}(M_p)} = \boldsymbol{\Psi}^{\text{r-pool}}$ and $\boldsymbol{u}^{\text{opt}(M_p)} = \boldsymbol{u}^{\text{pool}}$
for $i = M_p : -1 : M+1$ **do**
  $\boldsymbol{c} = \underset{\boldsymbol{c}}{\text{argmin}} \left\| \boldsymbol{u}^{\text{opt}(i)} - \boldsymbol{\Psi}^{\text{opt}(i)} \boldsymbol{c} \right\|_2^2$
  Initiate S to be a set including N_{sub} randomly selected integers between 1 and i
  for $j = 1 : N_{sub}$ **do**
    $\boldsymbol{c}^{(S_j)} = \boldsymbol{c} + \frac{u_{(S_j)}^{\text{opt}(i)} - \left\langle \boldsymbol{\Psi}_{(S_j)}^{\text{opt}(i)}, \boldsymbol{c} \right\rangle}{\left\| \boldsymbol{\Psi}_{(S_j)}^{\text{opt}(i)} \right\|_2^2} \boldsymbol{\Psi}_{(S_j)}^{\text{opt}(i)^T}$ ▷ Update $\boldsymbol{c}$ assuming the S_jth candidate is newly added to the pool
    $e^{(S_j)} = \left\| \boldsymbol{u}^{\text{pool}} - \boldsymbol{\Psi}^{\text{r-pool}} \boldsymbol{c}^{(S_j)} \right\|_2^2$
  end for
  $j^* = \underset{j}{\text{argmax}} \{ e^{(S_j)} \}$
  $\boldsymbol{\Psi}^{\text{opt}(i-1)} = \boldsymbol{\Psi}^{\text{opt}(i)}$ and $\boldsymbol{u}^{\text{opt}(i-1)} = \boldsymbol{u}^{\text{opt}(i)}$
  $\boldsymbol{\Psi}_{(S_{j^*})}^{\text{opt}(i-1)} = [\,]$ and $\boldsymbol{u}_{(S_{j^*})}^{\text{opt}(i-1)} = [\,]$; ▷ Permanently remove row S_j^*
end for

4. Numerical results

To demonstrate the advantage of the proposed approach, a set of four numerical examples is considered in this section: (i) the borehole function, (ii) the solution to a stochastic diffusion problem, (iii) the integral of Ackley function, and (iv) effective fracture toughens of layered media. This would allow for a comprehensive evaluation of the proposed sampling strategy for a variety of models with different combinations of dimensionality and order. Additionally, to ensure the robustness of the results, we report averages over 50 independent trials for each problem. We use relative error as the measure to report approximation accuracy which is computed as the relative difference between the exact and approximated high-fidelity QoI for a set of 10,000 high-fidelity samples. Finally, to better understand the benefits (i.e., in terms of computational cost) gained in each of the steps, we report the results for GKA with and without basis selection.

4.1. Borehole function

In this example, we illustrate the accuracy improvement achieved by the proposed approach in approximation of Borehole function. Borehole function models water flow through a borehole and has been used as a demonstrative example in several computational studies [34, 43]. The high-fidelity and low-fidelity borehole functions are expressed as:

$$u^h(\boldsymbol{\Xi}) = \frac{2\pi T_u (H_u - H_l)}{\ln(\frac{r}{r_w})(1 + \frac{2LT_u}{\ln(\frac{r}{r_w}) r_w^2 K_w} + \frac{T_u}{T_l})}, \tag{12}$$

$$u^l(\boldsymbol{\Xi}) = \frac{5T_u (H_u - H_l)}{\ln(\frac{r}{r_w})(1.5 + \frac{2LT_u}{\ln(\frac{r}{r_w}) r_w^2 K_w} + \frac{T_u}{T_l})}, \tag{13}$$

with $\boldsymbol{\Xi} = (T_u, H_u, H_l, r, r_w, L, K_w, T_l)$ the random vector containing the independent set of input random variables; see Table 1.

Table 1: Input random variables for the Borehole function

Random Variable	Definition	Probability Distribution
T_u	Transmissivity of upper aquifer (m^2/yr)	U[63700, 115600]
H_u	Potentiometric head of upper aquifer (m)	U[990, 1100]
H_l	Potentiometric head of lower aquifer (m)	U[700, 820]
r	Radius of influence (m)	U[100, 50000]
r_w	Radius of borehole (m)	U[0.05, 0.15]
L	Length of borehole (m)	U[1120, 1680]
K_w	Hydraulic conductivity of borehole (m/yr)	U[9855,12045]
T_l	Transmissivity of lower aquifer (m^2/yr)	U[63.1, 116]

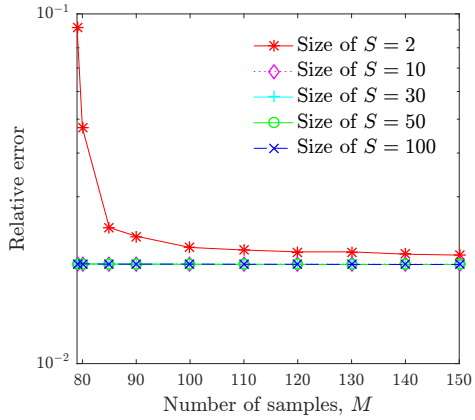

(a) Approximation accuracy for different sizes of S using GKA w/ basis selection, $Mp = 1500$

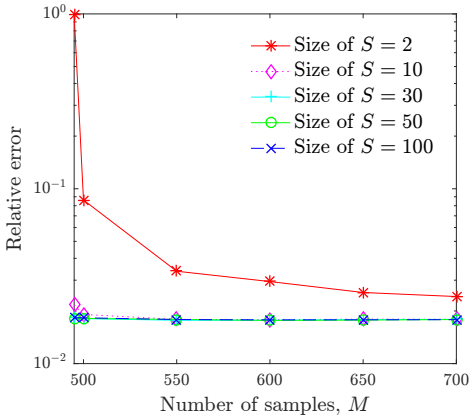

(b) Approximation accuracy for different sizes of S using GKA w/o basis selection, $Mp = 1500$

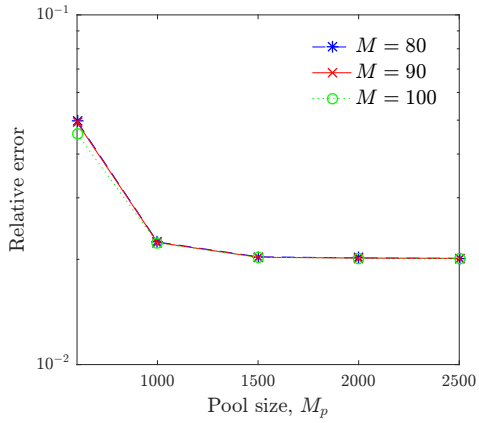

(c) Approximation accuracy for different sizes of M_p using GKA w/ basis selection, $S = 10$

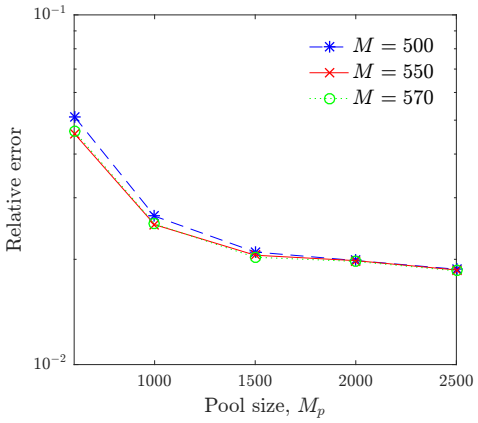

(d) Approximation accuracy for different sizes of M_p using GKA w/o basis selection, $S = 10$

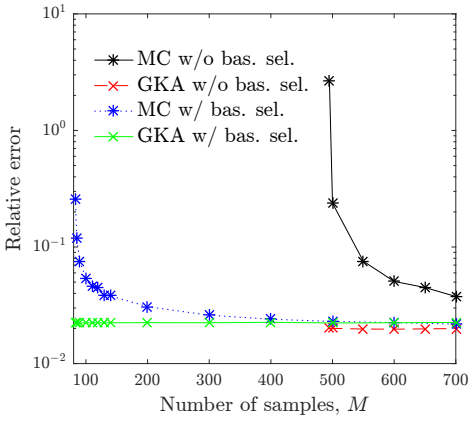

(e) Approximation accuracy using GKA and MC random sampling when $Mp = 1500$ and $S = 10$

Figure 1: Approximation accuracy for high-fidelity Borehole function under different settings of GKA and high-fidelity sample sizes.

We use Legendre polynomials up to fourth-order to approximate the eight-dimensional high-fidelity Borehole function. The full basis has a cardinally of 495, which is reduced to 79 through basis function selection. Figure 1 shows the approximation accuracy for high-fidelity Borehole function using GKA with and without basis selection under different scenarios. Specifically, Figures 1a and 1b demonstrate the impact of search subset size on approximation accuracy using GKA with and without basis selection, respectively. To generate these figures a pool of 1500 low-fidelity samples was used. It can be seen that, regardless of whether

basis function selection is performed or not, increasing the size of search subset beyond 10 does not improve the approximation accuracy any further. This is an observation of paramount importance as it implies the greedy Kaczmarz algorithm with subset updating (Algorithm 2) allows for performing the search for high-yield high-fidelity samples at substantially reduced computational cost.

We next fix the size of search subset at 10 samples and vary the size of low-fidelity sample pool, M_p, to evaluate its impact on approximation accuracy. Figures 1c and 1d show the results. Expectedly, it is observed that, for all sizes of high-fidelity samples and regardless of whether basis function selection is performed or not, approximation accuracy converges when the size of candidate pool is about three times of the cardinality of full basis set, K.

Finally, Figure 1e compares the improvement achieved by GKA with the case where naive Monte Carlo sampling is used for selecting high-fidelity samples. To generate Figure 1e, the sizes of search subset and candidate pool were set at 10 and 1500 samples, respectively. As expected, the saturation error without basis selection is smaller than the saturation error with basis selection. The results highlight (i) the significance of basis function selection in substantially reducing the number of required high-fidelity simulations for accurate construction of polynomial approximation (ii) the utility of the proposed GKA in further reducing the number of required high-fidelity simulations particularly for low values of M, that is for severely limited computational budget (iii) most importantly, that the number of high-fidelity simulations required to arrive at a low relative error is about the size of the coefficient vector (full or reduced). This signifies a remarkable performance for scenarios where limited budget (number of high-fidelity simulations) may hamper the process of uncertainty quantification.

4.2. Stochastic diffusion problem

In this example, we consider stochastic diffusion over a one dimensional physical domain, mathematically described by

$$\begin{aligned} -\frac{\partial}{\partial x}\left(a(x,\boldsymbol{\Xi})\frac{\partial u}{\partial x}(x,\boldsymbol{\Xi})\right) &= 2, \quad x \in (0,1), \\ u(0,\boldsymbol{\Xi}) = 0, \quad u(1,\boldsymbol{\Xi}) &= 0, \quad \boldsymbol{\Xi} \in [-1,1]^{20}. \end{aligned} \tag{14}$$

We assume that the diffusion coefficient, $a(x,\boldsymbol{\Xi})$, takes the following analytical form

$$a(x,\boldsymbol{\Xi}) = 1 + \sum_{k=1}^{20} \frac{1}{k\pi}\cos(2\pi k x)\Xi_k. \tag{15}$$

with $\boldsymbol{\Xi}$ uniformly distributed over the hypercube $[-1,1]^{20}$, and consider the solution of diffusion problem at $u(0.5,\boldsymbol{\Xi})$ to be the quantity of interest. We then obtain low-fidelity and high-fidelity evaluations of the QoI using 1×20 and a 1×1000 meshes, respectively, and aim to construct a second-order polynomial approximation of the high-fidelity QoI over the random input space. We use finite difference and the three-stage Lobatto IIIa formula (bvp4c in MATLAB) to solve (14).

Here, the accuracy of the second-order PCE surrogate is limited and a higher order PCE is needed for more accurate approximation of the high fidelity solution of the diffusion problem. We, however, note that the main goal of the proposed sampling strategy is to do the best under fixed and limited budget. In this particular case our sampling approach (or rather its resulting sample locations) provides *the most accurate* PCE given a fixed budget that allows up to second-order approximation.

Figure 2 depicts the accuracy of PCE approximation for high-fidelity solution of the stochastic diffusion problem using GKA with and without basis function selection under different scenarios. Particularly, Figures 2a and 2b demonstrate the impact of size of search subset on approximation accuracy, when the subset is randomly selected from a pool of 1000 low-fidelity samples. It is observed that the approximation accuracy does not improve any further by increasing the size of search subset beyond 30 samples.

Figures 2c and 2d, examine the impact of pool size, M_p, on approximation accuracy, when the size of search subset is fixed at 30 samples. It is again seen that (i) the approximation accuracy converges when the pool size is about three times the cardinality of full basis set, $K = 231$, (ii) remarkably, the number of high-fidelity simulations required (the main driver of the computational cost) is about the size of the coefficient vector whether or not basis reduction is performed.

Lastly, Figure 2e compares the improvement achieved by GKA over random Monte Carlo sampling, when the sizes of the search subset and the pool are set at 30 and 1000, respectively. Here, the process of basis function selection reduces the cardinality of the basis from 231 to 53, thereby substantially reducing the

number of required high-fidelity samples for PCE construction. The accuracy of PCE approximation is then significantly improved through the application of GKA for the identification of high-yield sample locations for high-fidelity simulation.

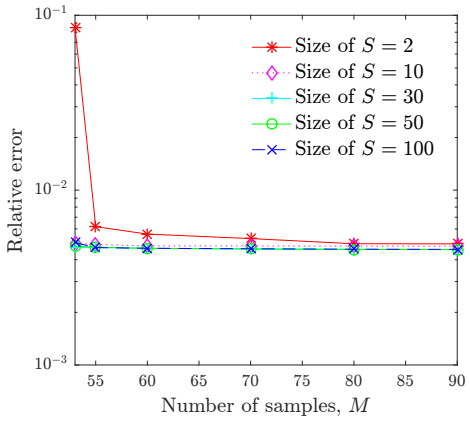

(a) Approximation accuracy for different sizes of S using GKA w/ basis selection, $Mp = 1000$

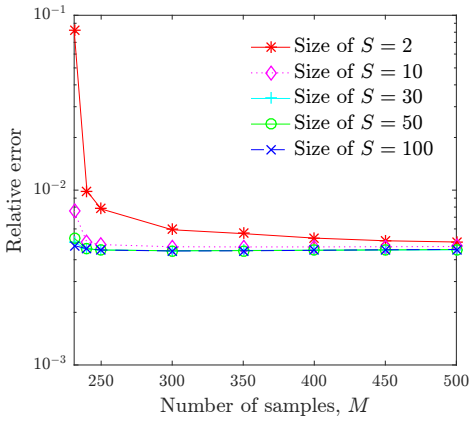

(b) Approximation accuracy for different sizes of S using GKA w/o basis selection, $Mp = 1000$

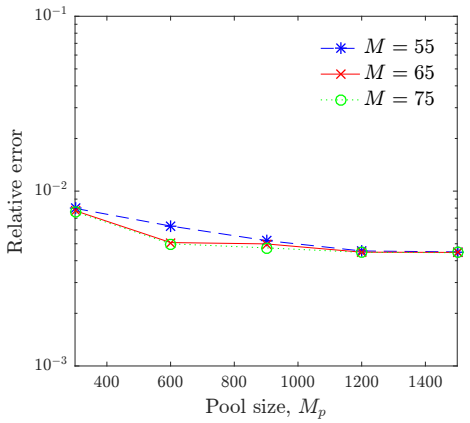

(c) Approximation accuracy for different sizes of M_p using GKA w/ basis selection, $S = 30$

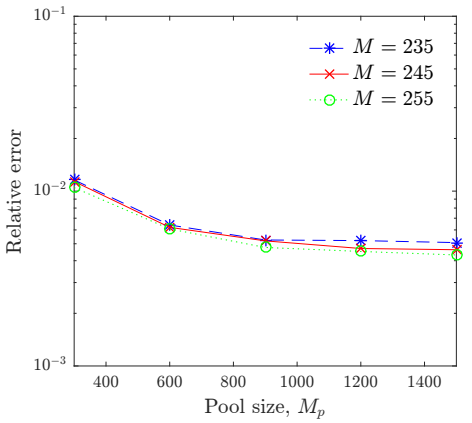

(d) Approximation accuracy for different sizes of M_p using GKA w/o basis selection, $S = 30$

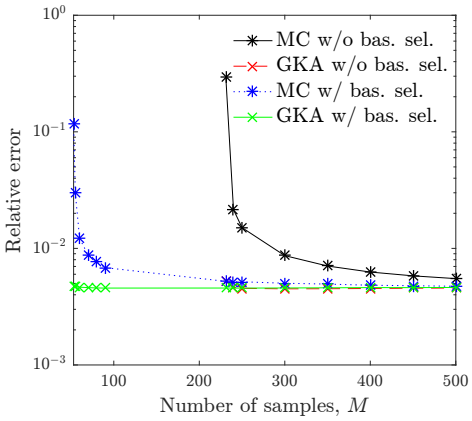

(e) Approximation accuracy using GKA and MC random sampling when $Mp = 1000$ and $S = 30$

Figure 2: Comparison of approximation accuracy for high-fidelity solution of stochastic diffusion problem under different settings of GKA and high-fidelity sample sizes.

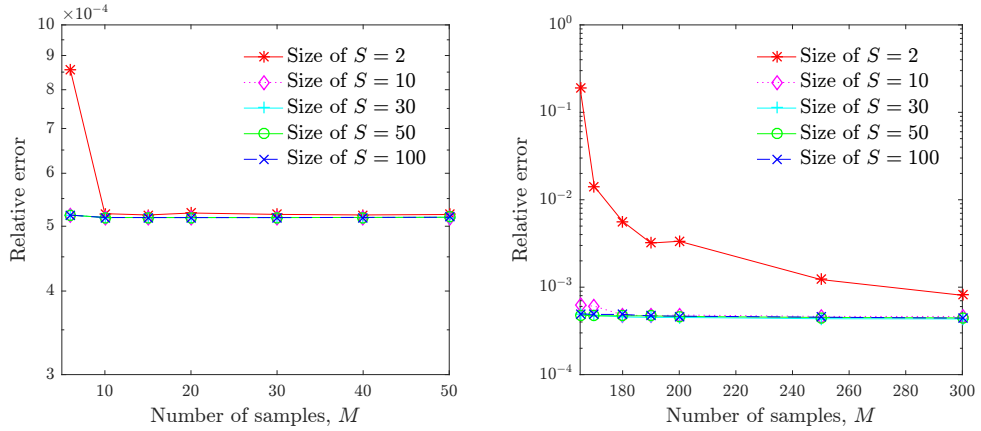

(a) Approximation accuracy for different sizes of S using GKA w/ basis selection, $Mp = 1000$

(b) Approximation accuracy for different sizes of S using GKA w/o basis selection, $Mp = 1000$

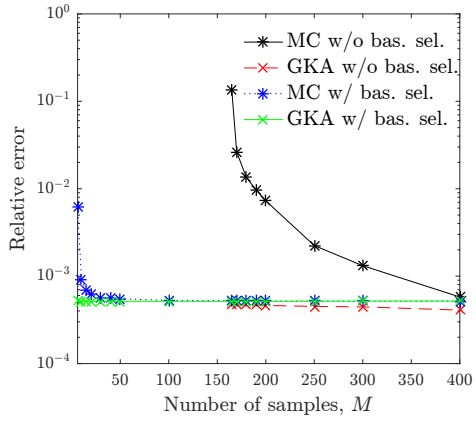

(c) Approximation accuracy using GKA and MC random sampling when $Mp = 1500$ and $S = 30$

Figure 3: Comparison of approximation accuracy for the integral of Ackley function under different settings of GKA and high-fidelity sample sizes.

4.3. Ackley function

Next, we consider the highly irregular Ackley function, which in 2D is expressed as:

$$f(x,y) = -20\exp(-0.2\sqrt{0.5(x^2+y^2)}) - \exp(0.5(\cos(2\pi x)+\cos(2\pi y))) + \exp(1) + 20. \tag{16}$$

We randomize the function in Equation (16) using three random variables:

$$\begin{aligned} f(x,y,\boldsymbol{\Xi}) = &-20(1+0.2\Xi_1)\exp(-0.2(1+0.3\Xi_2)\sqrt{0.5(x^2+y^2)}) \\ &- \exp(0.5(\cos(2\pi(1+0.2\Xi_3)x)+\cos(2\pi(1+0.3\Xi_3)y))) + \exp(1) + 20, \end{aligned} \tag{17}$$

with $\boldsymbol{\Xi}$ is uniformly distributed on $[-1,1]^3$ and define the quantity of interest to be the integral of (17) over $[-5,5]$:

$$u(\boldsymbol{\Xi}) := \int_{-5}^{5}\int_{-5}^{5} f(x,y,\boldsymbol{\Xi})\mathrm{d}x\mathrm{d}y \tag{18}$$

We evaluate the low-fidelity and high-fidelity QoI by performing trapezoidal numerical integration over a 126×126 and a 1001×1001 mesh, respectively, and aim to construct an eight-order polynomial approximation of the high-fidelity QoI over the random input space.

Figure 3 demonstrates the accuracy of PCE approximation for high-fidelity numerical integration of Ackley function using GKA with and without basis function selection under different settings. Figures 3a and 3b evaluate the impact of search subset size on approximation accuracy, when the subset is randomly selected from a pool of 1000 low-fidelity samples. As in the preceding sections, the results indicate that

accurate approximations can be achieved by randomly selecting fewer than 30 samples of the low-fidelity pool in each iteration of GKA.

An examination of the impact of pool size, M_p, on approximation accuracy and the relationship between the number of required high-fidelity simulations and the size of coefficient vector reveals a trend fully consistent with the observations in Sections 4.1 and 4.2 which we do not repeat here for the sake of brevity. A comparison of the proposed multi-fidelity approach with random Monte Carlo sampling is included in Figure 3c where we see selecting the locations of high-fidelity samples using GKA significantly improves the approximation accuracy. Finally, Figure 4 depicts three snapshots of GKA sample selection process from the starting point, i.e., the low-fidelity sample pool containing 1000 uniformly distributed samples to when only 165 high-yield samples are retained. Expectedly, it is observed that the condition number of the measurement matrix with the GKA identified samples is substantially lower compared to the measurement matrix with random samples. In this example, the 165 high-yield samples form a measurement matrix with condition number of 250, while measurement matrices with random samples have a condition number larger than 2×10^5 on average. The non-uniform distribution of the samples at the end as well as the improved condition number for the measurement matrix is strong evidence of GKA's ability to select samples where they are needed which in the case of integral of the Ackley function appears in the form of emptying the middle and focusing on the boundaries.

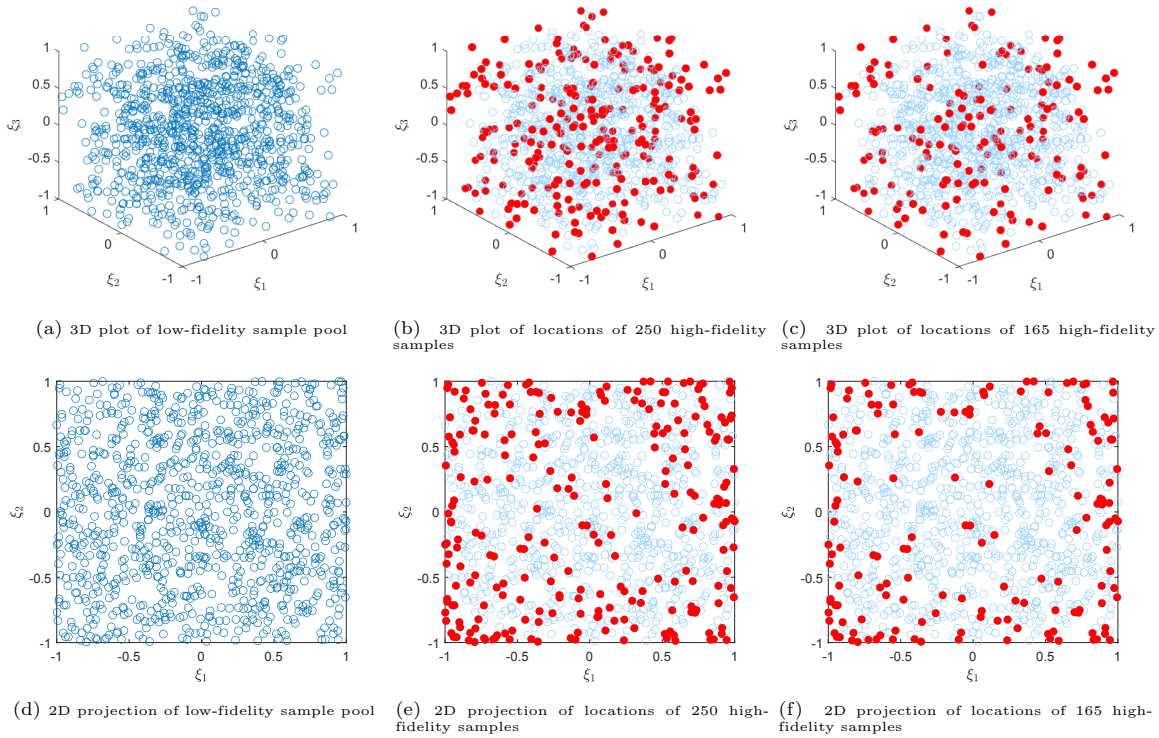

(a) 3D plot of low-fidelity sample pool (b) 3D plot of locations of 250 high-fidelity samples (c) 3D plot of locations of 165 high-fidelity samples

(d) 2D projection of low-fidelity sample pool (e) 2D projection of locations of 250 high-fidelity samples (f) 2D projection of locations of 165 high-fidelity samples

Figure 4: Locations of high-fidelity samples for accurate approximation of the integral of Ackley function

Figures 5a and 5b depict the approximation accuracy of the PCE surrogates for the integral of Ackley function using GKA without and with basis selection for different polynomial orders. It is evident that, regardless of the polynomial order chosen and whether basis selection is performed or not, the proposed strategy already results in the highest achievable accuracy for the PCE surrogates right around when the number of high-fidelity samples is equal to the cardinality of basis set.

Figure 5c shows the impact of truncation tolerance on the approximation accuracy of GKA with basis selection for a fixed polynomial order, $k = 8$. As expected, it is observed that selecting a large truncation tolerance can result in poor approximation accuracy. It is also observed that reducing the truncation tolerance beyond a threshold results in an accuracy practically identical to that of the full basis approximation. A too small of a truncation error, however, can lead to retaining many insignificant basis functions, thereby unnecessarily increasing the computational cost due to the need for more high-fidelity simulations.

Here, we do not provide a formal analysis of the impact of truncation tolerance when the proposed

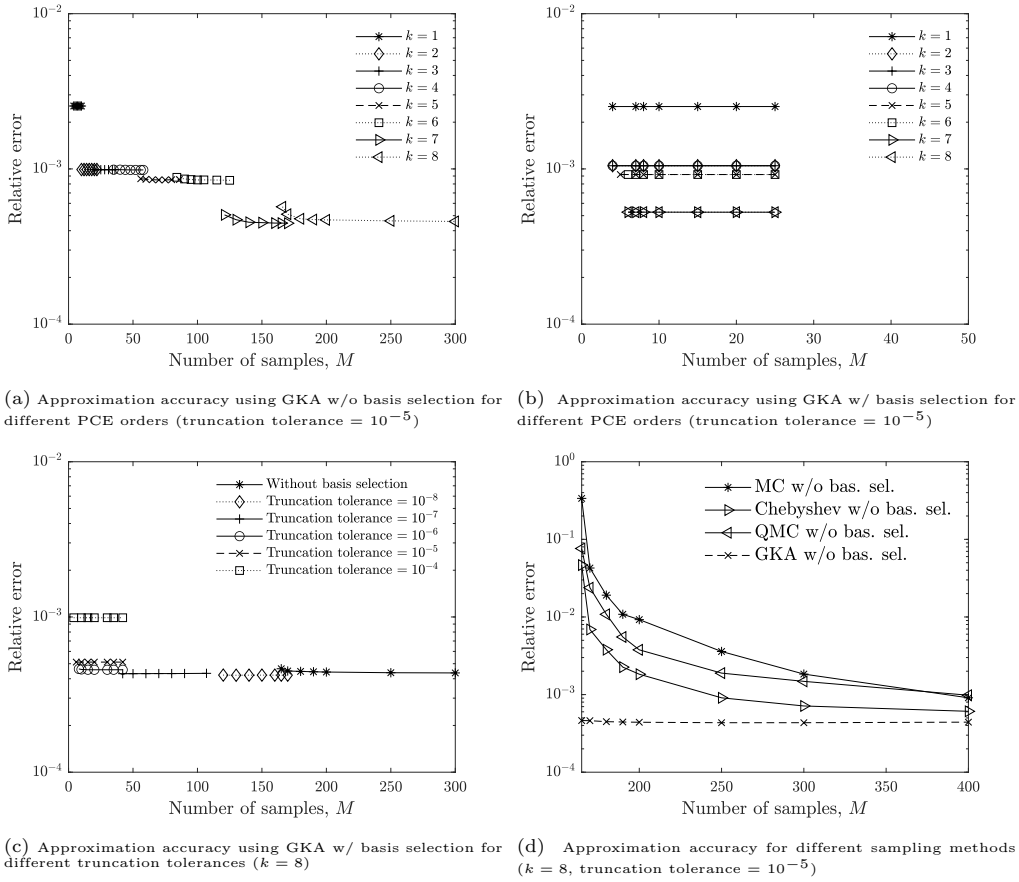

(a) Approximation accuracy using GKA w/o basis selection for different PCE orders (truncation tolerance = 10^{-5})

(b) Approximation accuracy using GKA w/ basis selection for different PCE orders (truncation tolerance = 10^{-5})

(c) Approximation accuracy using GKA w/ basis selection for different truncation tolerances ($k = 8$)

(d) Approximation accuracy for different sampling methods ($k = 8$, truncation tolerance = 10^{-5})

Figure 5: Comparison of approximation accuracy of the PCE surrogates constructed for the integral of Ackley function using GKA for different polynomial orders and truncation tolerances

strategy is used in conjunction with basis reduction or the impact of the PCE degree on the location of high-yield samples since (i) basis reduction is merely a choice and the propsoed strategy can be adopted with or without basis reduction (ii) the strength of the proposed strategy lies in its ability to locate the most informative sampling locations for high-fidelity simulations that provide highly accurate approximations with remarkable efficiency; that is with the number of samples equal to the caridinality of adopted basis.

As a final demonstration of the superior performance of the proposed approach, we compare its accuracy with Quasi-Monte Carlo sampling (using the Sobol sequence) and Chebyshev sampling, which is expected to outperform uniform sampling when $d < k$ and replicates coherence-optimal sampling when Legendre polynomial order is significantly larger than the dimensionality of the uncertain input [11]. As can be clearly seen in Figure 5d, selecting the locations of high-fidelity samples using GKA substantially improves the approximation accuracy of the constructed PCE, specially when the number of samples is closer to the cardinality of the basis set.

4.4. Effective fracture energy of layered media

As the final example, we look at effective fracture toughness for layered composites with elastic and fracture properties heterogeneity. The geometry, elastic moduli E_i, and fracture energies $\mathcal{G}_c^i$ of the constituents are illustrated in Figure 6, with E_i/E_0 and $\mathcal{G}_c^i/\mathcal{G}_c^0$ uniformly distributed on $[1, 3]$ and $[1, 2.25]$ respectively. The notched sample is subjected to an incremental linear displacement field at two boundaries with maximum strain of 18%. The Poisson's ratio is 1/3.

Fracture simulations are performed via a potential-of-mean-force approach [44, 45, 46], using lattice element method (LEM). The domain is discretized into material points interacting with each other via interaction potentials, similar to the ones used in atomistic and meso-scale simulations. The total energy of the system is the sum of ground state energy ϵ_i for each material point i, as well as central force U_{ij}^n and

bending potentials U_{ij}^b between the neighboring points i and j:

$$U_{tot} = \sum_{i}^{N} -\epsilon_i + \sum_{ij}^{N} \left(U_{ij}^n + U_{ij}^b\right) \tag{19}$$

The deformation field at every equilibrium state, minimizes the potential energy of the entire system, i.e. $\Pi = U_{tot} - W$ with W representing the external work. The interaction forces $\vec{F}_i^j$ are obtained from the gradient of force-field and are used to calculate the virial stresses. Fracture is modeled as an irreversible process of energy release between equilibrium states by breakage of bonds between the material points using Griffith energy criteria and the potential energy release rate is calculated from

$$\mathcal{G} = -\frac{\partial \Pi}{\partial \Gamma} \tag{20}$$

with $d\Gamma$ the surface created due to fracture. The QoI is the maximum energy release rate representing the effecti
material p

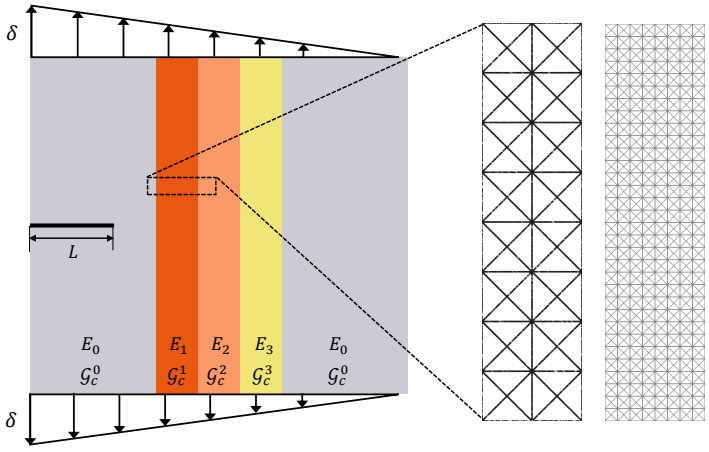

Figure 6: A multi-phase layered material subject to linear prescribed displacement

The results of high-fidelity simulations are illustrated in Figure 7 for a single realization of material properties. The first figure displays the variation of total potential energy of the system as a function of crack length. The red dots in this figure along with the three remaining figures show total potential energy and stress contours as crack reaches the phase interfaces when it propagates across the layered composite.

We aim to develop a polynomial surrogate for the high-fidelity QoI. The accuracy of the polynomial surrogate for the low-fidelity simulation reaches a plateau at fourth order i.e., the relative error converges to 0.04. The polynomial order is, hence, set at four. As in the preceding sections, the low-fidelity simulations are leveraged to select the most informative sample locations for the high-fidelity simulation. Figure 8 compares the approximation accuracy of polynomial surrogate for the high-fidelity simulation using random sampling and the proposed GKA approach. It is evident that the proposed approach substantially improves the approximation accuracy, particularly at low values of M corresponding to low computational budget. To better highlight the computational benefit of the proposed GKA, let us look at the approximation accuracy for a fixed computational budget. The computation times for the low- and high-fidelity simulations are 30 and 5000 seconds, respectively (simulations are run on cluster nodes that have cores with Intel (R) Xeon(R) Gold 6126 CPU @2.60 GHz). Consider a scenario in which the computational budget allows 1,130,000 seconds as elapsed computation time. This budget is enough for running 226 high-fidelity simulations or running 220 high-fidelity simulations plus 1000 low-fidelity simulations that can be used as the sample pool for the GKA algorithm. Using the entire budget with random sampling and high-fidelity simulations results in a surrogate with a relative error of about 0.4, while distributing the budget to select the samples using GKA will lead to a surrogate with a relative error of about 0.2. This is a significant improvement in approximation accuracy for a fixed computational budget.

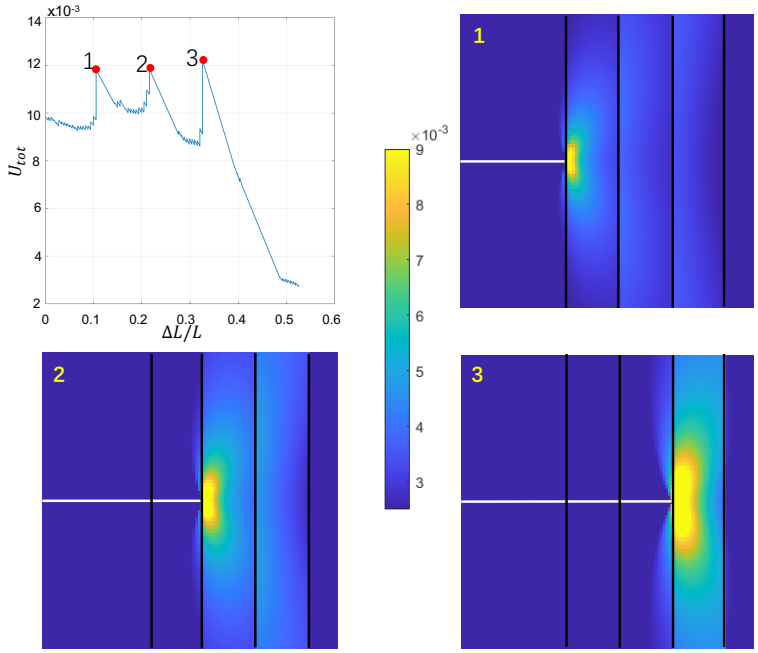

Figure 7: High-fidelity simulation results; Top left figure shows the potential energy in terms of crack length with red dots representing the interfaces, the remaining figures show stress contours respectively at points 1,2, and 3.

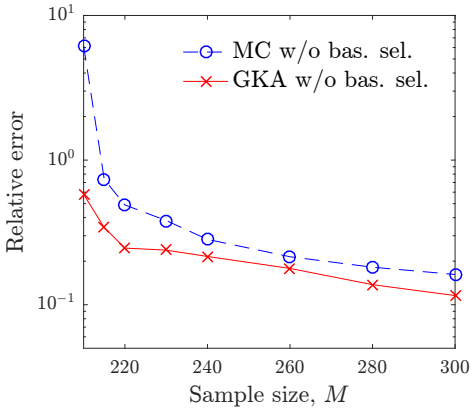

Figure 8: Comparison of approximation accuracy for effective fracture energy of layered media using MC and GKA.

4.5. A note on computational benefit and cost

The main computational benefit of the proposed GKA is due to the fact that it reduces the number of required high-fidelity simulations for surrogate construction. The ultimate computational saving, however, depends on the relative computational cost of the low- and high-fidelity simulations. It was observed in Section 4 that accurate approximations are achievable when the number of high-fidelity simulations is set as low as the cardinality of the (reduced) basis set and the size of the low-fidelity sample pool is about three times the cardinality of the full basis set. Assuming that the computational complexity of the low-fidelity simulation is substantially lower than the high-fidelity one, it can be concluded that the proposed GKA approach reduces the computational cost of PCE construction by half compared to the available multi-fidelity approaches [34, 35], that require the number of high-fidelity samples to be twice the cardinality of basis set.

The main computational cost in Algorithm 1, at each iteration, pertains to solving the least squares problem to update the coefficient vector once the most insignificant sample is identified and removed from the pool. On the other hand, identifying the least contributive sample, in each iteration, has minimal computational cost as the Kaczmarz algorithm only involves vector-by-vector multiplications. Moreover, we search among a rather small subset instead of the whole pool of sample candidates, which further improves

the computational efficiency. Ultimately, the computational cost of Algorithm 1 depends on the number of times that least square problem must be solved, which is equal to $M_p - M$. Therefore, considering the complexity of $\mathcal{O}(R^2 M_p)$ for the least square, Algorithm 1 has complexity of $\mathcal{O}(R^2 M_p^2)$. Overall, for scenarios where the cardinalities of (reduced) basis set and consequently M_p are not very large, it is fair to consider the computational cost of GKA to be negligible compared to other costs, e.g. low- and high-fidelity simulations, etc. Furthermore, for non-sparse high-dimensional problems when R is extremely large, Kaczmarz algorithm that imposes $\mathcal{O}(R)$ operations can be used to replace direct least square solution in line 3 of Algorithm 1 for estimating the coefficient vector $\boldsymbol{c}$ (readers can refer to [42] for more details). This would reduce the overall computation complexity to $\mathcal{O}(RM_p)$.

Figure 9 shows the elapsed computational time for selecting the location of 200 high-fidelity samples, out of a pool of sample candidates, for accurate approximation of the integral of the Ackley function. The following points deserve attention: (i) the proposed GKA is very fast as it takes less than a minute to select the 200 samples from a pool of 2000 sample candidates, (ii) performing basis function selection prior to sample selection substantially reduces the computational cost, (iii) expectedly, as the size of the pool increases it takes more time to search the pool and select the locations of high-fidelity samples. However, as shown by the numerical results, initiating a pool where the number of sample candidates is about three times the cardinality of full basis provides a benchmark beyond which little or no gain in accuracy is observed.

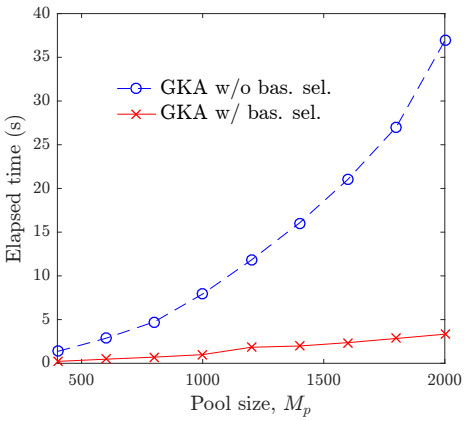

Figure 9: Elapsed time for selecting the locations of high-fidelity samples using GKA for the approximation of the integral pf Ackley function, when size of search subset is fixed at 30 samples.

5. Concluding remarks

In this paper, we presented an efficient framework for non-intrusively constructing PCE surrogates in a multi-fidelity setting. The proposed framework searches among a rather large pool of low-fidelity samples to select those high-yield samples that lead to most accurate PCE approximation for the high-fidelity model. The search relies on the Kaczmarz updating scheme and consists of iteratively identifying and removing the least influential sample from the pool defined as the sample, which if newly added to the pool, results in the least amount of improvement in the accuracy of the PCE approximation for the low-fidelity model. The efficiency provided by the Kaczmarz algorithm is complemented with a subset updating strategy that circumvents the need for looping over all samples in the pool. Several numerical examples are provided to demonstrate the performance of the proposed multi-fidelity framework. The results show a substantial improvement in accuracy can be achieved by selecting the locations of high-fidelity samples using the proposed multi-fidelity approach. They also indicate a remarkable computational gain, practically delivering accurate results with the number of the high-fidelity simulations about the cardinality of the basis. This makes the proposed approach an ideal means for efficient uncertainty quantification on fixed and limited budget especially if the performance is boosted via basis reduction which is shown to have a tangible positive impact on the overall efficiency.

Acknowledgements

The authors would like to acknowledge support from NSF grants CMMI-1351742, CMMI-1401575, CMMI-1826155 and CMMI-2047832. The authors also wish to thank the two anonymous reviewers for their constructive comments and suggestions.

References

[1] J. S. Liu, Monte Carlo strategies in scientific computing, Springer Science & Business Media, 2008.

[2] R. E. Caflisch, Monte Carlo and quasi-Monte Carlo methods, Acta numerica 7 (1998) 1–49.

[3] R. G. Ghanem, P. D. Spanos, Stochastic finite elements: a spectral approach, Courier Corporation, 2003.

[4] D. Xiu, G. E. Karniadakis, The wiener–askey polynomial chaos for stochastic differential equations, SIAM journal on scientific computing 24 (2) (2002) 619–644.

[5] D. Xiu, Numerical methods for stochastic computations: a spectral method approach, Princeton university press, 2010.

[6] O. Le Maître, O. M. Knio, Spectral methods for uncertainty quantification: with applications to computational fluid dynamics, Springer Science & Business Media, 2010.

[7] P. R. Conrad, Y. M. Marzouk, Adaptive smolyak pseudospectral approximations, SIAM Journal on Scientific Computing 35 (6) (2013) A2643–A2670.

[8] V. Barthelmann, E. Novak, K. Ritter, High dimensional polynomial interpolation on sparse grids, Advances in Computational Mathematics 12 (4) (2000) 273–288.

[9] G. T. Buzzard, Global sensitivity analysis using sparse grid interpolation and polynomial chaos, Reliability Engineering & System Safety 107 (2012) 82–89.

[10] L. Guo, A. Narayan, T. Zhou, Constructing least-squares polynomial approximations, SIAM Review 62 (2) (2020) 483–508.

[11] J. Hampton, A. Doostan, Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression, Computer Methods in Applied Mechanics and Engineering 290 (2015) 73–97.

[12] P. G. Constantine, E. Dow, Q. Wang, Active subspace methods in theory and practice: applications to kriging surfaces, SIAM Journal on Scientific Computing 36 (4) (2014) A1500–A1524.

[13] R. Tipireddy, R. Ghanem, Basis adaptation in homogeneous chaos spaces, Journal of Computational Physics 259 (2014) 304–317.

[14] Y. Shin, D. Xiu, On a near optimal sampling strategy for least squares polynomial regression, Journal of Computational Physics 326 (2016) 931–946.

[15] Y. Shin, D. Xiu, Nonadaptive quasi-optimal points selection for least squares linear regression, SIAM Journal on Scientific Computing 38 (1) (2016) A385–A411.

[16] X. Ma, N. Zabaras, An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations, Journal of Computational Physics 228 (8) (2009) 3084–3113.

[17] J. Foo, G. E. Karniadakis, Multi-element probabilistic collocation method in high dimensions, Journal of Computational Physics 229 (5) (2010) 1536–1557.

[18] A. Doostan, H. Owhadi, A non-adapted sparse approximation of pdes with stochastic inputs, Journal of Computational Physics 230 (8) (2011) 3015–3034.

[19] J. D. Jakeman, A. Narayan, T. Zhou, A generalized sampling and preconditioning scheme for sparse approximation of polynomial chaos expansions, SIAM Journal on Scientific Computing 39 (3) (2017) A1114–A1144.

[20] L. Yan, L. Guo, D. Xiu, Stochastic collocation algorithms using 1-minimization, International Journal for Uncertainty Quantification 2 (3) (2012).

[21] E. J. Candès, M. B. Wakin, An introduction to compressive sampling, IEEE signal processing magazine 25 (2) (2008) 21–30.

[22] J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies, Journal of Computational Physics 280 (2015) 363–386.

[23] N. Alemazkoor, H. Meidani, A near-optimal sampling strategy for sparse recovery of polynomial chaos expansions, Journal of Computational Physics 371 (2018) 137–151.

[24] N. Alemazkoor, H. Meidani, Divide and conquer: An incremental sparsity promoting compressive sampling approach for polynomial chaos expansions, Computer Methods in Applied Mechanics and Engineering 318 (2017) 937–956.

[25] J. D. Jakeman, M. S. Eldred, K. Sargsyan, Enhancing 1-minimization estimates of polynomial chaos expansions using basis selection, Journal of Computational Physics 289 (2015) 18–34.

[26] N. Alemazkoor, H. Meidani, A preconditioning approach for improved estimation of sparse polynomial chaos expansions, Computer Methods in Applied Mechanics and Engineering 342 (2018) 474–489.

[27] L. Le Gratiet, C. Cannamela, Cokriging-based sequential design strategies using fast cross-validation techniques for multi-fidelity computer codes, Technometrics 57 (3) (2015) 418–427.

[28] L. Parussini, D. Venturi, P. Perdikaris, G. E. Karniadakis, Multi-fidelity gaussian process regression for prediction of random fields, Journal of Computational Physics 336 (2017) 36–50.

[29] X. Song, L. Lv, W. Sun, J. Zhang, A radial basis function-based multi-fidelity surrogate model: exploring correlation between high-fidelity and low-fidelity models, Structural and Multidisciplinary Optimization 60 (3) (2019) 965–981.

[30] C. Piazzola, L. Tamellini, R. Pellegrini, R. Broglia, A. Serani, M. Diez, Uncertainty quantification of ship resistance via multi-index stochastic collocation and radial basis function surrogates: A comparison, in: AIAA AVIATION 2020 FORUM, 2020, p. 3160.

[31] A. Narayan, C. Gittelson, D. Xiu, A stochastic collocation algorithm with multifidelity models, SIAM Journal on Scientific Computing 36 (2) (2014) A495–A521.

[32] X. Zhu, A. Narayan, D. Xiu, Computational aspects of stochastic collocation with multifidelity models, SIAM/ASA Journal on Uncertainty Quantification 2 (1) (2014) 444–463.

[33] L. W.-T. Ng, M. Eldred, Multifidelity uncertainty quantification using non-intrusive polynomial chaos and stochastic collocation, in: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA, 2012, p. 1852.

[34] P. S. Palar, T. Tsuchiya, G. T. Parks, Multi-fidelity non-intrusive polynomial chaos based on regression, Computer Methods in Applied Mechanics and Engineering 305 (2016) 579–606.

[35] S. Salehi, M. Raisee, M. J. Cervantes, A. Nourbakhsh, An efficient multifidelity 1-minimization method for sparse polynomial chaos, Computer Methods in Applied Mechanics and Engineering 334 (2018) 183–207.

[36] S. Karczmarz, Angenaherte auflosung von systemen linearer glei-chungen, Bull. Int. Acad. Pol. Sic. Let., Cl. Sci. Math. Nat. (1937) 355–357.

[37] H.-J. Bungartz, M. Griebel, Sparse grids, Acta numerica 13 (2004) 147–269.

[38] A. Cohen, G. Migliorati, Optimal weighted least-squares methods, The SMAI journal of computational mathematics 3 (2017) 181–203.

[39] A. Doostan, R. G. Ghanem, J. Red-Horse, Stochastic model reduction for chaos representations, Computer methods in applied mechanics and engineering 196 (37-40) (2007) 3951–3966.

[40] B. Sudret, Global sensitivity analysis using polynomial chaos expansions, Reliability engineering & system safety 93 (7) (2008) 964–979.

[41] R. Sznajder, Kaczmarz algorithm revisited, Czasopismo Techniczne (2016).

[42] Y. Shin, D. Xiu, A randomized algorithm for multivariate function approximation, SIAM Journal on Scientific Computing 39 (3) (2017) A983–A1002.

[43] S. Xiong, P. Z. Qian, C. J. Wu, Sequential design and analysis of high-accuracy and low-accuracy computer codes, Technometrics 55 (1) (2013) 37–46.

[44] H. Laubie, S. Monfared, F. Radjaï, R. Pellenq, F.-J. Ulm, Effective Potentials and Elastic Properties in the Lattice-Element Method: Isotropy and Transverse Isotropy, Journal of Nanomechanics and Micromechanics (2017). doi:10.1061/(ASCE)NM.2153-5477.0000125.

[45] H. Laubie, S. Monfared, F. Radjai, R. Pellenq, F.-J. Ulm, Disorder-induced stiffness degradation of highly disordered porous materials, Journal of the Mechanics and Physics of Solids 106 (2017) 207–228.

[46] X. Wang, M. Botshekan, F.-J. Ulm, M. Tootkaboni, A. Louhghalam, A hybrid potential of mean force approach for simulation of fracture in heterogeneous media, Computer Methods in Applied Mechanics and Engeering (2021) 114084.