

Uncertainty quantification of discontinuous outputs via a non-intrusive bifidelity strategy

Original

Uncertainty quantification of discontinuous outputs via a non-intrusive bifidelity strategy / Canuto, Claudio; Pieraccini, Sandra; Xiu, Dongbin. - In: JOURNAL OF COMPUTATIONAL PHYSICS. - ISSN 0021-9991. - ELETTRONICO. - 398:(2019). [10.1016/j.jcp.2019.108885]

Availability:

This version is available at: 11583/2719883 since: 2020-07-03T16:15:59Z

Publisher:

ELSEVIER

Published

DOI:10.1016/j.jcp.2019.108885

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Uncertainty quantification of discontinuous outputs via a non-intrusive bifidelity strategy

Claudio Canuto^a, Sandra Pieraccini^b, Dongbin Xiu^c

^a*Dipartimento di Scienze Matematiche, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy*

^b*Dipartimento di Ingegneria Meccanica e Aerospaziale, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy*

^c*Department of Mathematics, The Ohio State University, Columbus, Ohio, USA*

Abstract

A non-intrusive bifidelity strategy [1] is applied to the computation of statistics of a quantity of interest (QoI) which depends in a non-smooth way upon the stochastic parameters. The procedure leverages the accuracy of a high-fidelity model and the efficiency of a low-fidelity model, obtained through the use of different levels of numerical resolution, to pursue a high quality approximation of the statistics with a moderate number of high-fidelity simulations. The method is applied first to synthetic test cases with outputs exhibiting either a continuous or a discontinuous behaviour, then to the realistic simulation of a flow in an underground network of fractures [2], whose stochastic geometry outputs a non-smooth QoI. In both applications, the results highlight the efficacy of the approach in terms of error decay versus the number of computed high-fidelity solutions, even when the QoI lacks smoothness. For the underground simulation problem, the observed gain in computational cost is at least of one order of magnitude.

Keywords: Multifidelity strategy; discontinuous outputs; uncertainty quantification; stochastic discrete fracture networks.

1. Introduction

The combined use of multiple computational models to simulate a physical situation of interest has become a common practice in applied Science and Engineering, and it is gaining increasing popularity among practitioners. Typically, one disposes of a model providing the desired accuracy in the simulation but at an expensive computational cost, especially if the model needs to be invoked many times with different inputs; this will be referred to as the ‘high-fidelity’ model. In addition, one or more models are available, that yield lower levels of accuracy but at cheaper prices; these will be the ‘low-fidelity’ models. For instance, one can go from high- to low-fidelity models by reducing the details in the mathematical description of the phenomena of interest; another popular and handy option consists of using the same mathematical model, but reducing the quality of its numerical discretization, e.g. by coarsening the computational mesh. Multifidelity methods aim at combining a high-fidelity model with one or more low-fidelity models (if only one low-fidelity model is invoked, one speaks of a bifidelity method), in such a way that the resulting accuracy is comparable to that of the high-fidelity model but the resulting cost is comparable to that of the low-fidelity model(s). This is possible when the response of the model exhibits a low-rank dependence upon the input data, that can already be captured sufficiently well through the information provided by the low-fidelity models.

Multifidelity strategies find wide applications in such fields as Uncertainty Quantification, Statistical Inference, Optimization (see, e.g., [3] and the references therein). In particular, in non-intrusive uncertainty quantification techniques, which are the focus of the present paper, the role of the low-fidelity computations

Email addresses: claudio.canuto@polito.it (Claudio Canuto), sandra.pieraccini@polito.it (Sandra Pieraccini), xiu.16@osu.edu (Dongbin Xiu)

may be (at least) two-fold: generate in a cheap way a surrogate of the high-fidelity output for a given (arbitrary) value of the stochastic variable (e.g., in order to compute statistical quantities of interest); and, search for the (hopefully, few) values of the stochastic variable at which the expensive high-fidelity computation has to be performed. An algorithm that implements such an approach was proposed in [1] and further refined and improved in [4–6]. The method is based on the computation of the low-fidelity solution for a set Γ of stochastic parameters, dense enough to catch the stochastic behaviour of the quantity of interest (QoI); a greedy procedure is used to iteratively select a moderate number of significant stochastic parameters, and the high-fidelity solution is computed only for these selected parameters. For any value of the stochastic parameter, an interpolation rule is constructed using the low-fidelity model, and then used to build a surrogate of the high-fidelity solution. Theoretical and practical error estimates are provided in the cited papers, and a wealth of numerical experiments clearly indicates that the bi- or tri-fidelity strategy is more efficient than classical collocation methods in different situations. However, the quantities of interest considered in these experiments depend on the stochastic variables in a smooth way.

The purpose of the present paper is to investigate the performance of the bifidelity method presented in [1] in the non-smooth case, i.e., when the QoI exhibits jumps or gradient discontinuities whose location cannot be predicted or easily detected. This is a situation that occurs in various applications, such as, e.g., in the presence of geometric uncertainties: a small variation of the stochastic parameters may induce an abrupt change in the topological structure of the domain, which is then reflected by a jump in the values of an associated QoI. We will give a significant example of such a behaviour, related to the stochastic simulation of subsurface networks of fractures. The non-smoothness of the QoI prevents an effective use of a standard stochastic collocation approach on sparse grids, since the error decay in the numerical integrations would be exceedingly slow. To avoid this drawback, one could resort to multi-element techniques (see e.g. [7]), however incurring in the non-trivial problem of properly identifying the appropriate partition of the parameter space; indeed, the localization of discontinuity interfaces through, e.g., edge detection methods (see, e.g., [8, 9]) is not easy to be applied in the realistic problem tackled herein, especially in the case of high stochastic dimensions. In general, the topological structure of the singularity set may be very complicated, with non-predictable regularity or shape. Application of techniques which rely on specific parametrizations of the discontinuity interface such as, e.g., [10] appears to be not a suitable choice in our application. A viable option is to resort to some fast realization of a Monte Carlo technique, whose convergence is not affected by possible non-smoothness of the integrand. The effectiveness of a Multilevel Monte Carlo method (which can be indeed casted in a multifidelity framework, see [3]) for the same physical application considered in the sequel has been presented in [11].

In principle, for the bifidelity strategy introduced in [1], the presence of discontinuities in the QoI should not represent an obstacle to achieving a fast decay of the error. Indeed, the error estimate presented therein shows that this error can be bounded in terms of the n -width of the solution manifold and the ‘interpolation’ error committed by generating the high-fidelity surrogates. Both errors can be made small by properly selecting the values of the stochastic parameters at which the high-fidelity samples are computed: if their distribution reflects the local structure of the solution manifold in a faithful way, errors are expected to be small. In the proposed strategy, the selection phase, realized by the greedy procedure based on the low-fidelity model, turns out to be successful in correctly identifying the location and the density of parameter values according to the local variations of the QoI. Thus, we anticipate that the main message of the present paper is that the considered bifidelity procedure works well even when it is applied to non-smooth QoI’s. However, more precise and rigorous analysis on this methodology is highly challenging in a general setting, and will be left as a future research direction.

The content of the paper is as follows. In Sect. 2, we recall the main structure of the bifidelity algorithm, and we briefly comment on certain implementation issues. In Sect. 3 we test our procedure in some ‘academic’ situations, which allow us to investigate its behaviour in a completely controlled setting; the similarity of response in the cases of smooth or non-smooth QoI’s is clearly highlighted. Sect. 4 contains the main results of the paper. In a well-assessed numerical model describing a flow in a discrete fracture network [2, 12], stochasticity is introduced by assuming that some fractures have diameters that are independent random variables; the QoI is the total flux leaving the network through some part of the boundary. Different levels of mesh refinement inside the fractures produce different approximations of the QoI, in terms of

accuracy and cost. We apply the bifidelity strategy to this problem, for several choices of pairs of high-/low-fidelity models, and for increasing stochastic dimension, and we analyze its performances. Finally, Sect. 5 summarizes the conclusions of our investigation.

2. Methodology

For the description of the bifidelity approximation, we closely follow [1, 4]. As a general framework, we consider a problem whose solution u is defined on a bounded space domain $\mathcal{D} \subset \mathbb{R}^D$ and depends on a stochastic parameter $z \in \mathcal{Z} \subset \mathbb{R}^d$, which represents the value assumed by a random variable Z defined on a given probability space $(\Omega, \mathcal{A}, \mathbb{P})$. A quantity of interest (QoI) is considered, which is a function of the solution u :

$$Q = \hat{Q}(u) : \mathcal{D} \times \mathcal{Z} \rightarrow \mathbb{R}.$$

Let V^L and V^H denote two finite-dimensional Hilbert spaces in which low- and high-fidelity approximations of $u(x, z)$ are respectively sought, for each $z \in \mathcal{Z}$, with $\dim V^L \ll \dim V^H$. Such approximations will be denoted by $u^L(z)$ and $u^H(z)$, respectively. It is assumed that the task of computing a high-fidelity approximation is significantly more expensive than computing a low-fidelity approximation.

Within this framework, we consider a discrete subset $\Gamma \subset \mathcal{Z}$ large enough to catch the stochastic behaviour of Q ; we denote by $M = \#\Gamma$ the cardinality of Γ . For each $z \in \Gamma$, the low-fidelity approximation $u^L(z)$ is computed, forming a set denoted $u^L(\Gamma)$. The choice of the set Γ depends on the sampling method adopted (namely, stochastic collocation, Monte Carlo, etc.).

A subset $\gamma_N \subset \Gamma$ is then selected, with cardinality $N = \#\gamma_N$, $N \ll M$. As before, the set of low-fidelity simulations computed for $z \in \gamma_N$ is denoted $u^L(\gamma_N)$, and the vector space spanned by elements in $u^L(\gamma_N)$ is denoted $U^L(\gamma_N)$. Analogous notation is used for the high-fidelity approximations. A key point is given by the selection of γ_N . In [1], a greedy approach is proposed, in which, at each step k , starting from a given set γ_{k-1} , a new node is chosen, and then added to γ_{k-1} , by maximizing the distance between the solution corresponding to the new node, and the vector space spanned by the solutions corresponding to nodes in γ_{k-1} . Namely, at step k , one set

$$\gamma_k := \gamma_{k-1} \cup \{z_{i_k}\} \tag{1}$$

where

$$z_{i_k} = \operatorname{argmax}_{z \in \Gamma} \operatorname{dist}(u^L(z), U^L(\gamma_{k-1})). \tag{2}$$

This step, as mentioned in [1, Lemma 3.1], can be tackled in several ways, all involving rather standard linear algebra operations. All the approaches rely on the construction of the matrix $\mathbf{V} \in \mathbb{R}^{n^L \times M}$, where $n^L = \dim V^L$, whose i -th column corresponds to the DOFs of $u^L(z_i)$, $i = 1, \dots, M$, and the Gramian matrix $\mathbf{G} \in \mathbb{R}^{n^L \times n^L}$ corresponding to the scalar product in V^L ; the outcome of the approaches is a permutation matrix \mathbf{P} such that the nodes detected to form γ_N correspond to the first N rows of \mathbf{PZ} , where $\mathbf{Z} \in \mathbb{R}^{M \times d}$ is the matrix whose rows are the parameters $z_i \in \Gamma$. The suggested approaches are the following ones:

- a) compute the Cholesky decomposition of \mathbf{G} , $\mathbf{G} = \mathbf{H}\mathbf{H}^T$, and then perform a column-pivoted QR decomposition of the matrix $\mathbf{V}^T\mathbf{H}$, namely $\mathbf{V}^T\mathbf{H} = \mathbf{Q}\mathbf{R}\mathbf{P}^T$;
- b) perform a full-pivoted LU decomposition of $\mathbf{V}^T\mathbf{G}\mathbf{V}$, namely $\mathbf{V}^T\mathbf{G}\mathbf{V} = \mathbf{P}^T\mathbf{L}\mathbf{U}\mathbf{P}$;
- c) perform a pivoted Cholesky decomposition of $\mathbf{V}^T\mathbf{G}\mathbf{V}$, namely $\mathbf{V}^T\mathbf{G}\mathbf{V} = \mathbf{P}^T\mathbf{L}\mathbf{L}^T\mathbf{P}$.

Note that with each one of the proposed approaches one always has $\gamma_{N_1} \subset \gamma_{N_2}$ for $N_1 < N_2$. Computational and memory storage considerations make the third approach the one adopted in [1, 4]. It is worth mentioning that the Cholesky decomposition process (which indeed coincides with the adaptive cross approximation method [13] when applied to symmetric matrices) need not be fully performed, as it can be interrupted at N -th step, once N is chosen. We also remark that, although the matrix \mathbf{V} is assumed to have full rank, in order to avoid numerical singularity, the Cholesky decomposition process may get interrupted at an earlier stage, thus revealing the numerical rank of \mathbf{V} and \mathbf{L} .

After the selection step is performed, the reconstruction phase comes into play. The idea behind the reconstruction phase is to build, for any $z \in \Gamma$, the projection of $u^L(z)$ onto $U^L(\gamma_N)$, namely

$$\mathcal{P}_{U_N^L}[u^L(z)] = \sum_{k=1}^N c_k u^L(z_{i_k}), \quad z_{i_k} \in \gamma_N, \quad (3)$$

where the coefficients c_k are given by the solution of the linear system

$$\tilde{\mathbf{G}}\mathbf{c} = \mathbf{g}, \quad (4)$$

being $\tilde{\mathbf{G}}$ the truncated Gramian matrix and the elements of \mathbf{g} given by $g_k = \langle u^L(z), u^L(z_{i_k}) \rangle$, $k = 1, \dots, N$. Once the vector of coefficients \mathbf{c} is computed, the bifidelity reconstruction $u_{r,N}^H(z)$ is built mimicking (3) on $U^H(\gamma_N)$, namely:

$$u_{r,N}^H(z) := \sum_{k=1}^N c_k u^H(z_{i_k}), \quad z_{i_k} \in \gamma_N. \quad (5)$$

We remark that the linear system (4) may suffer from ill-conditioning, especially if near rank deficiency is encountered in the selection phase. This phenomenon will clearly emerge from the numerical experiments presented in the paper.

Note that model (5) is defined for all $z \in \mathcal{Z}$; if $z \in \Gamma$, it is enough to compute $u^L(z)$ and then proceed as indicated above. However, since $u^L(z)$ has not been involved in the greedy selection phase, then the accuracy of the projection will depend on how the set Γ is representative of the behaviour of the parameter z on the whole set \mathcal{Z} .

3. Synthetic test cases

In this section we discuss some synthetic test cases introduced in order to assess the potential of the multifidelity approach for nonsmooth problems. We consider the following framework. Let a function $u(x; z)$ be given, depending on the space variable $x \in \mathcal{D} \subset \mathbb{R}^D$ and on a stochastic parameter $z \in \mathcal{Z} \subset \mathbb{R}^d$, with $|\mathcal{D}|, |\mathcal{Z}| < \infty$. Assuming $u \in L^2(\mathcal{D} \times \mathcal{Z})$ with respect to the Lebesgue measure, we consider as a quantity of interest

$$Q(z) := \|u(\cdot; z)\|_{L^2(\mathcal{D})}. \quad (6)$$

We investigate the use of the bifidelity approach for approximating the mean value of Q . The common setting of all test cases reported in this section is the following: we approximate u , supposed to be continuous in the x variable, by interpolating it either by global polynomials or by piecewise multi-linear functions. We consider test cases corresponding to two different situations: a continuous QoI in the stochastic space, and a discontinuous one; in the former case, we use a mildly oscillating function u (case A), and a highly oscillating function (case B).

3.1. Framework 2D-1D

Hereafter, we consider several cases, in which the space variable x is 2D, whereas the stochastic variable z is 1D.

3.1.1. Continuous QoI

We start by discussing two cases in which the QoI is a continuous, smooth function of z ; these tests are introduced for the sake of reference. We take u as

$$u(x, y; z) = \sin(a\pi(xy + bz(1 - x^2 + y^2))), \quad (x, y) \in [-1, 1]^2, \quad z \in \mathcal{Z} := [-1, 1]. \quad (7)$$

Here a, b are two parameters, chosen as follows:

$$\text{Case A: } a = 0.8, \quad b = 1, \quad \text{Case B: } a = 4, \quad b = 1.2.$$




Figure 1: Framework 2D-1D. Case A: function $u(x, y; z)$ for $z = -0.5$ (left) and interpolating polynomials of degree $n = 3$ (middle) and $n = 7$ (right).




Figure 2: Framework 2D-1D. Case B: function $u(x, y; z)$ for $z = -0.5$ (left) and interpolating polynomials of degree $n = 7$ (middle) and $n = 19$ (right).

Case B clearly corresponds to a function which is in general displaying more oscillations in the x space, for a fixed stochastic parameter z , than case A. A typical example is reported in Figs. 1 and 2 for the two cases, with the same fixed value of z (namely, $z = -0.5$).

The same figures report some interpolating polynomials built on tensorized Gauss-Lobatto nodes. Note that, as expected, in case A few nodes in each direction ($n = 7$) are enough to catch the correct behaviour of u , whereas, in case B, $n = 7$ nodes are yet not enough; $n = 19$ nodes are needed for a good agreement between u and its interpolating polynomial.

We consider a high-fidelity approximation u^H obtained as the global interpolating polynomial built on the tensorized Gauss-Lobatto grid with n^H nodes in each direction (using $n^H = 15$ for case A and $n^H = 30$ for case B), whereas the low-fidelity approximation u^L is obtained by one of the following approximations:

1. u^L is the global polynomial interpolating u on the tensorized Gauss-Lobatto grid with $n^L < n^H$ nodes in each direction;
2. u^L is a continuous, piecewise bilinear function interpolating u on the tensorized uniform grid with n^L nodes in each direction.

In all cases, each approximating low-fidelity (high-fidelity, resp.) function is identified by a vector of $(n^L)^2$ ($(n^H)^2$, resp.) elements containing either the function values at the nodes (case 2. from the previous list) or the coefficients of the discrete Legendre transform (high-fidelity and case 1., in which we express the functions in terms of the Legendre polynomials).

We have considered a set Γ of $M = 1000$ random stochastic parameters with uniform distribution. We have applied the multifidelity selection algorithm recalled in Sect. 2 in order to select a subset $\gamma_N \subset \Gamma$ made of N parameters (several values of N have been used, see later), and then we have computed the corresponding high-fidelity approximations on γ_N . Then, for each given value $z \in \Gamma$ of the stochastic parameter, we reconstruct a high-fidelity approximation $u_{r,N}^H(z)$ starting from $u^L(z)$ and from the set $u^H(\gamma_N)$. For each of

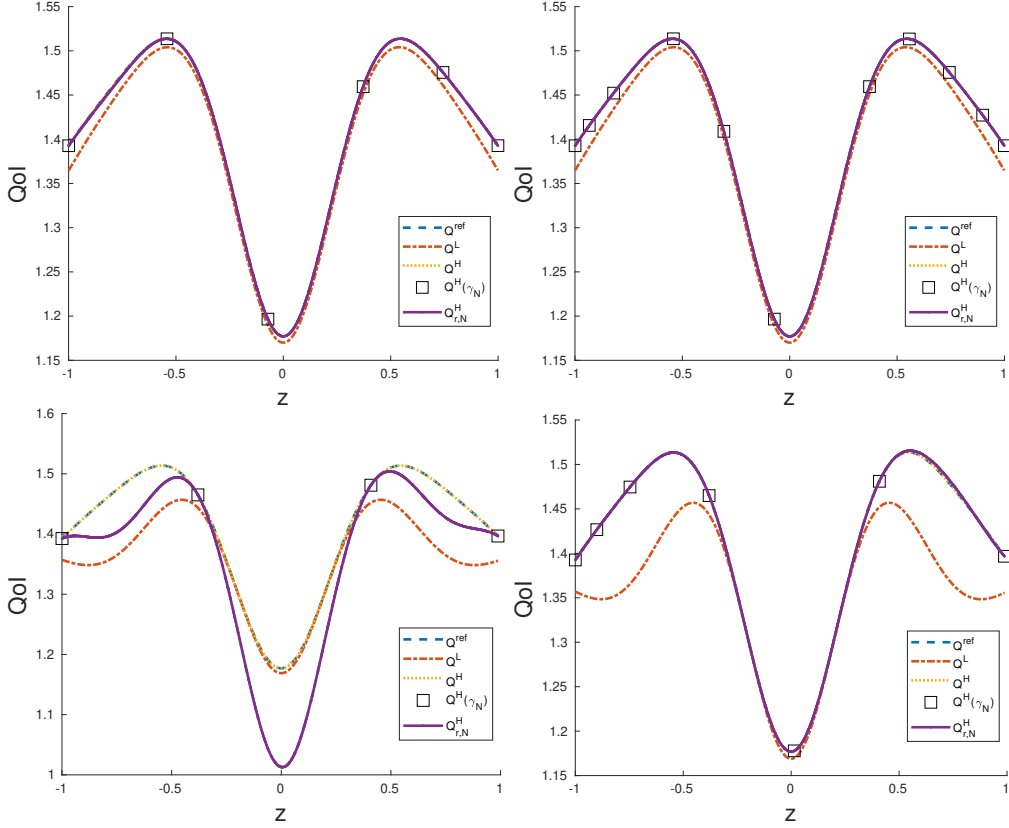


Figure 3: Framework 2D-1D. Case A: different approximations of the QoI versus z . Low-fidelity approximations by piecewise bilinear functions with $n^L = 20$ (top), or by global polynomials with $n^L = 5$ (bottom). Left to right: $N = 6$ and $N = 11$ (top); $N = 4$ and $N = 7$ (bottom)

these approximations, we compute the QoI by using a suitable quadrature rule in \mathcal{D} . We have therefore at hand the following quantities:

- $Q^L(z)$, the low-fidelity approximation of the QoI;
- $Q^H(z)$, the *true* high-fidelity approximation;
- $Q^H_{r,N}(z)$, the reconstructed value for the high-fidelity approximation;
- $Q^{\text{ref}}(z)$, the “exact” reference value, namely, a good approximation of the exact QoI using a large number (approximately $2n^H$) of Gauss-Lobatto quadrature nodes in each direction for computing (6).

The aforementioned quantities are plotted versus z in Fig. 3 (case A) and Fig. 4 (case B). Different values of N are considered in the plots. We also highlight in boxes the *pivot* values of $Q^H(z)$ that are used by the algorithm for the reconstruction (namely, $Q^H(z)$ for $z \in \gamma_N$). Note that Q^H is always graphically superimposed to Q^{ref} , and $Q^H_{r,N}$ is essentially superimposed to Q^{ref} for the largest considered values of N . Obviously, larger values of N are needed in case B than in case A.

In Figs. 5 and 6 we report, for cases A and B, respectively, the relative error on the mean QoI versus N (*mean error* for short in the figure labels); the error is computed as follows:

$$\text{err}_N = \frac{|\bar{Q}^H_{r,N} - \bar{Q}^{\text{ref}}|}{|\bar{Q}^{\text{ref}}|},$$

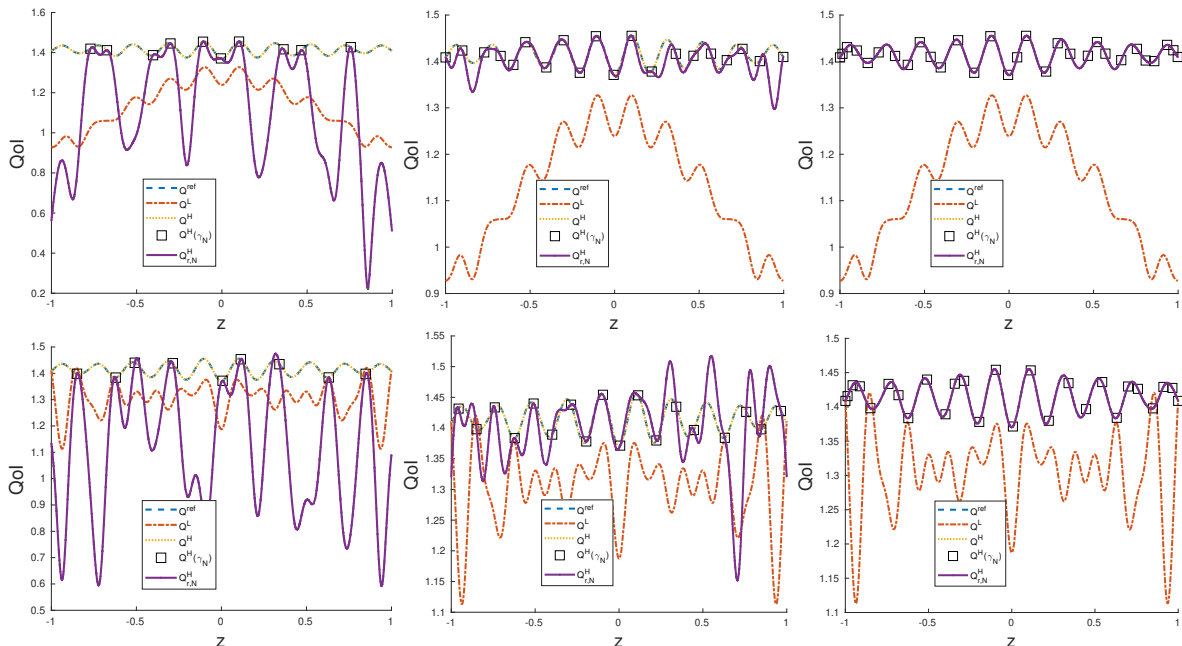


Figure 4: Framework 2D-1D. Case B: different approximations of the QoI versus z . Low-fidelity approximations by piecewise bilinear functions with $n^L = 20$ (top), or by global polynomials with $n^L = 9$ (bottom). Left to right: $N = 10, 20, 30$ (top) and $N = 9, 18, 27$ (bottom).

where

$$\bar{Q}_{r,N}^H = \frac{\sum_{z \in \Gamma} Q_{r,N}^H(z)}{\#\Gamma}, \quad \bar{Q}^{\text{ref}} = \frac{\sum_{z \in \Gamma} Q^{\text{ref}}(z)}{\#\Gamma}.$$

In the same figure we also report, for the sake of reference, the relative error on the mean QoI which is obtained by using only high-fidelity approximations (labeled “High fid.” in the legend), namely the value

$$\text{err}_* = \frac{|\bar{Q}^H - \bar{Q}^{\text{ref}}|}{|\bar{Q}^{\text{ref}}|}, \quad \bar{Q}^H = \frac{\sum_{z \in \Gamma} Q^H(z)}{\#\Gamma}.$$

Exponential convergence is apparent, for both types of low-fidelity approximations (piecewise bilinears or global polynomials), although the decay is somehow less regular in the latter case. Furthermore, the blue, red and yellow curves are nearly superimposed as long as they exist, indicating that the different resolutions of the low-fidelity approximations have a minor influence on the behaviour of the error. The break-down of the error curves is due to a procedure for preventing the onset of numerical rank deficiency in the selection step of the algorithm; note that this occurs at values of N that are significantly smaller than the dimension of the low-fidelity space V^L , indicating the onset of numerical linear dependence among the remaining low-fidelity functions.

3.1.2. Discontinuous QoI

We introduce now a test case (case C) with the same setting as in Sect. 3.1.1, but presenting a discontinuity in the QoI with respect to the stochastic variable. Namely, we still consider u given by (7), but the values of the parameters a, b are now piecewise constant, rather than constant, as functions of z ; this creates a jump in the space of stochastic parameters. To be precise, the following values are used in (7):

$$a = \begin{cases} 0.8 & z < 0.8, \\ 4 & z \geq 0.8, \end{cases} \quad b = \begin{cases} 1 & z < 0.8, \\ 1.2 & z \geq 0.8. \end{cases}$$




Figure 5: Framework 2D-1D. Case A: relative error on the mean QoI versus N . Low-fidelity approximations by piecewise bilinear functions (left), or by global polynomials (right).




Figure 6: Framework 2D-1D. Case B: relative error on the mean QoI versus N . Low-fidelity approximations by piecewise bilinear functions (left), or by global polynomials (right).

The resulting approximations of the QoI, namely $Q^L(z)$, $Q_{r,N}^H(z)$, $Q^H(z)$ and $Q^{\text{ref}}(z)$, are reported in Fig. 7. The mean errors, computed as in the continuous case, are reported in Fig. 8. Note the clear separation of scales introduced by the discontinuity, and the fact that the distribution of pivot values of Q is well adapted to the local structure of the target function: we observe a coarse distribution in the (larger) interval where the function is locally smoother, and a dense distribution in the (smaller) interval where the function is locally more oscillating. The error plots confirm the exponential decay of the errors even in the presence of discontinuities in the stochastic parameter space.

3.2. Framework 2D-2D

We consider now several test cases in which both the space and the stochastic variables are 2D. Based on the results of the previous subsection, indicating that using piecewise bilinear functions or global polynomials in the low-fidelity approximations yields comparable results, hereafter we confine ourselves to using piecewise bilinear interpolations for constructing the low-fidelity approximations u^L , while we keep considering higher-order global interpolation polynomials for the high-fidelity approximations u^H ; in all cases we use tensorized grids in the square $[-1, 1]^2$.

3.2.1. Continuous QoI

We begin with a QoI which is a continuous function of the stochastic parameters. The function u is now defined as

$$u(x, y; z) = \sin(\pi(x + a(z_1xy + bz_2(1 - x^2 + y^2))), \quad (x, y) \in [-1, 1]^2, \quad z = (z_1, z_2) \in \mathcal{Z} := [-1, 1]^2. \quad (8)$$




Figure 7: Framework 2D-1D. Case C: different approximations of the QoI versus z . Low-fidelity approximations by piecewise bilinear functions with $n^L = 20$ (top), or by global polynomials with $n^L = 9$ (bottom). Left to right: $N = 6, 12, 18$.

We consider the following two cases:

$$\text{Case D: } a = 0.8, b = 1,$$

$$\text{Case E: } a = 4, b = 1.2.$$

In Fig. 9, left, we show the set γ_N , i.e., the values of z (highlighted by square symbols) at which the pivot values of Q^H are computed. The relative errors on the mean QoI versus N are reported in Fig. 9, right. Similar results are reported for Case E in Fig. 10. In both cases, the distribution of pivot points appears to be coherent with the local structure of the target function. The exponential decay of the errors is confirmed; obviously, different values of N are needed to achieve the same accuracy when varying the local smoothness of the function.

3.2.2. Discontinuous QoI

Hereafter, we consider two cases in which, similarly to Sect. 3.1.2, a discontinuity in the space of stochastic parameters is introduced by letting the parameters depend upon z . Given

$$u(x, y; z) = \sin(\pi(x + a(z_1xy + bz_2(1 - x^2 + y^2)))) + c, \quad (x, y) \in [-1, 1]^2, \quad z \in \mathcal{Z} := [-1, 1]^2, \quad (9)$$

our test cases correspond to the following choices of parameters a, b, c :

- Case F: $a = 1.6, b = 1$,

$$c = \begin{cases} 1 & \text{if } z \in \mathcal{B}_F \\ 2 & \text{if } z \notin \mathcal{B}_F \end{cases}$$

$$\text{where } \mathcal{B}_F = \{z \in [-1, 1]^2 : z_1^2 + z_2^2 \leq \frac{1}{4}\}.$$

- Case G: $b = 1, c = 0$,

$$a = \begin{cases} 4 & \text{if } z \in \mathcal{B}_G \\ 0.8 & \text{if } z \notin \mathcal{B}_G \end{cases}$$

$$\text{where } \mathcal{B}_G = \{z \in [-1, 1]^2 : (z_1 - 1)^2 + (z_2 + 1)^2 \leq 1\}.$$




Figure 8: Framework 2D-1D. Case C: relative error on the mean QoI versus N . Low-fidelity approximations by piecewise bilinear functions (left), or by global polynomials (right).

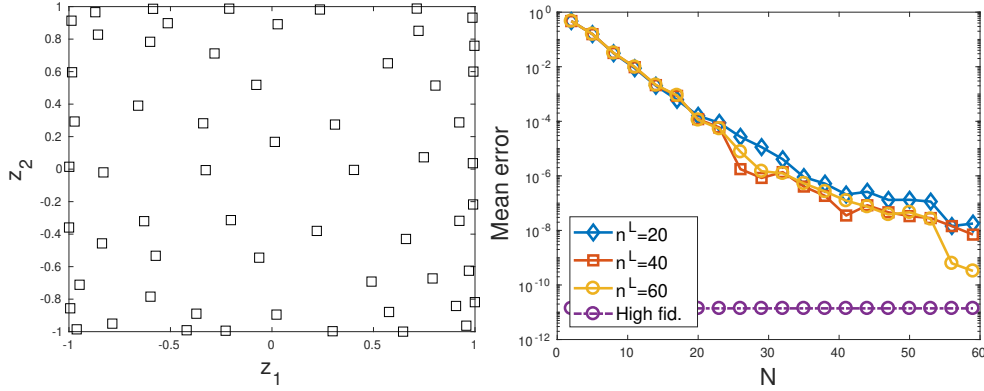


Figure 9: Framework 2D-2D, Case D. Left: pivot values of z ($N = 59$); right: relative error on the mean QoI versus N

The exact QoI's corresponding to the two cases are reported in Fig. 11, which clearly highlights the discontinuity interfaces. Fig. 12 reports, for case F, the values of z at which the pivot QoI's are computed, i.e., the set γ_N (left), and the relative errors on the mean QoI versus N (right). Similar results are reported in Fig. 13 for case G. As for case C, but even more evidently, the distribution of pivot values appears to be insensitive to the presence of a discontinuity in the stochastic parameter space; in particular, note that looking at Fig. 12, left, it is impossible to detect the position of the discontinuity curve in the middle of the square. What really matters is the “local smoothness”, or local oscillatory structure, of the target function, as clearly documented by the distribution of pivot values shown in Fig. 13, left. On the other hand, the exponential decay of the error is achieved also for these cases (note that in Fig. 12, right, the target error “High fid.” is significantly smaller than in all other cases, probably due to some cancellation effect).

To shed more light on the effect of discontinuities in parameter space, we focus again on case G and we show in Figure 14 the plots of the low-fidelity functions $u^L(z)$ for two representative values of z , namely, the first $z \in \gamma_N \setminus \mathcal{B}_G$ (left plot) and the first $z \in \gamma_N \cap \mathcal{B}_G$ (right plot) selected by the greedy algorithm (1)-(2); the differences in structure are manifest. While the reconstruction (5) involves in the expansion *all* the selected high-fidelity functions u^H , associated to parameters z_{i_k} lying both inside and outside \mathcal{B}_G , the projection (3)-(4) is able to give the appropriate weight to each family of basis functions, thus properly reproducing the local structure of the function to be reconstructed. This explains why the reconstruction is so accurate for all regions in parameter space. An example is provided by Figure 15, where we plot for two selected values of $z \notin \gamma_N$ (one outside \mathcal{B}_G , the other inside \mathcal{B}_G) the reconstructed functions $u_{r,N}^H(z)$ for $N = 50$ and $N = 100$, as well as the target function $u^H(z)$. A direct inspection of the entries of vector




Figure 10: Framework 2D-2D. Case E. Left: pivot values of z ($N = 198$); right: relative error on the mean QoI versus N




Figure 11: Framework 2D-2D. Discontinuous case: QoI versus z . Case F (left) and Case G (right).

c (see equation (4)) highlights that the basis functions contributing more to the reconstruction of $u^H(z)$ correspond to values of the stochastic parameter belonging to the same region to which z belongs.

4. Application to Discrete Fracture Networks

4.1. Problem description

We consider here an application arising in the field of underground flow simulations in fractured media [14–16]. We consider an impervious rock matrix crossed by a set of intersecting fractures, as depicted in Fig. 16. Each fracture is usually represented as an ellipse or a polygon, and endowed with a symmetric positive definite tensor (the so-called fracture transmissivity tensor) describing the hydrogeological properties of the fracture. We consider the flow in such network of fractures as ruled by the Darcy law on each fracture, with flux exchange occurring through fracture intersections. The primal variable is the pressure on each fracture, usually called hydraulic head. A related quantity is the co-normal derivative of the hydraulic head: computed at the intersection between two fractures, it describes the flow exchange between the two fractures; computed at a boundary edge of a fracture, it represents the flow exiting the network through such boundary edge. With reference to Fig. 16, we focus on a problem in which an inlet flow is considered on the left-hand side of the network, whereas an outlet flow is considered on the right-hand-side of the network (namely, pressure is set equal to 10 on the left, equal to 0 on the right). On the right of the network, a large fracture can be seen, which may represent an outcrop. We are interested in measuring the flux outflowing

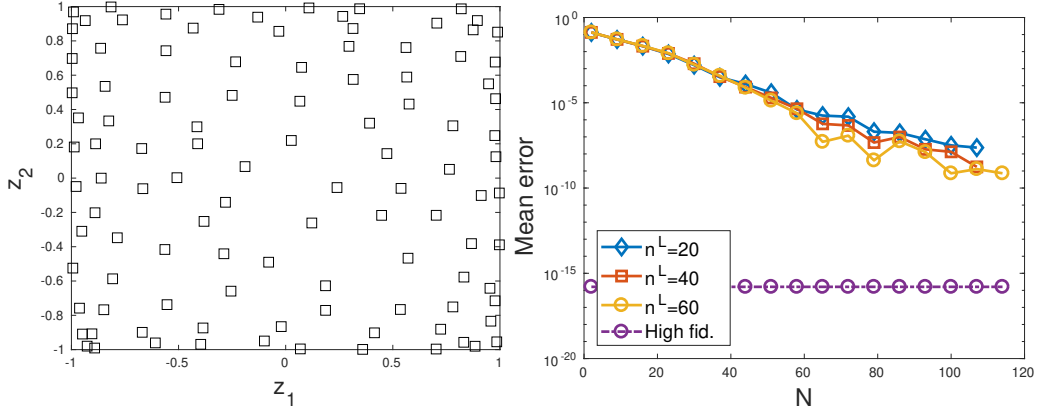


Figure 12: Framework 2D-2D. Case F. Left: pivot values of z ($N = 107$); right: relative error on the mean QoI versus N

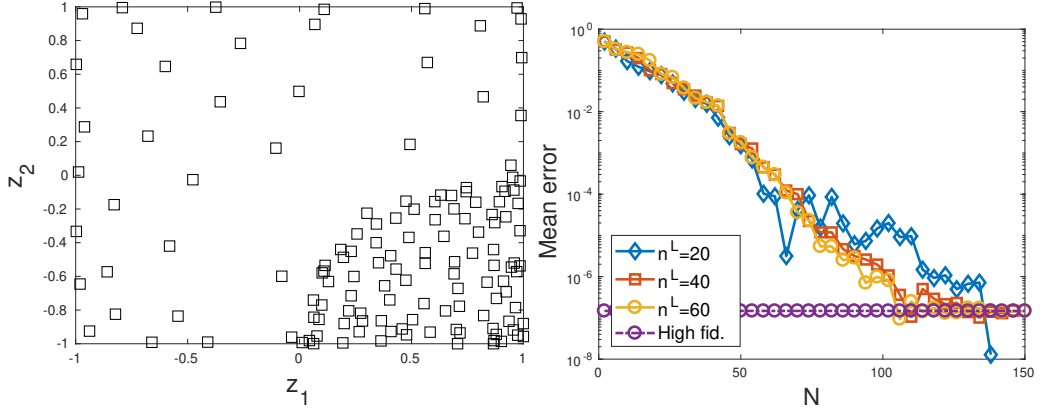


Figure 13: Framework 2D-2D. Case G. Left: pivot values of z ($N = 138$); right: relative error on the mean QoI versus N

from the network through this fracture. So, the QoI is here the total flux entering this fracture through its intersection with other fractures, which in turn exits the network from its non-insulated boundary edges.

Within this framework, it is a crucial point the lack of full deterministic description of all (hydrogeological and geometrical) features of the fractures; probabilistic distributions of such properties are only available. As such, uncertainty quantification is a crucial issue. A quite interesting and challenging situation is the one in which the geometry is described by stochastic parameters, as nonsmoothness in the QoI is likely to arise, see [11].

We consider in this section a realistic network in which fractures are assumed to be circular (represented as octagons). The overall network approximately counts 700 fractures. Among them, we select $d \geq 1$ fractures as having a radius which is a random variable with a given distribution; d represents therefore the stochastic dimension. When $d > 1$, we will assume the random variables to be independent. Let R denote the random variable describing the radius of a selected fracture; we assume that R has a power-law distribution with cut-off. This is a common distribution adopted for fracture radii in geomechanics, see [17, 18]. To be precise, let the probability density function of R be of the form

$$f_R(r) \propto r^{-\alpha}$$

with $\alpha > 1$. We impose a lower bound $r_L > 0$ on the values of R , in such a way that f is well-defined (coherently with the need of having non-vanishing radii), and an upper bound r_U , introduced in order to




Figure 14: Framework 2D-2D. Case G. Examples of basis function $u^L(z)$, $z \in \gamma_N$. Left: first selected value $z \notin \mathcal{B}_G$; right: first selected $z \in \mathcal{B}_G$.

avoid excessively large fractures with respect to the domain size. By enforcing the normalization

$$\int_{r_L}^{r_U} f_R(r) dr = 1,$$

one has

$$f_R(r) = \frac{1 - \alpha}{r_U^{1-\alpha} - r_L^{1-\alpha}} r^{-\alpha}.$$

The previous relations are often written in terms of $\gamma = \alpha - 1 > 0$. Using the inversion method (see, e.g., [19, Proposition 2.11]), the random variable R can be sampled starting from a uniformly distributed sample $u \in [0, 1]$ following

$$r = r_L \left(1 + u \left(\left(\frac{r_U}{r_L} \right)^{-\gamma} - 1 \right) \right)^{-\frac{1}{\gamma}}. \quad (10)$$

Values used for r_L , r_U and γ are $r_L = 15$, $r_U = 560$, $\gamma = 2.5$ in a domain of size 1000 [17]. In Fig. 17, we report the density function $f_R(r)$ (left plot), the cumulative distribution function $F_R(r)$ (center plot), and the behaviour of radii sampled from (10) (right plot). The random variable R plays therefore the role of the random variable Z introduced in Section 2.

4.2. Numerical treatment

The flow problem is modeled by Darcy law [2] and is numerically solved by linear finite elements on each fracture [20]. As a general framework, we consider as low- and high-fidelity approximations the solutions obtained with coarse and fine meshes, respectively. We follow the approach of having on each fracture approximately the same number of elements, and consider several possible mesh levels: $\ell = 1$ corresponds to approximately 16 elements on each fracture; at its opposite, $\ell = 7$ corresponds to approximately 1024 elements on each fracture. The finest mesh ($\ell = 7$) is only used for obtaining the reference solution. Mesh levels from 1 to 3 are used for the low-fidelity approximations; mesh levels from 4 to 6 are used for the high-fidelity approximations (see Table 1 for a summary). In what follows, for the sake of brevity, we will use the notation LF= ℓ (HF= ℓ , resp.) as a shorthand for the low-fidelity (high-fidelity, resp.) solution being obtained with mesh level ℓ . We recall that the QoI is here the integral of the co-normal derivative of the hydraulic head on all the intersections of a selected boundary fracture F with other fractures; these intersections are called traces of F . With reference to the notation introduced in Sec. 2, for each sample of the stochastic variable, the vectors u^L and u^H are obtained collecting the DOFs for the hydraulic head on F and the DOFs corresponding to the discretization of the co-normal derivative of the head on the traces of F . The overall number of DOFs obtained, corresponding to $\dim V^L$ and $\dim V^H$, respectively, is listed in Table 1. Note that this information only refers to the fracture F on which the QoI is measured. The total

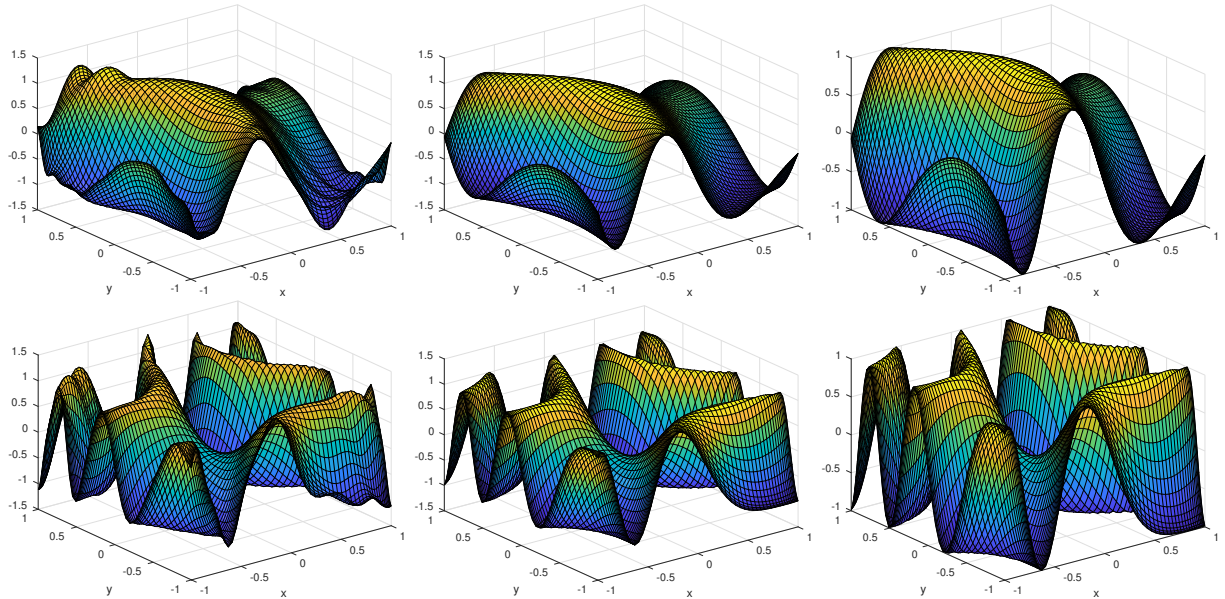


Figure 15: Framework 2D-2D. Case G. Functions $u_{r,50}^H(z)$ (left), $u_{r,100}^H(z)$ (middle), and $u^H(z)$ for two selected z values. Top: $z \notin \mathcal{B}_G$; bottom: $z \in \mathcal{B}_G$.

number of DOFs required for the computation of the solution is much larger (approximately 700 times such value).

In the sequel, we investigate the behaviour of the quantity of interest while increasing the stochastic dimensionality of the problem, i.e., the number of fractures whose radius is randomly varied.

mesh level	approx # of elements per fracture	n^L, n^H	usage
1	16	18	LF
2	32	34	LF
3	64	58	LF
4	128	117	HF
5	256	216	HF
6	512	431	HF
7	1024	852	reference

Table 1: Description of mesh levels

4.3. Case $d = 1$

In Fig. 18 we report, in blue color, the QoI versus the radius of F evaluated at 10000 random points, in order to have a dense set of samples. For the application of the bifidelity algorithm, we have considered a set Γ made of $M = 1500$ randomly selected radii. The non-smooth behaviour is clear, suggesting that both jumps and gradients discontinuities are present. In the same figure we also report, in red color and with labels on the right y axis, the number of fracture intersections versus the samples of the stochastic variable. Most points at which the QoI is non-smooth correspond to a change in the number of fractures intersections. This is clearly seen, for example, at the leftmost jump in the blue plot, which occurs at $r \simeq 40$ and do correspond to a (albeit tiny) jump in the red curve, corresponding to an increment of just one fracture intersection. The same can be easily seen for the jump occurring at $r \simeq 10^2$, which again do correspond to an increment of just one trace. The reason for this behavior is that, while changing the radius




Figure 16: Geometry of the network




Figure 17: Power-law with cut-off: density function (left), cumulative distribution function (center), sampling of r with u (right)

of a fracture, a new intersection may appear/disappear, and this may impact on the network connectivity and flux directionality, thus introducing non-smoothness in the QoI. On the other hand, in some situations (e.g., the leftmost jump in the number of fracture intersections) the new fracture intersections may have no consequence on the regularity of the QoI. This may be due to the fact that the new fracture intersections possibly involve a fracture which is situated in a “dead-end” branch of the network.

In Fig. 19 we report, for a fixed LF level and a fixed HF level, the values of the low-fidelity quantity of interest $Q^L(u)$ versus the samples of the random variable u , the true high-fidelity values $Q^H(u)$, and the reconstructed high-fidelity values $Q_{r,N}^H(u)$ obtained from N high-fidelity samples. We report the results for different values of N , also highlighting the QoI’s generated by the high-fidelities selected by the algorithm (black boxes labeled *pivot HF* in the figure); the corresponding values of the stochastic parameter form the set γ_N . Focusing on the leftmost portion of the plots, where the QoI is constant, it is worth noting that the high-fidelity reconstruction misses the true high-fidelity value until a pivot HF is selected in that region (compare the plots for $N = 1$ or 5 to those for $N = 6$ or 12); furthermore, just one pivot HF is needed and selected in that region, showing the good adaptivity properties of the algorithm.

In Fig. 20 we report the errors versus N in computing the mean value of the QoI by averaging the reconstructed values over 1500 random samples. The horizontal line represents the “true” high-fidelity error




Figure 18: Case $d = 1$. Blue dots (left y axis): QoI versus fracture radius (10,000 samples with LF=1); red dots (right y axis): total number of fracture intersections in the network versus fracture radius.

obtained in the approximation of the mean value, i.e., the difference between the mean values of Q^H and Q^{ref} , where the latter quantity is obtained using the reference level $\ell = 7$; again, mean values are computed by a Monte Carlo method with the same 1500 random samples used above. Note that an exponential decay of the error is apparent even in the present setting, which is by far more complex than the one considered in Sect. 3; remarkably, the presence of many discontinuities in the QoI has no negative influence on the convergence histories.

In the selection phase, all the three strategies mentioned in Sec. 2 have been tested. We found essentially the same behaviour with the three approaches, at least far from rank deficiency. Indeed, the low-fidelity vectors tend to form poorly linearly independent sets, and numerical rank of the matrix $\mathbf{V} = (u_1^L, \dots, u_M^L)$ is revealed relatively soon. However, while approaching rank deficiency, few minimal differences in the set of selected high-fidelities were observed. All the results here reported have been obtained adopting the QR strategy, item a).

As far as the reconstruction phase is concerned, the solution of the linear systems (4) is needed to obtain the coefficients for building the reconstructions. We observed that a severe ill-conditioning in $\tilde{\mathbf{G}}$ appears as rank deficiency is approached in the selection phase. In order to improve as much as possible the quality of the numerical solution, we tested the Cholesky decomposition, the QR factorization, and the SVD decomposition of $\tilde{\mathbf{G}}$ for treating very small singular values. The best quality results were obtained with the QR decomposition, although only marginally better than the others. All the results here reported are obtained with this latter approach. The quick growth of ill-conditioning is illustrated in the bottom right panel of Fig. 20, which depicts the spectral condition number of the factor R of the QR decomposition of $\tilde{\mathbf{G}}$.

4.4. Case $d = 2$

We follow the same sequence of plots as in the $d = 1$ case. In Fig. 21 we report different views of the QoI versus the samples of the stochastic variables; similarly to what is done in Fig. 18, in the bottom right plot we report the number of fracture intersections versus the fracture radii. The non-smooth dependence of these quantities upon the two independent random variables is apparent. With the available number of sampling it would be rather difficult to identify a relatively regular manifold of discontinuity, separating smooth variations of the QoI. Going towards very high stochastic dimension, such detection would become in practice unfeasible. Fig. 22 reports the low-fidelity QoI's, the high-fidelity QoI's, and the reconstructed high-fidelity QoI's, along with the selected pivot HF; note that again the latter concentrate in the regions of larger variations of the QoI. Finally, Fig. 23 reports the errors versus N for increasing values of the HF level, along with the condition number of the factor R in the QR factorization of $\tilde{\mathbf{G}}$. The convergence

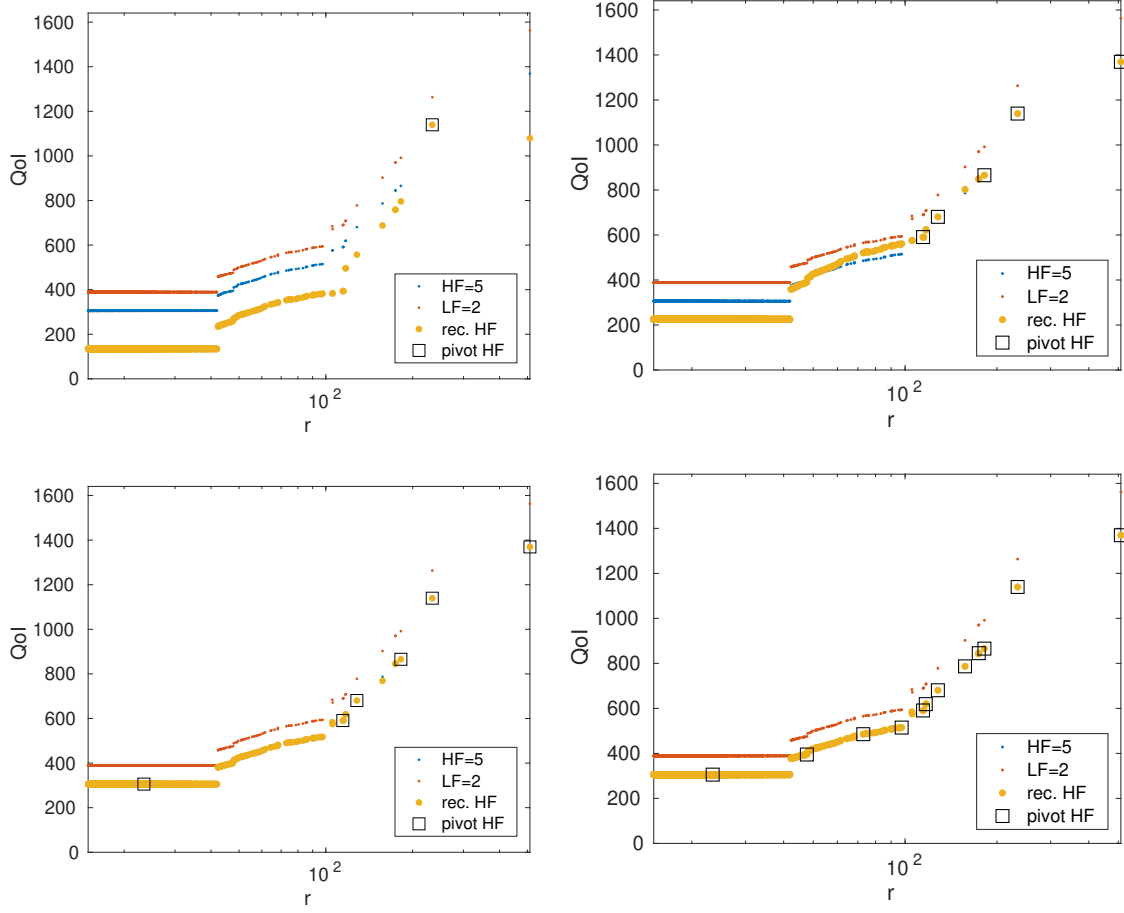


Figure 19: Case $d = 1$. QoI: LF with $\ell = 2$; HF with $\ell = 5$; reconstructed QoI and pivotal HF used for reconstruction. Left to right, top to bottom: $N = 1, 5, 6, 12$.

histories confirm the exponential decay of the error, although they exhibit a much more pronounced oscillatory behaviour than in the $d = 1$ case; several outliers appear as well.

4.5. Cases $d = 4$ and $d = 10$

For these cases we just report the errors and the conditioning of R (Figs. 24 and 25). These results confirm the trends emerged in the analysis of the previous cases. By comparing the cases $d = 2$, $d = 4$ and $d = 10$, one notes that the number of pivot high-fidelities needed to approach the target accuracy (about $N = 15, 20$) is rather independent of the stochastic dimension d and the high-fidelity level HF. However, while increasing HF error plots become more oscillatory.

An interesting observation which can be drawn by the analysis of these results is related to the numerical rank revealed during the selection phase. Indeed, the numerical rank increases with stochastic dimension, as can be noted by the fact that the points at which the plots, for a fixed LF level, get interrupted correspond to larger and larger values of N . Coherently, the rates at which the conditioning of the factor R deteriorates are reduced (compare the bottom right panels in Figs. 20, 23, 24 and 25). Recalling that the low-fidelities are computed on the exit fracture F , we think that the increased linear independency of the low-fidelities is related to the increased stochasticity of the network, which produces a larger variability on the solutions on F , resulting in an increased dimension of the space $U^L(\Gamma)$. This influence could be possibly further increased by increasing the number of stochastic fractures.




Figure 20: Case $d = 1$. Error versus N for increasing HF levels. Bottom-right: $K_2(R)$ versus N .

We end with some comments on the computational costs. We remark that in the application considered within this section, the cost for the construction and evaluation of model (3)-(5), once the needed samples are made available, is negligible, with respect to the total cost of computing the samples. Let us denote by \bar{c}_L and \bar{c}_H the average cost for a single low- and high-fidelity computation, respectively. The total cost C_B for the construction of the bi-fidelity model is therefore approximately $C_B = M\bar{c}_L + N\bar{c}_H$. The gain with respect to the total cost C_H which corresponds to the use of high-fidelity approximations only, which is $C_H = M\bar{c}_H$, clearly depends on the ratio between \bar{c}_L and \bar{c}_H . Let Δ denote the difference between the high- and low-fidelity levels. For $\Delta = 3$, as an average, we have $\bar{c}_H \simeq 15\bar{c}_L$; for $\Delta = 4$, we have $\bar{c}_H \simeq 30\bar{c}_L$, with no significant differences among the several values of d . The bifidelity procedure, considering, e.g., an average over $M = 1500$ samples, and a reconstruction made with $N = 30$ high-fidelities, yields an overall computational gain $C_H/C_B \simeq 12$ for $\Delta = 3$ and $C_H/C_B \simeq 18$ for $\Delta = 4$.

5. Conclusions

We have considered the problem of computing statistics of a quantity of interest (QoI) which depends in a non-smooth way upon the stochastic parameters, as it may occur e.g. in the presence of geometric uncertainties. To this end, we have applied a non-intrusive bifidelity strategy originally proposed in [1], which leverages the accuracy of a high-fidelity model and the efficiency of a low-fidelity model. The two models are herein obtained through the use of different levels of numerical resolution. Firstly, the method




Figure 21: Case $d = 2$. Top and bottom left plots: QoI obtained with 10000 samples with $LF=1$ (three different views) versus fracture radii r_1 and r_2 . Dot colouring proportional to z -value. Bottom right: total number of fracture intersections in the network versus fracture radii

has been tested on a set of synthetic cases, with analytically known outputs depending on certain physical variables and stochastic parameters; these are constructed to display either a smooth behaviour or a jump discontinuity in the stochastic parameter space. Upon varying the chosen high- and low-fidelity levels, in both cases we report a similar performance in terms of quality of the approximation; in particular, the algorithm is able to select the few parameters at which the high-fidelity computation has to be performed in strict accordance with the local behaviour of the QoI. The error decay is exponential with respect to the cardinality of the set of active high-fidelities, independently of the presence of singularities in the QoI. Next, we consider a real-life application originating in the framework of underground flow simulations, where the QoI represented by an outflow depends stochastically on certain parameters related to the geometry of the spatial domain. Even in this complex situation, the algorithm does not suffer from the presence of jumps and gradient discontinuities in the QoI, and produces approximations of its expected value that quickly approach the one computed by the high-fidelity model only, with a mild dependence on the stochastic dimension. The observed gain in computational cost is at least of one order of magnitude.

Acknowledgments

Research performed in the framework of the Italian MIUR Award “Dipartimento di Eccellenza 2018-2022” (E11G18000350001) to the Department of Mathematical Sciences, Politecnico di Torino.

The first and second authors are members of the INdAM-GNCS research group, which granted partial support to this research.

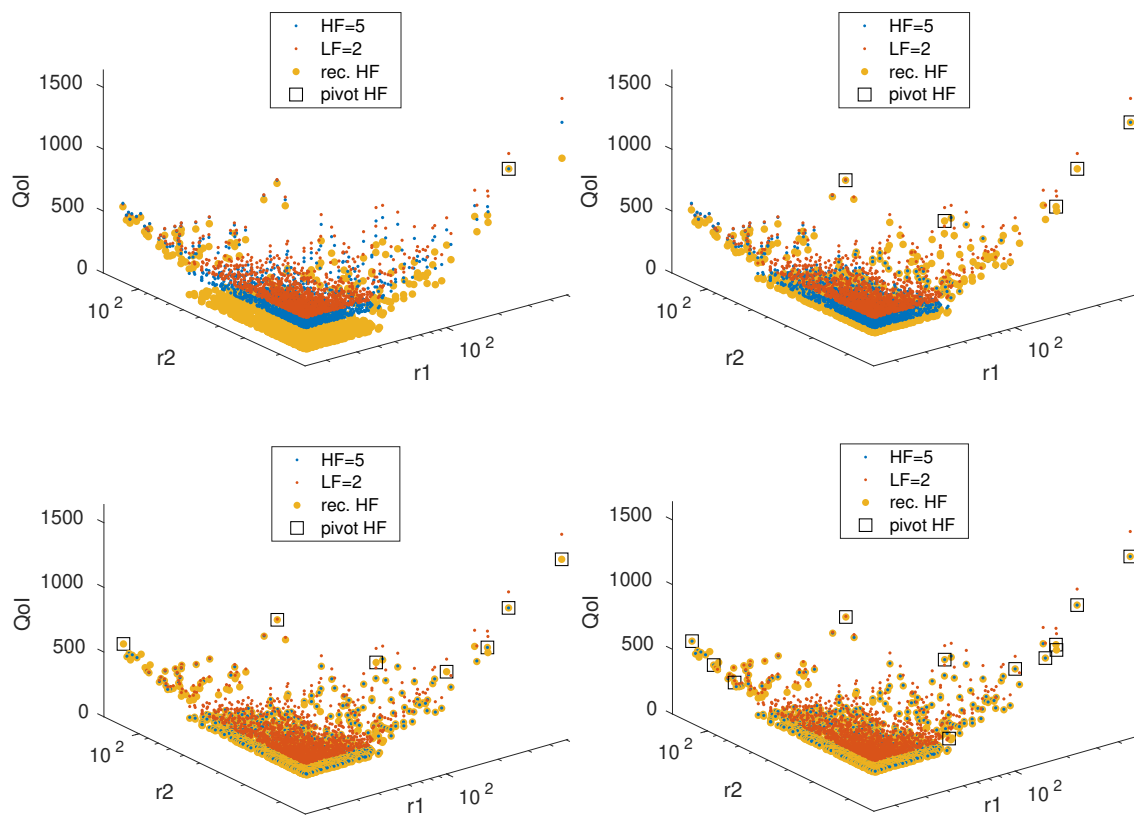


Figure 22: Case $d = 2$. QoI versus fracture radii r_1 and r_2 : LF with $\ell = 2$; HF with $\ell = 5$; reconstructed QoI and pivotal HF used for reconstruction. Left to right, top to bottom: $N = 1, 5, 7, 12$.

References

- [1] A. Narayan, C. Gittelsohn, D. Xiu, A stochastic collocation algorithm with multifidelity models, *SIAM J. Sci. Comput.* 36 (2) (2014) A495–A521.
- [2] S. Berrone, S. Pieraccini, S. Scialò, A PDE-constrained optimization formulation for discrete fracture network flows, *SIAM J. Sci. Comput.* 35 (2) (2013) B487–B510. doi:<http://dx.doi.org/10.1137/120865884>.
- [3] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods, in uncertainty propagation, inference, and optimization, *SIAM Review* 60 (3) (2018) 550–591.
- [4] X. Zhu, A. Narayan, D. Xiu, Computational aspects of stochastic collocation with multifidelity models, *SIAM/ASA J. Uncertainty Quantification* 2 (1) (2014) 444–463.
- [5] J. Hampton, H. R. Fairbanks, A. Narayan, A. Doostan, Practical error bounds for a non-intrusive bi-fidelity approach to parametric/stochastic model reduction, *Journal of Computational Physics* 368 (2018) 315 – 332. doi:<https://doi.org/10.1016/j.jcp.2018.04.015>.
URL <http://www.sciencedirect.com/science/article/pii/S0021999118302298>
- [6] X. Zhu, E. M. Linebarger, D. Xiu, Multi-fidelity stochastic collocation method for computation of statistical moments, *Journal of Computational Physics* 341 (2017) 386 – 396. doi:<https://doi.org/10.1016/j.jcp.2017.04.022>.
URL <http://www.sciencedirect.com/science/article/pii/S0021999117302930>
- [7] J. D. Jakeman, A. Narayan, D. Xiu, Minimal multi-element stochastic collocation for uncertainty quantification of discontinuous functions, *Journal of Computational Physics* 242 (2013) 790–808.
- [8] J. Y. R. Archibald, A. Gelb, Polynomial fitting for edge detection in irregularly sampled signals and images, *SIAM J. Numer. Anal.* 43 (2005) 259–279.
- [9] A. Gorodetsky, Y. Marzouk, Efficient localization of discontinuities in complex computational simulations, *SIAM Journal on Scientific Computing* 36 (6) (2014) A2584–A2610.
- [10] G. Zhang, C. G. Webster, M. Gunzburger, J. Burkardt, Hyperspherical sparse approximation techniques for high-dimensional discontinuity detection, *SIAM Review* 58 (3) (2016) 517–551.

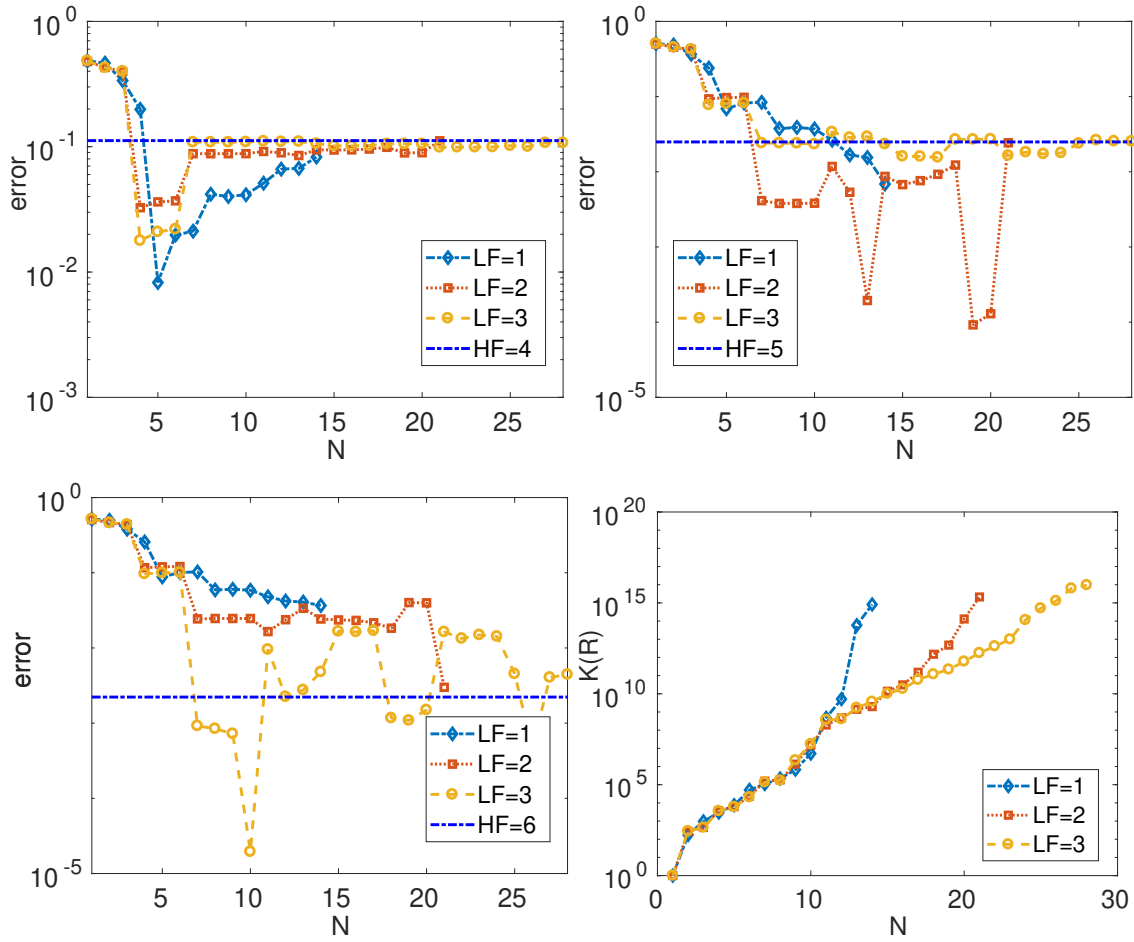


Figure 23: Case $d = 2$. Error versus N for increasing HF levels. Bottom-right: $K_2(R)$ versus M .

- [11] S. Berrone, C. Canuto, S. Pieraccini, S. Scialò, Uncertainty quantification in discrete fracture network models: Stochastic geometry, *Water Resources Research* 54 (2) (2018) 1338–1352.
- [12] S. Berrone, S. Pieraccini, S. Scialò, Towards effective flow simulations in realistic discrete fracture networks, *J. Comput. Phys.* 310 (2016) 181–201. doi:http://dx.doi.org/10.1016/j.jcp.2016.01.009.
- [13] E. E. Tyrtushnikov, Incomplete cross approximation in the mosaic-skeleton method, *Computing* 64 (4) (2000) 367–380.
- [14] B. Berkowitz, Characterizing flow and transport in fractured geological media: a review, *Adv. Water Resour.* 25 (8) (2002) 861–884.
- [15] K. T. B. MacQuarrie, K. U. Mayer, Reactive transport modeling in fractured rock: A state of the science review, *Earth-Science Reviews* 72 (2005) 189–227.
- [16] S. P. Neuman, Trends, prospects and challenges in quantifying flow and transport through fractured rocks, *Hydrogeol. J.* 13 (1) (2005) 124–147.
- [17] J. D. Hyman, G. Aldrich, H. Viswanathan, N. Makedonska, S. Karra, Fracture size and transmissivity correlations: Implications for transport simulations in sparse three-dimensional discrete fracture networks following a truncated power law distribution of fracture size, *Water Resources Research*.
- [18] Svensk Kärnbränslehantering AB, Data report for the safety assessment, SR-site, Tech. Rep. TR-10-52, SKB, Stockholm, Sweden (2010).
- [19] D. Xiu, *Numerical methods for stochastic computations*, Princeton University Press, 2010.
- [20] S. Berrone, S. Pieraccini, S. Scialò, F. Vicini, A parallel solver for large scale DFN flow simulations, *SIAM J. Sci. Comput.* 37 (3) (2015) C285–C306. doi:http://dx.doi.org/10.1137/140984014.

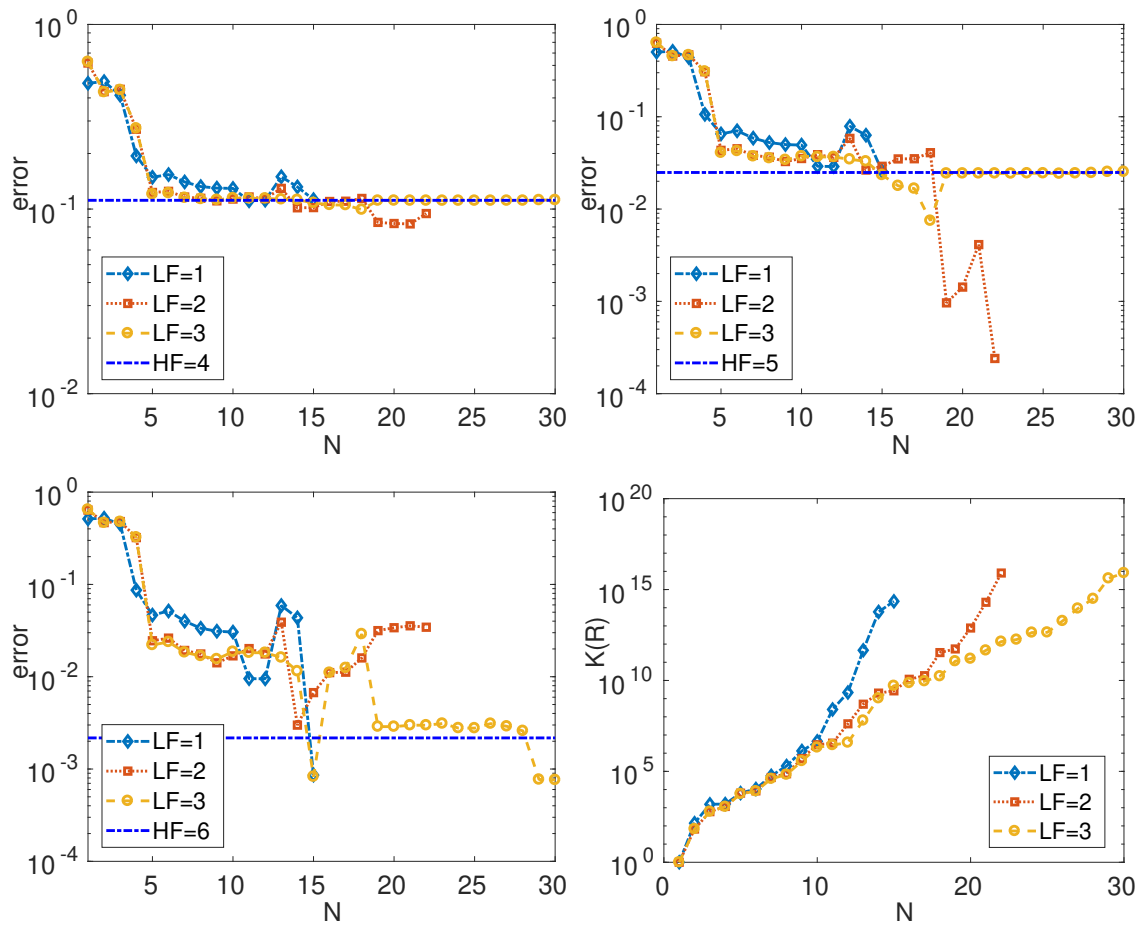


Figure 24: Case $d = 4$. Error versus N for increasing HF levels. Bottom-right: $K_2(R)$ versus M .




Figure 25: Case $d = 10$. Error versus N for increasing HF levels. Bottom-right: $K_2(R)$ versus M .