



Technische Universität München
Lehrstuhl für Wissenschaftliches Rechnen

Context-aware Model Hierarchies for Higher-dimensional Uncertainty Quantification

Ionuț-Gabriel Farcaș

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität
München zur Erlangung des Akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzende(r): Prof. Dr. Francisco Javier Esparza Estaun

Prüfer der Dissertation: 1. Prof. Dr. rer. nat. Hans-Joachim Bungartz

2. Prof. Karen E. Willcox, Ph.D.
The University of Texas at Austin, USA

Die Dissertation wurde am 06.03.2020 bei der Technischen Universität München
eingereicht und durch die Fakultät für Informatik am 18.05.2020 angenommen.

I	Introduction and Background	1
1	Introduction	3
1.1	Opening	3
1.2	Computational challenges in uncertainty quantification	4
1.3	Our contributions in brief	9
1.4	Thesis outline	11
2	Theoretical background	15
2.1	Setup	15
2.2	Assumptions	16
2.3	Outer-loop scenarios	17
2.4	Uncertainty propagation and Bayesian inference: an overview	17
2.4.1	Uncertainty propagation	18
2.4.2	Bayesian inference for parameter estimation	19
2.5	Approximation with dimension-adaptive sparse grids	21
2.5.1	Approximation operators	24
2.5.2	Weighted (L)-Leja sequences	28
2.5.3	Subspace or dimension-adaptivity	30
2.6	Approximation with spatially adaptive sparse grids	34
2.6.1	Interpolation with hierarchical bases with local support	34
2.6.2	Local or spatial adaptivity	37
2.7	Multilevel estimation of quantities of interest	38
2.8	Multifidelity sampling	41
2.8.1	Standard Monte Carlo sampling	42
2.8.2	Variance reduction via control variates	44
2.8.3	Multifidelity Monte Carlo sampling	45
2.9	Sobol' indices for global sensitivity analysis	48
2.10	Model hierarchies in uncertainty quantification	51
3	Applications background	55
3.1	Fluid-structure interaction problems	55
3.1.1	Overview	56
3.1.2	Modelling fluid-structure interaction	56
3.1.3	The fluid-structure interaction simulation codes	59
3.1.4	The multi-challenge	60
3.2	Plasma microturbulence analysis	60
3.2.1	Overview	61
3.2.2	The gyrokinetic approach and its numerical simulation	61
3.2.3	The plasma microturbulence code GENE	63
II	Context-aware sparse grid and multilevel approximations for uncertainty propagation	67
4	Sensitivity-driven dimension-adaptive sparse grid approximations	69
4.1	Uncertainty in plasma microturbulence analysis	70

CONTENTS

4.2	Sensitivity-driven dimension-adaptive sparse grid algorithm	71
4.2.1	Setup	73
4.2.2	Unnormalized sensitivity indices	73
4.2.3	Sensitivity scores	76
4.2.4	Maximum sensitivity score	79
4.2.5	Sensitivity-driven dimension-adaptive sparse grid algorithm	80
4.2.6	Computational cost	83
4.2.7	Illustrative examples	83
4.3	Numerical results: computed dominant amplitude eigenmode	85
4.3.1	Modified Cyclone Base Case	86
4.3.2	Realistic test case: overview	92
4.3.3	Realistic test case with three stochastic parameters	92
4.3.4	Realistic test case with 12 stochastic inputs	98
4.4	Numerical results: truncating particular microinstabilities	101
4.4.1	Realistic test case with three stochastic parameters	102
5	Multilevel adaptive spectral projections with online dimensionality reduction	111
5.1	Quantifying uncertainty in fluid-structure interaction simulations	112
5.2	Stochastic dimensionality reduction	113
5.3	Multilevel adaptive spectral projection with online dimensionality reduction	114
5.3.1	Spatially adaptive sparse pseudo-spectral approximations	118
5.3.2	Multilevel spectral projections	122
5.3.3	Online stochastic dimensionality reduction	123
5.3.4	Computational cost	128
5.3.5	Termination of the multilevel algorithm	128
5.4	Numerical results	129
5.4.1	Damped oscillator	130
5.4.2	Fluid flow over an elastic vertical flap	137
III	Context-aware multifidelity Monte Carlo sampling for uncertainty propagation	145
6	Context-aware model reduction for multifidelity Monte Carlo sampling	147
6.1	Reduced order models for uncertainty propagation	148
6.2	Context-aware multifidelity Monte Carlo sampling	149
6.2.1	Setup	150
6.2.2	Preliminaries	151
6.2.3	Context-aware multifidelity Monte Carlo sampling algorithm	155
6.3	Numerical results	167
6.3.1	Heat conduction in a two-dimensional domain	168
6.3.2	ASDEX Upgrade experiment	178
IV	Context-aware multilevel decompositions for Bayesian inference	185
7	Multilevel adaptive sparse grid Leja approximations for Bayesian inference	187
7.1	Posterior-focused surrogate models in Bayesian inversion	188
7.2	Multilevel adaptive sparse Leja algorithm	191
7.2.1	Multilevel setup	193

7.2.2	Level $\ell(1)$	194
7.2.3	Levels $\ell(j)$ with $j \geq 2$	195
7.2.4	Computational cost	200
7.2.5	Termination of the multilevel algorithm	202
7.3	Adaptive strategies for surrogate constructions	203
7.3.1	Standard dimension-adaptive interpolation and quadrature	203
7.3.2	Directional variance dimension-adaptive sparse interpolation	203
7.3.3	Illustrative examples	207
7.4	Numerical results	209
7.4.1	Simple quadrature showcase	210
7.4.2	Source inversion with one source in a 2D spatial domain	213
7.4.3	Source inversion with two sources in a 2D spatial domain	217
7.4.4	Higher-dimensional problem in a 3D spatial domain	222
V	Conclusion	229
8	Summary and conclusion	231
8.1	What was achieved	231
8.2	Outlook	233

Acknowledgements

First of all, I would like to express my gratitude towards my supervisor, Prof. Hans-Joachim Bungartz. Thank you for your guidance, support, words of encouragement, for giving me the freedom to pursue my ideas, and to make my own mistakes! During my doctoral studies, I was fortunate to make not one, not two, but three research stays. I would therefore like to thank my three hosts. Thank you, Prof. Peherstorfer for hosting me at the University of Wisconsin at Madison, thank you, all the staff from the Institute for Pure and Applied Mathematics in Los Angeles, and thank you Prof. Marzouk for hosting me at the Massachusetts Institute of Technology. In addition, I want to thank my mentors, Dr. Tobias Neckel, Dr. Tobias Görler and Prof. Thomas Wick. Dr. Neckel, thank you for introducing me to Uncertainty Quantification, for guiding me, for always making time to discuss with me, and for your words of encouragement! Dr. Görler, thank you for introducing me to Plasma Physics and for having the patience to allow me to learn some of this fascinating new field. Prof. Wick, thank you for all your advice and for hosting my first invited talk in your group! Moreover, I would like to thank Prof. Karen Willcox for agreeing to be the second reviewer of this thesis. Hans, Karen, it is an honour and a privilege to have you as my reviewers! I also want to express my gratitude towards Prof. Willcox and Prof. Frank Jenko for agreeing to guide the next step of my academic development. I look forward to working with and learning from you.

This work can be characterized by many words, but an important one is interdisciplinarity. In the papers that I published or submitted during my doctoral studies, I had eight co-authors from the fields of Informatics, Mathematics and Plasma Physics! Moreover, the list of past, present and future collaborators extends even further. Written in no particular order, I want to express my gratitude to Hans-Joachim Bungartz, Tobias Neckel, Benjamin Uekermann, Florian Künzner, Friedrich Menhorn, Ivana Jovanovic and Paul Cristian Sârbu from the chair of Scientific Computing in Computer Science at TUM, to Frank Jenko and Tobias Görler from the Max Planck Institute for Plasma Physics in Garching, to Elisabeth Ullmann and (now, Dr.) Jonas Latz from Mathematics at TUM, to Daniel Straub, Iason Papaioannou, Wolfgang Betz, Felipe Uribe and Elizabeth Bismut from the Engineering Risk Analysis Group at TUM, to Benjamin Peherstorfer from the Courant Institute of Mathematical Sciences at the New York University and to Youssef Marzouk from the Massachusetts Institute of Technology. Also, I want to thank Frank Jenko, Tobias Görler, Jonas Latz, Friedrich Menhorn and Paul Cristian Sârbu for proofreading some of the chapters.

I would like to express my deepest gratitude to the German Academic Exchange Service (DAAD) for offering me the financial support not only during the first three years of my doctoral studies, but also during my Master's studies in Germany! I also want to acknowledge the financial support provided by TUM International Graduate School of Science and Engineering (IGSSE). Besides many other things, two of my three research stays would not have been possible without the financial support of IGSSE!

Dedicating oneself to a goal requires sacrifice, not only from the person directly involved, but also from other people. Thus, I want to thank my family and friends for their support and understanding in the past four years.

Lastly, and mostly, I would like to thank my wife, Lavinia, my black swan! Thank you for your support and patience throughout my PhD, and especially during the months leading up to the completion of this thesis. She is my most trustworthy proofreader: every typo or misshapen paragraph that you do not find in this work (including this paragraph!), is due to her.

Abstract

Quantifying uncertainty in predictive simulations for real-world problems is of paramount importance, and far from trivial, mainly due to the large number of stochastic parameters and significant computational requirements. Prominent approaches to overcome, at least to a certain extent, these challenges are: (i) dimension-adaptive sparse grid approximations, (ii) standard multilevel decompositions and (iii) static multifidelity sampling methods. However, dimension-adaptivity is typically based on global refinement indicators that do not distinguish between individual stochastic inputs nor fully exploit the information about their coupling. In addition, in Bayesian inference, sparse grid approximations are typically performed with respect to the prior measure, hence, when the observation data are informative, many grid points are placed in the region of low posterior probability. Furthermore, standard multilevel decompositions based on telescoping sums can still become prohibitive when the number of uncertain inputs or the evaluation cost of the forward model increase. Finally, in static multifidelity sampling methods, the surrogates are assumed to be readily available, and thus they cannot be improved or refined, which is generally the case for data-driven reduced models.

To this end, we formulate four novel context-aware algorithms based on model hierarchies arising in sparse grid approximations, multilevel decompositions and multifidelity sampling, aimed to enable an efficient quantification of uncertainty in complex, computationally expensive problems. All four approaches treat the forward model as a black box. The first algorithm is based on dimension-adaptive sparse grid approximations, the novelty being a refinement indicator computed using sensitivity information, with the goal of exploiting the anisotropic coupling of the stochastic inputs and the lower intrinsic stochastic dimensionality. The lower intrinsic dimensionality is exploited in the second proposed approach as well, in which we enhance standard multilevel decompositions with a novel online dimensionality reduction step, performed at no additional computational cost. In the third and last contribution for uncertainty propagation, we formulate a new multifidelity sampling algorithm in which we construct context-aware reduced models such that there is a trade off between improving the accuracy of the surrogates and sampling them. As representative examples of complex, real-world applications, we test these three methods in fluid-structure interaction and plasma microinstability simulations, in which we consider stochastic scenarios with dimensionality between two and 12. Finally, we formulate a novel deterministic, multilevel, adaptive approach based on sparse grid approximations to address the computational challenges appearing in Bayesian inference. We perform a sequential update of the reference information which allows to place sparse grid points in the region of high posterior probability and thus reduce the cost of finding accurate surrogates in Bayesian inverse problems. As an example of a more complex application, we test the proposed approach in an inverse problem with an elliptic forward model defined on a three-dimensional spatial domain, in which we infer eight stochastic parameters. Our results clearly show that the four proposed algorithms are more efficient than standard approaches and that they are able to cope with the challenges of quantifying uncertainty in higher-dimensional, complex problems.

Zusammenfassung

Die Quantifizierung der Unsicherheit in prädiktiven Simulationen für reale Probleme ist von größter Wichtigkeit und bei weitem nicht trivial, vor allem wegen der großen Anzahl stochastischer Parameter und der erheblichen Rechenanforderungen. Prominente Ansätze, um diese Herausforderungen zumindest teilweise zu bewältigen, sind: (i) dimensionsadaptive dünne Gitterapproximationen, (ii) standardmäßige Mehrebenen-Methoden und (iii) statische Multifidelitäts-Methoden. Die Dimensionsadaptivität basieren jedoch in der Regel auf globalen Verfeinerungsindikatoren, die nicht zwischen einzelnen stochastischen Parametern unterscheiden und die Informationen über ihre Interaktion nicht vollständig ausnutzen. Darüber hinaus werden bei der Bayes'schen Inferenz dünne Gitterapproximationen in Bezug auf die vorherige Messung durchgeführt. Wenn daher die Beobachtungsdaten informativ sind, werden viele Gitterpunkte in den Bereich niedriger posteriorer Wahrscheinlichkeit gelegt. Darüber hinaus können standardmäßige Mehrebenen-Methoden auf der Grundlage von Teleskopsummen immer noch untragbar werden, wenn die Anzahl unsicherer Eingabeparameter oder die Bewertungskosten des Vorwärtsmodells zunehmen. Schließlich wird bei statischen multifidelität-Methoden angenommen, dass die Surrogate leicht verfügbar sind und daher nicht verbessert oder verfeinert werden können, was im Allgemeinen bei datengetriebenen reduzierten Modellen der Fall ist.

Zu diesem Zweck formulieren wir vier neuartige kontextsensitive Algorithmen auf der Grundlage von Modellhierarchien, die bei der Näherung mit dünnen Gittern, bei Mehrebenen-Methoden und bei Multifidelitäts-Methoden entstehen, um eine effiziente Quantifizierung der Unsicherheit bei komplexen, rechenintensiven Problemen zu ermöglichen. Alle vier Ansätze behandeln das Modell als Black-Box. Der erste Algorithmus basiert auf dimensionsadaptiven Näherungen mit dünnem Gitter, wobei die Neuerung ein Verfeinerungsindikator ist, der unter Verwendung von Sensitivitätsinformationen berechnet wird, mit dem Ziel, die anisotrope Interaktion der stochastischen Eingabeparameter und die niedrigere intrinsische stochastische Dimension auszunutzen. Die niedrigere intrinsische stochastische Dimensionalität wird auch im zweiten vorgeschlagenen Algorithmus ausgenutzt, bei dem wir standardmäßige Mehrebenen-Methoden mit einem neuen Online-Schritt zur Verringerung der stochastischen Dimensionalität verbessern, der ohne zusätzlichen Rechenaufwand durchgeführt wird. Im dritten und letzten Beitrag zur Vorwärtsunsicherheitsquantifizierung formulieren wir einen neuen multifidelität-Algorithmus, in dem wir kontextsensitive reduzierte Modelle konstruieren, das es einen Kompromiss zwischen der Verbesserung der Genauigkeit der Surrogate und deren Abtastung gibt. Als repräsentative Beispiele für komplexe reale Anwendungen testen wir diese drei Methoden in Fluid-Struktur-Wechselwirkungs- und Plasma-Mikroinstabilitätssimulationen, in denen wir stochastische Szenarien mit Dimensionen zwischen zwei und zwölf betrachten. Schließlich formulieren wir einen neuartigen deterministischen, mehrebenen, adaptiven Ansatz, der auf dünnen Gitterapproximationen basiert, um die in der Bayes'schen Inferenz auftretenden rechnerischen Herausforderungen zu verringern. Wir führen eine sequentielle Aktualisierung der Referenzinformationen durch, die es ermöglicht, dünne Gitterpunkte im Bereich hoher posteriorer Wahrscheinlichkeit zu platzieren und damit die Kosten für das Auffinden genauer reduzierter Modelle in Bayes'schen inversen Problemen zu senken. Als Beispiel für eine komplexere Anwendung testen wir den vorgeschlagenen Ansatz in einem inversen Problem mit einem elliptischen Modell, das in einer dreidimensionalen räumlichen Domäne definiert ist, in der wir acht stochastische Parameter ableiten. Unsere Ergebnisse zeigen deutlich, dass die vier vorgeschlagenen Algorithmen effizienter als Standardansätze sind und die Herausforderungen der Quantifizierung der Unsicherheit in höherdimensionalen, komplexen Problemen bewältigen können.

PART I

INTRODUCTION AND BACKGROUND

"Distinguishing the signal from the noise requires both scientific knowledge and self-knowledge: the serenity to accept the things we cannot predict, the courage to predict the things we can, and the wisdom to know the difference."

Nate Silver

"...as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns - the ones we don't know we don't know."

Donald Rumsfeld

1

Introduction

1.1 Opening

Science is fundamentally the process of understanding the world around us, from subatomic particles and micro-organisms, to galaxies and the entire visible universe, with the ultimate goal of using this knowledge for reliable predictions to the benefit of mankind. The scientific process has been revolutionized through computing: numerical simulations in which mathematical models are used to derive data, i.e., the *deductive* approach, allow us to investigate phenomena much more complex than theoretical analysis alone permits. Moreover, data-driven computing, i.e., the *inductive* approach in which data are used to derive models allows us to more effectively explore, understand, and use data resulting from experiments, observations, and simulations. Despite the progress of the predictive capacity of inductive and deductive computing, several aspects of this process still need to be accounted for to ensure that mathematical models gives useful and reliable information (see the first quote above).

In recent years, it has become well established that a proper quantification of uncertainty is an important step towards predictive numerical simulations of real-world phenomena. Whether stemming from measurement errors, incomplete knowledge or inherent variabilities, uncertainty is intrinsic to most real-world problems and it needs to be accounted for *ab initio*. To give the reader a first idea about what quantifying uncertainty in numerical simulations means, let us take a closer look at the second quote from above. In physics, for example, there are universal constants such as the speed of light, c , or the gravitational acceleration, g , which have known, deterministic values. These constants are examples of *known knowns*, which, in general, refer to anything that describes the phenomena under consideration which is known with certainty. However, measurement errors, incomplete knowledge or inherent variabilities lead to *known unknowns*, i.e., descriptors that are known, but which no longer have fixed values. Treating and incorporating known unknowns in computer simulations is the goal of uncertainty quantification (UQ), which is also the main theme of this thesis. Finally, looking at the last part of the second quote, *unknown unknowns* are generally more difficult to deal with. Examples include scientific theories yet to be discovered or specific scenarios for which current theories fail, both leading to new knowledge, known unknowns or known knowns.

Our focus in this work is on UQ in complex, real-world applications such as multi-physics or plasma microturbulence simulations. Such applications are very challenging both mathematically and computationally, and our main goal is to design novel algorithms that enable an efficient UQ in these problems. As we will see throughout this work, the leitmotif of our algorithms is that even though most existing UQ approaches become infeasible in problems with a large number of uncertain parameters and/or computationally expensive simulation codes due to, e.g., the so-called "*curse of dimensionality*", in practical applications, we typically benefit from the "blessing of low rank" or "the blessing of sparsity", or more generally, "the blessing of structure", which makes UQ computationally tractable.

We briefly discussed why an *ab initio* treatment of uncertainty is needed to obtain reliable numerical simulations. Next, we provide the reader with an introduction into the two UQ branches of interest in this work, uncertainty propagation and Bayesian inference, we briefly overview some of the existing techniques and discuss the computational challenges arising in UQ, which serve as motivation for the novel approaches proposed in this thesis.

1.2 Computational challenges in uncertainty quantification

In uncertainty propagation, we are given mathematical models of real-world phenomena and a set of uncertain input parameters, the task being to quantify and reduce the uncertainty in outputs or observables of this model. Thus, the simulation is performed in the direction input-model-output. The uncertain inputs are generally modelled probabilistically via a (multivariate) probability measure, which is assumed to be available a priori. This input probability measure is then sampled, e.g., using pseudo-random numbers or deterministic interpolation/-collocation/quadrature points.

Remark: In this work, we use the term “sample” to refer broadly to points that fall within the support of the underlying probability measure. This includes the “traditional” samples, i.e., (pseudo-) random numbers, but also deterministic, e.g., collocation or quadrature points.

Each sample gives rise to an input parameter set for which the underlying model is solved. Thus, once the input measure is sampled, the uncertainty propagation problem becomes *deterministic*. The ensemble of model evaluations is afterwards used to compute quantities such as expectation, variance, failure probability, or sensitivity indices.

Remark: If the forward model is a garbage in - garbage out mapping, then UQ algorithms will yield nonsensical results; see, e.g., [174] for a more detailed discussion. The goal of forward UQ is to quantify and reduce uncertainties in physical phenomena for which the associated model can be trusted to yield meaningful results, and not to ascertain the validity of the underlying model. There are, of course, stochastic techniques for model validation, calibration or selection, mainly based on the Bayesian framework (see, for example, Chapters 6 – 8 from [164]), but these are outside of the scope of uncertainty propagation.

A prominent field of study in which uncertainty propagation is relevant is computational fluid dynamics. As an example, [197] studied the effects of input uncertainty in the free stream Mach number, M_∞ , angle of attack, α , and thickness-to-cord ratio, t/c , into the lift, drag and pitching moment coefficients in the transonic RAE 2822 airfoil test case. Another application, also used in one of the contribution chapters of this thesis is fluid-structure interaction, i.e., problems in which one or more structures interact with an internal or surrounding fluid flow. Representative examples range from human (patho)physiology, e.g., the study of blood flow, calcification, and aneurysms, see e.g. [113], to aeronautics. In general, these problems feature several physical parameters which suffer from uncertainty as they typically stem from measurements: both fluid and structure density, the viscosity of the fluid, or the structural elasticity and plasticity parameters, possibly different in every spatial dimension. This therefore requires a proper quantification of uncertainty in numerical simulations. Quantifying uncertainty prior to numerical simulations is also needed in the simulation of plasma microturbulence, which we employ in two of our four contribution chapters. In fusion plasmas, the microturbulence is driven by the free energy provided by the (unavoidably) steep temperature and density gradients of the particle species in the plasma. Unfortunately, these gradients and other parameters influencing the underlying microinstabilities are affected by uncertainties due to, e.g., measurement noise.

We next overview three classes of methods for uncertainty propagation: (i) intrusive approaches, (ii) non-intrusive deterministic approaches and (iii) non-intrusive methods based on sampling. **Intrusive approaches** were historically the first to be developed in the context of UQ and they were made popular in the book by Ghanem and Spanos [73]. Generally speaking, intrusive methods aim to decrease the computational cost of uncertainty propagation by finding *reduced models* in which the forward solution is *projected* onto a lower dimensional space/manifold. In their standard form, these methods require that the underlying model operators are available such that the model and the associated simulation code are modified.

The most prominent intrusive approach in UQ is the *stochastic Galerkin* projection, first proposed in [73]; see also [29, 38, 73, 163, 202] for more details. Therein, ideas from finite element (FE) analysis were borrowed to compute the coefficients of spectral or polynomial chaos (PC) expansions. Spectral or PC expansions represent a random variable as a series of orthogonal polynomials that span the underlying probabilistic space. The idea emerged in 1938 with Wiener’s seminal work on homogeneous chaos [192]. In [24], Cameron and Martin proved optimal convergence of the Hermite PC in terms of Gaussian random variables. Furthermore, in 2002 Xiu and Karniadakis [201] extended Hermite PC to generalized PC (gPC) which allows optimal approximations in terms of orthogonal polynomial bases from the Askey scheme. Later on gPC was further extended to arbitrary orthogonal polynomial bases (see, e.g., [194]). Some other methods for projection-based model reduction for uncertainty propagation include the reduced basis approach [25], empirical interpolation or combinations of reduced bases and sparse grids (see [29] for an overview). Since intrusive methods require the modification of the simulation code, their practicality deteriorates as the complexity of either of the two increases. For example, the application scenarios considered in this work, fluid-structure interaction simulations and plasma microturbulence analysis, have complex mathematical models and simulation codes. For these types of applications, intrusive approaches are impractical. Thus, in the remainder of this thesis, we focus exclusively on approaches that require only independent evaluations of the given simulation code, i.e., non-intrusive approaches.

Non-intrusive methods treat the underlying simulation code as a *black box*, in the sense that the model operators and the simulation code do not need to be modified. Furthermore, simulations associated to different *samples* of the uncertain input can be performed *independently* of each other, in *parallel*, with little or no communication overhead. In addition, since most existing simulation codes for complex applications are already parallelized, non-intrusive methods allow for multiple layers of parallelism: each (individual) simulation can be performed in parallel on P_1 processes (parallel layer 1), groups of P_2 simulations corresponding to P_2 input samples can be performed on $P_2 \cdot P_1$ processes in total (parallel layer 2) and so on. This clearly makes non-intrusive methods for uncertainty propagation computationally more advantageous than standard intrusive approaches.

Deterministic non-intrusive methods employ samples with fixed coordinates in the input probability space. Examples of deterministic samples include interpolation or quadrature points, such as Gaussian, Clenshaw-Curtis or Leja points. One of the most prominent non-intrusive method is the *pseudo-spectral approach*, see, e.g., [199], in which quadrature is employed to assess the coefficients of gPC approximations. An alternative approximation scheme based on Lagrange interpolation, called *stochastic collocation*, was proposed in [5]. The samples used in standard pseudo-spectral or stochastic collocation approaches stem typically from *full-tensor grids*: if N samples are used in one dimension, then standard tensorization leads to $\underbrace{N \cdot N \cdot \dots \cdot N}_{d_{\text{sto}} \text{ times}} = N^{d_{\text{sto}}}$ samples, where d_{sto} denotes the number of stochastic parameters. This number becomes prohibitive even for relatively small N or d_{sto} . For example, if $d_{\text{sto}} = 6$ and $N = 10$ points per dimension are used, then $N^{d_{\text{sto}}} = 10^6$, which means that we need to perform

one million simulations! This exponential increase in the number of degrees of freedom with the dimensionality in standard full-tensor approaches is a manifestation of the *curse of dimensionality* (see, e.g., [20]), which prevents most standard deterministic non-intrusive methods to be applicable for problems with a large number of uncertain parameters or computationally expensive simulation codes.

Overcoming or delaying the curse of dimensionality is one of the greatest challenges in UQ in particular and in scientific computing in general. Approximations based on sparse grids [20] have been well established as suitable counter-measures. Loosely speaking, sparse grid approximations reduce computational effort by weakening the assumed coupling between the input dimensions (see, e.g., [35]). In this way, the exponential increase of the number of degrees of freedom with the dimensionality decreases to a subexponential, e.g., logarithmic dependency in the hierarchical formulation [20,206]. Starting with works such as [139,200], standard sparse grid approximations have been extensively used in uncertainty propagation. However, although sparse grids succeed in delaying the curse of dimensionality, in their standard form they can still be computationally prohibitive especially when the given simulation code is expensive.

In recent years, significant efforts have been put into designing enhanced approximation strategies for computationally expensive UQ problems. Spatial adaptivity was used to enhance sparse grid stochastic collocation formulated in terms of hierarchical basis functions with local support, for example, in [62,119]. To accelerate stochastic collocation formulated in terms of global polynomials, a previously assembled lower fidelity interpolant was used in [65] to predict the solution of the stochastic model at each collocation point. A multilevel stochastic collocation approach based on hierarchies given by anisotropic combinations of physical domain discretization grids and sparse grids for stochastic space sampling was proposed in [176]. Therein, the hierarchy is given by two parameters. This was extended in [85] to a multiindex stochastic collocation method in which the hierarchy is given by more than two parameters, e.g., time step discretization, physical space discretization parameter and sparse grid level for stochastic space sampling. Hierarchies given by models with different fidelities were used in a multifidelity stochastic collocation approach in [133,137]. Strategies to improve standard sparse grid pseudo-spectral projections were investigated as well. For example, in [37] the equivalence between stochastic collocation and spectral projection constructed using Gaussian nodes was exploited to formulate sparse pseudo-spectral projections free of internal aliasing errors. This approach was generalized to arbitrary point sets and was rigorously analysed in [35] where several adaptive strategies have also been proposed. Furthermore, in [193], an enhanced strategy for adaptive pseudo-spectral projection was proposed for problems with two sets of inputs: deterministic design parameters and stochastic input parameters.

In **non-intrusive sampling approaches** for uncertainty propagation the input samples are (pseudo-)random numbers. By far the most popular sampling technique is the Monte Carlo (MC) algorithm [130]. In its standard form, however, MC is not applicable in higher-dimensional, computationally expensive uncertainty propagation problems because of its slow convergence. Various efforts have been made to accelerate standard MC sampling. Heinrich in [91] and Giles in [74] proposed the multilevel Monte Carlo (MLMC) method, which is nowadays one of the most popular techniques for uncertainty propagation (see also, e.g., [6,33,34] and the references therein). An adaptive version of MLMC was proposed in [42], where the levels giving rise to the model hierarchy were treated as *continuous* random variables. For a detailed review of MLMC and its variants, we refer the reader to [75]. Multiindex Monte Carlo, the stochastic variant of multiindex stochastic collocation of [85], was proposed in [86] to extend MLMC via model hierarchies determined by more than two indices. The idea of using model hierarchies to accelerate standard MC was generalized in [145] to the multifidelity MC (MFMC) sampling algorithm in which surrogates or *low-fidelity* models e.g., data-driven or reduced physics models,

1.2. COMPUTATIONAL CHALLENGES IN UNCERTAINTY QUANTIFICATION

are used together with the high-fidelity models to reduce the variance and computational cost of standard MC sampling. MFMC was analysed in [142]. MFMC sampling was also employed in other settings, for example, in rare events simulation [106, 143], or in the estimation of variance and Sobol' indices for global sensitivity analysis [148]. A comprehensive review of multifidelity methods is available in [146]. Another approach to reduce the variance of standard MC estimators, based on regression, was proposed in [132].

Remark: The majority of approaches for uncertainty propagation are in the realm of intrusive or non-intrusive methods based on projection, deterministic approximation and sampling, or pseudo-random sampling. Nevertheless, there exist alternative methodologies as well. For example, in [104], probabilistic cellular automata were used for uncertainty propagation. However, in this work, we concentrate only on non-intrusive approaches for uncertainty propagation, therefore we do not review further alternative techniques.

In uncertainty propagation, the input probability measure is assumed to be available a priori, the task being to map the parameter space to the observation space via a given forward operator. In some situations, input parameters cannot be observed directly, but are nevertheless required for model-based predictions. A standard way to deal with these cases is to estimate these parameters from indirect observations, a procedure known as an *inverse problem*. Examples include computer tomography in medical diagnosis, in which, given measurements of X-ray intensity going through tissues and a decay model of X-rays, the tissue density is inferred. Another example is wave scattering in geosciences in which the unknown wave speed is inferred from seismogram data and the wave equation. In ice dynamics, again relevant in geosciences for the prediction of, e.g., future sea level rise, the basal boundary friction condition is inferred indirectly from surface velocity data and a flow model given as the solution to incompressible nonlinear Stokes equations (see [94]). Inference is relevant in other fields as well, such as visual perception. Human visual perception is an inherently uncertain process. Despite our subjective experience, the way we perceive our environments is not a one-to-one mapping of “raw” sensory input. The same sensory information, e.g., colour, is more reliable in some contexts, for example, during daylight than in others, such as in the dark. Despite this, our brains have developed into so-called “Bayes optimal” perceptual systems and are able to account for information uncertainties when transforming our sensory input into coherent perceptions (see [103]).

In most parameter estimation problems, the given measurements are noisy and their number is insufficient to identify a unique associated parameter value. Moreover, even if a unique deterministic value can be obtained, this solution does not incorporate a quantification of the parameters' uncertainty due to, e.g., measurement or model errors. To obtain a unique solution¹ which also incorporates a quantification of the parameters uncertainty, we adopt a probabilistic formulation of the inverse problem in which the model parameters are treated as random variables, and the task is to infer their probability measure. The most widely used approach for probabilistic inversion is the Bayesian framework, which we also employ in this work.

The Bayesian approach to inverse problems [99, 171, 175] provides a consistent mechanism to combine noisy or incomplete data with prior knowledge, e.g., expert opinion, available data, or simply the “best available guess” on the input parameters, and to quantify the uncertainty in the parameter estimate. The prior knowledge is incorporated into a probability distribution over the parameter space, which is usually referred to as the *prior (measure)*. The Bayesian solution to the parameter estimation problem is then a *posterior (measure)* arising from conditioning the prior knowledge (measure) on the observations. However, obtaining the posterior measure is usually not the end of the story, as often, if not always, a postprocessing of the posterior solution

¹We will make clear what a unique solution means in Section 2.4.2.

is necessary. Examples include determining the maximum a-posteriori estimate (MAP) point or computing the expectation, covariance or even higher statistical moments of the posterior. In addition, the posterior measure could be used as input in an uncertainty propagation problem, and thus samples from the posterior are needed in this setting as well. Most often, however, the posterior is intractable in the sense that it does not admit closed form analytic expressions, making direct sampling impossible, hence approximations have to be used in practice.

We briefly review three types of approximations in Bayesian inference: (i) sampling-based approximations, (ii) variational approaches, which reformulate the sampling problem as an optimization problem, and (iii) sampling-free methods, which are of interest in this work. **Sampling-based** posterior approximations include methods such as Markov chain Monte Carlo (MCMC) [18] or sequential Monte Carlo (SMC) [40]. These algorithms were successfully employed in many problems, and historically, methods such as MCMC were among the first to enable the practical use of Bayesian inversion. However, since each sample entails the evaluation of the forward map, the total cost of standard sampling methods becomes prohibitive if the forward model is computationally expensive or if the number of parameters to be inferred is large. To this end, enhanced strategies for high-dimensional Bayesian inversion typically exploit the presence of low-dimensional structure. For example, in most problems, the posterior is a low-dimensional update of the prior, in the sense that the likelihood is influential, relative to the prior, only on a low-dimensional subspace (see [205]). In linear inverse problems with Gaussian prior, [168] showed that low-rank approximations of the prior-preconditioned Hessian of the log-likelihood yield an optimal approximation of the posterior expectation and covariance. In the nonlinear case, heuristics that exploit the geometry of the log-likelihood to accelerate MCMC sampling were proposed in [12, 39, 76]. In [112], a novel multilevel sequential scheme was proposed to decrease the cost of posterior sampling.

An alternative research direction that gained much attention recently is **variational inference**. Drawing inspiration from optimal transport (see [190]), variational inference methods reformulate the sampling as an optimization problem that approximates the posterior distribution with a transformed prior by minimizing their Kullback-Leibler divergence. The resulting mapping is called a *transport map* (see [121] for a general introduction). To improve standard transport mapping, [167] established a link between the Markov properties of the posterior distribution and the existence of low-dimensional couplings, induced by transport maps that are sparse or decomposable. In addition, [14] used a composition of low-dimensional transport maps for high-dimensional Bayesian inference. Moreover, in [43], invertible neural networks and optimal transport were merged to efficiently sample the posterior distribution. Furthermore, the recently developed Stein variational methods seek a composition of simple transport maps parametrized by kernel functions, such as radial bases. In particular, Stein variational methods based on gradient descent [117] or Newton [41] optimization achieve fast convergence, but only in relatively low dimensions. As an improvement, [31] proposed the projected Stein variational Newton method. An alternative two-stage-approach was proposed in [144], in which multifidelity methods and transport maps are combined to speed up MCMC sampling.

Sampling-free approaches offer an alternative way to address complex Bayesian inversion. Specifically, in these methods, either a surrogate of the forward map, negative log-likelihood or directly of the posterior is created, or integration involving the posterior is performed via (enhanced) deterministic techniques such as adaptive sparse grid quadrature. For example, surrogates based on gPC [51, 115, 122–124, 169, 204] or sparse grids with spatial adaptivity [120, 207] exploit the smoothness of the parameters-to-observables map. The sparsity w.r.t. the prior density was exploited via adaptive sparse grid quadrature in [160–162], whereas [30] took advantage of the geometry of the posterior around the MAP point to formulate dimension-independent adaptive sparse grid quadrature schemes. Other efficient emulators of the forward

map or of the negative log likelihood are based on Gaussian process regression [100,149,172]. In addition, [61,64,111,116,138] employed model reduction to speed-up the inference, while [26–28] proposed using sparse grids in combination with reduced bases. Moreover, in cases when the response surface is not very smooth, [128] proposed an approach based on piecewise polynomial approximations together with Voronoi tessellations of the parameter space. Quasi-Monte Carlo [44–46,67,159] is in principle a sampling-free method which does not rely on surrogates, however, it requires a smooth mapping, and it often employs randomization.

Remark: The above mentioned approaches for parameter inference begin from a prior measure on the parameters which is updated to a posterior measure - the solution to the inverse problem - via a likelihood function. Nevertheless, there also exist alternative approaches for parameter inference. For instance, [22,23] proposed a so-called consistent framework based on an updating scheme that finds a posterior probability measure. This produces a push-forward measure when propagated through the deterministic model. This measure matches exactly the observed probability measure on the data. An alternative idea was outlined in [125–127], where conditional expectation was leveraged to formulate a sampling-free nonlinear Bayesian updating scheme in the form of a filter. In this work, we employ the Bayesian framework defined in terms of a prior, a likelihood and a posterior and thus we do not review further alternative approaches.

The major challenge in uncertainty propagation is dealing with computationally expensive forward models with a large(r) number of uncertain inputs. In Bayesian inference, additionally to the aforementioned challenges, we also have to find suitable strategies to approach inverse problems with concentrated posteriors or to find accurate surrogates at a low cost even when the number of parameters to be inferred is small. We take these challenges as motivation to formulate novel strategies suitable for quantifying uncertainty in higher-dimensional, computationally expensive problems. Next, we give a brief summary of the contributions of this work.

1.3 Our contributions in brief

Motivated by the computational challenges appearing in higher-dimensional uncertainty propagation and Bayesian inference, we formulate four novel algorithms, three for uncertainty propagation and one for Bayesian inverse problems. We base our proposed approaches on approximation with sparse grids, multilevel decompositions and multifidelity sampling methods. As we will discuss throughout this work, all three approaches employ *model hierarchies*. In sparse grid approximation, the given forward model is used in all computations, the hierarchy appearing in the stochastic space, where different levels of approximation are employed. Going one step further, in multilevel decompositions, we have a *within model* hierarchy, in which the given forward model is discretized on different grids, and these discretizations are combined with a (complementary) sampling hierarchy of the stochastic space. Finally, multifidelity sampling methods employ a *between models* hierarchy, that is, heterogeneous models with different fidelities are used to estimate statistics of high-fidelity model outputs. The proposed approaches are also *context-aware*, in the sense that to overcome the computational challenges of higher-dimensional UQ, they exploit the *structure* of the underlying problem. By structure, we refer to properties such as anisotropic coupling of the stochastic inputs, lower stochastic dimensionality or properties of reduced models, such as accuracy and cost rates. Next, we briefly summarize the four proposed algorithms:

1. *Sensitivity-driven dimension-adaptive sparse grid algorithm.* This approach employs *hierarchies* appearing in dimension-adaptive sparse grid approximations. We use adaptive

sparse grids to efficiently assess quantities of interest such as expectation, standard deviation or sensitivity indices of output of interest in uncertainty propagation settings. As approximation operators, we consider pseudo-spectral projections and interpolation. In this way, besides computing quantities of interest, our approach also yields a reduced model of the underlying forward model. Our novelty is a *context-aware* sensitivity-driven refinement indicator for dimension-adaptivity, with the goal to explore and exploit the anisotropic coupling and lower intrinsic dimensionality of the stochastic input space;

2. *Multilevel adaptive spectral projection with online dimensionality reduction.* This method is based on *model hierarchies* appearing in sparse grid approximations with spatial adaptivity and in multilevel decompositions. Specifically, at each level in the multilevel decomposition, we formulate a novel *context-aware* approach to estimate spectral projection coefficients via spatially adaptive sparse grid interpolation and one-dimensional quadrature. The multilevel decomposition is used to assess the global spectral projection coefficients, which are used to estimate quantities of interest such as expectation, standard deviation or sensitivity indices. The major novelty of our multilevel approach is that we endow it with the possibility to perform *context-aware* online stochastic dimensionality reduction at no additional computational cost: if a stochastic input is rendered unimportant, based on its sensitivity index, we replace it with a deterministic value and we thus reduce the stochastic dimensionality online;
3. *Context-aware multifidelity sampling algorithm.* This approach is based on *model hierarchies* used in multifidelity Monte Carlo sampling. The major novelty is that we construct *context-aware* reduced models: surrogates are constructed for being explicitly used together with the high-fidelity model to estimate statistics of the high-fidelity model outputs. We show that there is a trade-off between exploring and exploiting the surrogates and that in consequence, low-fidelity models can be too accurate for multifidelity methods, which is in contrast with traditional model reduction. In addition, to show the generality of the proposed approaches in this work, we employ the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm as a reduced model in one of the text case;
4. *Multilevel adaptive sparse Leja algorithm.* The last proposed approach employs *hierarchies* given by (adaptive) sparse grid approximations and multilevel decompositions. The major contribution of this algorithm is the construction of posterior-focused surrogates from sequentially updated reference densities, at no additional cost. To further reduce the computational cost, we employ *context-aware*, adaptive sparse grid approximations, including a dimension-adaptive strategy enhanced with a filtering step based on sensitivity information about the uncertain inputs w.r.t. the sequentially updated reference density. We note that in contrast with the majority of existing approaches for Bayesian inference, our proposed method is fully deterministic in the sense that no sampling is used either to construct surrogates nor for sampling from the posterior solution.

We employ the proposed approaches for uncertainty propagation in complex applications such as plasma microturbulence analysis and fluid-structure interaction. The proposed approach for Bayesian inference is tested in a problem with a forward model defined in a 3D physical domain. We will show that our algorithms outperform existing approaches by orders of magnitude.

We provide a visual summary of the four algorithms in Figure 1.1. Besides the fact that all approaches are context-aware and based on model hierarchies, there are also several other connections between them. For example, both Multilevel adaptive spectral projection with online dimensionality reduction and the Multilevel adaptive sparse Leja algorithms employ multilevel

decompositions and adaptive sparse grid approximations with dimension or spatial adaptivity. In addition, in the Multilevel sparse Leja algorithm, we employ sparse grids with dimension-adaptivity, which are also at the core of the Sensitivity-driven dimension-adaptive sparse grid algorithm. In both algorithms, we enhance standard dimension-adaptivity to further exploit the structure of the underlying stochastic problem. Finally, in the Context-aware multifidelity sampling algorithm, we employ the Sensitivity-driven dimension-adaptive sparse grid algorithm as a reduced model in one of the numerical experiments. In summary, the four proposed algorithms are the pieces of a puzzle which, as a whole, aims to tackle the challenges of quantifying uncertainty in a broad spectrum of computationally expensive, higher-dimensional problems.




Figure 1.1: Visual summary of the four proposed context-aware approaches in this work.

1.4 Thesis outline

The remainder of this work is organized as follows. In Chapter 2, we overview the fundamental concepts and methodologies on which we will build our four proposed algorithms:

- in Section 2.1, we introduce some of our notation and discuss uncertainty modelling;
- we state the assumptions used throughout this work in Section 2.2;
- in Section 2.3, we give a brief overview of outer-loop applications, which include uncertainty propagation and inverse problems;
- a more detailed introduction to uncertainty propagation and the Bayesian approach to inverse problems is given in Section 2.4;
- in Section 2.5, we overview *dimension-adaptive sparse grid approximations*;
- in Section 2.6, we summarize *spatially adaptive sparse grid approximations*;


- *multilevel decompositions* are presented in Section 2.7;
- in Section 2.8, we present *multifidelity sampling*;
- Section 2.9 is concerned with Sobol' indices for global sensitivity analysis.

We end this tour-de-force with showing, in Section 2.10, that intuitively, both multilevel decompositions and multifidelity sampling can be seen as (abstract) combination schemes summarized in Section 2.5. We will introduce most of our notation in Chapter 2. We remark that in the literature on uncertainty propagation, Bayesian inference, sparse grid approximations, multilevel decompositions, multifidelity sampling and sensitivity analysis, we generally encounter rather different notation standards. Here, we regard these concepts in a unified framework, therefore we introduce a unified notation, while maintaining, where possible, the standardized notation.


The algorithms formulated in this work are aimed for quantifying uncertainty in higher-dimensional, computationally expensive and practically relevant problems. As representatives of such problems, we consider fluid-structure interaction, an example of multi-physics applications, and *plasma microturbulence analysis*. We summarize these two applications in Chapter 3.

In Parts II – Parts IV, we present in detail the four proposed algorithms:

- in Part II, we present the first two contributions for uncertainty propagation. In Chapter 4, we focus on the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm, which is based on model hierarchies appearing in dimension-adaptive sparse grids and Sobol' indices for global sensitivity analysis (see the diagram below). We apply the new approach in two plasma microturbulence analysis test cases with three, eight or 12 uncertain inputs.



Chapter 5 is dedicated to our second contribution for uncertainty propagation, the Multilevel adaptive spectral projection with online stochastic dimensionality reduction. Here, we employ model hierarchies appearing in spatially adaptive sparse grid approximations and multilevel decompositions, and Sobol' indices for global sensitivity analysis; see the illustration below. We test the new approach in a damped oscillator problem, first considering five and then six uncertain inputs, and in a fluid-structure interaction test case with five stochastic parameters.



- in Part III, we present our third and final contribution for uncertainty propagation. In Chapter 6, we switch focus from model hierarchies appearing deterministic approximations, such as those based on sparse grids, and formulate a novel Context-aware multifidelity Monte Carlo sampling algorithm. We test the algorithm in a thermal block problem with nine uncertain parameters and in a plasma microturbulence analysis test case with 12 uncertain inputs. In the second test case, we employ the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm as a reduced model; see the depiction below.

Chapter 6

Section 2.4.1 Section 2.8 Chapter 4 Section 3.2

- finally, in Part IV, we present our fourth and last contribution in this work. We switch focus from uncertainty propagation to Bayesian inference and formulate, in Chapter 7, a novel Multilevel adaptive sparse Leja algorithm for Bayesian inverse problems. The proposed approach depends on multilevel decompositions, model hierarchies appearing in dimension-adaptive approximations as well as Sobol' indices for global sensitivity analysis. We depict these dependencies below. We apply the new algorithm in Bayesian inverse problems with elliptic forward models defined in 2D and 3D spatial domains, in which we infer two or eight uncertain inputs.

Chapter 7

Section 2.4.2 Section 2.5 Section 2.7 Section 2.9

We end this thesis in Part V with a summary and outlook for future research.

This work comprises aspects of applied mathematics, computer science, probability theory, statistics, and physics. Nevertheless, it was written to be, as much as possible, self contained. A reader familiar with either one of the aforementioned fields or any other related field should be able to read this work without the need to overview additional materials. We provide, however, relevant references which the reader can use for more details about the discussed topics. Next, we provide the reader with the theoretical background, including the three approximation strategies on which we will build our algorithms, and introduce most of our notation.

"If you wish to make an apple pie from scratch, you must first invent the universe."

Carl Sagan

2

Theoretical background

In this chapter, we overview the fundamental concepts and methodologies used in this work. In Section 2.1, we introduce some of our notation and discuss uncertainty modelling. We state the assumptions used throughout this work in Section 2.2. In Section 2.3, we give a brief overview of outer-loop applications, which include uncertainty propagation and inverse problems, the two UQ branches considered in this work. We provide a more detailed introduction to uncertainty propagation and summarize the Bayesian approach to inverse problems in Section 2.4. In the remainder of this chapter, we present the methodologies on which we will build our four proposed context-aware approaches. Our algorithms are based on (i) approximations with adaptive sparse grids (Sections 2.5 and 2.6), (ii) multilevel decompositions (Section 2.7) and (iii) multifidelity sampling (Section 2.8). In addition, another branch of UQ of which we make use to formulate three of our four algorithms is global sensitivity analysis via Sobol' indices, summarized in Section 2.9. We end this tour-de-force with showing, in Section 2.10, that intuitively, both multilevel decompositions and multifidelity sampling can be seen as combination techniques, summarized in Section 2.5. Furthermore, we will also discuss that the three aforementioned approximation approaches are based on model hierarchies: in sparse grid approximations, we have an approximation hierarchy in the stochastic domain, in multilevel decomposition we have a within-model hierarchy, while in multifidelity sampling we have between-models hierarchy. We introduce most of our notation in this chapter. We remark that in the literature on uncertainty propagation, Bayesian inference, sparse grid approximations, multilevel decompositions, multifidelity sampling and sensitivity analysis, different notation standards are generally employed. Since in this work, we regard these concepts in a unified framework, we introduce a unified notation, but we also try to maintain, where possible, the standardized notation.

2.1 Setup

Our goal in this work is to quantify uncertainty in complex, real-world phenomena such as fluid-structure interaction and plasma microturbulence analysis. These problems are typically governed by a model, \mathcal{F} , specified by a (complex and nonlinear) ODE/PDE system, stemming from first principles. The solution to this model is very rarely, if ever, available analytically, hence we resort to numerical approximations. Let us denote the discretized version of the underlying continuous model by \mathcal{F}_h . Let $\mathcal{E}(\mathcal{F} - \mathcal{F}_h)$ denote the approximation error, where \mathcal{E} denotes a conveniently chosen error indicator, e.g., a norm. The discretization parameter, h , can comprise, for example, the time step size and/or the spatial discretization parameter, e.g., mesh size used to obtain \mathcal{F}_h . Throughout this work, \mathcal{F}_h is assumed to be sufficiently accurate to yield results with any desired accuracy, i.e., $\mathcal{E}(\mathcal{F} - \mathcal{F}_h)$ can be made arbitrarily small. Thus, we refer to this model as the *high-fidelity model* (see, e.g., [146]). Moreover, we assume that \mathcal{F}_h is available as a *black box* in the sense that in numerical simulations, we do not need to modify the mathematical operators describing \mathcal{F} and thus also the underlying simulation code.

The high-fidelity model, \mathcal{F}_h , depends on *deterministic inputs*, \mathbf{x} , and *stochastic inputs*, $\boldsymbol{\theta}$, i.e., the high-fidelity model is $\mathcal{F}_h \equiv \mathcal{F}_h(\mathbf{x}, \boldsymbol{\theta})$. The deterministic inputs comprise time, spatial coordinates, initial, boundary condition etc., that is, everything describing the underlying problem which is a priori known with certainty. Since \mathbf{x} is fixed, we omit it and denote the high-fidelity model as $\mathcal{F}_h \equiv \mathcal{F}_h(\boldsymbol{\theta})$. We further assume \mathcal{F}_h is a bounded and measurable function w.r.t. the Lebesgue measure and the Borel σ -algebra on \mathbb{R} .

Uncertainty enters \mathcal{F}_h via $\boldsymbol{\theta} := (\theta_1, \theta_2, \dots, \theta_{d_{\text{sto}}}) \subset \mathbb{R}^{d_{\text{sto}}}$, where $d_{\text{sto}} \in \mathbb{N}$ is referred to as the *stochastic dimensionality*. We model $\boldsymbol{\theta}$ as a *multivariate* random vector defined in a probability space (Θ, A, P) , where Θ is the event space, which is equipped with σ -algebra A and probability measure P . Moreover, $\boldsymbol{\theta}$ is a *continuous* random vector characterized by a *probability density function*, $\boldsymbol{\pi}$, with support \mathbf{X} . Thus, $\boldsymbol{\theta} \in \mathbf{X}$. Additionally, the components θ_i of $\boldsymbol{\theta}$ are assumed to be L^2 random variables, i.e., they have finite expectation

$$\mathbb{E}[\theta_i] := \int_{X_i} \theta_i \pi_i(\theta_i) d\theta_i$$

and finite variance

$$\text{Var}[\theta_i] := \mathbb{E}[\theta_i^2] - (\mathbb{E}[\theta_i])^2,$$

where X_i is the domain and π_i is the (marginal) density of θ_i .

Since the stochastic input is a continuous random variable with multivariate density $\boldsymbol{\pi}$, in computer simulations we need therefore discrete realizations, that is, samples. For simplicity, throughout this work, we use the term sample to refer broadly to points that fall within the support of $\boldsymbol{\pi}$. This includes the “traditional” samples, i.e., (pseudo-) random numbers, but also deterministic, e.g., collocation or quadrature points.

We employ the following notation conventions in the remainder of this thesis. Non-scalar or tensorized quantities, or quantities with a product structure are denoted by *italic* letters in **bold**, e.g., $\boldsymbol{\theta}$. In contrast, univariate, scalar or multivariate quantities that do not have a product structure are denoted only by *italic* letters, e.g., θ_i . Throughout this work, we denote by $\boldsymbol{\theta}$ the stochastic parameters of the high-fidelity model and by \mathbf{X} their support. Moreover, we use the letter π or $\boldsymbol{\pi}$, sometimes with super or subscripts, to refer to probability density functions. Comparisons between vectors are understood component-wisely, that is, $\boldsymbol{\theta} = \boldsymbol{\omega}$ or $\boldsymbol{\theta} \leq \boldsymbol{\omega}$ means, respectively, that $\theta_i = \omega_i$ or $\theta_i \leq \omega_i$ for all $i = 1, 2, \dots$. The remaining of our notation is introduced in the following sections.

In this section, we introduced some of our notation, we introduced concepts such as high-fidelity model and discussed how we model uncertainty. In the next section, we specify the assumptions we make about the stochastic space, \mathbf{X} , and its associated density function, $\boldsymbol{\pi}$.

2.2 Assumptions

We make the following assumptions regarding the stochastic space, \mathbf{X} , and the probability density function characterizing the uncertain inputs, $\boldsymbol{\pi}$, which are needed to formulate the four proposed approaches in this work.

- Assumption 2.1: The parameter space is a tensor *product* space, i.e.,

$$\mathbf{X} = \bigotimes_{i=1}^{d_{\text{sto}}} X_i,$$

where $X_i \subset \mathbb{R}$ for $i = 1, 2, \dots, d_{\text{sto}}$.

- Assumption 2.2: The probability density $\boldsymbol{\pi}$ has a *product* structure, i.e.,

$$\boldsymbol{\pi}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} \pi_i(\theta_i),$$

where $\pi_i: X_i \rightarrow \mathbb{R}$ are its *marginals*, where $i = 1, 2, \dots, d_{\text{sto}}$.

Assumption 2.1 can always be satisfied by embedding a non-tensorized parameter space into a hyperspace of suitable dimension. Assumption 2.2 is fulfilled if the components of $\boldsymbol{\theta}$ are stochastically independent. If Assumption 2.2 is not satisfied, we could use a transformation such as a copula [170] or a transport map [121]. Any additional assumptions are only needed in specific chapters (see, for example, Chapter 6).

2.3 Outer-loop scenarios

In this work, we are interested in uncertainty propagation and Bayesian inference in computationally expensive, higher-dimensional problems. Uncertainty propagation and inference are examples of so-called *outer-loop scenarios* (see, for example, [146]). Other examples of outer-loop scenarios are optimization or control.

Broadly speaking, outer-loop scenarios are applications that require an ensemble of solutions of the underlying model for different input instances. That is, they are, at least conceptually, *loops* wrapped around the given model. An intuitive example is numerical quadrature, in which the underlying integrand is evaluated at quadrature nodes. In our context, the outer-loop is put around the high-fidelity model, hence yielding an ensemble of high-fidelity evaluations. We remark that applications such as (intrusive) stochastic Galerkin projections [73], are not, in general, outer-loops: in these methods, the solution corresponds to a *modified model*, which is typically solved *serially*, as a *monolithic* application. Instances of the stochastic input, $\boldsymbol{\theta}$, are input into the high-fidelity model. For each such instance, a high-fidelity output is computed and the ensemble of these outputs is used to estimate the desired outer-loop result. For example, in uncertainty propagation or Bayesian inference, we typically estimate the mean value of the high-fidelity output. In general, although this is not restrictive, the outer-loop result is a scalar quantity. We give an intuitive illustration of outer-loops in Figure 2.1.

In computer simulations, outer-loops seem to be, intuitively, implemented as for or do loops. However, this is not necessarily the case. For example, if we revisit the aforementioned quadrature example, assuming we have N processes available and N quadrature nodes in total, the evaluation of the integrand can be done in parallel using one process per node. In addition, the size of the outer-loop is not always available a priori. This is the case, e.g., when employing adaptivity or other type of sequential construction of the outer-loop.

Up to this point, we provided the reader with an intuitive description of uncertainty propagation and Bayesian inference. In the following, we give a more formal definition of the two outer-loop scenarios, including mathematical details.

2.4 Uncertainty propagation and Bayesian inference: an overview

In Section 2.4.1, we overview uncertainty propagation. A summary of Bayesian inference are presented in Section 2.4.2.




Figure 2.1: Throughout this work, we assume we are given a high-fidelity model which can yield results with any desired accuracy, but which is computationally too expensive for outer-loop scenarios such as uncertainty propagation or Bayesian inference. To make such outer-loops computationally feasible, we design context-aware algorithms based on model hierarchies. Our goal is therefore two-fold: we aim to exploit the properties of either the underlying high-fidelity model to create accurate surrogates in single fidelity settings or of the surrogates themselves to create context-aware low-fidelity models in multifidelity settings.

2.4.1 Uncertainty propagation

Given a stochastic vector $\boldsymbol{\theta} \in \mathbf{X}$ with independent components $\theta_1, \theta_2, \dots, \theta_{d_{\text{sto}}}$ and a well-defined, well-posed forward mapping \mathcal{F}_h , uncertainty propagation is concerned with quantifying the effect of input uncertainty in specific outputs of the forward model. That is, the uncertainty is propagated through \mathcal{F}_h in the direction input \rightarrow model \rightarrow output, i.e., we have a forward problem. Well-posedness here is understood in the sense of Hadamard [84], that is

1. a solution exists,
2. it is unique,
3. the solution depends continuously on the data.

All forward models in this work are well-posed. Moreover, they are treated as black boxes.

In uncertainty propagation, we are typically not interested in the full output given by \mathcal{F}_h , but rather in a more specific so-called *output of interest*. For example, in a fluid flow problem, the output of the deterministic simulation typically comprises the pressure and velocity fields in the entire flow domain, whereas in uncertainty propagation, the output of interest might be the pressure at only a user-defined location in the domain. We note, however, that especially when the forward model is treated as a black box, we generally still compute the same output as in a deterministic simulation, from which we extract the output of interest.

As we briefly touched upon in the introduction, there are two main categories of approaches for uncertainty propagation, *intrusive* and *non-intrusive*. On the one hand, standard intrusive methods require changing the underlying model operators and with that, the simulation code as well. In this work, we are concerned with black box uncertainty propagation in problems with complex mathematical models, which usually also have complex simulation codes. Thus, we do not consider intrusive UQ methods. Non-intrusive methods, on the other hand, such as MC sampling or collocation-based approaches, are based on ensembles of independent, black

2.4. UNCERTAINTY PROPAGATION AND BAYESIAN INFERENCE: AN OVERVIEW

box high-fidelity evaluations. Our focus in this thesis is exclusively on non-intrusive algorithms for uncertainty propagation. In the remainder of this work, by uncertainty propagation we refer to only non-intrusive uncertainty propagation.

In more formal terms, uncertainty propagation is generally performed as follows. Via a suitable sampling procedure, we draw samples from $\boldsymbol{\pi}$ and we solve the high-fidelity forward model, \mathcal{F}_h , for each input sample. At the end of one simulation, an observation operator, \mathbb{O} , is employed to determine the output of interest, $o(\boldsymbol{\theta})$. The propagation of samples from the input distribution through \mathcal{F}_h gives rise to a probability density of $o(\boldsymbol{\theta})$, π_o as well. Finally, the ensemble of simulations are used to assess quantities such as expectation

$$\mathbb{E}[o] := \int_{\mathbf{X}} o(\boldsymbol{\theta})\boldsymbol{\pi}(\boldsymbol{\theta})d\boldsymbol{\theta},$$

variance or standard deviation

$$\text{Var}[o] := \mathbb{E}[o^2] - (\mathbb{E}[o])^2, \quad \text{Std}[o] := \sqrt{\text{Var}[o]},$$

sensitivity indices, which we discuss in Section 2.9. We call such quantities *quantities of interest*, denoted by \mathcal{Q} . Note that the output of interest is itself a random variable as well with a density π_o . Since in the context of uncertainty propagation, the output of interest, o , is specified as $o(\boldsymbol{\theta}) := \mathcal{F}_h(\boldsymbol{\theta})$, we generally refer to the quantities of interest as depending on the high-fidelity model, i.e., we use, for example, $\mathbb{E}[o] := \mathbb{E}[\mathcal{F}_h]$. We provide a visual illustration of the uncertainty propagation pipeline in Figure 2.2.




Figure 2.2: Visual illustration of uncertainty propagation. The uncertain input parameter, $\boldsymbol{\theta}$, is a vector with independent continuous components, characterized by a probability density function, $\boldsymbol{\pi}$, which is known a priori. Via a suitable sampling procedure, we draw samples from $\boldsymbol{\pi}$ and we solve the high-fidelity forward model, \mathcal{F}_h , for each input sample. At the end of one simulation, an observation operator, \mathbb{O} , is employed to determine the output of interest, $o(\boldsymbol{\theta})$. The propagation of samples from the input distribution through \mathcal{F}_h gives rise to a probability density of $o(\boldsymbol{\theta})$, π_o as well. Note that the output of interest in an uncertainty propagation problem can be a subset or even different from the output of the equivalent deterministic simulation. Finally, the ensemble of simulations are used to assess quantities of interest \mathcal{Q} such as expectation, standard deviation, or sensitivity measures of the output of interest.

2.4.2 Bayesian inference for parameter estimation

The Bayesian approach to inverse problems provides a consistent and optimal probability formulation for learning parameters of science and engineering models from observational or experimental data under uncertainty. In this framework, the mathematical model of the process under consideration and the observational data are used to update the prior belief on the model parameters to a posterior distribution, the solution of the Bayesian inverse problem. The posterior

distribution provides a probabilistic description of the model parameters and thus a quantification of uncertainty. In the following, we overview inverse problems in general, we briefly state why these problems are generally ill-posed, and then we summarize how the Bayesian approach can be used to cure the ill-posedness of these problems. For a more detailed description of Bayesian inverse problems, we refer the reader to [99, 171, 175].

Let $\mathbf{Y} = \mathbb{R}^{N_{\text{obs}}}$ be a separable Banach space that denotes the data space. Since the parameter and data spaces, \mathbf{X} and \mathbf{Y} , are *finite-dimensional*, allows us to work with densities w.r.t. the Lebesgue measure. Noisy observations $\mathbf{y} \in \mathbf{Y}$ are performed, for which we employ a statistical, additive noise model. With this model \mathbf{y} is a realisation of the random variable $\mathcal{F}_h(\boldsymbol{\theta}^{\text{true}}) + \boldsymbol{\eta}$, where we assume that $\boldsymbol{\eta} \sim \mathcal{N}(0, \Gamma)$ is non-degenerate Gaussian noise with symmetric and positive definite covariance matrix $\Gamma > 0$, and $\boldsymbol{\theta}^{\text{true}} \in \mathbf{X}$ is the true parameter. The task of the inverse problem is to identify the parameter $\boldsymbol{\theta}^{\text{true}}$, i.e., solve, for $\boldsymbol{\theta}$, the equation

$$\mathcal{F}_h(\boldsymbol{\theta}) + \boldsymbol{\eta} = \mathbf{y}. \quad (2.1)$$

Problem (2.1) is typically ill-posed in the sense of Hadamard [84], due to noise, and since the low-dimensional data space is often not sufficiently rich to allow the identification of a unique parameter in the high-dimensional space \mathbf{X} . The ill-posedness can be cured, for example, by using deterministic regularization, such as Tikhonov regularization [95]. This approach, however, does not provide any probabilistic description of the inferred parameters $\boldsymbol{\theta}$ and thus no quantification of their uncertainty. A probabilistic approach to cure the ill-posedness of (2.1), which we also adopt in this work, is the Bayesian approach [171]. Next, we introduce our notation and overview the Bayesian approach to inverse problems.

Assume that the \mathbf{X} -valued random variable $\boldsymbol{\theta}$ is distributed according to a *prior measure* μ_0 with Lebesgue density π_0 on the parameter space \mathbf{X} . Further, we assume that $\boldsymbol{\theta}$ is stochastically independent of the noise $\boldsymbol{\eta}$. The density π_0 reflects our knowledge about $\boldsymbol{\theta}$ before we make an observation \mathbf{y} . Intuitively, the prior π_0 regularizes the inverse problem, restricting the space in which the sought parameters lie. Note that since π_0 is the given density on the parameter space, \mathbf{X} , it has a product structure by Assumption 2.2. The information provided by \mathbf{y} is modelled by the (data) likelihood. Since the noise $\boldsymbol{\eta}$ is additive and Gaussian by assumption, the likelihood reads

$$\begin{aligned} L(\boldsymbol{\theta}|\mathbf{y}) &:= \exp(-\Phi(\boldsymbol{\theta}; \mathbf{y})), \\ \Phi(\boldsymbol{\theta}; \mathbf{y}) &:= \frac{1}{2} \left\| \Gamma^{-1/2}(\mathbf{y} - \mathcal{F}_h(\boldsymbol{\theta})) \right\|^2. \end{aligned} \quad (2.2)$$

The function Φ is called *potential* or *negative log-likelihood*. The solution of the Bayesian inverse problem is the *posterior (measure)* $\mu^{\mathbf{y}}$, i.e., the conditional measure of $\boldsymbol{\theta}$ given that the event $\{\mathcal{F}_h(\boldsymbol{\theta}) + \boldsymbol{\eta} = \mathbf{y}\}$ occurred. The posterior measure $\mu^{\mathbf{y}}$ has also a density $\pi^{\mathbf{y}}$ which can be computed using Bayes's formula:

$$\pi^{\mathbf{y}}(\boldsymbol{\theta}) = \frac{L(\boldsymbol{\theta}|\mathbf{y})\pi_0(\boldsymbol{\theta})}{Z(\mathbf{y})}, \quad \boldsymbol{\theta} \in \mathbf{X}, \mathbf{y} \in \mathbf{Y}, \quad (2.3)$$

$$Z(\mathbf{y}) = \int_{\mathbf{X}} L(\boldsymbol{\theta}|\mathbf{y})\pi_0(\boldsymbol{\theta})d\boldsymbol{\theta}, \quad (2.4)$$

provided that $0 < Z(\mathbf{y}) < \infty$. In the given setting, i.e., non-degenerate Gaussian additive noise, finite dimensional data space, one can show that $Z(\mathbf{y})$ is always finite and bounded away from 0. This implies existence of the posterior measures, see [110]. The work [110] also establishes that Bayesian inverse problems of this type are always *well-posed*.

2.5. APPROXIMATION WITH DIMENSION-ADAPTIVE SPARSE GRIDS

Finally, let again $o : \mathbf{X} \rightarrow \mathbb{R}$ to denote an output of interest depending on θ . We are interested in computing quantities of interest \mathcal{Q} w.r.t. o . Since θ is a random variable, one is typically interested in the forward propagation of uncertainties through the action of o . We provide a visual illustration of Bayesian parameter inference Figure 2.3.




Figure 2.3: Visual illustration of Bayesian parameter inference. Given the high-fidelity forward model, \mathcal{F}_h , prior information in terms of a prior density, π_0 , and measurement data, \mathbf{y} (all three depicted in black) affected by measurement noise η , assumed to be Gaussian, the Bayesian approach finds a posterior density, $\pi^{\mathbf{y}}$, which characterizes the uncertain inputs, θ . Let again o denote an output of interest depending on θ . Once the posterior density is available, we are typically interested in computing quantities of interest \mathcal{Q} w.r.t. o . We depict the posterior density and all other quantities that are obtained after the posterior is available with red.

A common task in Bayesian inference is to evaluate integrals of o w.r.t. the posterior density:

$$\mathcal{Q} := \mathbb{E}_{\pi^{\mathbf{y}}}[o] = \int_{\mathbf{X}} o(\theta)\pi^{\mathbf{y}}(\theta)d\theta, \quad (2.5)$$

where, for clarity, we explicitly wrote the integration weight function, $\pi^{\mathbf{y}}$, in a subscript. We can write the expected value in (2.5) in terms of a ratio of two expected values w.r.t. the prior:

$$\mathbb{E}_{\pi^{\mathbf{y}}}[o] = \frac{1}{Z(\mathbf{y})} \int o(\theta)L(\theta|\mathbf{y})\pi_0(\theta)d\theta = \frac{\mathbb{E}_{\pi_0}[o(\cdot)L(\cdot|\mathbf{y})]}{\mathbb{E}_{\pi_0}[L(\cdot|\mathbf{y})]}. \quad (2.6)$$

Integrals as in (2.6) are generally approximated via numerical quadrature w.r.t. the prior density, since the prior is typically more accessible compared to the posterior. Specifically, one typically uses either MC quadrature via i.i.d. sampling or other numerical quadrature rules such as quasi-Monte Carlo [44] or sparse grid quadrature [160,162]. The efficient estimation of integrals such as (2.5) via posterior-focused adaptive sparse grid surrogates is addressed by our proposed Multilevel sparse Leja algorithm, detailed in Chapter 7.

In this section, we summarized uncertainty propagation and Bayesian inference. In the next four sections, we overview the three classes of model hierarchies which will be used to formulate our proposed approaches for uncertainty propagation and Bayesian inference in computationally expensive, higher-dimensional problems.

2.5 Approximation with dimension-adaptive sparse grids

In this section, we summarize sparse grid approximations with subspace or dimension-adaptivity, which are at the core of two of our proposed algorithms: the Sensitivity-driven dimension-

adaptive sparse grid algorithm for uncertainty propagation, detailed in Chapter 4, and the Multilevel adaptive sparse Leja algorithm used in the context of Bayesian inversion, presented in Chapter 7. Dimension-adaptive sparse grids are characterized by using *global* approximation operators, that is, approximations supported in the entire domain of the underlying problem. In the literature, these types of approximations are associated with several names and derivations, such as *Smolyak's algorithm* [165] or the *combination technique* [70, 71, 83]. In our context, we construct the approximation for $\mathcal{F}_h(\boldsymbol{\theta})$. Recall that by Assumptions 2.1 and 2.2, the stochastic space, \mathbf{X} , has a tensorized structure, whereas the input density, $\boldsymbol{\pi}$, is the product of its marginals π_i for $i = 1, 2, \dots, d_{\text{sto}}$. Throughout this section, our notation is similar to [35, 71].

Loosely speaking, dimension-adaptive sparse grids are a methodology to construct multivariate approximations by weakening the assumed coupling between the input parameters (see, e.g., [35]). Let $f : X_i \rightarrow \mathbb{R}$ denote a univariate, scalar-valued function. Moreover, let $\mathcal{U}^i[f]$ be a sequence of 1D linear continuous operators depending on the marginals π_i and let $\mathcal{U}_k^i[f]$ be approximations such that

$$\|\mathcal{U}^i[f] - \mathcal{U}_k^i[f]\| \xrightarrow{k \rightarrow \infty} 0, \quad i = 1, 2, \dots, d_{\text{sto}},$$

in a suitable norm $\|\cdot\|$, where $1 \leq k \in \mathbb{N}$ is referred to as *level*. $\mathcal{U}_k^i[f]$ is typically obtained from discrete evaluations of $\mathcal{U}^i[f]$ at $m(k)$ points in X_i , where $m(k) : \mathbb{N} \rightarrow \mathbb{N}$ is called the *level-to-nodes* function. For example, if $\mathcal{U}^i[f]$ denotes integration w.r.t. π_i , then the corresponding approximation $\mathcal{U}_k^i[f]$ could be a quadrature scheme with $m(k)$ nodes and weights that depend on π_i . In general, $\mathcal{U}^i[f]$ denotes a continuous operator that stems from theory but which is not available computationally, while $\mathcal{U}_k^i[f]$ denotes its discrete (convergent) approximation which can be used in numerical computations. We assume that $\mathcal{U}^i[f], \mathcal{U}_k^i[f]$ are *global* operators, that is, their support is the entire domain X_i .

The key idea behind formulating sparse grid approximations is to make use of the fact that $\mathcal{U}_k^i[f]$ converges to $\mathcal{U}^i[f]$ as $k \rightarrow \infty$ and write $\mathcal{U}^i[f]$ as a *telescoping series* of the form

$$\mathcal{U}^i[f] = \mathcal{U}_1^i[f] + (\mathcal{U}_2^i[f] - \mathcal{U}_1^i[f]) + (\mathcal{U}_3^i[f] - \mathcal{U}_2^i[f]) + \dots = \sum_{k=1}^{\infty} \mathcal{U}_k^i[f].$$

We denote by $\Delta_k^i[f] := \mathcal{U}_k^i[f] - \mathcal{U}_{k-1}^i[f]$ the so-called *hierarchical surpluses*, with $\mathcal{U}_0^i[f] := 0$. Using this notation, the above telescoping series becomes

$$\mathcal{U}^i[f] = \sum_{k=1}^{\infty} \Delta_k^i[f]. \quad (2.7)$$

Remark: (i) The telescoping series (2.7) is a *multilevel* representation of the operator $\mathcal{U}^i[f]$ depending on the level k .

(ii) If the point set used to determine $\mathcal{U}_{k-1}^i[f]$ is a subset of the set of $m(k)$ points used to assess $\mathcal{U}_k^i[f]$, $k \geq 2$, that is, if the point sets at two adjacent levels are *nested*, the linearity of $\mathcal{U}_k^i[f]$ implies that the $m(k-1)$ evaluations of f used to determine $\mathcal{U}_{k-1}^i[f]$ can be reused when computing $\mathcal{U}_k^i[f]$. In other words, to determine $\mathcal{U}_k^i[f]$, it is sufficient to evaluate f at only $m(k) - m(k-1)$ points.

Before extending the above formulas to d_{sto} dimensions, we introduce the definitions for *accuracy* and *half accuracy approximation sets* from [35] which we need later on.

Definition: ([35, Definition 2.1]). For a given operator $\mathcal{U}^i[f]$ and approximation $\mathcal{U}_k^i[f]$, we define the *exact approximation set* by $\mathcal{W}(\mathcal{U}_k^i[f]) := \{f : \mathcal{U}_k^i[f] = \mathcal{U}^i[f]\}$. Moreover, the *half exact approximation set* is defined as $\mathcal{W}^2(\mathcal{U}_k^i[f]) := \{f : \mathcal{U}_k^i[f^2] = \mathcal{U}^i[f^2]\}$.

The definition of the exact approximation set is straightforward and it is used to establish the functions for which sparse grid approximations are exact. In general, these functions are polynomials. The half exact approximation set seems less intuitive, but it will be useful when we discuss pseudo-spectral approximations in Section 2.5.1.

In this work, we are interested in problems that depend on $d_{\text{sto}} \geq 2$ stochastic parameters. Since the input probability density π is assumed to have a product structure, a natural way to lift the 1D operators to d_{sto} dimensions is via *tensorization*, i.e.,

$$\mathbf{u}^{d_{\text{sto}}}[\mathcal{F}_h] = \left(\mathcal{U}^1 \otimes \mathcal{U}^2 \otimes \dots \otimes \mathcal{U}^{d_{\text{sto}}} \right) [\mathcal{F}_h] = \left(\bigotimes_{i=1}^{d_{\text{sto}}} \mathcal{U}^i \right) [\mathcal{F}_h]. \quad (2.8)$$

Plugging the telescoping series (2.7) into the tensorized representation (2.8), we obtain

$$\begin{aligned} \mathbf{u}^{d_{\text{sto}}}[\mathcal{F}_h] &= \left(\sum_{k_1=1}^{\infty} \Delta_{k_1}^1 \otimes \sum_{k_2=1}^{\infty} \Delta_{k_2}^2 \otimes \dots \otimes \sum_{k_{d_{\text{sto}}}=1}^{\infty} \Delta_{k_{d_{\text{sto}}}}^{d_{\text{sto}}} \right) [\mathcal{F}_h] \\ &= \left(\sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}}} \Delta_{k_1}^1 \otimes \Delta_{k_2}^2 \otimes \dots \otimes \Delta_{k_{d_{\text{sto}}}}^{d_{\text{sto}}} \right) [\mathcal{F}_h] = \sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}}} \Delta_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h], \end{aligned} \quad (2.9)$$

where $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}}) \in \mathbb{N}^{d_{\text{sto}}}$ denotes a *multiindex* and

$$\Delta_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h] = \sum_{\mathbf{z} \in \{0,1\}^{d_{\text{sto}}}} (-1)^{|\mathbf{z}|_1} \mathbf{u}_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h],$$

where $\mathbf{u}_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h] = \left(\bigotimes_{i=1}^{d_{\text{sto}}} \mathcal{U}_{k_i}^i \right) [\mathcal{F}_h]$ and $|\mathbf{z}|_1 := \sum_{i=1}^{d_{\text{sto}}} z_i$.

Formula (2.9) is a representation of $\mathbf{u}^{d_{\text{sto}}}[\mathcal{F}_h]$ with an infinite number of terms. To make (2.9) suitable for numerical simulations, we need to truncate it, i.e., we need to restrict the multiindices \mathbf{k} to a finite set. To this end, consider a finite multiindex set $\mathcal{K} \subset \mathbb{N}^{d_{\text{sto}}}$ and define

$$\mathbf{u}_{\mathcal{K}}^{d_{\text{sto}}}[\mathcal{F}_h] = \sum_{\mathbf{k} \in \mathcal{K}} \Delta_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h]. \quad (2.10)$$

Since (2.10) involves tensorizations of univariate difference operators (2.7) the multiindex set \mathcal{K} must be constructed such that the summation in (2.10) telescopes correctly. Such suitable multiindex sets \mathcal{K} are called *admissible* or *downward closed* (see [71]). In particular, for an admissible set \mathcal{K} it holds that $\mathbf{k} \in \mathcal{K} \Rightarrow \mathbf{k} - \mathbf{e}_i \in \mathcal{K}$ for $i = 1, 2, \dots, d_{\text{sto}}$, where \mathbf{e}_i denotes the i th unit vector in $\mathbb{R}^{d_{\text{sto}}}$. Note that (2.10) can be rewritten as a *combination scheme* (see, e.g., [83]),

$$\mathbf{u}_{\mathcal{K}}^{d_{\text{sto}}}[\mathcal{F}_h] = \sum_{\mathbf{k} \in \mathcal{K}} a_{\mathbf{k}} \mathbf{u}_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h], \quad (2.11)$$

where $a_{\mathbf{k}} = \sum_{\mathbf{z} \in \{0,1\}^{d_{\text{sto}}}} (-1)^{|\mathbf{z}|_1} \chi_{\mathcal{K}}(\mathbf{k} + \mathbf{z})$ and $\chi_{\mathcal{K}}$ is the characteristic function on \mathcal{K} , i.e., $\chi_{\mathcal{K}}(\mathbf{k}) = 1$ if $\mathbf{k} \in \mathcal{K}$ and $\chi_{\mathcal{K}}(\mathbf{k}) = 0$ otherwise (see, e.g., [155]). The combination formula (2.11) tells us that the sum of multivariate hierarchical surpluses expressed in (2.10) can be rewritten as a linear combination of *full-grid* approximations. The trick is that these full grids are *anisotropic*, with an overall number of points considerably smaller than a full isotropic grid. We provide some examples in Section 2.5.2.

We now briefly touch upon the exactness of multivariate approximations of the form (2.10), (2.11). The following theorem of Conrad and Marzouk [35] states that if the underlying 1D approximation operators have nested exact sets, then the exact set of the resulting multivariate approximation contains the union of the exact 1D sets.

Theorem 2.5.1 ([35, Theorem 3.2]) Assume that for the approximation $\mathcal{U}_{\mathcal{K}}^{d_{\text{sto}}}[\mathcal{F}_h]$ defined in (2.10), (2.11), the exact sets of the underlying 1D approximations are nested, i.e., $k \leq k' \Rightarrow \mathcal{W}(\mathcal{U}_k^i) \subseteq \mathcal{W}(\mathcal{U}_{k'}^i)$ for $i = 1, 2, \dots, d_{\text{sto}}$. Then, we have

$$\mathcal{W}(\mathcal{U}_{\mathcal{K}}^{d_{\text{sto}}}[\mathcal{F}_h]) \supseteq \bigcup_{\mathbf{k} \in \mathcal{K}} \mathcal{W}(\Delta_{\mathbf{k}}^{d_{\text{sto}}}[\mathcal{F}_h]) \supseteq \bigcup_{\mathbf{k} \in \mathcal{K}} \left(\bigotimes_{i=1}^{d_{\text{sto}}} \mathcal{W}(\Delta_{k_i}^i) \right) [\mathcal{F}_h].$$

To fully define the approximation in (2.10) or (2.11) we therefore need three ingredients: (i) the linear operators $\mathcal{U}^i[f]$, (ii) a point set to construct the approximation $\mathcal{U}_k^i[f]$ from $m(k)$ evaluations of $\mathcal{U}^i[f]$, and (iii) the (finite) multiindex set \mathcal{K} . Without loss of generality, the operators $\mathcal{U}^i[f]$ in this work are identical for all input directions i , which is generally the case when employing sparse grid approximations. Hence we omit the superscript and use the notation $\mathcal{U}[f]$ and $\mathcal{U}_k[f]$, respectively. Moreover, we also omit the superscript d_{sto} in (2.10) or (2.11); unless specified otherwise, all multivariate quantities in this work are d_{sto} -dimensional. We summarize the employed approximation operators in Section 2.5.1. Section 2.5.2 reviews weighted (L)-Leja sequences, which we use to construct the approximations. Finally, in Section 2.5.3, we summarize the dimension-adaptive algorithm [71,90] used to determine the multiindex set, \mathcal{K} , and thus to fully specify the sparse grid approximation (2.10).

2.5.1 Approximation operators

In this work, we employ three types of global approximations to define dimension-adaptive sparse grid approximations (2.10), (2.11): quadrature (*qu*), pseudo-spectral projection (*psp*) and interpolation with global Lagrange polynomials (*in.g*). The abbreviations in parenthesis are used as superscripts to distinguish between the three operators. In addition, we also use the superscript *op* to refer to any of the three approximation operations.

Interpolatory quadrature

Let $f : X_i \rightarrow \mathbb{R}$ denote a one-dimensional integrable function w.r.t. the marginal density π_i . We are interested in the numerical approximation of weighted integrals of the form

$$\mathcal{U}[f] := \int_{X_i} f(\theta) \pi_i(\theta) d\theta$$

via *interpolatory quadrature schemes*.

Interpolatory quadrature rules employ the *approximate then integrate* principle to approximate (weighted) integrals: first, the integrand is approximation with a polynomial, and this polynomial is then integrated *exactly*, up to the employed numerical precision, using nodes and weights associated to the integration weight function. Thus, interpolatory quadrature schemes are characterized by a *degree-of-exactness*, which is the maximum polynomial degree which can be integrated exactly. The 1D schemes are defined as

$$\mathcal{U}_k^{\text{qu}}[f] := \sum_{n=1}^{m(k)} f(\theta_n) w_n, \tag{2.12}$$

where $\{\theta_n, w_n\}_{n=1}^{m(k)}$ are the quadrature nodes and weights associated to π_i . Note that since π_i is a density function, $\int_{X_i} f(\theta) \pi_i(\theta) d\theta = \mathbb{E}[f]$, and the weights w_n are normalized, i.e., $\sum_{n=1}^{m(k)} w_n = 1$. The exact set of $\mathcal{U}_k^{\text{qu}}[f]$ in (2.12) is then $\mathcal{W}(\mathcal{U}_k^{\text{qu}}[f]) = \{\sum_{m=0}^p \theta^m : p = 0, 1, \dots, DE_{m(k)}\}$, where $DE_{m(k)}$ denotes the degree-of-exactness associated to the underlying $m(k)$ nodes and weights.

Example: For Gaussian quadrature, $DE_{m(k)} = 2m(k) - 1$, whereas for Clenshaw-Curtis [32] or Leja quadrature (see, e.g., [134]), $DE_{m(k)} = m(k) - 1$. Even though this seems to indicate that Gaussian quadrature is more accurate as compared to the other rules, works such as [53, 180] established that schemes with lower degree-of-exactness can be compatibly accurate, or sometimes even more accurate than Gaussian quadrature when the integrand is nonlinear. For linear or quasi-linear integrands, Gaussian quadrature is, by construction, more accurate.

Next, we present the d_{sto} -variate integration operator and its quadrature approximation for a multiindex $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}})$. Assume that $\mathcal{F}_h : \mathbf{X} \rightarrow \mathbb{R}$ is integrable w.r.t. the multivariate probability density $\boldsymbol{\pi}$. The d_{sto} -variate integral operator reads

$$\mathcal{U}[\mathcal{F}_h] = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \boldsymbol{\pi}(\boldsymbol{\theta}) d\boldsymbol{\theta}$$

and its approximation via quadrature is

$$\mathcal{U}_{\mathbf{k}}^{\text{qu}}[\mathcal{F}_h] := \left(\bigotimes_{i=1}^{d_{\text{sto}}} \mathcal{U}_{k_i}^{\text{qu}} \right) [\mathcal{F}_h] = \sum_{n_1=1}^{m(k_1)} \sum_{n_2=1}^{m(k_2)} \dots \sum_{n_{d_{\text{sto}}}=1}^{m(k_{d_{\text{sto}}})} \mathcal{F}_h(\theta_{n_1}, \theta_{n_2}, \dots, \theta_{n_{d_{\text{sto}}}}) w_{n_1} w_{n_2} \dots w_{n_{d_{\text{sto}}}}.$$

The associated degree-of-exactness follows from the ones of the one-dimensional schemes and it reads $\mathcal{W}(\mathcal{U}_{\mathbf{k}}^{\text{qu}}[\mathcal{F}_h]) = \{\sum_{\mathbf{m}=0}^{\mathbf{p}} \boldsymbol{\theta}^{\mathbf{m}} : |\mathbf{p}|_1 \leq \sum_{i=1}^{d_{\text{sto}}} DE_{m(k_i)}\}$.

Pseudo-spectral projection

Let \mathbb{P}_{P_k} denote the (vector) space of univariate polynomials of degree $P_k \in \mathbb{N}$. In addition, denote by $H_i = L^2(X_i; \pi_i)$ the separable weighted Hilbert space of univariate square-integrable functions $f : X_i \rightarrow \mathbb{R}$ with inner product

$$f, g \in H_i, \quad \langle f, g \rangle := \int_{X_i} f(\theta) g(\theta) \pi_i(\theta) d\theta. \quad (2.13)$$

The 1D pseudo-spectral projection (PSP) operator is defined as an expansion of the form:

$$\mathcal{U}_{\mathbf{k}}^{\text{PSP}} : H_i \rightarrow \mathbb{P}_{P_k}, \quad \mathcal{U}_{\mathbf{k}}^{\text{PSP}}[f] := \sum_{p=0}^{P_k} c_p \Phi_p(\theta), \quad (2.14)$$

where $\{\Phi_p\}_{p=0}^{P_k}$ are polynomials orthogonal w.r.t. the inner product (2.13), i.e., $\langle \Phi_p, \Phi_q \rangle = \gamma_p \delta_{pq}$. In UQ, approximations such as (2.14) in which the polynomial basis is orthogonal w.r.t. the input density are called *generalized polynomial chaos* approximations [201]. For example, if π_i is the standard uniform density defined on $X_i = [-1, 1]$, the associated polynomials are the Legendre polynomials, whereas the Hermite polynomials are associated to the standard Gaussian density with $X_i = \mathbb{R}$. In this work, we refer to approximation in (2.14) as PSP. To simplify the notation, in the following we assume that the orthogonal basis in (2.14) is normalized, i.e., $\Phi_p := \Phi_p / \sqrt{\gamma_p}$, that is, the polynomial bases are *orthonormal*.

The PSP coefficients c_p are found by imposing the residual of (2.14) to be orthonormal to the space spanned by the polynomial basis (see, e.g., [193, 201]), i.e.,

$$\left\langle f - \sum_{p'=0}^{P_k} c_{p'} \Phi_{p'}, \Phi_p \right\rangle = 0 \Leftrightarrow c_p = \langle f, \Phi_p \rangle := \int_{X_i} f(\theta) \Phi_p(\theta) \pi_i(\theta) d\theta =: \mathbb{E}[f \Phi_p]. \quad (2.15)$$

Since (2.15) can rarely be evaluated analytically, we typically resort to numerical quadrature for an approximation, i.e., $c_p = \mathcal{U}_k^{\text{qu}}[f\Phi_p]$.

Assume that the quadrature scheme $\mathcal{U}_k^{\text{qu}}[\cdot]$ has degree-of-exactness $DE_{m(k)}$. In general, the PSP basis maximum degree, P_k , and the number of quadrature nodes and weights, $m(k)$, used to assess the PSP coefficients are chosen independently. However, Constantine, Eldred and Phipps observed in [37], in numerical experiments involving Gaussian quadrature that unless the maximum PSP basis degree is chosen depending on the degree-of-exactness of $\mathcal{U}_k^{\text{qu}}[\cdot]$, detrimental errors ensue. These errors are essentially due to the inability of the quadrature scheme to integrate exactly the higher-order basis polynomials. Congrad and Marzouk extended this result to arbitrary quadrature schemes in [35] and proved that choosing P_k independent of $DE_{m(k)}$ can lead to an $O(1)$ quadrature error (called *internal aliasing error*) which can have drastic effects on the accuracy of PSP. It was shown that P_k has to be chosen depending on $\mathcal{W}^2(\mathcal{U}_k^{\text{qu}}[\cdot])$, i.e., the half accuracy approximation set of the employed quadrature rule. Specifically, P_k should be such that the *product* of two PSP basis polynomials, which appears in the definition of the inner product, is integrated exactly by $\mathcal{U}_k^{\text{qu}}[\cdot]$, i.e.,

$$\langle \Phi_p, \Phi_q \rangle \stackrel{!}{=} \mathcal{U}_k^{\text{qu}}[\Phi_p \Phi_q] = \delta_{pq}, \quad p, q = 0, 1, \dots, P_k. \quad (2.16)$$

With this in mind, we follow [35, 37] and choose P_k in (2.14) such that (2.16) is satisfied, i.e.,

$$P_k = \lfloor DE_{m(k)}/2 \rfloor. \quad (2.17)$$

If standard Gaussian quadrature is used, i.e., $DE_{m(k)} = 2m(k) - 1$, then $P_k = m(k) - 1$, whereas when quadrature schemes having $DE_{m(k)} = m(k) - 1$ are employed, then $P_k = \lfloor (m(k) - 1)/2 \rfloor$.

The extension to d_{sto} directions is straightforward. Let $\mathbf{H} = \bigotimes_{i=1}^{d_{\text{sto}}} H_i$ denote the multivariate separable Hilbert space of square integrable functions and assume that $\mathcal{F}_h : \mathbf{H} \rightarrow \mathbb{R}$ is such a function. Further, let $\mathbb{P}_{\mathbf{P}_k}$ be the (vector) space of multivariate polynomials of degree $\mathbf{P}_k := (P_{k_1}, P_{k_2}, \dots, P_{k_{d_{\text{sto}}}})$ for a given multiindex $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}})$.

Remark: In UQ settings, $\mathcal{F}_h \in \mathbf{H}$ is equivalent to saying that $\text{Var}[\mathcal{F}_h] < \infty$, i.e., the variance of \mathcal{F}_h is finite. This assumption is valid for all problems considered in this work.

Let $\mathbf{p} = (p_1, p_2, \dots, p_{d_{\text{sto}}})$. We have that

$$\mathcal{U}_{\mathbf{k}}^{\text{PSP}} : \mathbf{H} \rightarrow \mathbb{P}_{\mathbf{P}_k}, \quad \mathcal{U}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h] := \sum_{\mathbf{p}=\mathbf{0}}^{\mathbf{P}_k} c_{\mathbf{p}} \Phi_{\mathbf{p}}(\boldsymbol{\theta}), \quad (2.18)$$

where $\Phi_{\mathbf{p}}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}$ with $\langle \Phi_{\mathbf{p}}(\boldsymbol{\theta}), \Phi_{\mathbf{q}}(\boldsymbol{\theta}) \rangle = \delta_{\mathbf{p}\mathbf{q}} = \prod_{i=1}^{d_{\text{sto}}} \delta_{p_i q_i}$. Moreover,

$$c_{\mathbf{p}} = \langle \mathcal{F}_h, \Phi_{\mathbf{p}} \rangle := \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \Phi_{\mathbf{p}}(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} =: \mathbb{E}[\mathcal{F}_h \Phi_{\mathbf{p}}]. \quad (2.19)$$

To evaluate the PSP coefficients in (2.19), we resort to quadrature, i.e., $c_{\mathbf{p}} = \mathcal{U}_{\mathbf{k}}^{\text{qu}}[\mathcal{F}_h \Phi_{\mathbf{p}}]$. Extending (2.17) to d_{sto} dimensions yields a multivariate PSP degree

$$\mathbf{P}_k = (\lfloor DE_{m(k_1)}/2 \rfloor, \lfloor DE_{m(k_2)}/2 \rfloor, \dots, \lfloor DE_{m(k_{d_{\text{sto}}})}/2 \rfloor).$$

One of the advantages of using PSP is that after the expansion coefficients were assessed, the entire postprocessing can be performed depending solely on them. Thus, [201] showed that $\mathbb{E}[\mathcal{F}_h]$ and $\text{Var}[\mathcal{F}_h]$ can be estimated as

$$\hat{\mathbb{E}}[\mathcal{F}_h] = c_{\mathbf{0}} \quad (2.20)$$

$$\widehat{\text{Var}}[\mathcal{F}_h] = \sum_{\mathbf{p}:|\mathbf{p}|_1 \geq 1}^{P_{\mathbf{k}}} c_{\mathbf{p}}^2, \quad \widehat{\text{Std}}[\mathcal{F}_h] = \sqrt{\sum_{\mathbf{p}:|\mathbf{p}|_1 \geq 1}^{P_{\mathbf{k}}} c_{\mathbf{p}}^2}. \quad (2.21)$$

Additionally, in Section 2.9, we show that the PSP coefficients can be also used to estimate Sobol' indices for global sensitivity analysis.

Interpolation with global Lagrange polynomials

Let $C^0(X_i)$ denote the space of continuous univariate functions $f : X_i \rightarrow \mathbb{R}$ and consider again \mathbb{P}_{P_k} , the vector space of univariate polynomials of degree $P_k \in \mathbb{N}$. A popular interpolation strategy in UQ is Lagrange interpolation (see, e.g., [60, 85, 134, 176]), which we use in this work. Note that in the UQ literature, the interpolation of the forward model usually bears the name of *stochastic collocation*. The 1D Lagrange interpolation operator at level k reads

$$\mathcal{U}_k^{\text{in-g}} : C^0(X_i) \rightarrow \mathbb{P}_{P_k}, \quad \mathcal{U}_k^{\text{in-g}}[f] := \sum_{n=1}^{m(k)} f(\theta_n) L_n(\theta), \quad (2.22)$$

where $\{\theta_n\}_{n=1}^{m(k)}$ are interpolation knots and $L_n(\theta) := \prod_{0 \leq m \leq P_k, m \neq n} \frac{\theta - \theta_m}{\theta_n - \theta_m}$ are canonical Lagrange polynomials satisfying $L_n(\theta_m) = \delta_{nm}$. With $m(k)$ interpolation nodes in (2.22) we can represent polynomials of degrees in the range $0, 1, \dots, m(k) - 1$. Thus, $P_k = m(k) - 1$ and

$$\mathcal{U}_k^{\text{in-g}}[f] := \sum_{p=0}^{P_k} f(\theta_p) L_p(\theta).$$

Remark: Recall that when employing Gaussian quadrature for PSP, $P_k := m(k) - 1$, as for Lagrange interpolation. Thus, Lagrange interpolation with Gaussian nodes is equivalent to PSP plus Gaussian quadrature. This equivalence was exploited in [37] to show how to choose the PSP degrees such that the underlying PSP space is free of internal aliasing error.

For improved numerical stability, we implement the 1D Lagrange interpolation in terms of the so-called *first form of the barycentric formula* (see, e.g., [11]).

Let again $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}})$ be a multiindex and let $C^0(\mathbf{X})$ be the space of continuous d_{sto} -variate functions defined on \mathbf{X} . Assume that $\mathcal{F}_h \in C^0(\mathbf{X})$. Further, consider $\mathbb{P}_{\mathbf{P}_{\mathbf{k}}}$, the vector space of multivariate polynomials of degree $\mathbf{P}_{\mathbf{k}} := (m(k_1) - 1, m(k_2) - 1, \dots, m(k_{d_{\text{sto}}}) - 1)$. The d_{sto} -variate Lagrange interpolation operator associated to \mathbf{k} reads

$$\mathcal{U}_{\mathbf{k}}^{\text{in-g}} : C^0(\mathbf{X}) \rightarrow \mathbb{P}_{\mathbf{P}_{\mathbf{k}}}, \quad \mathcal{U}_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}_h] := \sum_{\mathbf{p}=0}^{\mathbf{P}_{\mathbf{k}}} \mathcal{F}_h(\boldsymbol{\theta}_{\mathbf{p}}) \mathbf{L}_{\mathbf{p}}(\boldsymbol{\theta}),$$

where $\mathbf{L}_{\mathbf{p}}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} L_{p_i}(\theta_i)$ is the multivariate Lagrange basis.

Remark: To be able to employ quadrature, we need integrable functions. PSP requires square-integrability, whereas interpolation in terms of Lagrange bases assumes continuity of the target function. These assumptions are necessary for the underlying theory to hold true. From a practical point-of-view, however, to be able to make use of these approximations in high(er)-dimensional settings, we need smoother solutions of the underlying model, and moreover, solutions with certain *structure*, such as anisotropic coupling of the inputs or a lower intrinsic stochastic dimensionality. Otherwise, as the dimensionality of the UQ problem increases, the accuracy of these approaches will deteriorate and will make them infeasible.

2.5.2 Weighted (L)-Leja sequences

In this section, we summarize weighted (L)-Leja sequences, which we use to formulate all dimension-adaptive sparse grid approximations in this work. Since these approximations are constructed using tensorizations of one dimensional operators, it is sufficient to show how Leja points are constructed in 1D w.r.t. the marginal density π_i for $i = 1, 2, \dots, d_{\text{sto}}$.

The popular point sets used for PSP, interpolation and quadrature are Gaussian points, such as Gauss-Legendre or Gauss-Hermite [199], and Clenshaw-Curtis points [32]. However, these points can only be used for specific probability densities. For example, Gauss-Hermite points are associated with Gaussian densities, and Gauss-Legendre and Clenshaw-Curtis points are used for uniform densities. Moreover, although all of them have excellent approximation properties, their number usually increases exponentially with the level. In addition, only some, e.g., Clenshaw-Curtis points, are nested (recall Section 2.5.1).

To this end, we want dimension-adaptive sparse grid approximations capable of handling arbitrary probability density functions. Moreover, we also want these approximations to be suitable for computationally expensive problems. Thus, we desire a point set that is

1. *nested*, so that we can reuse computations from previous levels;
2. *grows slowly with the level k* , so that the total number of grid points is not very large;
3. has *good approximation* properties.

A point set having all the aforementioned properties is the weighted (L)-Leja sequence (see e.g. [96, 134]) which we employ in this work.

The univariate weighted (L)-Leja sequences [96] are constructed recursively as follows:

$$\begin{aligned} \theta_1 &= \operatorname{argmax}_{\theta \in X_i} \pi_i(\theta), \\ \theta_n &= \operatorname{argmax}_{\theta \in X_i} \pi_i(\theta) \prod_{m=1}^{n-1} |(\theta - \theta_m)|, \quad n = 2, 3, \dots \end{aligned} \tag{2.23}$$

Note that this point sequence is in general not uniquely defined, because the optimization problem (2.23) might have multiple solutions. In that case, we simply pick one of the maximizers. We illustrate the weighted (L)-Leja construction (2.23) in Figure 2.4. On the left-hand side we depict the (L)-Leja points for the uniform density on $X_i = [0, 1]$ using $\theta_1 = 0.5$. In the right plot we depict the points for the standard Gaussian density ($X_i = \mathbb{R}$) using $\theta_1 = 0$.

The construction (2.23) indicates that weighted (L)-Leja points form an interpolatory sequence, i.e., only one extra point is needed to increase the interpolation degree by one. Thus, for interpolation, we employ a level-to-nodes mapping $m(k) = k$. Since (L)-Leja points are nested, the number of points between adjacent levels is $m(k) - m(k-1) = 1$, i.e., only one extra function evaluation is needed as we increase the level. For PSP and quadrature, on the other hand, we use $m(k) = 2k - 1$, i.e., we add two extra (L)-Leja points for all levels $k \geq 2$. This choice is motivated as follows: when the input density is symmetric, (2.23) places the first (L)-Leja point in the center of the domain, while the next two points are symmetric w.r.t. the first. In this way, adding only one (L)-Leja point at a time will lead to a zero quadrature weight at level $k = 2$, causing adaptive algorithms to stop prematurely.

In Figure 2.5, we compare one-dimensional (L)-Leja points for interpolation, i.e., growing as $m(k) = k$, constructed w.r.t. the standard uniform density $\pi_i = 1_{[0,1]}$, with linearly increasing Gauss-Legendre points, that is, $m(k) = 2k - 1$, defined for the same density function. On the left, we depict the Leja points, while the Gauss-Legendre points are visualized on the right,

2.5. APPROXIMATION WITH DIMENSION-ADAPTIVE SPARSE GRIDS




Figure 2.4: Weighted (L)-Leja sequence of points defined in (2.23). From top to bottom we depict the function $\pi_i(\theta) \prod_{m=1}^{n-1} |(\theta - \theta_m)|$ and the corresponding Leja points (dots) for $n = 2, 3, 4, 5$. On the left-hand side $X_i = [0, 1]$, $\pi_i = 1_{[0,1]}$ and $\theta_1 = 0.5$. In the right part $X_i = \mathbb{R}$, $\pi_i(\theta) = \frac{1}{\sqrt{2\pi}} \exp(-\theta^2/2)$ and $\theta_1 = 0$.




Figure 2.5: Left: One-dimensional (L)-Leja points for interpolation at levels $k = 3, 4$ and 5 for a uniform weight, i.e., $\pi_i = 1_{[0,1]}$. (L)-Leja points are nested and increase as $m(k) = k$. Right: One-dimensional Gauss-Legendre for the standard uniform density $\pi_i = 1_{[0,1]}$. These points increase linearly as $m(k) = 2k - 1$, but they are not nested (besides the point 0.5).

for levels $k = 3, 4$ and 5 . The (L)-Leja points are nested and have arbitrary granularity, which makes them an excellent choice for higher-dimensional approximations based on tensorizations. Moreover, as it was shown in, e.g., [134], these points also have an excellent approximation accuracy. In contrast, the linearly increasing Gauss-Legendre points, although having very good approximation properties as well, they are not nested (besides the point 0.5). Standard choices of point sets which are nested, such as Clenshaw-Curtis points, increase exponentially with the level, which can lead to high cardinalities if the dimensionality or the sparse grid level are large. Moreover, all standard point sets are restricted to specific density functions, whereas weight (L)-Leja points can be constructed for arbitrary weight functions. More details on (L)-Leja points and their properties can be found in [96, 134] and the references therein.

We now look at the number of (L)-Leja points in d_{sto} dimensions. Consider a multiindex $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}})$. Since this point sequence is *nested*, we are only interested in the number of *surplus* points corresponding to \mathbf{k} . Let this number be $\Delta L_{\mathbf{k}}^{\text{op}}$. Further, let $\delta m(k_i) = m(k_i) - m(k_{i-1})$ such that $\delta m(1) = 0$, and define $\Delta m(k_i)$ as

$$\Delta m(k_i) = \begin{cases} 1 & \text{if } m(k_i) = m(k_{i-1}) \text{ or } k_i = 1 \\ \delta m(k_i) & \text{else.} \end{cases}$$

Hence, if $m(k_i) = m(k_{i-1})$, then $\delta m(k_i) = 0$ and $\Delta m(k_i) = 1$. With this notation, it follows that $\Delta L_{\mathbf{k}}^{\text{op}} = \prod_{i=1}^{d_{\text{sto}}} \Delta m(k_i)$. On the one hand, for 1D interpolation, (L)-Leja points increase as $m(k) = k$, thus $m(k) - m(k-1) = 1$, from which we have that $\Delta m(k_i) = 1$ and

$$\Delta L_{\mathbf{k}}^{\text{in-g}} = 1$$

regardless of \mathbf{k} . For quadrature and PSP, on the other hand, $m(k) = 2k - 1$ and therefore $m(k) - m(k-1) = 2$ in 1D. Thus, $\Delta m(k_i)$ can take two values: $\Delta m(k_i) = 1 = 2^0$ or $\Delta m(k_i) = 2 = 2^1$. If we multiply these two values, we have $2 = 1 \cdot 2 = 2^{0+1} = 2^{(0+2)/2}$. We deduce that

$$\Delta L_{\mathbf{k}}^{\text{qu}} = \Delta L_{\mathbf{k}}^{\text{psp}} = 2^{(\sum_{i=1}^{d_{\text{sto}}} \delta m(k_i))/2}.$$

This implies that $1 \leq \Delta L_{\mathbf{k}}^{\text{qu}}, \Delta L_{\mathbf{k}}^{\text{psp}} \leq 2^{d_{\text{sto}}}$ since $2^{(\sum_{i=1}^{d_{\text{sto}}} 2)/2} = 2^{d_{\text{sto}}}$ when $k_i \neq 1$ for all $i = 1, 2, \dots, d_{\text{sto}}$. For example, if $d_{\text{sto}} = 2$ and $\mathbf{k} = (2, 2)$, then $L_{(2,2)}^{\text{qu}} = L_{(2,2)}^{\text{psp}} = 2^{(2+2)/2} = 4$, whereas if $d_{\text{sto}} = 5$ and $\mathbf{k} = (2, 1, 3, 4, 1)$, then $L_{(2,1,3,4,1)}^{\text{qu}} = L_{(2,1,3,4,1)}^{\text{psp}} = 2^{(2+0+2+2+0)/2} = 2^3 = 8$. To summarize, for a given multiindex $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}})$, (L)-Leja points for interpolation have the finest granularity $\Delta L_{\mathbf{k}}^{\text{in-g}} = 1$ regardless of multiindex \mathbf{k} , whereas for quadrature and PSP, we have that $1 \leq \Delta L_{\mathbf{k}}^{\text{qu}}, \Delta L_{\mathbf{k}}^{\text{psp}} \leq 2^{d_{\text{sto}}}$.

2.5.3 Subspace or dimension-adaptivity

To fully define (2.10) or (2.11), three ingredients are needed. First, we need the linear operators $\mathcal{U}^i[f]$, which we reviewed in Section 2.5.1. The choice of the operator depends on the task at hand. Further, we need a point set to construct the approximation $\mathcal{U}_{\mathbf{k}}^i[f]$; in Section 2.5.1 we summarized weighted (L)-Leja points, which we employ in this work. This second step is more critical for the overall computational cost, because, as we discussed, nested point sets are usually more cost-efficient for expensive problems. To fully specify a sparse grid approximation (2.9), we need the third and last ingredient, the finite multiindex set, \mathcal{K} . The choice \mathcal{K} has influence on both the overall cost and on *how* the structure of the underlying problem is exploited. By the structure of the underlying problem we especially refer to the *coupling* and *importance* of the underlying stochastic inputs. A context-aware construction of \mathcal{K} is at the core of the

2.5. APPROXIMATION WITH DIMENSION-ADAPTIVE SPARSE GRIDS

Sensitivity-driven dimension-adaptive sparse grid algorithm in Chapter 4 and of one of the adaptive strategies used in the Multilevel adaptive sparse Leja algorithm in Chapter 7.

Recall that \mathcal{K} must be admissible or downward closed. The standard, a priori choice reads

$$\mathcal{K} = \{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}} : |\mathbf{k}|_1 \leq K_{\text{max}}^{\text{op}} + d_{\text{sto}} - 1\} \quad (2.24)$$

for a user-defined level $K_{\text{max}}^{\text{op}} \in \mathbb{N}$; see Figures 2.6 and 2.7 for a visual depiction for $d_{\text{sto}} = 2$ and $K_{\text{max}}^{\text{op}} = 6$. This choice is based on a priori known properties of the underlying function, such as smoothness in terms of, for example, number of bounded derivatives [20, 181], or as it was pointed out in [179], alignment with the standard Cartesian axes. However, when these properties are not fully known beforehand or when the problem has a richer structure, such as anisotropy, which is typically the case in most problems (see, e.g., [35, 193]), constructing \mathcal{K} *dynamically* using a posteriori estimates is preferred.

To this end, [71, 90] proposed constructing the multiindex set \mathcal{K} via *dimension* or *subspace* adaptivity. In the following, we briefly summarize only the basic idea of the algorithm and refer the reader to, e.g., [35, 71, 90, 193] for more details. In dimension-adaptivity, \mathcal{K} is split in two sets, the *old index set*, \mathcal{O} , and the *active set*, \mathcal{A} such that $\mathcal{K} = \mathcal{O} \cup \mathcal{A}$ is *admissible* and \mathcal{A} is *\mathcal{O} -admissible*, i.e., $\forall \mathbf{k} \in \mathcal{A}, k_i \geq 2, \mathbf{k} - \mathbf{e}_i \in \mathcal{O}, i = 1, 2, \dots, d_{\text{sto}}$. The active set \mathcal{A} is the one that drives the adaptive process. In the first step, $\mathcal{O} = \{\mathbf{1}\}$, because the number of points corresponding to $\mathbf{1}$ is $1^{d_{\text{sto}}} = 1$, and $\mathcal{A} = \{\mathbf{1} + \mathbf{e}_i, i = 1, 2, \dots, d_{\text{sto}}\}$. In the remaining steps, the algorithm employs the following principle: if the subspace associated to a multiindex $\mathbf{k} \in \mathcal{A}$ is significant for the current approximation, its adjacent neighbours are likely to be significant as well. To this end, in each refinement step, a refinement indicator $\epsilon(\mathbf{k})$ is computed for each multiindex $\mathbf{k} \in \mathcal{A}$. The multiindex with the largest refinement indicator is moved to \mathcal{O} . Then, all its forward neighbours that preserve the admissibility of \mathcal{K} are added to \mathcal{A} . In addition, a surrogate for the global error indicator $\rho = \sum_{\mathbf{k} \in \mathcal{A}} \epsilon(\mathbf{k})$ is computed at each refinement step. The algorithm typically stops if $\rho < \text{tol}^{\text{op}}$ for a user-defined tolerance tol^{op} , if $\mathcal{A} = \emptyset$ or if a user-defined maximum level $K_{\text{max}}^{\text{op}}$ is reached.

The essential ingredient in the described algorithm is the refinement indicator, $\epsilon(\cdot)$. A standard choice (see [35, 71] for examples) is of the form

$$\epsilon(\mathbf{k}) := \mathcal{I}(\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h], \mathcal{C}_{\mathbf{k}}^{\text{op}}), \quad (2.25)$$

where \mathcal{I} is a function depending on the multivariate hierarchical surplus and on the cost $\mathcal{C}_{\mathbf{k}}^{\text{op}}$ of assessing $\Delta_{\mathbf{k}}^{\text{op}}$, typically the number of grid points corresponding to $\Delta_{\mathbf{k}}^{\text{op}}$.

We summarize standard dimension-adaptive refinement in Algorithm 2.1. The inputs are the user-defined tolerance tol^{op} and the maximum reachable level, $K_{\text{max}}^{\text{op}}$, the forward model \mathcal{F}_h w.r.t. which we perform the approximation and the multivariate density characterizing the stochastic input, $\boldsymbol{\pi}$. Since this algorithm is employed in stochastic computations, everything is computed w.r.t. $\boldsymbol{\pi}$. At the end, the multiindex set, $\mathcal{K} = \mathcal{O} \cup \mathcal{A}$ is determined (line 20) and the sparse grid approximation (2.10) is returned (line 21). We illustrate the difference between static, a priori chosen multiindex sets \mathcal{K} and multiindex sets determined dynamically via *dimension-adaptivity* in the following example.

Example: Consider two functions, $\mathcal{F}, \mathcal{G} : [0, 1]^2 \rightarrow \mathbb{R}$ that depend on a two-dimensional uniformly distributed input in $[0, 1]^2$. For simplicity, we employ only sparse grid interpolation.

We compare a priori chosen multiindex sets with multiindex sets found via dimension-adaptivity. For dimension adaptivity, we consider the refinement indicator [35]

$$\mathcal{I}(\Delta_{\mathbf{k}}^{\text{in-g}}[\cdot], \mathcal{C}_{\mathbf{k}}^{\text{in-g}}) = \left\| \Delta_{\mathbf{k}}^{\text{in-g}}[\cdot] \right\|_{L^2} / \Delta L_{\mathbf{k}}^{\text{in-g}} = \left\| \Delta_{\mathbf{k}}^{\text{in-g}}[\cdot] \right\|_{L^2},$$

Algorithm 2.1 Standard dimension-adaptive sparse grid approximation

```

1: procedure STDADAPTIVITY( $tol^{op}$ ,  $K_{\max}^{op}$ ,  $\mathcal{F}_h$ ,  $\boldsymbol{\pi}$ )
2:    $\mathbf{1} = (1, 1, \dots, 1)$ ,  $\mathcal{O} = \emptyset$ ,  $\mathcal{A} = \{\mathbf{1}\}$ 
3:    $a = \boldsymbol{\Delta}_{\mathbf{1}}^{op}[\mathcal{F}_h]$ 
4:   Compute  $\epsilon(\mathbf{1}) = \mathcal{I}(\boldsymbol{\Delta}_{\mathbf{1}}^{op}[\mathcal{F}_h], \mathcal{C}_{\mathbf{1}}^{op})$  based on  $\boldsymbol{\pi}(\boldsymbol{\theta})$ 
5:    $\rho = \epsilon(\mathbf{1})$ 
6:   while  $\rho \geq tol^{op}$  or  $\mathcal{A} \neq \emptyset$  or  $\max(\mathcal{K}) \leq K_{\max}^{op}$  do
7:     Select  $\mathbf{k}$  from  $\mathcal{A}$  with the largest  $\epsilon(\mathbf{k})$ 
8:      $\mathcal{A} = \mathcal{A} \setminus \{\mathbf{k}\}$ ,  $\mathcal{O} = \mathcal{O} \cup \{\mathbf{k}\}$ 
9:      $\rho = \rho - \epsilon(\mathbf{k})$ 
10:    for  $i \leftarrow 1, 2, \dots, d_{sto}$  do
11:       $\mathbf{r} = \mathbf{k} + \mathbf{e}_i$ 
12:      if  $\mathbf{r} - \mathbf{e}_q \in \mathcal{O}$  for all  $q = 1, 2, \dots, d_{sto}$  then
13:         $\mathcal{A} = \mathcal{A} \cup \{\mathbf{r}\}$ 
14:         $a = a + \boldsymbol{\Delta}_{\mathbf{r}}^{op}[\mathcal{F}_h]$ 
15:        Compute  $\epsilon(\mathbf{r}) = \mathcal{I}(\boldsymbol{\Delta}_{\mathbf{r}}^{op}[\mathcal{F}_h], \mathcal{C}_{\mathbf{r}}^{op})$  based on  $\boldsymbol{\pi}(\boldsymbol{\theta})$ 
16:         $\rho = \rho + \epsilon(\mathbf{r})$ 
17:      end if
18:    end for
19:  end while
20:   $\mathcal{K} = \mathcal{O} \cup \mathcal{A}$ 
21:  return  $a$ 
22: end procedure

```

where the last inequality is due to $\Delta L_{\mathbf{k}}^{\text{in-g}} = 1$ for interpolation. Let us denote by $\mathcal{K}_n^{\text{std}}$ the a priori chosen multiindex set and by $\mathcal{K}_n^{\text{adapt}}$ the one obtained at the end of adaptive refinement, where $n = 1$ for \mathcal{F} and $n = 2$ for \mathcal{G} .

We consider

$$\mathcal{F}(\boldsymbol{\theta}) = \sin(4\theta_1) + \sin\left(\frac{3}{2}\theta_2\right) \quad (2.26)$$

that is, a function with an additive structure in which the first input has a larger weight than the second input. We define \mathcal{G} to have a complementary structure, i.e., the stochastic input $\boldsymbol{\theta}$ is additive, having the same weights as in the definition of \mathcal{F} , i.e.,

$$\mathcal{G}(\boldsymbol{\theta}) = \sin\left(4\theta_1 + \frac{3}{2}\theta_2\right). \quad (2.27)$$

Note that both functions are in $C^\infty([0, 1]^2)$, thus interpolation can be used to accurately approximate them.

We depict the results for \mathcal{F} in Figure 2.6 and the results corresponding to \mathcal{G} in Figure 2.7. In the right plot in each figure, we depict the underlying function. In the left figures, we depict $\mathcal{K}_{1,2}^{\text{adapt}}$ and $\mathcal{K}_{1,2}^{\text{std}}$ as follows. The black squares represent the subspaces associated to the multiindices $\mathbf{k} \in \mathcal{K}_{1,2}^{\text{adapt}}$, whereas the black together with the red squares comprise the subspaces associated to $\mathcal{K}_{1,2}^{\text{std}}$ of level $K_{\max}^{\text{in-g}}$ equal to the maximum level reached by the adaptive algorithm; in both examples, $K_{\max}^{\text{in-g}} = 6$. Moreover, in each square, we show the numerical estimate of $\left\| \boldsymbol{\Delta}_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}] \right\|_{L^2}^2$. We see that for the additive function \mathcal{F} , the adaptive algorithm detects that the two stochastic directions are *independent* due to the additive form of \mathcal{F} and thus it invests effort only in the subspaces associated to each individual directions. That is, subspaces associated

2.5. APPROXIMATION WITH DIMENSION-ADAPTIVE SPARSE GRIDS

to multiindices $(1, k_2)$ for the first direction and $(k_1, 1)$ for the second direction; the multiindex $(2, 2)$ is also added because it is always within the forward neighbour set of the first refined multiindex (recall the description of the dimension-adaptive algorithm). We see that the values of $\left\| \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}] \right\|_{L^2}^2$ are large in subspaces associated to $\mathcal{K}_1^{\text{std}}$. In the remaining subspaces which complete $\mathcal{K}_1^{\text{std}}$, $\left\| \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}] \right\|_{L^2}^2 = 0$ because here there is no contribution to the approximation of \mathcal{F} . Thus, a standard sparse grid approach employs many unnecessary subspaces, whereas the greedy adaptive approach better exploits the structure of the underlying problem. The structure of \mathcal{G} is more complex because it is neither additive nor multiplicative; see Figure 2.7, right. This is reflected in the left plot in Figure 2.7. We see, however, that in this case, too, although the sparse grid approximation of \mathcal{G} is more involved, dimension-adaptivity better exploits the anisotropic coupling of the stochastic inputs. Specifically, it detects that the first input, θ_1 , is more important than θ_2 since it has a larger multiplication weight. Moreover, since the interaction between the two stochastic inputs is nonnegligible, the dimension-adaptive algorithm also adds several subspaces with both indices different from one.

In this example, we saw that dimension-adaptivity can exploit the anisotropic coupling of stochastic inputs to create a multiindex set that reflects this anisotropy. When the problem has a simple structure, such as \mathcal{F} , the adaptive construction of the multiindex set can lead to significant cost reduction as compared to a static choice of \mathcal{K} . Even when the problem has a more complex structure, as \mathcal{G} does, the adaptive algorithm can exploit, for example, the anisotropic coupling of the input parameters, which is usually characteristic to most problems. Of course, the adaptive algorithm is based on heuristics and it is not guaranteed that it will always be beneficial. However, in a very large number of numerical test cases (see [35, 71, 90, 134, 193] and the references therein) it was shown that dimension-adaptivity leads to cost reduction.




Figure 2.6: $\mathcal{F}(\theta) = \sin(4\theta_1) + \sin(\frac{3}{2}\theta_2)$. In the left plot, we show the adaptive multiindex set, $\mathcal{K}_1^{\text{adapt}}$ using black squares, and the associated a priori chosen $\mathcal{K}_1^{\text{std}}$ of level $K_{\text{max}}^{\text{in-g}} = 6$ (back and red squares). Moreover, we also depict the value of $\left\| \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}] \right\|_{L^2}^2$ in each subspace. In the right figure, we depict \mathcal{F} .




Figure 2.7: $\mathcal{G}(\theta) = \sin(4\theta_1 + \frac{3}{2}\theta_2)$. In the left plot, we show the adaptive multiindex set, $\mathcal{K}_2^{\text{adapt}}$ using black squares, and the associated a priori chosen $\mathcal{K}_2^{\text{std}}$ of level $K_{\text{max}}^{\text{in-g}} = 6$ (back and red squares). Moreover, we also depict the value of $\|\Delta_k^{\text{in-g}}[\mathcal{G}]\|_{L^2}^2$ in each subspace. In the right figure, we depict \mathcal{G} .

2.6 Approximation with spatially adaptive sparse grids

In Section 2.5, we summarized sparse grid approximations with dimension adaptivity, which are defined depending on global operators. In this section, we present another type of sparse grid approximations in terms of *hierarchical basis functions with local support*. For hierarchical sparse grids, it is natural to perform *local* or *spatial* adaptivity, thus they are called *spatially adaptive sparse grid approximations*. We employ them in the second proposed algorithm, the Multilevel adaptive spectral projection with online dimensionality reduction in Chapter 5. In Section 2.6.1, we review interpolation with basis functions with local support, such as linear hat functions or piecewise polynomials. To distinguish this approximation from the ones in Section 2.5.1, we employ the superscript *in-l*, where we use *l* to refer to *local*. We summarize the considered refinement strategies in Section 2.6.2. Our notation is similar to [20, 147].

2.6.1 Interpolation with hierarchical bases with local support

Interpolation with hierarchical bases is typically defined for functions defined on the unit hypercube, $[0, 1]^{d_{\text{sto}}}$. We outline first the setup for the 1D domain. Let $k, v \in \mathbb{N}$ denote the level and spatial position, respectively. The starting point is a grid of Newton-Cotes (dyadic) nodes $u_{k,v} = v h_k \in [0, 1]$, $h_k := 2^{-k}$, $v \geq 1$, and standard linear hat basis functions $\varphi_{k,v}(u)$ centered at $u_{k,v}$, with support $[u_{k,v} - h_k, u_{k,v} + h_k]$, $\varphi_{k,v}(u) = \varphi(2^k u - v)$, where $\varphi(u) = \max(1 - |u|, 0)$. Note that this construction leads to no boundary points. Furthermore, if $v = 1$, the number of grid points is $m(k) = 1$, whereas if $v > 1$, we have $m(k) = 2^{v-1} + 1$. By construction, the Newton-Cotes nodes are *nested*. We visualize 1D piecewise linear (hat) functions at levels $k = 1, 2, 3$ in the left plot in Figure 2.8.

Remark: The Newton-Cotes nodes defined above increase geometrically with the level k . In contrast, the weighted (L)-Leja points defined in Section 2.5.2 increase linearly, i.e., $m(k) = k$

2.6. APPROXIMATION WITH SPATIALLY ADAPTIVE SPARSE GRIDS

for interpolation and $m(k) = 2k - 1$ for quadrature and PSP. However, weighted (L)-Leja are globally defined and thus employed in dimension-adaptivity, whereas the Newton-Cotes allow for local adaptivity, as we will see next.




Figure 2.8: Standard (left) and modified (right) piecewise linear (hat) basis functions at levels $k = 1, 2, 3$. The standard basis functions are used for Newton-Cotes grids with boundary points, whereas the modified basis are used for grids with only interior points.

The extension to d_{sto} -dimensions is done via *tensorization*. This hence leads to

$$\varphi_{\mathbf{k}, \mathbf{v}}(\mathbf{u}) = \prod_{i=1}^{d_{\text{sto}}} \varphi_{k_i, v_i}(u_i),$$

where $\mathbf{k} = (k_1, k_2, \dots, k_{d_{\text{sto}}}) \in \mathbb{N}^{d_{\text{sto}}}$ and $\mathbf{v} = (v_1, v_2, \dots, v_{d_{\text{sto}}}) \in \mathbb{N}^{d_{\text{sto}}}$ are multiindices. Furthermore, let $W_{\mathbf{k}} = \text{span}\{\varphi_{\mathbf{k}, \mathbf{v}} | \mathbf{v} \in \Upsilon_{\mathbf{k}}\}$ denote the *hierarchical increment* space, where $\Upsilon_{\mathbf{k}} = \{\mathbf{v} \in \mathbb{N}^{d_{\text{sto}}} : 1 \leq v_i \leq 2^{k_i} - 1, v_i \text{ odd}, i = 1, 2, \dots, d_{\text{sto}}\}$. Given a properly defined multiindex set \mathcal{K} , the sparse grid interpolant $\mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h]$ of the high-fidelity model \mathcal{F}_h reads

$$\mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h] = \sum_{\mathbf{k} \in \mathcal{K}, \mathbf{v} \in \Upsilon_{\mathbf{k}}} \alpha_{\mathbf{k}, \mathbf{v}} \varphi_{\mathbf{k}, \mathbf{v}}(\mathbf{u}), \quad (2.28)$$

where $\alpha_{\mathbf{k}, \mathbf{v}}$ are the so-called *hierarchical surpluses*, computed as

$$\alpha_{\mathbf{k}, \mathbf{v}} = \mathcal{F}_h(\mathbf{u}_{\mathbf{k}, \mathbf{v}}) - \mathcal{U}_{\mathcal{K} \setminus \{\mathbf{k}\}}^{\text{in},1}[\mathcal{F}_h(\mathbf{u}_{\mathbf{k}, \mathbf{v}})], \quad (2.29)$$

where $\mathbf{u}_{\mathbf{k}, \mathbf{v}} := (u_{k_1, v_1}, u_{k_2, v_2}, \dots, u_{k_{d_{\text{sto}}}, v_{d_{\text{sto}}}})$. Note that (2.28) is defined in terms of Newton-Cotes grids with interior points only. Thus, when using standard basis functions, for (2.28) to hold, we need \mathcal{F}_h to vanish at the boundary of $[0, 1]^{d_{\text{sto}}}$. Recall that the standard choice of the multiindex set \mathcal{K} is (2.24), i.e.,

$$\mathcal{K} := \{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}} : |\mathbf{k}|_1 \leq K_{\text{max}}^{\text{in},1} + d_{\text{sto}} - 1\},$$

for $K_{\text{max}}^{\text{in},1} \in \mathbb{N}$, which leads to the standard sparse grid space $V_{K_{\text{max}}^{\text{in},1}} := \bigotimes_{\mathbf{k} \in \mathcal{K}} W_{\mathbf{k}}$ with mesh width $h_{\text{in},1} := 2^{-K_{\text{max}}^{\text{in},1}}$. In [20], it was shown that if \mathcal{F}_h has *bounded mixed derivatives*, i.e.,

$$\mathcal{F}_h \in H_2^{\text{mix}}([0, 1]^{d_{\text{sto}}}) := \{f : [0, 1]^{d_{\text{sto}}} \rightarrow \mathbb{R} : D^{\mathbf{k}} f \in L^2([0, 1]^{d_{\text{sto}}}), |\mathbf{k}|_{\infty} \leq 2, f_{\partial[0, 1]^{d_{\text{sto}}}} = 0\},$$

where the mixed derivative $D^{\mathbf{k}}f := \partial^{|\mathbf{k}|}f / \partial u_1^{k_1} \dots \partial u_{d_{\text{sto}}}^{k_{d_{\text{sto}}}}$, then we have an approximation error $\|\mathcal{F}_h - \mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h]\|_{L^2} \in O\left(h^2 (K_{\text{max}}^{\text{in},1})^{d_{\text{sto}}-1}\right)$. Moreover, the cost in terms of number of grid points is in $O\left(2^{K_{\text{max}}^{\text{in},1}} (K_{\text{max}}^{\text{in},1})^{d_{\text{sto}}-1}\right)$. In contrast, when using a full grid interpolant the approximation error is in $O\left(h_{\text{in},1}^2\right)$, with an associated cost in $O\left(2^{K_{\text{max}}^{\text{in},1} d_{\text{sto}}}\right)$. We see therefore that (2.28) leads to an approximation that has similar accuracy as a full grid approximation, but with a significantly reduced total number of grid points.

If the underlying high-fidelity model solution has more smoothness, i.e., bounded higher-order mixed derivatives, [19] showed that the presented methodology can be extended to using higher-order piecewise polynomial bases. That is, if $\mathcal{F}_h \in H_{p+1}^{\text{mix}}([0, 1]^{d_{\text{sto}}})$ for some $p \geq 2$, then $\|\mathcal{F}_h - \mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h]\|_{L^2} \in O\left(h_{\text{in},1}^{p+1} (K_{\text{max}}^{\text{in},1})^{d_{\text{sto}}-1}\right)$. In the numerical experiments in Chapter 5, we consider polynomial basis functions of degree two. We depict polynomial bases of degree $p = 2$ for levels $k = 1, 2, 3$, in Figure 2.9.




Figure 2.9: Standard (left) and modified (right) piecewise polynomial basis functions of degree $p = 2$ at levels $k = 1, 2, 3$. The standard basis functions are used for grids comprising boundary points, whereas the modified basis are used for grids with only interior points.

Remark: We saw that the accuracy of the interpolation approximation (2.28) depends on the smoothness of the underlying high-fidelity model, as measured in terms of bounded mixed derivatives. Geometrically, as it was discussed, for example, in [179], this means that sparse grid interpolation is suitable with functions *aligned* with the standard Cartesian system (see [179] for a more detailed discussion; see also the discussion from Chapter 7, Section 7.2.3).

We simplify the notation by writing (2.28) more compactly as

$$\mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h] = \sum_{\mathbf{k} \in \mathcal{K}} \alpha_{\mathbf{k}} \varphi_{\mathbf{k}}(\mathbf{u}), \quad (2.30)$$

since $\Upsilon_{\mathbf{k}}$ in (2.28) depends on the multiindex \mathbf{k} as well.

To make (2.30) useful in arbitrary settings, we need two more ingredients. First, since the presented construction leads to no Newton-Cotes points on the boundary of $[0, 1]^{d_{\text{sto}}}$, we need to

extend it to incorporate boundary information, too. Our solution is to follow [147] and modify the standard basis functions so as to linearly extrapolate the boundary information. Modified linear and polynomial bases of degree $p = 2$ can be visualized in the right plots in Figures 2.8 and 2.9. Second, for stochastic domains \mathbf{X} different than $[0, 1]^{d_{\text{sto}}}$, we need a transformation

$$\mathbf{T} : [0, 1]^{d_{\text{sto}}} \rightarrow \mathbf{X}$$

with which the surpluses (2.29) are computed as

$$\alpha_{\mathbf{k}} = \mathcal{F}_h(\mathbf{T}(\mathbf{u}_{\mathbf{k}})) - \mathcal{U}_{\mathcal{K} \setminus \{\mathbf{k}\}}^{\text{in},1}[\mathcal{F}_h(\mathbf{T}(\mathbf{u}_{\mathbf{k},v}))]. \quad (2.31)$$

We provide examples of such transformations in Chapter 5. In this section, we defined interpolation with hierarchical bases with local support defined in terms of Newton-Codes nodes. To fully define this approximation in d_{sto} dimensions, we need to specify the multiindex set \mathcal{K} . We do this via spatially adaptive refinement, which we summarize next.

2.6.2 Local or spatial adaptivity

To address the computational challenges of interpolating computationally expensive forward models, we employ local or spatially adaptive refinement (see [120, 147]), an intrinsic property of sparse grid interpolation in terms of hierarchical bases with local support. Hierarchical surpluses (2.31) are indicators of local interpolation error. Hence, we employ refinement criteria based on surplus values, which allow us to exploit the *local* structure of the underlying problem. For a broader overview, we consider two refinement strategies. Since in this work, our focus is not on developing novel spatially adaptive algorithms, we only briefly summarize the two considered criteria and refer the reader to [120, 147] for more details and discussions.

Let $\bar{\mathcal{K}} = \{\mathbf{k} : \mathbf{u}_{\mathbf{k}} \text{ is refinable}\}$ be the set containing all levels corresponding to grid points that can be refined. The first refinement criterion employs the *maximum absolute value* of the local interpolation error, which is given by the *hierarchical surplus* (2.31), i.e.,

$$\max_{\mathbf{k} \in \bar{\mathcal{K}}} |\alpha_{\mathbf{k}}|. \quad (2.32)$$

The second criterion, used, for example, in [62], employs additionally the *maximum volume* contained by the *hierarchical basis functions*, i.e.,

$$\max_{\mathbf{k} \in \bar{\mathcal{K}}} |\alpha_{\mathbf{k}} \mathbb{E}[\varphi_{\mathbf{k}}]|,$$

where $\mathbb{E}[\varphi_{\mathbf{k}}] := \int_{[0,1]^{d_{\text{sto}}}} \varphi_{\mathbf{k}}(\mathbf{u}) d\mathbf{u}$. For example, if $\varphi_{\mathbf{k}}$ are hat functions, then $\mathbb{E}[\varphi_{\mathbf{k}}] := 2^{-|\mathbf{k}|_1}$. Note that if not all hierarchical parents exist in the refined grid, we ensure that they are added as well (see [147]).

Remark: Local adaptivity can be used also for functions which do not satisfy the regularity assumptions of the standard, static formulation. For example, spatially adaptive sparse grid interpolation can be employed for discontinuous functions or functions with sharp gradients. Of course, accurate approximations of such functions might inquire a possibly large computational cost. Moreover, the accuracy rates known for functions satisfying the standard regularity assumptions will likely deteriorate.

In this work, we use spatial adaptivity as follows. We begin with a standard grid of level K_{init} . Since we employ modified basis functions, to ensure that we obtain model evaluations in the interior of domain, we usually begin with $K_{\text{init}} = 3$ (see Figures 2.8 and 2.9). We then prescribe a list of three parameters $r := (N_{\text{adapt}}; \text{ref_perc}; \text{ref_pol})$, where N_{adapt} represents the

percentage ref_perc of the grid points to be refined locally using the refinement policy ref_pol . Thus, in a refinement step, we refine N_{adapt} times ref_perc of the current grid points with the *largest* refinement indicators using refinement policy ref_pol . We illustrate spatial adaptivity based on the maximum absolute hierarchical surplus in the following example.

Example: In this example, we illustrate spatial adaptivity using ref_pol based on the maximum absolute hierarchical surplus (2.32). As bases, we employ modified piecewise linear hat functions. We consider the same functions as in the example to illustrate dimension-adaptivity in Section 2.5.3, \mathcal{F} , defined in (2.26), and \mathcal{G} in (2.27). We therefore let again \mathcal{F} Note that since the domain \mathbf{X} of the two functions is $[0, 1]^2$, the transformation \mathbf{T} is the identity.

We set $K_{init} = 3$ and refine only the grid point with the largest indicator. To illustrate the local structure of the two functions, we prescribe $N_{adapt} = 50$ for \mathcal{F} , which is a simple, additive function, and $N_{adapt} = 800$ for \mathcal{G} , which has a more complex structure. We depict the results for \mathcal{F} in Figure 2.10 and the results corresponding to \mathcal{G} in Figure 2.11. In the left figures, we depict the adapted sparse grids. We see that the local structure of the two functions is reflected in the two sparse grids. In Figure 2.10, left, we see a (quasi-)symmetric grid. which resembles the surface of \mathcal{F} . Moreover, due to the simple structure of this function, we have only points with coordinates in $\{0.25, 0.5, 0.75\}$. The grid corresponding to \mathcal{G} is depicted in the left plot in Figure 2.11. Observe that, as expected, the local adaptive algorithm refined more in the upper right part of the domain, where the \mathcal{G} has a more complex structure.

This example underlines how spatially adaptive refinement can be used to exploit local structure of functions. Moreover, local adaptivity can be used even for discontinuous functions or functions with sharp gradients, but at a possibly nonnegligible computational cost.




Figure 2.10: $\mathcal{F}(\boldsymbol{\theta}) = \sin(4\theta_1) + \sin(\frac{3}{2}\theta_2)$. In the left plot, we show the resulting sparse grid after 20 adaptive steps in which we refine the grid point with the largest absolute hierarchical surplus. In the right figure, we depict the surface of \mathcal{F} .

2.7 Multilevel estimation of quantities of interest

In Sections 2.5 and 2.6, we summarized approximation with global, dimension-adaptive sparse grids or local, spatially adaptive grids. In this section, we introduce the second class of approximations relevant in this work, the *multilevel estimation of quantities of interest*. Standard




Figure 2.11: $\mathcal{G}(\boldsymbol{\theta}) = \sin(4\theta_1 + \frac{3}{2}\theta_2)$. In the left plot, we show the resulting sparse grid after 20 adaptive steps in which we refine the grid point with the largest absolute hierarchical surplus. In the right figure, we depict the surface of \mathcal{G} .

multilevel decompositions based on telescoping sums are used to formulate our Multilevel adaptive spectral projection with online dimensionality reduction approach in Chapter 5. Moreover, the Multilevel adaptive sparse Leja algorithm proposed in Chapter 7 is based on a (nonstandard) multilevel decomposition as well.

Multilevel decompositions are among the most popular methods in UQ. Inspired by multigrid solvers, Heinrich, in [91] and Giles, in [74] proposed the Multilevel MC (MLMC) sampling algorithm to decrease the variance of standard MC sampling; see also [6, 33, 34, 42, 75]. Multilevel decompositions were also used in deterministic quadrature in e.g., [81]. Furthermore, the works [176, 187] employed multilevel decompositions in stochastic collocation to decrease the cost of standard single level collocation approaches. In the following, we provide a generic summary of multilevel decompositions for UQ without going into the particularities of either of the variants. For more details, we refer the reader to the aforementioned references.

Let again $\mathcal{Q}[\mathcal{F}_h]$ be a quantity of interest in either uncertainty propagation or Bayesian parameter inference, which depends on forward model, \mathcal{F}_h . In standard multilevel estimations, $\mathcal{Q}[\mathcal{F}_h]$ is assumed to be *linear*, such as the expectation operator. Examples of nonlinear quantities of interest include, for example, the variance; there are multilevel formulations for estimating nonlinear quantities of interest (see, e.g., [13]). Nevertheless, in this work we are only concerned with linear, continuous quantities of interest.

Since $\mathcal{Q}[\mathcal{F}_h]$ can rarely be evaluated analytically, we resort to numerical approximations. Therefore, let $\mathcal{Q}_r[\mathcal{F}_h]$ denote an approximation of $\mathcal{Q}[\mathcal{F}_h]$ depending on a parameter $r \in \mathbb{R}$. For example, if $\mathcal{Q}[\mathcal{F}_h]$ is the expectation of the output of interest and if sparse grid quadrature is employed for its approximation, then r could be either the maximum grid level (if static grids are used) or the tolerance (when adaptive refinement is employed). To simplify the notation, we denote $\mathcal{Q}_{h,r} := \mathcal{Q}_r[\mathcal{F}_h]$.

Let $J \geq 2$ be an integer denoting the number of *levels* and let $j = 1, 2, \dots, J$. In addition, consider two (increasing) hierarchies of model discretization parameters and approximations of the quantity of interest, respectively, i.e.,

$$h_1 \leq h_2 \leq \dots \leq h_J, \quad r_1 \leq r_2 \leq \dots \leq r_J.$$

With this notation, the goal of standard *single level* approximations is to compute \mathcal{Q}_{h_J, r_J} ; in this setting, the high-fidelity model is \mathcal{F}_{h_J} . However, when the evaluation cost of \mathcal{F}_{h_J} is large, computing \mathcal{Q}_{h_J, r_J} becomes prohibitive. Thus, to reduce this cost without significantly deteriorating the overall accuracy, we employ standard multilevel decompositions, as follows.

The high-fidelity model, \mathcal{F}_{h_J} , is written as a *telescoping sum*, i.e.,

$$\mathcal{F}_{h_J} = \sum_{j=1}^J (\mathcal{F}_{h_j} - \mathcal{F}_{h_{j-1}}), \quad (2.33)$$

where $\mathcal{F}_{h_0} := 0$. Multilevel decompositions rely on the following observation: as j increases, $\mathcal{F}_{h_j} - \mathcal{F}_{h_{j-1}}$ decreases as well, and hence the computational effort to approximate the quantity of interest w.r.t. $\mathcal{F}_{h_j} - \mathcal{F}_{h_{j-1}}$ decreases as well. The question then is how to best exploit (2.33) to compute the quantity of interest. To preserve the accuracy of single level approximations and to insure that most computational effort is invested into the coarsest model discretization and the least effort into the finest discretization, the quantity of interest w.r.t. $\mathcal{F}_{h_j} - \mathcal{F}_{h_{j-1}}$ is assessed as $\mathcal{Q}_{r_{J-j+1}}[\mathcal{F}_{h_j} - \mathcal{F}_{h_{j-1}}] = \mathcal{Q}_{r_{J-j+1}}[\mathcal{F}_{h_j}] - \mathcal{Q}_{r_{J-j+1}}[\mathcal{F}_{h_{j-1}}] = \mathcal{Q}_{h_j, r_{J-j+1}} - \mathcal{Q}_{h_{j-1}, r_{J-j+1}}$, i.e., using approximations of *complementary* levels. To avoid carrying these long subscripts, we further simplify the notation by writing $\mathcal{Q}_{h_s, r_m} = \mathcal{Q}_{s, m}$ with the convention that the first subscript refers to the forward model's discretization parameter and the second subscript to the quantity of interest discretization parameter. Therefore, the *multilevel* estimation of the quantity of interest reads

$$\mathcal{Q}_{J, J} \approx \sum_{j=1}^J \mathcal{Q}_{J-j+1}[\mathcal{F}_j - \mathcal{F}_{j-1}] = \sum_{j=1}^J (\mathcal{Q}_{J-j+1}[\mathcal{F}_j] - \mathcal{Q}_{J-j+1}[\mathcal{F}_{j-1}]) = \sum_{j=1}^J (\mathcal{Q}_{j, J-j+1} - \mathcal{Q}_{j-1, J-j+1}), \quad (2.34)$$

We visualize the multilevel decomposition (2.34) for $J = 3$ in the left plot in Figure 2.12.

The multilevel decomposition (2.34) is computationally efficient if the approximation levels h_j and h_{j-1} are *nested* for $j = 1, 2, \dots, J$. But this is not always true, since there are situations in which these approximations are non-nested, e.g., if unstructured FE are used to discretize the physical space of the underlying model. However, if the approximations used to compute the quantity of interest are *nested*, which is the case, for example, when sparse grid quadrature in terms of weighted (L)-Leja points is employed (cf. Section 2.5), we can exploit the linearity of the quantity of interest to arrive at (see [81, Theorem 6.1])

$$\mathcal{Q}_{J, J} \approx \sum_{j=1}^J \mathcal{Q}_j[\mathcal{F}_{J-j+1}] - \mathcal{Q}_{j-1}[\mathcal{F}_{J-j+1}] = \sum_{j=1}^J (\mathcal{Q}_{J-j+1, j} - \mathcal{Q}_{J-j+1, j-1}). \quad (2.35)$$

We obtained (2.35) by rearranging the terms in (2.34). For example, for $J = 3$, we have that

$$\mathcal{Q}_{1,3} + \mathcal{Q}_{2,2} - \mathcal{Q}_{1,2} + \mathcal{Q}_{3,1} - \mathcal{Q}_{2,1} = \mathcal{Q}_{3,1} + \mathcal{Q}_{2,2} - \mathcal{Q}_{2,1} + \mathcal{Q}_{1,3} - \mathcal{Q}_{1,2}.$$

When the approximations used to compute the quantity of interest are nested, it is sufficient to compute only $\mathcal{Q}_{J-j+1, j}$ in each term in (2.35), since $\mathcal{Q}_{J-j+1, j-1} \subset \mathcal{Q}_{J-j+1, j}$. We depict the multilevel decomposition (2.35) for $J = 3$ in Figure 2.12, right.

Remark: Observe that standard multilevel decompositions (2.34), (2.35) resemble the sparse grid combination technique approximation (2.11): using (2.33), we write the model as a telescoping sum, which we then plug in (2.34) or (2.35) to assess the quantity of interest. This connection was studied in, e.g., [81], and we also exploit it in this work (see also Section 2.10).

To fully define (2.34) or (2.35), we need the values of J and r_1, r_2, \dots, r_J . To this end, standard multilevel methods proceed, in summary, as follows. The goal is to assess (2.34) with a prescribed accuracy, ε . To this end, the approximation error of (2.34) is theoretically established, usually in a norm. An upper bound of this error is derived, consisting of two terms: a term due to using J levels and another term due to the multilevel decomposition involving r_1, r_2, \dots, r_J . It is assumed that a cost rate of evaluating \mathcal{F}_{h_j} and a decay rate of $\mathcal{Q}_{j, J-j+1} - \mathcal{Q}_{j-1, J-j+1}$ are theoretically available. These rates will then enter the two aforementioned terms. The prescribed accuracy, ε , is split equally between the two aforementioned error terms. Moreover, the cost of the multilevel decomposition is also determined. To find J and r_1, r_2, \dots, r_J , the two terms are minimized numerically via, e.g., Lagrange multipliers to find the optimal estimates of J and r_1, r_2, \dots, r_J . For more details and examples of complexity theorems for multilevel methods, we refer the reader to [6, 33, 34, 42, 74, 75, 81, 176, 187].




Figure 2.12: Multilevel decomposition of the quantity of interest for $J = 3$. On the x-axis, we have the three approximation levels of the forward operator, h_1, h_2 and h_3 . The three approximation levels of the quantity of interest, r_1, r_2 and r_3 , are displayed on the y-axis. In each square, we depict the combinations of the two levels used to obtain each of the five terms in the expression of \mathcal{Q}_3 , as well as their sign. In the left plot, we depict the combination (2.34) which assumes nested levels in the first direction, i.e., $\mathcal{Q}_{3,3} \approx \mathcal{Q}_{1,3} + \mathcal{Q}_{2,2} - \mathcal{Q}_{1,2} + \mathcal{Q}_{3,1} - \mathcal{Q}_{2,1}$. On the right, we visualize the combination (2.35) which assumes nested levels in the second direction, i.e., $\mathcal{Q}_{3,3} \approx \mathcal{Q}_{3,1} + \mathcal{Q}_{2,2} - \mathcal{Q}_{2,1} + \mathcal{Q}_{1,3} - \mathcal{Q}_{1,2}$.

2.8 Multifidelity sampling

The third and last class of approximations relevant for this work are *multifidelity sampling* schemes. We begin with brief summary of standard MC sampling in Section 2.8.1. In Section 2.8.2, we review control variates, which are used to reduce the variance of standard MC samplers. Finally, Section 2.8.3 summarizes the standard multifidelity MC algorithm proposed in [145].

2.8.1 Standard Monte Carlo sampling

One of the most popular sampling techniques in uncertainty propagation is the MC algorithm developed by von Neumann, Metropolis and Ulam in the 1940s (see [130]). Broadly speaking, in its standard form, MC is a *stochastic quadrature* algorithm with *uniform weights* and (*pseudo*-) *random nodes*.

Assume that the task is to evaluate the expected value $\mu_{\mathcal{F}} := \mathbb{E}[\mathcal{F}]$ of the continuous, nondiscretized version of the high fidelity model, \mathcal{F} w.r.t. to the input density, $\boldsymbol{\pi}$. We want to compute this estimate using MC sampling. Let $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_M \sim \boldsymbol{\pi}$ denote $M \in \mathbb{N}$ independent and identically distributed (i.i.d.) samples drawn from the input density. The standard MC expectation estimator $\hat{\mu}_{\mathcal{F}} \approx \mu_{\mathcal{F}}$ reads

$$\hat{\mu}_{\mathcal{F}} := \frac{1}{M} \sum_{n=1}^M \mathcal{F}(\boldsymbol{\theta}_n). \quad (2.36)$$

By the strong law of large numbers, we have that $\hat{\mu}_{\mathcal{F}} \xrightarrow{M \rightarrow \infty} \mu_{\mathcal{F}}$, i.e., the MC estimator is *ergodic*. However, in computations, only a finite number of evaluations of \mathcal{F} are possible. This approximation introduces an error, which we want to determine. Specifically, we are interested in the *mean-squared error* (MSE) of the estimator (2.36), which we denote by $\text{MSE}(\hat{\mu}_{\mathcal{F}})$.

Lemma: 2.8.1 (see [152]). $\text{MSE}(\hat{\mu}_{\mathcal{F}}) = \text{Bias}^2(\hat{\mu}_{\mathcal{F}}) + \text{Var}[\hat{\mu}_{\mathcal{F}}]$, where $\text{Bias}(\hat{\mu}_{\mathcal{F}}) := \mu_{\mathcal{F}} - \mathbb{E}[\hat{\mu}_{\mathcal{F}}]$ is the estimator's bias and $\text{Var}[\hat{\mu}_{\mathcal{F}}] := \mathbb{E}[\hat{\mu}_{\mathcal{F}}^2] - \mathbb{E}[\hat{\mu}_{\mathcal{F}}]^2$ represents its variance.

Proof: We have

$$\begin{aligned} \text{MSE}(\hat{\mu}_{\mathcal{F}}) &:= \mathbb{E}[(\mu_{\mathcal{F}} - \hat{\mu}_{\mathcal{F}})^2] = \mathbb{E}[\mu_{\mathcal{F}}^2 - 2\mu_{\mathcal{F}}\hat{\mu}_{\mathcal{F}} + \hat{\mu}_{\mathcal{F}}^2] \\ &= \mathbb{E}[\mu_{\mathcal{F}}^2] - \mathbb{E}[2\mu_{\mathcal{F}}\hat{\mu}_{\mathcal{F}}] + \mathbb{E}[\hat{\mu}_{\mathcal{F}}^2] = \mu_{\mathcal{F}}^2 - 2\mu_{\mathcal{F}}\mathbb{E}[\hat{\mu}_{\mathcal{F}}] + \mathbb{E}[\hat{\mu}_{\mathcal{F}}^2] \\ &= (\mu_{\mathcal{F}}^2 - 2\mu_{\mathcal{F}}\mathbb{E}[\hat{\mu}_{\mathcal{F}}] + \mathbb{E}[\hat{\mu}_{\mathcal{F}}]^2) + (\mathbb{E}[\hat{\mu}_{\mathcal{F}}^2] - \mathbb{E}[\hat{\mu}_{\mathcal{F}}]^2) \\ &= (\mu_{\mathcal{F}} - \mathbb{E}[\hat{\mu}_{\mathcal{F}}])^2 + (\mathbb{E}[\hat{\mu}_{\mathcal{F}}^2] - \mathbb{E}[\hat{\mu}_{\mathcal{F}}]^2) = \text{Bias}^2(\hat{\mu}_{\mathcal{F}}) + \text{Var}[\hat{\mu}_{\mathcal{F}}]. \quad \blacksquare \end{aligned}$$

Corollary: 2.8.2 If the estimator $\hat{\mu}_{\mathcal{F}}$ is unbiased, then $\text{MSE}(\hat{\mu}_{\mathcal{F}}) = \text{Var}[\hat{\mu}_{\mathcal{F}}] = \text{Var}[\mathcal{F}]/M$.

Proof: If the estimator is unbiased, then $\mathbb{E}[\hat{\mu}_{\mathcal{F}}] = \mu_{\mathcal{F}}$ which means that $\text{Bias}^2(\hat{\mu}_{\mathcal{F}}) = 0$ in Lemma 2.8.1. Moreover, $\text{Var}[\hat{\mu}_{\mathcal{F}}] = \text{Var}[\frac{1}{M} \sum_{n=1}^M \mathcal{F}(\boldsymbol{\theta}_n)] = \frac{1}{M^2} \sum_{n=1}^M \text{Var}[\mathcal{F}] = \frac{1}{M} \text{Var}[\mathcal{F}]$. \blacksquare

From Corollary 2.8.2 we have that $\text{MSE}(\hat{\mu}_{\mathcal{F}}) = \text{Var}[\mathcal{F}]/M$, thus the root MSE is in $O(1/\sqrt{M})$. This statement is generally regarded as “the Monte Carlo (expectation) estimator has an MSE independent of d_{sto} and it thus breaks the curse of dimensionality”. We want to point out that the above statement should be taken with care because, although the rate $1/\sqrt{M}$ is indeed independent of d_{sto} , the constant in $O(1/\sqrt{M})$ might not always be. As it was noted elsewhere, see, e.g., [114], for the MSE derived in Corollary 2.8.2 to be independent of the dimensionality d_{sto} , we need $\text{Var}[\mathcal{F}]$ to be independent of d_{sto} as well, which is not always the case. We illustrate this point with the following example, also considered in [114].

Example: Consider the function $\mathcal{F}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} a \sin(2\pi\theta_i)$, where $a > 0$. A simple calculation shows that $\mu_{\mathcal{F}} = 0$ and $\text{Var}[\mathcal{F}(\boldsymbol{\theta})] = (a^2/2)^{d_{\text{sto}}}$. From Corollary 2.8.2, the MSE of an unbiased MC estimator of $\mu_{\mathcal{F}}$ would be

$$\text{MSE}(\hat{\mu}_{\mathcal{F}}) = \frac{(a^2/2)^{d_{\text{sto}}}}{M}. \quad (2.37)$$

The above MSE is dependent on d_{sto} except when $a = \sqrt{2}$, in which case $a^2/2 = 1$ and $\text{MSE}(\hat{\mu}_{\mathcal{F}}) = 1/M$. In contrast, when $a > \sqrt{2}$, $\text{MSE}(\hat{\mu}_{\mathcal{F}})$ increases exponentially with the dimensionality. This means that to obtain the same MSE in $d_{\text{sto}} + 1$ dimensions as in d_{sto} dimensions, we need more computational effort. Moreover, if $a < \sqrt{2}$, then the MSE is still dependent on d_{sto} but decreases exponentially with the dimensionality, meaning that in this case, the curse of dimensionality is broken. We illustrate, in Figure 2.13, this behaviour of the MSE (2.37) for different values of a . In applications, we wish the underlying integrand to be




Figure 2.13: Behaviour of $\text{MSE}(\hat{\mu}_{\mathcal{F}}) = (a^2/2)^{d_{\text{sto}}}/M$ corresponding to the MC estimation of the mean of $\mathcal{F}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} a \sin(2\pi\theta_i)$ for different values of a when the dimensionality is 1, 2, 4 and 8. In the left plot, $a = \sqrt{2}$ and thus $(a^2/2)^{d_{\text{sto}}}/M = 1/M$, i.e., the MSE is independent of the dimensionality. In the center, $a = \sqrt{2} + 0.2$ which means that the MSE increases with the dimensionality, i.e., the associated MC sampler suffers from the curse of dimensionality. The curse of dimensionality is broken in the right plot where $a = \sqrt{2} - 0.2$ because in this case the MSE decreases with the dimensionality.

smooth such that standard quadrature techniques are not (too) drastically affected by the curse of dimensionality. Exploiting the regularity of the integrand is central to approaches such as sparse grid [20] or quasi-Monte Carlo (see, e.g., [67]) quadrature.

We assumed in the the above derivations that the high-fidelity model \mathcal{F} is available analytically. In practical applications, however, the high-fidelity model depends on an approximation parameter, h , i.e., $\mathcal{F}_h \approx \mathcal{F}$ (see Section 2.1). Therefore, $\hat{\mu}_{\mathcal{F}} \approx \hat{\mu}_{\mathcal{F}_h} = \frac{1}{M} \sum_{n=1}^M \mathcal{F}_h(\boldsymbol{\theta}_n)$, which introduces a bias term

$$\text{Bias}(\hat{\mu}_{\mathcal{F}_h}) = \mu_{\mathcal{F}} - \mathbb{E}\left[\frac{1}{M} \sum_{n=1}^M \mathcal{F}_h(\boldsymbol{\theta}_n)\right] = \mathbb{E}[\mathcal{F}] - \frac{1}{M} \sum_{n=1}^M \mathbb{E}[\mathcal{F}_h] = \mathbb{E}[\mathcal{F}] - \mathbb{E}[\mathcal{F}_h] = \mathbb{E}[(\mathcal{F} - \mathcal{F}_h)].$$

in the MSE of the estimator. We note that what was presented in this section can be analogously derived for higher-order moments such as the variance [152].

To summarize, we have a bias term in the MSE of MC sampling whenever the underlying model \mathcal{F} is discretely approximated depending on a parameter h . This bias term is controlled by the underlying approximation scheme. In this work, \mathcal{F}_h is assumed to be the high-fidelity model, therefore it is sufficiently accurate such that this bias term is (arbitrarily) small. The other term in the MSE, $\text{Var}[\mathcal{F}]/M$, stems from sampling error and it can reduced by either

increasing the number of samples, M , or decreasing the variance $\text{Var}[\mathcal{F}]$. Especially in the *pre-asymptotic* regime, i.e., when M is not large enough to dominate $\text{Var}[\mathcal{F}]$, reducing the variance of the standard MC estimator may lead to significant computational cost reduction. Next, we summarize how control variates can lead to variance reduction in MC sampling.

2.8.2 Variance reduction via control variates

Assume that the task at hand is to estimate the mean value $\mu_{\mathcal{Y}} := \mathbb{E}[\mathcal{Y}]$ of a continuous random variable \mathcal{Y} using M i.i.d. samples y_1, y_2, \dots, y_M via MC sampling, i.e.,

$$\hat{\mathcal{Y}}_M = \frac{1}{M} \sum_{n=1}^M y_n.$$

Since we have access to \mathcal{Y} , there is no bias term in the MSE and hence, from Corollary 2.8.2, we have that the MSE of $\hat{\mathcal{Y}}_M$ is $\text{Var}[\mathcal{Y}]/M$. To decrease this MSE we have two options. First, we could increase the denominator, M , but this requires more computational effort, making this option infeasible when obtaining y_i is expensive for $i = 1, 2, \dots, M$.

The second option is to decrease the numerator, i.e., the variance $\text{Var}[\mathcal{Y}]$ via a so-called *variance reduction* technique. Example of variance reduction techniques include antithetic, stratified, Latin hypercube or importance sampling, and control variates (see, e.g., [152]). In this work, we employ control variates [135], which are used in multifidelity MC sampling [145], which we summarize in the next section.

Control variates (CV) assume that we have an auxiliary random variable, \mathcal{Z} , with known expectation $\mu_{\mathcal{Z}}$, such that \mathcal{Z} is *correlated* with \mathcal{Y} . Let $\gamma \in \mathbb{R}$ and define the CV estimator as

$$\hat{\mathcal{Y}}^{\text{CV}} = \hat{\mathcal{Y}}_M + \gamma(\mu_{\mathcal{Z}} - \hat{\mathcal{Z}}_M),$$

where $\hat{\mathcal{Z}}_M = \frac{1}{M} \sum_{n=1}^M z_n$ is an unbiased MC estimator of $\mu_{\mathcal{Z}}$. Note that $\hat{\mathcal{Y}}^{\text{CV}}$ is an unbiased estimator of $\mu_{\mathcal{Y}}$ since $\mathbb{E}[\hat{\mathcal{Y}}^{\text{CV}}] = \mathbb{E}[\hat{\mathcal{Y}}_M + \gamma(\mu_{\mathcal{Z}} - \hat{\mathcal{Z}}_M)] = \mathbb{E}[\hat{\mathcal{Y}}_M] + \gamma\mathbb{E}[(\mu_{\mathcal{Z}} - \hat{\mathcal{Z}}_M)] = \mathbb{E}[\hat{\mathcal{Y}}_M] = \mu_{\mathcal{Y}}$. Therefore, from Corollary 2.8.2 we have that the MSE of $\hat{\mathcal{Y}}^{\text{CV}}$ is $\text{Var}[\hat{\mathcal{Y}}^{\text{CV}}]$, which reads:

$$\text{Var}[\hat{\mathcal{Y}}^{\text{CV}}] = \text{Var}[\hat{\mathcal{Y}}_M + \gamma(\mu_{\mathcal{Z}} - \hat{\mathcal{Z}}_M)] = \text{Var}[\mathcal{Y}] + \gamma^2 \text{Var}[\mathcal{Z}] - 2\gamma \text{Cov}[\mathcal{Y}, \mathcal{Z}], \quad (2.38)$$

where

$$\text{Cov}[\mathcal{Y}, \mathcal{Z}] := \mathbb{E}[\mathcal{Y}\mathcal{Z}] - \mathbb{E}[\mathcal{Y}]\mathbb{E}[\mathcal{Z}]$$

is the *covariance* between \mathcal{Y} and \mathcal{Z} . Minimizing $\text{Var}[\hat{\mathcal{Y}}^{\text{CV}}]$ w.r.t. γ yields $\gamma^* = \text{Cov}[\mathcal{Y}, \mathcal{Z}]/\text{Var}[\mathcal{Z}]$. If we plug γ^* in (2.38), we obtain

$$\text{Var}[\hat{\mathcal{Y}}^{\text{CV}}] = \text{Var}[\mathcal{Y}] - \frac{\text{Cov}^2[\mathcal{Y}, \mathcal{Z}]}{\text{Var}[\mathcal{Z}]} = \text{Var}[\mathcal{Y}] \left(1 - \frac{\text{Cov}^2[\mathcal{Y}, \mathcal{Z}]}{\text{Var}[\mathcal{Z}]\text{Var}[\mathcal{Y}]} \right) = \text{Var}[\mathcal{Y}](1 - \rho_{\mathcal{Y}, \mathcal{Z}}^2), \quad (2.39)$$

where $\rho_{\mathcal{Y}, \mathcal{Z}} := \frac{\text{Cov}[\mathcal{Y}, \mathcal{Z}]}{\sqrt{\text{Var}[\mathcal{Z}]\text{Var}[\mathcal{Y}]}}$ is the *Pearson correlation coefficient* of \mathcal{Y} and \mathcal{Z} . A simple application of the Cauchy-Schwarz inequality yields that $\rho_{\mathcal{Y}, \mathcal{Z}} \in [-1, 1]$. If in (2.39), $\rho_{\mathcal{Y}, \mathcal{Z}} \approx -1$, then \mathcal{Y} and \mathcal{Z} are *negatively correlated*, whereas $\rho_{\mathcal{Y}, \mathcal{Z}} \approx 1$ when they are *positively correlated*. If $1 - \rho_{\mathcal{Y}, \mathcal{Z}}^2 < 1$, that is, if $\rho_{\mathcal{Y}, \mathcal{Z}}^2 > 0$, then the control variate estimator $\hat{\mathcal{Y}}^{\text{CV}}$ leads to variance reduction. Moreover, when $\rho_{\mathcal{Y}, \mathcal{Z}}^2 \approx 1$, $\text{Var}[\hat{\mathcal{Y}}^{\text{CV}}] \approx 0$, i.e., the variance and hence the MSE of the CV estimator vanishes, thus we can estimate its mean with a single sample. Thus, the more correlated the auxiliary random variable \mathcal{Z} is with \mathcal{Y} , the more variance reduction we get with a CV estimator. The worst case scenario is when $\rho_{\mathcal{Y}, \mathcal{Z}} \approx 0$, because in this case the CV estimator reduces to a standard MC estimator. In the next section, we summarize the multifidelity MC estimator proposed by Peherstorfer, Willcox and Gunzburger in [145].

2.8.3 Multifidelity Monte Carlo sampling

The standard multifidelity MC (MFMC) sampling algorithm proposed in [145] employs CV and multifidelity model hierarchies to decrease the variance and hence the MSE of a standard MC estimator of the expectation of the high-fidelity model. In MFMC notation, the high-fidelity model is typically denoted by $f^{(1)}$. Thus, for consistency, in the remaining of this section, we denote $f^{(1)} \equiv \mathcal{F}_h$. Our proposed Context-aware multifidelity sampling algorithm in Chapter 6 is an extension of the standard MFMC approach, summarized next.

MFMC exploits the fact that in most problems we have available *surrogate* or *low-fidelity models* which approximate the given high-fidelity model, $f^{(1)}$. Low-fidelity models are characterized by a lower accuracy but also a (significantly) lower computational cost as compared to the high-fidelity model (see Figure 2.15). MFMC takes advantage of the lower computational cost of the surrogates to reduce the estimator's MSE and thus to speed up an equivalent standard MC estimator. Examples of surrogates include *data-driven* low-fidelity models, such as interpolation, spectral projection or regression approximations, *projection-based* approaches, for example, reduced basis or proper orthogonal decomposition, or *reduced-physics* surrogates, such as the Reynolds averaged Navier-Stokes model in turbulence simulations. Moreover, standard MFMC does not require *certified* surrogates, in the sense that there is no need for bounds to approximation errors such as $\mathcal{E}(f^{(1)} - f^{(j)}) := |f^{(1)} - f^{(j)}|$. The accuracy of low-fidelity models is ascertained via their Pearson correlation coefficient w.r.t. $f^{(1)}$, which we present shortly.

More formally, the goal of MFMC is to estimate the expectation of the high-fidelity model, $\mathbb{E}[f^{(1)}]$, via sampling such that the MSE of the estimator is reduced as compared to an equivalent standard MC estimator. To this end, assume that we have available $m - 1$ low-fidelity models $f^{(2)}, f^{(3)}, \dots, f^{(m)}$, resulting in a total of m models $f^{(1)}, f^{(2)}, \dots, f^{(m)}$ with $f^{(1)} \equiv \mathcal{F}_h$. The MFMC estimator is defined as follows. Let $M_1 < M_2 < \dots < M_m$ be positive integers such that $M_1 \geq 1$. We draw $\Theta_{M_m} = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_{M_m}\}$ i.i.d. realizations of the stochastic input $\boldsymbol{\theta}$. Define the standard MC estimators

$$\hat{f}_{M_j}^{(j)} = \frac{1}{M_j} \sum_{n=1}^{M_j} f^{(j)}(\boldsymbol{\theta}_n), \quad j = 1, 2, \dots, m.$$

The MFMC estimator $\hat{f}_{\text{MFMC}}^{(1)}$ of $\mathbb{E}[f^{(1)}]$ is hence defined as

$$\hat{f}_{\text{MFMC}}^{(1)} = \hat{f}_{M_1}^{(1)} + \sum_{j=2}^m \gamma_j (\hat{f}_{M_j}^{(j)} - \hat{f}_{M_{j-1}}^{(j)}), \quad (2.40)$$

where $\gamma_i \in \mathbb{R}$ are the CV coefficients. Note that $\hat{f}_{M_j}^{(j)}$ is obtained by evaluating the model $f^{(j)}$ at the first M_j samples from Θ_{M_m} , whereas $\hat{f}_{M_{j-1}}^{(j)}$ is obtained by reusing the first M_{j-1} evaluations from $\hat{f}_{M_j}^{(j)}$; this makes the two estimators dependent. Thus, the MFMC estimator (2.40) makes use of *all* available m models (see Figure 2.14). In contrast, a standard MC estimator would employ only *one* model, either the high-fidelity or one of the low-fidelity surrogates.

In [145, Lemma 3.1], it was shown that the estimator (2.40) is unbiased, which follows from (i) $M_1 \geq 1$, i.e., each model, including the high-fidelity model, is evaluated at least once, and (ii) the linearity of the expectation operator. Therefore, in MFMC, the high-fidelity model $f^{(1)}$ is kept in the MFMC estimator (2.40) to ensure that this estimator is unbiased. From Corollary 2.8.2, we have that $\text{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) = \text{Var}[\hat{f}_{\text{MFMC}}^{(1)}]$. In [145], it was shown that

$$\text{Var}[\hat{f}_{\text{MFMC}}^{(1)}] = \frac{\sigma_1^2}{M_1} + \sum_{j=2}^m \left(\frac{1}{M_{j-1}} - \frac{1}{M_j} \right) (\gamma_1^2 \sigma_j^2 - 2\gamma_j \rho_j \sigma_1 \sigma_j), \quad (2.41)$$




Figure 2.14: MFMC sampling: the high fidelity model $f^{(1)}$ and $m - 1$ low-fidelity models $f^{(2)}, f^{(3)}, \dots, f^{(m)}$ are used *together* to estimate the mean of $f^{(1)}$, $\hat{f}_{\text{MFMC}}^{(1)}$, via sampling.

where $\sigma_j^2 = \text{Var}[f^{(j)}]$ and ρ_j is the Pearson correlation coefficient between $f^{(1)}$ and $f^{(j)}$, i.e.,

$$\rho_j = \frac{\text{Cov}[f^{(1)}, f^{(j)}]}{\sigma_1 \sigma_j}, \quad j = 2, 3, \dots, m.$$

We want to find M_j and γ_j for $j = 1, 2, \dots, m$ such that the MSE (2.41) is minimized. Let $\mathcal{C} = \sum_{j=1}^m M_j \mathcal{C}_j$ denote the total computational cost for computing (2.40), where \mathcal{C}_j denotes the runtime (in seconds) for one evaluation of $f^{(j)}$. For simplicity, denote $\mathbf{M} = (M_1, M_2, \dots, M_m)$ and $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots, \gamma_m)$. The following theorem states that under mild assumptions about the Pearson correlation coefficients and costs, M_j and γ_j can be found *analytically*.

Theorem 2.8.3 ([145, Theorem 3.4]) *Let $f^{(1)}, f^{(2)}, \dots, f^{(m)}$ be m models characterizing the phenomenon under consideration with associated Pearson correlation coefficients ordered as $1 = |\rho_1| > |\rho_2| > \dots > |\rho_m|$ such that $\rho_{m+1} = 0$, and with costs $\mathbf{C} = (\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_m)$ satisfying*

$$\frac{\mathcal{C}_{j-1}}{\mathcal{C}_j} > \frac{\rho_{j-1}^2 - \rho_j^2}{\rho_j^2 - \rho_{j+1}^2}, \quad j = 2, 3, \dots, m.$$

Furthermore, let $\mathbf{r}^* = (r_1^*, r_2^*, \dots, r_m^*)$ with $r_j^* = \sqrt{\frac{\mathcal{C}_1(\rho_j^2 - \rho_{j+1}^2)}{\mathcal{C}_j(1 - \rho_j^2)}}$ and

$$\boldsymbol{\gamma}^* = (\gamma_2^*, \dots, \gamma_m^*), \quad \mathbf{M}^* = (M_1^*, M_2^*, \dots, M_m^*),$$

where $\gamma_j^* = \frac{\rho_j \sigma_1}{\sigma_j}$, $M_1^* = \frac{p}{\mathcal{C}_1^T \mathbf{r}^*}$ and $M_j^* = r_j^* M_1^*$ for $j = 2, 3, \dots, m$.

For a user defined computational budget $p > 0$, the pair (\mathbf{M}^*, γ^*) is the global minimum to

$$\begin{aligned} & \underset{\mathbf{M}, \gamma \in \mathbb{R}^m}{\operatorname{argmin}} && \operatorname{Var}[\hat{f}_{\text{MFMC}}^{(1)}], \\ & \text{subject to} && M_{j-1} - M_j \leq 0 \\ & && \sum_{j=1}^m M_j \mathcal{C}_j = p, \end{aligned}$$

with the convention $M_0 := 0$.

Note that, for simplicity, the optimization problem in Theorem 2.8.3 is solved over the real numbers, which means that the solution $\mathbf{M}^* \in \mathbb{R}^m$. To have an integer number of samples, we simply round-up \mathbf{M}^* , i.e., $\bar{M}_j^* = \lceil M_j^* \rceil$. Since in computations p is typically large, rounding up \mathbf{M}^* does not have a significant impact on the overall results. It is important to note that p should be chosen such that at least one evaluation of $f^{(1)}$ is possible, otherwise the MFMC estimator (2.40) cannot be defined.

Theorem 2.8.3 states that if we have a *model hierarchy* $f^{(1)}, f^{(2)}, \dots, f^{(m)}$ such that the models are ordered by accuracy, i.e., $1 = |\rho_1| > |\rho_2| > \dots > |\rho_m|$, and, in addition,

$$\frac{\mathcal{C}_{j-1}}{\mathcal{C}_j} > \frac{\rho_{j-1}^2 - \rho_j^2}{\rho_j^2 - \rho_{j+1}^2}, \quad j = 2, 3, \dots, m.$$

then we can analytically find the optimal M_j and γ_j for $j = 1, 2, \dots, m$. Note that we do not always a priori know if the above conditions are satisfied for the available set of models. In that case, we need a *model selection* algorithm which orders the high- and low-fidelity models such that the above conditions are satisfied; see Algorithm 1 in [145].

With the optimal parameters from Theorem 2.8.3, we obtain

$$\operatorname{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) = \operatorname{Var}[\hat{f}_{\text{MFMC}}^{(1)}] = \underbrace{\frac{\operatorname{Var}[f^{(1)}]}{p}}_{\operatorname{MSE} \text{ single model}} \underbrace{\left(\sum_{j=1}^m \sqrt{\mathcal{C}_j (\rho_j^2 - \rho_{j+1}^2)} \right)^2}_{\text{variance reduction}}. \quad (2.42)$$

Thus, MFMC reduces the variance of standard MC sampling via an expression that depends only on the correlation coefficients and computational costs of the underlying hierarchy of models. Note then when the correlation coefficients ρ_j or the variances $\sigma_j^2 = \operatorname{Var}[f^{(j)}]$ are not available analytically, they are typically evaluated numerically via, e.g., sampling in a preprocessing step, which requires an additional computational cost. However, [145] showed that estimates obtained with a small number of samples are sufficient, hence this extra cost is typically not too large. If these costs, denoted by p_{preproc} , are large, we can subtract them from p and use the remaining budget, $\hat{p} := p - p_{\text{preproc}}$ for MFMC sampling. We provide a visual summary of MFMC in Figure 2.18.

Remark: Standard multilevel MC sampling [74, 75] is a particular instance of MFMC for grid-based hierarchies, in which the high-fidelity model is given by the finest while the lowest-fidelity surrogate is the coarsest discretization of the underlying model (see Section 2.7).

We note that MFMC can be employed in other settings as well, such as rare events simulation [106, 143] or in the estimation of variance and Sobol' indices for global sensitivity analysis [148].




Figure 2.15: MFMC sampling: high- and low-fidelity models are used together to find an estimate of the expectation of the high-fidelity model with a lower variance as compared to a standard MC sampler. The high-fidelity model is used in MFMC to ensure unbiasedness, whereas the low-fidelity models are leveraged for computational speed up.

2.9 Sobol' indices for global sensitivity analysis

In the previous four sections, we summarized the three main classes of approximations of interest in this work. Next, we overview Sobol' indices [166] for *variance-based global sensitivity analysis*. Sobol' indices are employed in the formulation of two of the proposed algorithms, the Sensitivity-driven dimension-adaptive sparse grid algorithm from Chapter 4 and the Multilevel adaptive spectral projection with online dimensionality reduction, presented in Chapter 5. Moreover, Sobol' indices are also used to design an enhanced adaptive strategy to be employed in the Multilevel adaptive sparse Leja algorithm in Chapter 7.

The goal of global sensitivity analysis is to assess the contribution of uncertain inputs and combinations thereof to the total output uncertainty. As the name suggests, in variance-based sensitivity analysis, the output measure of uncertainty is the *variance*. To ascertain the contribution of stochastic inputs or their interaction to the output uncertainty, and hence their importance, we employ Sobol' indices. In the following, we provide a brief overview of global sensitivity analysis and Sobol' indices and refer the reader to, e.g., [148, 156, 166, 173] for a more comprehensive overview of these topics. In the following, our notation is similar to [148].

Assume that the high-fidelity model, \mathcal{F}_h , is square-integrable w.r.t. the input density, $\boldsymbol{\pi}$, i.e., $\mathbb{E}[\mathcal{F}_h]$ and $\text{Var}[\mathcal{F}_h]$ are both finite. Recall also that by Assumption 2.2, the input density has a product structure, i.e., $\boldsymbol{\pi}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} \pi_i(\theta_i)$. We use the following notation. We denote $D := \{1, 2, \dots, d_{\text{sto}}\}$ and let $\boldsymbol{v} \subset D$ be a non-empty subset of indices from D . Furthermore, by $\boldsymbol{\pi}_{\boldsymbol{v}}$ we denote the product of marginals having indices in \boldsymbol{v} .

The starting point is to write \mathcal{F}_h as a sum of functions of subsets of its inputs

$$\mathcal{F}_h(\boldsymbol{\theta}) = \mathcal{F}_h^0 + \sum_{i=1}^{d_{\text{sto}}} \mathcal{F}_h^i(\theta_i) + \sum_{1 \leq i < j \leq d_{\text{sto}}} \mathcal{F}_h^{i,j}(\theta_i, \theta_j) + \dots \mathcal{F}_h^{1,2,\dots,d_{\text{sto}}}(\boldsymbol{\theta}) = \mathcal{F}_h^0 + \sum_{\mathbf{v} \subset D} \mathcal{F}_h^{\mathbf{v}}(\boldsymbol{\theta}_{\mathbf{v}}), \quad (2.43)$$

where \mathcal{F}_h^0 is a *scalar* function, $\mathcal{F}_h^i(\theta_i)$ are *univariate functions*, $\mathcal{F}_h^{i,j}(\theta_i, \theta_j)$ are *bivariate functions* and so on until $\mathcal{F}_h^{1,2,\dots,d_{\text{sto}}}(\boldsymbol{\theta})$, which involves all d_{sto} stochastic inputs. Observe that (2.43) is a *representation*, not an approximation of the high-fidelity model with a total of $2^{d_{\text{sto}}}$ terms. To make (2.43) unique, the following orthogonality condition is imposed

$$\int \mathcal{F}_h^{\mathbf{v}}(\boldsymbol{\theta}_{\mathbf{v}}) \boldsymbol{\pi}_i(\boldsymbol{\theta}_i) d\boldsymbol{\theta}_i = 0, \quad \forall i \in \mathbf{v}, \forall \mathbf{v} \subset D. \quad (2.44)$$

The decomposition (2.43) with the orthogonality condition (2.44) is called the *Sobol-Hoeffding* or *ANOVA* (ANalysis Of VAriance) decomposition. It ensures that $\mathcal{F}_h^0 := \mathbb{E}[\mathcal{F}_h]$ and allows the following decomposition of $\text{Var}[\mathcal{F}_h]$:

$$\text{Var}[\mathcal{F}_h] = \sum_{i=1}^{d_{\text{sto}}} \text{Var}[\mathcal{F}_h^i] + \sum_{1 \leq i < j \leq d_{\text{sto}}} \text{Var}[\mathcal{F}_h^{i,j}] + \dots \text{Var}[\mathcal{F}_h^{1,2,\dots,d_{\text{sto}}}] = \sum_{\mathbf{v} \subset D} \text{Var}[\mathcal{F}_h^{\mathbf{v}}]. \quad (2.45)$$

Note that the above variance representation comprises $2^{d_{\text{sto}}} - 1$ terms. From this variance decomposition, the *local* Sobol' indices for global sensitivity analysis are defined as

$$S_{\mathbf{v}} = \frac{\text{Var}[\mathcal{F}_h^{\mathbf{v}}]}{\text{Var}[\mathcal{F}_h]}. \quad (2.46)$$

These indices measure the relative contribution of the subset of inputs \mathbf{v} to the output variance, $\text{Var}[\mathcal{F}_h]$, which is equivalent to say that they measure the (stochastic) importance of the subset of inputs \mathbf{v} relative to the output uncertainty. From (2.45), it follows that $\sum_{\mathbf{v} \subset D} S_{\mathbf{v}} = 1$. Of particular interest are the *first order* local Sobol' indices

$$S_i = \frac{\text{Var}[\mathcal{F}_h^i]}{\text{Var}[\mathcal{F}_h]},$$

measuring the local importance of each individual parameter, i.e., $|\mathbf{v}|_1 = 1$. Besides local Sobol' indices, another useful measure of importance of stochastic parameters is the *total Sobol' index*, which ascertains the *total* contribution of an uncertain input to the output variance. That is, a total Sobol' index comprises the first order contribution and *all* interactions involving that stochastic parameter. Therefore, we have d_{sto} total Sobol' indices computed as

$$S_i^T = \sum_{\{\mathbf{v}: i \in \mathbf{v}\}} S_{\mathbf{v}}. \quad (2.47)$$

Note that contributions due to interactions between inputs are added multiple times. For example, if input i interacts with input j , $S_{i,j}$ is added to both S_i^T and S_j^T . Therefore, $\sum_{i=1}^{d_{\text{sto}}} S_i^T \geq 1$.

Remark: When $\sum_{i=1}^{d_{\text{sto}}} S_i^T \approx 1$, it follows that $S_{\mathbf{v}} \approx 0$ for all \mathbf{v} with $|\mathbf{v}| \geq 2$, i.e., the interactions between inputs are negligible. From (2.43), we therefore have that $\mathcal{F}_h(\boldsymbol{\theta}) \approx \mathcal{F}_h^0 + \sum_{i=1}^{d_{\text{sto}}} \mathcal{F}_h^i(\theta_i)$, i.e., the underlying stochastic model is linear. Hence, the Sobol' indices can also be used to establish the linearity of stochastic models in uncertainty propagation.

The variances used to assess Sobol' indices are generally computed using standard MC [156], multilevel [13] or MFMC sampling [148]. In this work, however, we do not employ sampling methods but rather exploit the link between Sobol' decompositions and PSP established in [173] and compute (2.46) and (2.47) using PSP coefficients; see also [1]. Let

$$\mathbf{u}^{\text{PSP}}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{P}^{\text{PSP}}} c_{\mathbf{p}} \Phi_{\mathbf{p}}(\boldsymbol{\theta})$$

denote a multivariate orthonormal PSP approximation (recall (2.18)), where \mathcal{P}^{PSP} of a set containing the (multivariate) PSP polynomial degrees. Examples of \mathcal{P}^{PSP} include a priori choices, such as a *total degree* set

$$\mathcal{P}^{\text{PSP}} = \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}} : |\mathbf{p}|_1 \leq P_{\text{max}}^{\text{PSP}}\}$$

or a *maximum degree* set

$$\mathcal{P}^{\text{PSP}} = \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}} : |\mathbf{p}|_{\infty} \leq P_{\text{max}}^{\text{PSP}}\},$$

where $P_{\text{max}}^{\text{PSP}}$ is a user-defined degree. Moreover, if dimension-adaptive PSP is employed (see Section 2.5), then

$$\mathcal{P}^{\text{PSP}} = \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}} : \mathbf{0} \leq \mathbf{p} \leq \mathbf{p}_{\text{max}}\},$$

where $\mathbf{p}_{\text{max}} = (k_{1,\text{max}} - 1, k_{2,\text{max}} - 1, \dots, k_{d_{\text{sto}},\text{max}} - 1)$ and $(k_{1,\text{max}}, k_{2,\text{max}}, \dots, k_{d_{\text{sto}},\text{max}})$ is the maximum reached multiindex by the dimension-adaptive algorithm. From (2.21), the variance of \mathcal{F}_h is assessed as $\text{Var}[\mathcal{F}_h] = \sum_{\mathbf{p} \in \mathcal{J}_{\text{var}}^{\text{PSP}}} c_{\mathbf{p}}^2$, where $\mathcal{J}_{\text{var}}^{\text{PSP}} = \{\mathbf{p} \in \mathcal{P}^{\text{PSP}} : |\mathbf{p}|_1 \geq 1\}$. To assess the numerators of the local Sobol' indices (2.46), *subsets* of PSP coefficients are employed, i.e.,

$$\text{Var}[\mathcal{F}_h^{\mathbf{v}}] = \sum_{\mathbf{p} \in \mathcal{J}_{\mathbf{v}}^{\text{PSP}}} c_{\mathbf{p}}^2, \quad \mathcal{J}_{\mathbf{v}}^{\text{PSP}} = \{\mathbf{p} \in \mathcal{P}^{\text{PSP}} : \mathbf{p}_{\mathbf{v}} \neq \mathbf{0} \wedge \mathbf{p}_{D \setminus \{\mathbf{v}\}} = \mathbf{0}\},$$

that is, only multivariate degrees with non-zero entries corresponding to the indices in \mathbf{v} are considered. For example, if $d_{\text{sto}} = 3$ and $\mathbf{p} = (1, 2, 0)$, then we can use \mathbf{p} to assess $S_{1,2}$. Therefore, the local Sobol' indices are estimated from PSP coefficients as

$$\hat{S}_{\mathbf{v}} = \frac{\sum_{\mathbf{p} \in \mathcal{J}_{\mathbf{v}}^{\text{PSP}}} c_{\mathbf{p}}^2}{\sum_{\mathbf{p} \in \mathcal{J}_{\text{var}}^{\text{PSP}}} c_{\mathbf{p}}^2}. \quad (2.48)$$

In addition, total Sobol' indices are estimated using (2.47) and (2.48). Before going further, we present an example in which we assess local Sobol' indices.

Example: Let us consider again the example at the end of Section 2.5.3. Therein, we considered two functions, $\mathcal{F}(\boldsymbol{\theta}) = \sin(4\theta_1) + \sin(\frac{3}{2}\theta_2)$, defined in (2.26), and $\mathcal{G} = \sin(4\theta_1 + \frac{3}{2}\theta_2)$, defined in (2.27), with $\boldsymbol{\theta}$ uniformly distributed in $[0, 1]^2$. To compute the Sobol' indices, we use a 2D PSP approximation with a statically chosen total degree set

$$\mathcal{P}^{\text{PSP}} = \{\mathbf{p} \in \mathbb{N}^2 : |\mathbf{p}|_1 \leq 9\}.$$

Since $d_{\text{sto}} = 2$, we estimate three local Sobol' indices, S_1 , S_2 and $S_{1,2}$.

For \mathcal{F} , we obtain

$$\hat{S}_1 = 0.7424, \quad \hat{S}_2 = 0.2750, \quad \hat{S}_{1,2} = 0.0000.$$

That is, the local contribution due to the first input is the most important, the contribution of θ_2 is second, while the Sobol' index associated to their interaction is zero. These values reflect the additive structure of \mathcal{F} : θ_1 is more important than θ_2 since it has a larger multiplication weight, and θ_1 and θ_2 do not interact. Recall that this structure was also exploited by dimension-adaptivity; see Figure 2.6.

Since \mathcal{G} is nonlinear in θ_1 and θ_2 , we obtain three non-zero indices

$$\hat{S}_1 = 0.8533, \quad \hat{S}_2 = 0.0572, \quad \hat{S}_{1,2} = 0.0893.$$

We see again that θ_1 is the most important parameter, but in this case, the parameters interactions is nonnegligible. The fact the the first direction is more important than the second was exploited by dimension-adaptivity in the example at end of Section 2.5.3; recall Figure 2.7.

2.10 Model hierarchies in uncertainty quantification

We end the theoretical overview with revisiting three summarized classes of approximation techniques of interest in this work: (i) global (dimension-adaptive) sparse grid approximations (Section 2.5), (ii) multilevel estimation of quantities of interest, presented Section 2.7 and (iii) MFMC sampling, summarized in Section 2.8.3. Note that we also summarized approximations with spatially adaptive sparse grids in Section 2.6. However, we do not discuss this technique further here since its primary use in this work is in multilevel decompositions, in Chapter 5.

All three approaches are outer-loop scenarios since they involve ensembles of model evaluations at samples of the stochastic input, θ . Furthermore, all methods intrinsically employ model hierarchies and, in addition, multilevel and multifidelity decompositions can be viewed as “generalized” sparse grid approximations (2.10) or “generalized” combination schemes (2.11).

Assume, without loss of generality, that $d_{\text{sto}} = 2$ and that all three schemes are used to approximate the expected value of the forward model. This means that we need global quadrature operators in sparse grid approximations and that $\mathcal{Q}[\mathcal{F}_h] := \mathbb{E}[\mathcal{F}_h]$ in the multilevel decomposition(2.35). In addition, for simplicity, we do not employ dimension-adaptivity for the sparse grid approximation, but a statically defined multiindex set (2.24). For all three approaches, we consider three levels: $K_{\text{max}}^{\text{qu}} = 3$ in global sparse grid quadrature, $J = 3$ in the multilevel approximation (2.35) and $m = 3$ in the MFMC estimator (2.40).

First, we look at global sparse grid quadrature. With the employed setup, the multiindex set is $\mathcal{K} = \{(1, 1), (1, 2), (2, 1), (1, 3), (2, 2), (3, 1)\}$. From (2.10) and (2.11), we obtain

$$\begin{aligned} \mathbb{E}[\mathcal{F}_h] &\approx \mathcal{U}_{\mathcal{K}}^{\text{qu}}[\mathcal{F}_h] = \Delta_{(1,1)}^{\text{qu}}[\mathcal{F}_h] + \Delta_{(1,2)}^{\text{qu}}[\mathcal{F}_h] + \Delta_{(2,1)}^{\text{qu}}[\mathcal{F}_h] + \Delta_{(1,3)}^{\text{qu}}[\mathcal{F}_h] + \Delta_{(2,2)}^{\text{qu}}[\mathcal{F}_h] + \Delta_{(3,1)}^{\text{qu}}[\mathcal{F}_h] \\ &= 0 \cdot \mathcal{U}_{(1,1)}^{\text{qu}}[\mathcal{F}_h] - 1 \cdot \mathcal{U}_{(1,2)}^{\text{qu}}[\mathcal{F}_h] - 1 \cdot \mathcal{U}_{(2,1)}^{\text{qu}}[\mathcal{F}_h] + 1 \cdot \mathcal{U}_{(1,3)}^{\text{qu}}[\mathcal{F}_h] \\ &\quad + 1 \cdot \mathcal{U}_{(2,2)}^{\text{qu}}[\mathcal{F}_h] + 1 \cdot \mathcal{U}_{(3,1)}^{\text{qu}}[\mathcal{F}_h]. \end{aligned} \tag{2.49}$$

The above formula introduces a model hierarchy, as follows. The same model, \mathcal{F}_h , is used at all times, but we evaluate it at the two sparse grid levels corresponding to the two stochastic inputs, θ_1 and θ_2 . Therefore, in global sparse grid approximations, the model hierarchy is given by the d_{sto} sparse grid levels corresponding to the d_{sto} stochastic inputs. For example, by increasing the level corresponding to θ_2 , we have the hierarchy $\Delta_{(1,1)}^{\text{qu}}[\mathcal{F}_h] \rightarrow \Delta_{(1,2)}^{\text{qu}}[\mathcal{F}_h] \rightarrow \Delta_{(1,3)}^{\text{qu}}[\mathcal{F}_h]$. We depict this example in Figure 2.16. In the left part, we visualize the combination of subspaces associated to the multiindices in \mathcal{K} (the latter equality in (2.49)). In the right figure, we make a “zoom-in” to the subspace associated to the multiindex (2, 2). Therein, $\Delta_{(2,2)}^{\text{qu}}$ is assessed using a tensorization of the quadrature nodes associated to (k_2, k_2) . Note that the aforementioned model hierarchy appears also when using dimension-adaptivity.

The multilevel decompositions (2.35) extends the hierarchy appearing in global sparse grids one step further by using different discretizations of the forward model together with different approximation levels in the stochastic space. Therefore, we have a grid-based model hierarchy or, what we call a *within-model* hierarchy. In this setting, the high-fidelity model is the given




Figure 2.16: Model hierarchy in global sparse grid quadrature. In the left plot, we depict a level three combination scheme (2.10), (2.11). On the right, we visualize the subspace associated to the multiindex (k_2, k_2) .

model employing the finest discretization in the problem domain (\mathcal{F}_{h_3} here). The three-level estimation (2.35) of the quantity of interest reads

$$\begin{aligned} Q_{3,3} &\approx Q_{3,1} + (Q_{2,2} - Q_{2,1}) + (Q_{1,3} - Q_{1,2}) \\ &= 0 \cdot Q_{1,1} - 1 \cdot Q_{2,1} - 1 \cdot Q_{1,2} + 1 \cdot Q_{1,3} + 1 \cdot Q_{2,2} + 1 \cdot Q_{3,1}. \end{aligned} \quad (2.50)$$

The work [81] established that in UQ, (2.35) is essentially a 2D combination scheme in an *abstract* 2D space, involving tensorizations of problem space discretization (first direction) and sampling (quadrature) of the stochastic space (second direction), which we can clearly see in the above formula (2.50). We visualize this analogy in Figure 2.17. In the left plot, we depict (2.50). In the right plot, we “zoom-in” into the subspace corresponding to $Q_{2,2}$. Therein, we have the second level discretization of the forward operator, which could be, e.g., a FE or finite difference scheme, and the second level approximation in the stochastic space, which could be, e.g., a (dimension-adaptive) sparse grid quadrature approximation.

Finally, let us look at MFMC sampling. The model hierarchy is given by the different *fidelities*, i.e., $f^{(3)} \rightarrow f^{(2)} \rightarrow f^{(1)}$, which generalizes the hierarchies appearing in global sparse grid and multilevel approximations. Therefore, MFMC makes use of a *between model* hierarchy. We can rewrite the MFMC estimator (2.40) for $m = 3$ as

$$\begin{aligned} \hat{f}_{\text{MFMC}}^{(1)} &= \hat{f}_{M_1}^{(1)} + \gamma_2(\hat{f}_{M_2}^{(2)} - \hat{f}_{M_1}^{(2)}) + \gamma_3(\hat{f}_{M_3}^{(3)} - \hat{f}_{M_2}^{(3)}) \\ &= 0 \cdot \hat{f}_{M_1}^{(3)} - \gamma_3 \cdot \hat{f}_{M_2}^{(3)} - \gamma_2 \cdot \hat{f}_{M_1}^{(2)} + \gamma_3 \cdot \hat{f}_{M_3}^{(3)} + \gamma_2 \cdot \hat{f}_{M_2}^{(2)} + 1 \cdot \hat{f}_{M_1}^{(1)}. \end{aligned} \quad (2.51)$$

We therefore obtain a representation very similar to the sparse grid combination scheme in (2.49). However, in (2.51) we have a *weighted combination* with weights γ_2, γ_3 and 1. Nevertheless, by redefining $\hat{f}'_{M_2}{}^{(3)} := \gamma_3 \cdot \hat{f}_{M_2}^{(3)}$ and the remaining weighted estimators analogously, we have indeed the standard combination scheme (2.49). Hence, we can interpret (2.51) as follows. The MFMC estimator $\hat{f}_{\text{MFMC}}^{(1)}$ is defined as a two-dimensional (weighted) combination scheme




Figure 2.17: Multilevel model hierarchy. In the left plot, we depict a level three multilevel approximation of the quantity of interest, which is equivalent to a 2D combination scheme (2.10), (2.11). On the right, we visualize the subspace associated to the multiindex (h_2, r_2) .

in an abstract 2D space, in which one direction represents the model fidelity and the other the number of MC samples used in the corresponding estimators. We depict this in Figure 2.18. In the left plot, we show how (2.51) is essentially a two-dimensional (weighted) combination scheme. By zooming-in the “subspace” in which the estimator $\gamma_2 \hat{f}_{M_2}^{(2)}$ is computed, we have the model with the second fidelity, $f^{(2)}$, a Gaussian process [149] in this example, which is evaluated at M_2 samples to obtain the (weighted) MC estimator $\gamma_2 \hat{f}_{M_2}^{(2)}$.

Remark: If $f^{(2)}$ would be, for example, a sparse grid-based surrogate, we could zoom-in further and see something similar to (2.49) and Figure 2.16. Thus, with this perspective, MFMC can comprise stacks of model hierarchy layers, each expressed as a (weighted) combination scheme.

Viewing the multilevel estimation of the quantity of interest and the MFMC estimator as (weighted) combination schemes does not change how these estimators are defined. However, thinking about them as combination schemes could facilitate, for example, the design of refinement strategies similar to the dimension-adaptive algorithm 2.1, which would be performed *directly* in the abstract 2D space in which these approximations are defined. Moreover, one can add multiple layers of refinement by, e.g., finding surrogates of the forward model via adaptivity and then computing the multilevel or multifidelity approximations using adaptivity as well. Note that at all these layers, the underlying models will be treated as black boxes. Furthermore, the combination technique perspective could make multilevel estimation applicable also in problems in which a thorough theoretical knowledge of the underlying problem is not available (cf. Section 2.7). In addition, when all these three approaches are employed in large scale applications simulated on supercomputers, the combination technique perspective facilitates also a straightforward incorporation of *fault tolerance*: when compute nodes on which simulations corresponding to one or several “rectangles” as in Figures 2.16, 2.17 or 2.18 fail, the *parallel fault tolerant combination scheme* can be employed (see [88, 89]). It is important to point out that viewing all multilevel and multifidelity schemes as (weighted) combination techniques was facilitated by the fact that we used these approaches to assess a linear quantity of interest, the




Figure 2.18: Multifidelity Monte Carlo model hierarchy. In the left plot, we depict an MFMC estimator with one high-fidelity model and two low-fidelity surrogates, which is equivalent to a 2D combination scheme (2.10), (2.11). On the right, we visualize the subspace associated to the multiindex $(f^{(2)}, M_2)$.

expectation of the high-fidelity model. If we want to estimate a nonlinear operator, such as the variance, this analogy breaks.

Remark: The global sparse grid quadrature approximation was chosen to be two-dimensional for convenience, to show the analogy with the multilevel and MFMC approximations. However, when viewing the multilevel or MFMC estimations as combination schemes, they are *always* 2D, as we discussed above.

To summarize, the three classes of model hierarchies considered in this work, global sparse grid approximations, multilevel decompositions and multifidelity sampling, can be viewed as (weighted) combination schemes. In sparse grid approximations, the same model is used in all computations, the hierarchy being given by using different sparse grid levels. Multilevel decompositions go one step further and introduce a grid-based hierarchy of models which is used together with a hierarchy of approximations of the quantity of interest. Finally, multifidelity approaches generalize the previous two approaches by employing a hierarchy of different models with different fidelities. In this work, we design four context-aware algorithms based on the aforementioned model hierarchies to address higher-dimensional, computationally expensive and practically relevant UQ problems. Before we begin presenting our proposed approaches and results, let us first review the two applications of interest in this work.

"...in real life mistakes are likely to be irrevocable. Computer simulation, however, makes it economically practical to make mistakes on purpose. If you are astute, therefore, you can learn much more than they cost. Furthermore, if you are at all discreet, no one but you need ever know you made a mistake. "

John H. Mcleod

3

Applications background

Our goal in this work is to formulate novel context-aware algorithms that overcome, at least to a certain extent, the challenges of quantifying uncertainty in higher-dimensional, computationally expensive, real-world problems. As representative examples of such applications, we consider fluid-structure interaction (FSI) and plasma microturbulence simulations, which we summarize next. Both applications are characterized by complex mathematical models and complex simulation codes, and feature parameter sets affected by uncertainty due to, e.g., measurement errors. Thus, performing realistic numerical simulations requires a proper treatment of uncertainty *ab initio*. We employ these two applications to test our proposed algorithms for uncertainty propagation in Chapters 4 – 6. In Chapter 7, where the focus is on Bayesian inference, as a representative example of a more complex application, we consider an inverse problem with an elliptic forward model defined on a three-dimensional domain. These models are relevant, e.g., in geotechnical engineering (steady-state groundwater flow problems), in which the elliptic forward model couples the permeability field with the hydrostatic pressure, see, e.g., [73].

The forward UQ algorithms proposed in this work are not directly designed for the two aforementioned applications in particular, but they are aimed for quantifying uncertainty in a broad spectrum of higher-dimensional, computationally expensive, real-world problems, of which FSI and plasma microturbulence are representative examples. Thus, our focus in the following is not on modelling. We present only a summary of the mathematical models used in these two applications and refer the reader to more detailed references. In Section 3.1, we summarize FSI and in Section 3.2, we summarize plasma microturbulence analysis. We try to keep the notation consistent with the standard literature.

3.1 Fluid-structure interaction problems

In this section, we briefly summarize FSI problems, which we use to test the proposed Multilevel spectral projection with online stochastic dimensionality reduction, detailed in Chapter 5. In Section 3.1.1, we provide the reader with example applications in which FSI is relevant and discuss why quantifying uncertainty in FSI simulations is necessary. In Section 3.1.2, we briefly summarize the models used in FSI. The codes that we use for FSI simulations are overviewed in Section 3.1.3. The fluid flow and structural mechanics models are solved numerically using the multi-physics code `Alya`, and their coupling is done via the partitioned approach implemented in the coupling library `preCICE`. We end, in Section 3.1.4, with a review of the mathematical and computational challenges of FSI, grouped under the umbrella of *multi-challenges*, and summarize the solutions used in this work. For a more detailed introduction to FSI, we refer the reader to [4, 8]. Moreover, the dissertation [69] also provides a good overview of the modelling and computational aspects of FSI.

3.1.1 Overview

FSI problems, in which one or more structures interact with an internal or surrounding fluid flow, play a prominent role in many domains and pose many challenges due to their non-linearity and multidisciplinary nature. Representative examples range from human (patho)physiology, that is, the study of blood flow, calcification, and aneurysms (see [113]), to aeronautics. For the latter, flow-induced vibrations have an important influence on the stability and durability of aircrafts [58]. Parachutes [177] or noise prediction [158] are other examples, to only name a few. We study FSI not only due to its practical relevance, but also as a representative of a challenging and computationally expensive multi-physics application. In general, FSI features several physical parameters which suffer from uncertainty as they typically stem from measurements: both fluid and structure density, the viscosity of the fluid, the structural elasticity and plasticity parameters, possibly different in every spatial dimension. This can sum up to several stochastic parameters, which, given the high computational cost of a single FSI simulation, makes the quantification of uncertainty computationally challenging.

3.1.2 Modelling fluid-structure interaction

In this work, we are concerned with quantifying uncertainty in FSI problems characterized by *incompressible*, *Newtonian* flows which interact with an *elastic* and *compressible* structure. Recall that Newtonian fluids, such as water or air, are characterized by a linear relationship between stress and strain rate and a vanishing strain rate for zero stress. The fluid flow is defined on a two-dimensional spatial domain, $\Omega_f \subset \mathbb{R}^2$. Moreover, the elastic structure is defined on a 2D domain $\Omega_s \subset \mathbb{R}^2$ as well. The time domain is $[0, t_{\text{end}}] \subset \mathbb{R}$, where $t_{\text{end}} > 0$. In the following, we briefly summarize the models used for incompressible flows and structural mechanics, the constitutive laws used in structural mechanics, and, finally, we discuss the coupling of incompressible flow and structural mechanics models into an FSI simulation.

Incompressible flows

The fluid flow, assumed to be incompressible and Newtonian, has a state characterized by the pressure, p_f , and the velocity field, $\mathbf{v}_f := (v_{f,x}, v_{f,y})$. The dynamics are generally modelled by the incompressible Navier-Stokes equations, written in an *Eulerian* framework: the state variables, p_f and \mathbf{v}_f , are functions of time, $t \in [0, t_{\text{end}}]$, and space, $(x, y) \in \Omega_f$, the task being to characterize their dynamics from the perspective of a *fixed* observer (a *fixed* frame of reference). The Navier-Stokes equations for incompressible flows comprise the *mass conservation* equation:

$$\nabla \cdot \mathbf{v}_f = 0, \quad (3.1)$$

and the *momentum conservation* equation:

$$\rho_f (\dot{\mathbf{v}}_f + (\mathbf{v}_f \cdot \nabla) \mathbf{v}_f) = -\nabla p_f + \mu_f \nabla^2 \mathbf{v}_f + \rho_f \mathbf{f}_f, \quad (3.2)$$

where μ_f is the *shear viscosity* of the fluid and \mathbf{f}_f comprises the distributed volume forces that act upon the fluid, such as gravity. Note that the Eulerian perspective entails having a *convective* term in the momentum conservation equation, i.e., the term $(\mathbf{v}_f \cdot \nabla) \mathbf{v}_f$ in (3.2). For a more detailed introduction for fluid dynamics, see, e.g., [59].

Structural mechanics

The structure is assumed to be elastic and compressible, with a configuration described in terms of its displacement, $\mathbf{u}_s := (u_{s,x}, u_{s,y})$, and velocity, $\mathbf{v}_s := \partial \mathbf{u}_s / \partial t := (\partial u_{s,x} / \partial t, \partial u_{s,y} / \partial t)$. The

structural dynamics equation is analogous to the momentum conservation equation (3.2) in fluid dynamics, but usually expressed using a Lagrangian perspective, that is, from the perspective of an observer following the movement of the material domain (a *moving* frame of reference). Mathematically, this means that the model characterizing the dynamics of the structure does not have a convective term:

$$\rho_s \frac{\partial^2 \mathbf{u}_s}{\partial t^2} = \nabla \cdot \mathbf{S} + \rho_s \mathbf{f}_s, \quad (3.3)$$

where ρ_s is the material density. Body forces are modelled by the *2nd Piola-Kirchhoff stress tensor*, \mathbf{S} , whereas body forces, such as gravity are contained in the distribution force, \mathbf{f}_s . For more details about structural mechanics, we refer the reader to [7, 16].

Constitutive Equations

In this work, we consider a *Saint-Venant-Kirchhoff* material model, characterized by the constitutive law

$$S_{ij} = \lambda E_{kk} \delta_{ij} + 2\mu E_{ij}, \quad (3.4)$$

where \mathbf{S} denotes the 2nd Piola-Kirchhoff stress tensor, \mathbf{E} denotes the *strain tensor*, $\boldsymbol{\delta}$ is the *unit tensor* and the two scalars, λ and μ denote the so-called *Lamé* constants. *Lamé* constants are related to material properties such as *Young's modulus*

$$E = \frac{\mu(3\lambda + 2\mu)}{\lambda + \mu}$$

and *Poisson's ratio*

$$\nu_s = \frac{\lambda}{2(\lambda + \mu)}.$$

Coupling of Fluids and Structures

The incompressible flow equations (3.1) - (3.2) are expressed using an Eulerian perspective. In contrast, structural dynamics (3.3) are expressed in a Lagrangian framework. Since the Eulerian and Lagrangian perspectives are incompatible, the coupling of these two models in an FSI problem is not possible in this form. To this end, the *arbitrary Lagrangian-Eulerian* (ALE) approach [48] was proposed. Briefly, in ALE, a third frame of reference is introduced, called a *referential domain*, which results in a mesh motion in addition to the material motion in the underlying spatial domain. This motion leads to a *convective velocity* which is the relative velocity between material (Lagrangian) and referential domains, as seen from the spatial (Eulerian) coordinate system. Denoting by \mathbf{v}_r the mesh velocity in the ALE framework and by $\mathbf{c}_f := \mathbf{v}_f - \mathbf{v}_r$ the resulting convective velocity, the incompressible Navier-Stokes equations written in the ALE framework (see, e.g., [69]) read:

$$\begin{aligned} \rho_f \nabla \cdot \mathbf{c}_f &= 0 \\ \rho_f (\dot{\mathbf{v}}_f + (\mathbf{c}_f \cdot \nabla) \mathbf{v}_f) &= -\nabla p_f + \mu_f \nabla \mathbf{v}_f + \rho_f \mathbf{f}_f. \end{aligned}$$

We next summarize how the fluid flow model and the structural mechanics model (3.3) with constitutive law (3.4) are coupled in the context of FSI. The coupling of the fluid domain Ω_f with boundary $\Gamma_f = \partial\Omega_f$ and of the structural domain Ω_s with boundary $\Gamma_s = \partial\Omega_s$ is done at their interface, usually called the *wet* surface, i.e.,

$$\Gamma_{fs} = \Gamma_f \cup \Gamma_s,$$

where we need to impose *kinematic* as well as *dynamic* interface conditions.

The kinematic interface conditions impose that displacements and velocities for the fluid and structure are the same at Γ_{fs} , i.e.,

$$\mathbf{u}_f = \mathbf{u}_s, \quad \mathbf{v}_f = \frac{\partial \mathbf{u}_s}{\partial t} \quad \text{at } \Gamma_{fs}. \quad (3.5)$$

That is to say, (i) the fluid and structure continuously fill their domains up to the wet surface, while neither overlaps nor gaps of material can occur (displacements are equal) and (ii) fluid molecules in contact with the structural surface are bound to it by molecular attraction forces (velocities are equal). Besides kinematic conditions, we also need dynamic interface conditions, that is, a balance of forces (expressed point-wisely) by the surface (Cauchy) stresses $\boldsymbol{\sigma}$:

$$\boldsymbol{\sigma}_f \cdot \mathbf{n}_f = -\boldsymbol{\sigma}_s \cdot \mathbf{n}_s \quad \text{at } \Gamma_{fs}, \quad (3.6)$$

where \mathbf{n}_f and \mathbf{n}_s and normal vectors with opposite directions.

To couple fluid flow and structural mechanics solvers into a single FSI simulations, there exists two main approaches: the *monolithic* approach and the *partitioned approach*. To provide the reader with a short description of these two approaches, we follow [69]. Briefly speaking, monolithic coupling approaches are driven by the idea of a closed discretization of the overall FSI system. Let \mathbf{y}_f and \mathbf{y}_s denote the flow and structural variables, respectively. The result of a monolithic discretization of an FSI problem is a *single* system, comprising the discrete flow variables, which can be abstractly written as

$$\mathbf{A}(\mathbf{y}_f, \mathbf{y}_s) = 0. \quad (3.7)$$

In contrast, the partitioned approach to FSI is based on a domain decomposition approach, in which the domain is split into a fluid and a structure component. Moreover, we assume that the ALE point-of-view is used for the fluid solver, thus the mesh motion scheme can be seen as a third component. As FSI is a surface-coupled problem, these three components are coupled by their common boundary values, that is, dynamic and kinematic boundary values, as described in (3.5) and (3.6). Let \mathbf{F} and \mathbf{S} generally denote the fluid flow and structure operators, with inputs \mathbf{f} and \mathbf{s} , respectively: \mathbf{f} denotes the vector of dynamic values, i.e., forces or stresses, whereas \mathbf{s} comprises displacements and velocities. With this notation, the partitioned approach to FSI translates into a *Dirichlet-Neumann* decomposition

$$\begin{aligned} \mathbf{F}(\mathbf{s}) &= \mathbf{f} \\ \mathbf{S}(\mathbf{f}) &= \mathbf{s}. \end{aligned} \quad (3.8)$$

We assume the dynamic mesh component to be formally included in the fluid flow operator, \mathbf{F} .

There are several numerical schemes to perform monolithic (3.7) or partitioned (3.8) FSI. Examples include *explicit* schemes, such as conventional serial or parallel staggering and their improved or generalized versions, multilevel-schemes, and *implicit* approaches, such as Schwarz procedures, constant or Aitken relaxation, vector extrapolation or Newton-Raphson methods. For more details on these methods, see [69] and the references therein.

Choosing between the monolithic and the partitioned approach to perform FSI simulations usually depends on the underlying problem. For example, works such as [183] established that the partitioned approach can be used to efficiently couple multi-physics solvers in strongly-coupled problems, such as hemodynamics, although the monolithic approach is usually employed in such problems. However, choosing between the monolithic and the partitioned approach is not our concern here, since we assume that the underlying simulation code is available as a

black-box, our goal being to efficiently quantify the input uncertainty. Therefore, in this work, we will use the partitioned approach to FSI, which is provided by the high-performance coupling library `preCICE`, summarized in Section 3.1.3. This concludes our brief summary of the models used for incompressible flows and structural mechanics. In the following, we present the two simulation codes used for FSI in this work.

3.1.3 The fluid-structure interaction simulation codes

The multi-physics simulation code `Alya`

The fluid flow and structural equations are solved numerically using the `Alya`¹ [188] multi-physics code. `Alya` system is a high-performance simulation code developed at the Barcelona Supercomputing Center in Spain and it is designed to efficiently solve complex coupled multi-physics / multi-scale / multi-domain problems such as incompressible/compressible flows, non-linear solid mechanics, chemistry, particle transport, heat transfer, turbulence modelling or electrical propagation. Inter-node parallelization is done via `MPI`² - for communication, and using `METIS`³ - for problem sub-division. Moreover, internally, `Alya` uses `OpenMP`⁴ directives to distribute the code loops amongst threads. Spatial discretization is done via FE while for time discretization, both explicit and implicit time integrators can be employed.

The three `Alya` modules that we are using in our simulations are `NASTIN`, for incompressible flow solvers, and `SOLIDZ`, the module for solid mechanics solvers. To allow for moving geometries, `Alya` employs the ALE approach. Finally, we also use the `PARALL` module, which provides the functionality to perform parallel simulations. For more details about the `Alya` code, we refer the reader to [188].

The coupling library `preCICE`

The `Alya` system, summarized in the previous section, is used to perform (parallel) fluid flow and structural mechanics simulations. To couple these two solvers into a single FSI simulation, we employed the multi-physics coupling library, `preCICE`⁵ [21]. For more details about the `Alya` - `preCICE` coupling, see [184].

`preCICE`, which stands for *Precise Code Interaction Coupling Environment*, is a coupling library for partitioned multi-physics simulations such as FSI or conjugate heat transfer. The main features of `preCICE` are:

- it employs fully parallel peer-to-peer communication: coupled solvers can directly exchange information without requiring a central instance. In this way, all coupling operations are executed directly on compute resources used by the solvers, thus reducing the communication overhead and enabling efficient simulations on parallel supercomputers;
- it is a library: `preCICE` was designed as a library, not as a framework, meaning that the coupled solvers call `preCICE`, not the other way around, making the coupling process minimally invasive;
- it comprises fast and robust coupling algorithms: `preCICE` employs state-of-the-art coupling algorithms, such as quasi-Newton schemes [129], which enable the portioned coupling even in strongly-coupled problems, such as hemodynamics;

¹<https://www.bsc.es/es/computer-applications/alya-system>

²<http://www.mpi-forum.org/>

³<http://glaros.dtc.umn.edu/gkhome/views/metis>

⁴<http://openmp.org/wp/>

⁵<https://www.precice.org/>

- it supports multi-coupling: `preCICE` can couple an arbitrary number of solvers, for example, fluid-structure-structure or fluid-structure-acoustics interaction.

To couple the fluid flow and structure solvers, we use a Dirichlet-Neumann implicit partitioned approach provided by `preCICE` and sub-iterate until convergence in every time step. A quasi-Newton scheme [129] is used to stabilize and accelerate the sub-iterations. For a more detailed overview of `preCICE`, the reader is referred to [21, 69, 183].

3.1.4 The multi-challenge

We saw that from a modelling perspective, multi-physics problems, such as FSI, are challenging because they bring together different models, e.g., fluid flow plus structural mechanics in FSI. Moreover, the governing conservation or constitutive laws of these problems are typically represented using different perspectives. We saw that fluid flows are often represented using an Eulerian frame of reference, whereas structural mechanics equations are generally written using the Lagrangian perspective. Furthermore, obtaining realistic multi-physics numerical simulations requires high-performance computing, hence adding another challenge to FSI. These challenges were discussed for example, in [101], and were grouped together as *multi-challenges*:

1. the *multi-physics* challenge, due to having multiple mathematical models;
2. the *multi-domain* challenge, since, generally, different frameworks are used to express the governing conservation or constitutive laws;
3. the *multi-core* challenge, due to the need of high-performance computing to obtain realistic numerical simulations.

In this work, the first two multi-challenges are addressed by coupling the fluid flow and structural mechanics models in an FSI problem using the partitioned approach and the ALE formalism to allow for moving geometries. The third multi-challenge is resolved by using high-performance simulation codes, `Alya` and `preCICE`.

Our goal in this work is to quantify uncertainty in complex and computationally expensive applications, such as FSI. This adds another multi-challenge to FSI,

4. the *multi-dimensional* challenge,

which is due to having, in general, higher-dimensional uncertain parameter sets. The multi-dimensional challenge is also related to the multi-core challenge, since non-intrusive UQ methods, which are the theme of this work, require ensemble of model evaluations. Our solution to the multi-dimensional challenge is presented in Chapter 5, where we formulate the Multilevel adaptive spectral projection with online stochastic dimensionality reduction.

In this section, we summarized FSI, the first of the two complex and computationally challenging applications of interest for this work. Next, we summarize our second application of interest, the simulation of plasma microturbulence.

3.2 Plasma microturbulence analysis

Another complex and practically relevant application considered in this work is the simulation of microturbulence in magnetized fusion plasmas. We use plasma microturbulence simulation scenarios to test the proposed Sensitivity-driven adaptive sparse grid algorithm in Chapter 4 and the Context-aware multifidelity sampling algorithm in Chapter 6. Section 3.2.1 overviews the

practical relevance of studying fusion plasmas and some challenges related to their mathematical modelling. In Section 3.2.2, we summarize the most popular approach for modelling and aspects of its numerical simulation, and we emphasize the need for UQ in such problems. Finally, we introduce the employed simulation code, GENE, in Section 3.2.3.

3.2.1 Overview

The simulation of plasma microturbulence is of high practical relevance for efforts such as the ITER⁶ experiment. The aim of the ITER is to create, for the very first time, a self-sustained (“burning”) plasma in the laboratory. This ambitious effort amounts to a milestone on the way towards harnessing green and practically unlimited energy from a fusion power plant. However, even though the mathematical theory that governs this process is well understood, fusion power plants are not a reality yet due to some physical and technological hurdles. For example, a physics obstacle is the fluctuation that appears at the small-scale, which causes energy loss rates despite sophisticated plasma confinements via strong and shaped magnetic fields. This so-called *microturbulence* is driven by the free energy given by the unavoidably steep plasma temperature and density gradients. The associated turbulent transport determines the energy confinement time which in turn is a key ingredient for creating a burning plasma in the laboratory. Any insight into the nature of microturbulence and ways for avoiding turbulence related confinement degradations are therefore crucial for the design of fusion power plants. Since the plasma temperature and density gradients, as well as further physical parameters affecting the underlying microinstabilities are subject to uncertainties, this calls for numerical simulations in which uncertainty is taken into account *ab initio*.

The study of turbulence is in general not an easy task, to say the least. In classical physics, as Richard Feynman famously said, turbulence remains one of the most important unsolved problems. In magnetically confined plasmas, that is, very hot but dilute ionized gases, the problem is further worsened by the low collisionality of the component species (deuterium ions, electrons etc.). This renders (macroscopic) fluid descriptions insufficient in many situations. Even more, a (microscopic) particle model, which would provide the most in-depth description of the plasma, is computationally infeasible even on the most powerful supercomputer at the time of writing this thesis, since a plasma usually consists of about $O(10^{20})$ particles. Thus, a descriptive-enough model which can lead to computationally feasible simulation codes is a *kinetic* model. A kinetic model for fusion plasmas is the 6D Vlasov model (three positions, three velocities in the state space), with one equation per component species, which are coupled via Maxwell’s equations. However, modelling the dynamics of the particles distribution function in six-dimensional state space remains computationally infeasible when the corresponding solutions are sought in complex geometries, which is typically the case for a fusion reactor such as a tokamak or stellerator⁷. This calls for reduced kinetic models, which we summarize next.

3.2.2 The gyrokinetic approach and its numerical simulation

The most popular theory for assessing plasma microturbulence is the so-called *gyrokinetic* theory [17, 107], which reduces the dimension of the state space from six (three positions, three velocities) to five (three positions, two velocities). Such a reduction makes sense because of a time-scale separation between the fast gyromotion of the particles around the magnetic field lines and typical turbulence time scales. Hence, the knowledge of the exact position of the par-

⁶ITER was initially an acronym for International Thermonuclear Experimental Reactor. Nowadays it is mainly referred to the Latin word “iter”, meaning “the way”; see <https://www.iter.org/>

⁷See <https://www.ipp.mpg.de/14731/anlagentypen> for further details on fusion devices.

ticles along their orbit is not relevant. Gyrokinetics therefore effectively removes the gyrophase information and yields a 5D system of equations, which is more suitable for nowadays computational resources and is considered a valid approach for a wide range of plasma parameters.

In the gyrokinetic theory, each particle species, s , is characterized by a distribution function

$$\pi_s(\mathbf{R}, v_{\parallel}, \mu_m),$$

where $\mathbf{R} = (R_x, R_y, R_z)$ describes the *position* of the gyrocenter in real space, whereas v_{\parallel} denotes the *parallel velocity* and $\mu_m := m_s v_{\perp}^2 / 2B$ is the *magnetic moment*, where m_s denotes the mass of species s , v_{\perp} is the perpendicular component of the velocity and B denotes the magnetic field. Examples of particle species, which we also use in our numerical simulations in Chapters 4 and 6, are deuterium ions and electrons. We note that other two velocity coordinates can be chosen to describe the distribution functions, which would, of course, modify the form of the equations that we summarize below. The dynamics of the distribution function π_s is modelled by the *gyrokinetic Vlasov equation* with collisions:

$$\dot{\pi}_s + \dot{\mathbf{R}} \cdot \nabla \pi_s + \dot{v}_{\parallel} \frac{\partial \pi_s}{\partial v_{\parallel}} + \dot{\mu}_m \frac{\partial \pi_s}{\partial \mu_m} = C(\pi_s, \pi_{s'}),$$

where $C(\pi_s, \pi_{s'})$ is a collision operator. The dynamics of the coordinates of π_s read

$$\begin{aligned} \dot{\mathbf{R}} &= v_{\parallel} \mathbf{b}_0 + \frac{B_0}{B_{0\parallel}^*} (\mathbf{v}_E + \mathbf{v}_{\nabla B} + \mathbf{v}_c) \\ \dot{v}_{\parallel} &= -\frac{\dot{\mathbf{R}}}{m_s v_{\parallel}} \cdot \left(q_s \nabla \bar{\Phi}_1 + \frac{q_s}{c} \frac{\partial \bar{A}_{1\parallel}}{\partial t} \mathbf{b}_0 + \mu_m \nabla (B_0 + \bar{B}_{1\parallel}) \right) \\ \dot{\mu}_m &= 0, \end{aligned}$$

where B_0 denotes the modulus of the equilibrium magnetic field vector \mathbf{B}_0 , $\mathbf{b}_0 := \mathbf{B}_0 / B_0$ is the corresponding unit vector, $B_{0\parallel}^* := \mathbf{b}_0 \cdot \mathbf{B}_0^*$ is the parallel component of $\mathbf{B}_0^* := \mathbf{B}_0 + \nabla \times (\mathbf{B}_0 v_{\parallel} / \Omega_s)$, where $\Omega_s := q_s B_0 / m_s c$ is the gyrofrequency of species s with mass m_s and charge q_s ; c is the speed of light. $\bar{B}_{1\parallel}$ denotes the gyro-averaged modulus of the parallel component of the perturbed magnetic field \mathbf{B}_1 , $\bar{\Phi}_1$ denotes the gyro-averaged perturbed part of the electrostatic potential, Φ , and $\bar{A}_{1\parallel}$ is the gyro-averaged modulus of the parallel component of the perturbed vector potential, A . The three characteristic drift terms entering the dynamics of \mathbf{R} satisfy

$$\begin{aligned} \text{Generalized } \mathbf{E} \times \mathbf{B} \text{ drift velocity: } \quad \mathbf{v}_E &= \frac{c}{B_0^2} \mathbf{B}_0 \times \nabla \bar{\chi}_1, \\ \text{Gradient-B drift velocity: } \quad \mathbf{v}_{\nabla B} &= \frac{\mu_m c}{q_s B_0^2} \mathbf{B}_0 \times \nabla B_0, \\ \text{Curvature drift velocity: } \quad \mathbf{v}_c &= \frac{v_{\parallel}^2}{F_s} (\nabla \times \mathbf{b}_0)_{\perp}, \end{aligned}$$

where $\bar{\chi}_1$ is the gyro-averaged perturbed part of the generalized potential, χ .

To obtain a complete description, we also need the *gyrokinetic field equations* for the perturbed electrostatic potential, Φ_1 , and for the perturbed electric potential, A_1 , obtained from Maxwell's equations, i.e.,

$$\begin{aligned} \text{Poisson's equation:} \quad -\nabla^2 \Phi_1 &= 4\pi \sum_s q_s n_{1s}, \\ \text{Ampère's law:} \quad -\nabla^2 A_1 &= \frac{4\pi}{c} \sum_s n_{1s} \mathbf{u}_{1s}, \end{aligned}$$

where n_{1s} is the 0th space moment and \mathbf{u}_{1s} the 1st order velocity moment of π_s .

Several numerical implementations have been developed over the last two decades and encouraging progress has been made [68, 107]. Earlier studies were often limited to restricted physics, e.g., adiabatic electrons and simplified geometries, which could only yield qualitative results and predictions. However, the tremendous progress in computing in the recent past enabled the (realistic) numerical simulation of plasma microturbulence using the gyrokinetic equations. Therefore, the complexity of the gyrokinetic simulations has dramatically increased and flagship codes aim for quantitative validation with various observables obtained from experimental measurements, see e.g., [191] and references therein. Some simulation codes aim for a full flux-driven setup, where profiles and turbulence are self-consistently developing in response to prescribed heat and particle source. However, these simulations are usually too costly for typical applications and therefore need to be performed with reduced physics.

An alternative scheme is to use the experimentally determined mean temperature and density profiles as well as the magnetic topology in a given time window as fixed physics inputs to the gyrokinetic codes and compute the resulting turbulent fluctuations. Naturally, all these physics inputs but also the experimental fluctuation observables are affected by measurement noise from, e.g., the experimental diagnostics used to acquire the data. Note that this can be quite troublesome since plasma profiles are often found to be quite chaotic, that is, a small increase in inputs such as the gradients may cause large differences in the resulting turbulent heat fluxes. The simulation parameters are taken then in an interval with bounds stemming from subtracting (left bound) and adding (right bound) the (assumed or experimentally measured) noise from/to the nominal mean value. After the intervals for all parameters are determined, the standard practice in the plasma physics community is to identify the most sensitive input parameters and use only these in subsequent simulations. Both steps are typically done via so-called *parameter scans*: equidistant points for all parameters' intervals are tensorized and for each such combination, a corresponding simulation is performed. After the most important parameters are identified, the simulations are performed using scans of only these parameters. It is easy to see that parameter scans suffer from the curse of dimensionality: a scan is essentially a fully tensorized grid with an exponentially increasing number of degrees of freedom with the dimensionality. Therefore, given the enormous computational efforts associated with high-dimensional parameter scans, the sensitivity analysis step is often performed without considering the nonlinearities in the Vlasov equation which can take up to 50% of the runtime.

To this end, in a first step, linear gyrokinetic simulations are typically performed. Such linear simulations are highly relevant to characterize the underlying microinstabilities such as *ion* or *electron temperature gradient* (ITG/ETG) driven modes, *trapped electron* modes (TEMs), *micro-tearing* modes (MTM), and many more. Determining their threshold values or transitions usually provides some guideline for input parameters scans in fully nonlinear simulations, which are computationally much more expensive, e.g., in the order of 10000 CPU hours/simulation. In this work, we restrict our focus on linear local (flux-tube) plasma microturbulence simulations, while nonlinear gyrokinetic simulations will be targeted in our future research. Given that standard parameter scans employed in the plasma physics community suffer from the curse of dimensionality, this motivates the development of modern UQ methods that exploit the structure of the underlying problem to reduce the overall computational cost, which we address in this work. We end this section with an overview of the employed simulation code.

3.2.3 The plasma microturbulence code GENE

The employed gyrokinetic solver is the GENE code, one of the first grid-based, i.e., Eulerian codes in this field, which has now been under continuous development for almost two decades [97].

GENE computes the time evolution of gyrocenter distribution functions on a fixed grid Ω in a 5D position-velocity state space. To discretize this grid we therefore need

$$N_x \times N_y \times N_z \times N_{v_{\parallel}} \times N_{\mu_m}$$

grid points in total, for all five coordinates of the distribution function, π_s . The underlying nonlinear gyrokinetic PDEs are solved via a mix of numerical methods also widely used in computational fluid dynamics, including finite difference, spectral, FE, and finite volume schemes. More details can be found in, e.g., [77]. GENE had been originally restricted to flux-tube simulation domains [9], i.e., thin magnetic field lines following boxes which allow for highly efficient simulations provided that the radial correlation lengths of the turbulent fluctuations are small compared to the profile scale lengths. With this setup the radial variations of the profiles as well as their gradients can be assumed to be constant across the simulation domain. Consequently, periodic boundary conditions in the directions perpendicular to the magnetic field line can be assumed which allows the usage of spectral methods; spectral methods can greatly simplify operators such as gyroaverages.

However, for applications in small devices or with steep profiles, locality cannot be assumed anymore. Moreover, profiles periodicity is lost as well by considering full radial profiles, at least in the radial direction. Numerically this translates into the need for alternatives to spectral methods. To this end, finite differences have been added as another option in GENE which allow for radially global simulations [77] or full flux-surface simulations [198], if toroidal instead of radial background variations are considered. The field-aligned non-orthogonal coordinate system has, however, been kept to exploit the high anisotropy of plasma turbulence which typically has correlation lengths of several meters along a magnetic field line but only of a few centimetres in the perpendicular plane. GENE simulations are parallelized by domain decomposition in all five dimensions, typically using MPI. Note, however, that in spite of all aforementioned efforts to speed up gyrokinetic simulations, fully nonlinear UQ simulations remain very challenging since a single such run may be computationally very expensive. That is why in this work, we focus on linear local (flux-tube) UQ simulations, while turbulence assessments will be targeted in our future research.

The form of the gyrokinetic equations implemented in GENE read

$$\dot{\boldsymbol{\pi}}_s = \mathbb{O}_{\text{lin}}(\boldsymbol{\pi}_s) + \mathbb{O}_{\text{nonlin}}(\boldsymbol{\pi}_s),$$

where $\boldsymbol{\pi}_s$ is a vector holding the linearized grid values of the transformed five-dimensional distribution function of species s on a discretization grid Ω . Moreover, \mathbb{O}_{lin} and $\mathbb{O}_{\text{nonlin}}$ are the linear and nonlinear parts of the gyrokinetic equations.

As stated above, in this work, we are interested in linear flux-tube simulations, which are simply obtained by neglecting the nonlinear part in the above equation, i.e.,

$$\dot{\boldsymbol{\pi}}_s = \mathbb{O}_{\text{lin}}(\boldsymbol{\pi}_s).$$

The discrete form of the operator \mathbb{O}_{lin} is a matrix, O_{lin} , i.e.,

$$\dot{\boldsymbol{\pi}}_s = O_{\text{lin}} \boldsymbol{\pi}_s.$$

O_{lin} typically up to approximately 30% non-zero entries and a rank of at least a few hundred thousand; for example, in one of our simulations, the rank was 552960. The microinstability mode is typically revealed by the so-called dominant eigenmode, i.e., the complex eigenvalue of O_{lin} corresponding to the first eigenvector. For an efficient computation of dominant eigenmodes, GENE employs solvers from the high-performance linear algebra libraries `Petsc/Slepc`⁸.

⁸<https://www.mcs.anl.gov/petsc/>, <http://slepc.upv.es/>

In this section, we provided the reader with an overview of the two applications of interest in this work, FSI and plasma microturbulence simulations. We consider FSI simulations defined on two-dimensional spatial domains, in which the fluid flow is governed by the incompressible Navier-Stokes equations, formulated in the ALE approach to allow for moving geometries. A non-linear Saint-Venant-Kirchhoff model governs the elastic structure. The flow and structure domains are discretized via FE. To solve numerically these two models, we use the code `Alya`. The FE meshes match at the boundary, where a balancing between stresses and displacements is enforced. To couple the fluid flow and structure solvers, we use a Dirichlet-Neumann implicit partitioned approach and sub-iterate until convergence in every time step; this functionality is provided by the coupling library `preCICE`. A quasi-Newton scheme is used to stabilize and accelerate the sub-iterations. We use FSI simulations to test the algorithm proposed in Chapter 6. Therein, we consider a stochastic scenario with five uncertain inputs.

The plasma microturbulence simulations are defined on a five-dimensional space, governed by the gyrokinetic Vlasov equation with collisions, coupled with Maxwell's equations. In this work, our attention falls on linear eigenvalue problems, which are used to characterize the nature of the microturbulence. To perform these numerical simulations, we employ the plasma physics code `GENE`. We employ two plasma microturbulence test cases with three, eight or 12 uncertain inputs in Chapters 4 and 5. In Chapter 7, where the focus is on Bayesian inference, as a representative example of a more complex application, we consider an inverse problem with an elliptic forward model defined on a 3D domain.

We stress again that the UQ algorithms proposed in this work are not directly designed for these applications in particular. Our algorithms are aimed for quantifying uncertainty in a broad spectrum of higher-dimensional, computationally expensive, real-world applications, of which FSI and plasma microturbulence are representative examples. In the next chapter, we present our first contribution, a novel refinement strategy based on sensitivity scores for dimension-adaptive sparse grid approximations for uncertainty propagation. So, let us begin!

PART II

CONTEXT-AWARE SPARSE GRID AND MULTILEVEL APPROXIMATIONS FOR UNCERTAINTY PROPAGATION

"Probability is not a mere computation of odds on the dice or more complicated variants; it is the acceptance of the lack of certainty in our knowledge and the development of methods for dealing with our ignorance."

Nassim Nicholas Taleb

4

Sensitivity-driven dimension-adaptive sparse grid approximations

In this chapter, we present our first contribution. Our goal is to exploit the anisotropic coupling of the stochastic inputs and the (lower) intrinsic dimensionality, which appear in most higher-dimensional uncertainty propagation problems, to drive the adaptive refinement. Thus, in this chapter we formulate a *structure-exploiting, context-aware* algorithm based on *model hierarchies* in sparse grid approximations aimed for quantifying uncertainty in higher-dimensional, computationally expensive problems. The proposed algorithm is based on dimension-adaptive refinement and it can be used in a broad spectrum of applications provided that only a few assumptions are fulfilled. As a practically relevant example, we consider the numerical simulation of plasma microturbulence.

Some of the ideas and results presented in this chapter are presented our work [54]. Therein, we formulated the Sensitivity driven dimension-adaptive sparse grid algorithm, presented in Section 4.2, in which $d_{\text{sto}} + 1$ sensitivity scores were computed, i.e., one for each input direction, resulting in d_{sto} scores, and the last score was computed for all stochastic input interactions. We applied the new algorithm to the two plasma microinstability test cases presented in Section 4.3.1 and Sections 4.3.3 and 4.3.4. Moreover, some of our preliminary results can be found in [78], where we used a dimension-adaptive sparse pseudo-spectral approximation method (adaptive SPAM) in the ASDEX Upgrade test case. The remaining parts presented in this chapter represent new results: we extended the sensitivity-driven algorithm to computing sensitivity scores for all stochastic input interactions (see Algorithm 4.2), which we applied in all our numerical results in Sections 4.3 and 4.4. In addition, another novelty of this chapter is that we perform uncertainty and sensitivity analyses in the the ASDEX Upgrade test case for another output of interest, i.e., the underlying microinstability mode. This UQ analysis offers novel insights into the impact of input uncertainties in plasma microinstability analysis.

This chapter is organized as follows. In Section 4.1, we discuss (i) the practical relevance of simulating microturbulence in fusion plasmas and (ii) how uncertainty is intrinsic in such phenomena, therefore motivating the need for numerical simulations in which uncertainty is incorporated *ab initio*. Performing uncertainty propagation and sensitivity analysis in such applications is challenging due to the large number of stochastic parameters and high associated computational costs. To address these challenges, we propose a *Sensitivity-driven dimension-adaptive sparse grid algorithm*. The proposed approach is a *deterministic, non-intrusive, context-aware* approach based on sparse grid *model hierarchies* in which sensitivity information about the *individual stochastic parameters* and their *coupling* is used to drive the adaptive process. Our algorithm is formulated as a methodology independent of approximation operators and point sets, provided that the sensitivity information needed for refinement can be obtained. In this work, we consider interpolation and PSP constructed using (L)-Leja points. The proposed sensitivity-driven algorithm is presented in details in Section 4.2. At the end of this section, we

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

employ the proposed approach in two illustrative examples with analytical solution and compare with the standard algorithm in terms of L^2 approximation error and cost in terms of number of (L)-Leja points. To test the power and usefulness of the proposed approach in real-world applications, we perform a comprehensive UQ study in plasma microturbulence simulations. In Section 4.3, we consider two test cases in which we perform the uncertainty propagation for the *computed dominant amplitude eigenmode*. The first test case is a modified version of a popular benchmark in the plasma microturbulence analysis in which we consider eight uncertain inputs. The second test case is a realistic problem based on a discharge [63] from the ASDEX fusion experiment. Therein, the input parameters stem from experimental measurements. Moreover, we break the uncertainty analysis into two parts, one in which we consider three uncertain parameters and another in which we extend the number to 12 stochastic inputs. In Section 4.4, we consider again the discharge from the ASDEX fusion experiment [63] with three uncertain parameters in which we perform the uncertainty analysis for particular microinstability amplitude eigenmodes.

Highlights and novelties

- We formulate a new Sensitivity-driven dimension-adaptive sparse grid algorithm for uncertainty propagation;
- Our approach is based on model hierarchies appearing in global sparse grid approximations, such as interpolation or PSP;
- The main novelty of this algorithm is a refinement indicator for dimension-adaptivity based on sensitivity scores;
- We apply the proposed approach in two plasma microturbulence test cases, in which we perform a comprehensive UQ analysis;
- We show that (i) the new method is more efficient than standard dimension-adaptivity and (ii) we highlight the benefits of dimension-adaptivity in UQ problems with anisotropically coupled inputs and lower intrinsic dimensionality;
- With the new approach, we needed at most 543 forward model runs to quantify uncertainty in a scenario with 12 uncertain inputs;
- Another novelty of this chapter is the undertaking, to the best of our knowledge, of one of the first UQ studies in plasma microinstability analysis;
- The performed uncertainty and sensitivity analyses aids the physical interpretation of the results and reveals insights about the underlying microturbulence.

4.1 Uncertainty in plasma microturbulence analysis

The simulation of microturbulence in magnetized plasmas is a problem of high practical relevance for efforts such as the ITER project, which aims to demonstrate a net gain from nuclear fusion. To make such an ambitious and complex endeavour possible, many technological and scientific hurdles need to be overcome. Notable obstacles are the small-scale fluctuations which cause energy loss rates despite sophisticated plasma confinements via strong and shaped magnetic fields. This microturbulence is driven by the free energy provided by the (unavoidably) steep temperature and density gradients of the particle species in the plasma. Unfortunately,

the measurement of these gradients and of other parameters influencing the underlying microinstabilities is affected by uncertainties due to, e.g., measurement noise. Moreover, the sources or sinks of amplitudes and profiles can also be affected by uncertainty. This therefore calls for a proper quantification of uncertainty *ab initio*.

To this end, in recent years, there has been a growing interest in quantifying uncertainty in the numerical simulation of fusion plasmas. For example, in [80, 178], a simple uniform, deterministic parameter scan was used in both linear and nonlinear gyrokinetic simulations. The focus in both works was on assessing the sensitivity to changes in the ions temperature gradient to validate the underlying simulation codes. In addition, in [140], a method to estimate the expectation and standard error of quantities such as spatially averaged density and temperature fluctuations or radial fluxes was proposed. The work [185] used nonintrusive stochastic collocation methods to quantify uncertainty in a drift-wave turbulence study from the CSDX linear plasma experiment. Moreover, the recent paper [186] assessed different methods for quantifying uncertainty of temporally varying quantities in nonlinear turbulent energy fluxes.

The existing approaches for uncertainty propagation in the numerical simulation of fusion plasmas typically sample the underlying stochastic space using one of the following two strategies. In so-called *parameter scans*, full, tensorized grids are employed, but these grids suffer from the curse of dimensionality. Another approach which was employed in more recent works (see, e.g., [185]), is based on a priori chosen, static sparse grids. However, with this approach the number of grid points can also grow prohibitively large with the dimensionality, especially when the underlying computer code is expensive to run. To enable the quantification of uncertainty in high(er)-dimensional plasma physics problems, new approaches are therefore needed. To this end, in this chapter, we formulate a novel structure-exploiting, context-aware sparse grid methodology, which is presented in details in the next section.

4.2 Sensitivity-driven dimension-adaptive sparse grid algorithm

Motivated by the challenges of quantifying uncertainty in complex, real-world applications such as plasma microturbulence analysis, our goal is to design an algorithm that:

1. *reduces* the overall *cost* of computing quantities of interest;
2. *exploits* the *structure* of the underlying (stochastic) problem;
3. *handles arbitrary* input probability *densities*;
4. samples the input probability space using *low-cardinality point sets*.

To achieve the first goal, we employ *model hierarchies*, specifically, global sparse grid approximations with dimension-adaptivity. Exploiting the structure of the underlying UQ problem depends on *how* the adaptive refinement is performed. To this end, in our formulation, we design *context aware* refinement strategies that exploit both the anisotropic coupling of the stochastic inputs and the intrinsic stochastic dimensionality. To achieve the third and fourth goal, we construct the sparse grids using (L)-Leja sequences (see Section 2.5.2).

Thus, we formulate a novel *deterministic, non-intrusive, context-aware* strategy based on dimension-adaptive sparse grids for uncertainty propagation called *Sensitivity-driven dimension-adaptive sparse grid algorithm*. We build our strategy on the dimension-adaptive algorithm of [71, 90] summarized in Section 2.5.3. Recall that in this approach, the essential ingredient is a *refinement indicator* $\epsilon(\cdot)$ whose value is used to guide the adaptive process. Specifically, the refinement indicator $\epsilon(\mathbf{k})$ is computed for each multiindex \mathbf{k} in the active set \mathcal{A} , and the subspace associated to the multiindex with the largest indicator is refined. Dimension-adaptivity

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

was shown to reduce the computational cost compared to standard approaches in a large number of UQ studies (see, for example, [26–28, 30, 35, 53, 160, 193]). However, standard refinement indicators are not computed using *stochastic information*, but rather using heuristics for the *deterministic approximation error*. Since in uncertainty propagation we usually have deterministic models parametrized in terms of stochastic input parameters, refinement indicators based on *stochastic information* might be preferred. Additionally, the employed heuristics in standard refinement indicators are generally based on *global information* in each subspace, usually a norm of the surplus Δ_k^{op} . In this way, neither information about the individual parameters nor about their coupling is exploited.

Consider, for example, that we have an uncertainty propagation problem with 10 stochastic inputs out of which only four are important. If we “zoom in” further, we see that only interactions between two of the four parameters are important. This means that we need to invest effort only in four out of 10 directions, and for interactions concerning two of the important parameters. Thus, we want to tune the dimension-adaptive algorithm to exploit this kind of anisotropic and lower-dimensional coupling between stochastic input parameters. To this end, in the proposed *sensitivity-driven adaptive algorithm*, we employ a *context-aware* refinement policy based on a *sensitivity scoring system*, as follows. We perform a Sobol’ decomposition in each candidate subspace for refinement to obtain the sensitivities of each input parameter and of parameters interaction. That is, we obtain information about the isotropy of the input parameters as well as about the stochastic dimensionality in the underlying subspace. We use this information to assess a corresponding *sensitivity score* such that the subspace with the largest score is refined. In this way, we ensure that more effort is invested only in subspaces that contribute significantly to the largest number of stochastic directions.

Remark: In an uncertainty propagation problem with d_{sto} stochastic inputs, we have d_{sto} individual sensitivities, one for each stochastic input, and $2^{d_{\text{sto}}} - d_{\text{sto}} - 1$ sensitivities due to all possible inputs interactions (see Section 2.9). Especially in settings in which d_{sto} is large, only a subset of the $2^{d_{\text{sto}}} - d_{\text{sto}} - 1$ interactions are expected to be significant. This is essentially the assumption behind sparse grid approximations: the stochastic parameters interact weakly at higher-order, i.e., sensitivities due to higher-order interactions are negligible.

The refinement process ends when a stopping criterion is met; we present the employing stopping conditions in the next section. At the end of the adaptive process, we implicitly obtain quantities of interest such as mean, standard deviation and Sobol’ indices for global sensitivity analysis of the output of interest, at no additional computational cost. We also have an approximation of the underlying forward model, but usually this is not the target of uncertainty propagation.

Next, we present in detail the algorithmic steps of our proposed Sensitivity-driven dimension-adaptive sparse grid approach. We build our method on the adaptive algorithm of [71, 90]. Our novelty is a context-aware refinement policy based on sensitivity information such that the important directions from a stochastic perspective are preferentially refined. To this end, in each subspace, we perform a Sobol’ decomposition [166] from which we obtain unnormalized Sobol’ indices. While obtaining the Sobol’ indices is straightforward for PSP, for interpolation we need to first perform a change of basis to an equivalent PSP basis and then we assess the desired Sobol’ indices from the coefficients in the new basis. We note that our algorithm is independent of the employed approximation operator, provided that sensitivity scores can be assessed. For a broader overview, we consider two types of approximation strategies, PSP and interpolation with Lagrange polynomials, which we construct using weighted (L)-Leja sequences.

4.2.1 Setup

The proposed approach in this chapter is based on model hierarchies appearing in global sparse grid approximations with dimension-adaptivity. The high-fidelity model is \mathcal{F}_h , which is treated like a black box. It depends on d_{sto} uncertain parameters $\boldsymbol{\theta} \in \mathbf{X}$ characterized by a probability density function, $\boldsymbol{\pi}$. Moreover, Assumptions 2.1 and 2.2 either hold true or suitable transformations are employed such that they are satisfied. We provide an intuitive visual description of the proposed algorithm in Figure 4.1.




Figure 4.1: Visual depiction of the proposed *Sensitivity-driven dimension-adaptive sparse grid algorithm* for uncertainty propagation. Context-aware adaptivity based on sensitivity scores is employed with the goal of exploiting the anisotropic coupling of the uncertain inputs and the lower intrinsic stochastic dimensionality. The high-fidelity model is treated like a black-box. At each adaptive step, this model is evaluated for each sample $\boldsymbol{\theta}$ from the subspaces added in the refinement process. These evaluations are then used to continue the adaptivity until a stopping criterion is met. At the end of the adaptive process, we can assess the mean, standard deviation and Sobol' indices for global sensitivity analysis of the output of interest at no additional cost. Moreover, although this is not the usual goal of uncertainty propagation, we also obtain a surrogate of the stochastic forward model. Note that the size of this outer-loop process is not known apriori, since we employ dimension-adaptive refinement based on a posteriori estimates.

4.2.2 Unnormalized sensitivity indices

We remind the reader the definition of multivariate PSP operators, summarized in Section 2.5.1. For a multiindex \mathbf{k} , the multivariate PSP operator is defined as (recall (2.18)):

$$\mathcal{U}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h] := \sum_{\mathbf{p}=0}^{\mathbf{P}_{\mathbf{k}}} c_{\mathbf{p}} \Phi_{\mathbf{p}}(\boldsymbol{\theta}),$$

where $\Phi_{\mathbf{p}}(\boldsymbol{\theta}) = \prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}$ is the L^2 -orthonormal basis associated with the density $\boldsymbol{\pi}$. Moreover, we evaluate the PSP coefficients via (L)-Leja quadrature, i.e., $c_{\mathbf{p}} = \mathcal{U}_{\mathbf{k}}^{\text{qu}}[\mathcal{F}_h \Phi_{\mathbf{p}}]$. For these operations, the 1D (L)-Leja points grow as $m(k_i) = 2m(k_i) - 1$ for $i = 1, 2, \dots, d_{\text{sto}}$. From (2.17), we have that $P_{k_i} = \lfloor (2k_i - 2)/2 \rfloor = k_i - 1$ and thus $\mathbf{P}_{\mathbf{k}} := (k_1 - 1, k_2 - 1, \dots, k_d - 1)$. To simplify the notation, let $\mathcal{P}_{\mathbf{k}} := \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}} : \mathbf{0} \leq \mathbf{p} \leq \mathbf{P}_{\mathbf{k}}\}$.

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

We are interested in the hierarchical PSP surplus associated to a $\mathbf{k} \in \mathcal{A}$. We have

$$\Delta_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \Delta c_{\mathbf{p}} \Phi_{\mathbf{p}}(\boldsymbol{\theta}), \quad (4.1)$$

where

$$\Delta c_{\mathbf{p}} := \sum_{z \in \{0,1\}^{d_{\text{sto}}}} (-1)^{|z|_1} c_{\mathbf{p}-z}, \quad (4.2)$$

with the convention $\Delta c_{\mathbf{0}} := c_{\mathbf{0}}$.

The PSP surplus (4.1) is defined in terms of multivariate L^2 -orthonormal polynomials w.r.t. $\boldsymbol{\pi}$. From Parseval's identity, we have that

$$\|\Delta_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h]\|_{L^2}^2 = \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \Delta c_{\mathbf{p}}^2 = \Delta c_{\mathbf{0}}^2 + \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}} \setminus \{\mathbf{0}\}} \Delta c_{\mathbf{p}}^2 = c_{\mathbf{0}}^2 + \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}} \setminus \{\mathbf{0}\}} \Delta c_{\mathbf{p}}^2. \quad (4.3)$$

Recall that the expectation and standard deviation of the underlying forward model can be directly assessed from the PSP coefficients (see (2.20), (2.21)). Thus, (4.3) is equivalent to

$$\|\Delta_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h]\|_{L^2}^2 = (\mathbb{E}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h])^2 + \Delta \text{Var}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h], \quad (4.4)$$

that is, the squared L^2 norm of the PSP surplus comprises a squared expectation contribution, which is usually small (see, e.g., [54, 193]), and the variance surplus, $\Delta \text{Var}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h]$. Therefore, in standard dimension-adaptivity as summarized in Algorithm 2.1, we see that it is natural to employ refinement indicators depending on the L^2 norm of the surpluses, (4.3), since this leads to the variance contribution of the subspaces associated to multiindices $\mathbf{k} \in \mathcal{A}$. This refinement indicator, however, represents *global* information in the underlying subspace and it therefore does not provide any information about the individual parameters, their interaction, or which of these are stochastically important. Note that we also use the L^2 of the surpluses in Section 7.3.1, in the context of adaptive sparse grid approximations for Bayesian inference.

To obtain the desired information, we decompose (4.4) further by exploiting the equivalence between PSP and Sobol' decompositions (see Section 2.9) to obtain the variance decomposition:

$$\Delta \text{Var}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h] = \sum_{i=1}^{d_{\text{sto}}} \Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i}[\mathcal{F}_h] + \sum_{i,j=1}^{d_{\text{sto}}} \Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i,j}[\mathcal{F}_h] + \dots + \Delta \text{Var}_{\mathbf{k}}^{\text{PSP},1,2,\dots,d_{\text{sto}}}[\mathcal{F}_h], \quad (4.5)$$

where

$$\Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{J}_i^{\text{PSP}}} \Delta c_{\mathbf{p}}^2, \quad \mathcal{J}_i^{\text{PSP}} := \{\mathbf{p} \in \mathcal{P}_{\mathbf{k}} \setminus \{\mathbf{0}\} : \mathbf{p}_i \neq 0 \wedge \mathbf{p}_j = 0, \forall j \neq i\},$$

$$\Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i,j}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{J}_{i,j}^{\text{PSP}}} \Delta c_{\mathbf{p}}^2, \quad \mathcal{J}_{i,j}^{\text{PSP}} := \{\mathbf{p} \in \mathcal{P}_{\mathbf{k}} \setminus \{\mathbf{0}\} : \mathbf{p}_i \neq 0 \wedge \mathbf{p}_j \neq 0 \wedge \mathbf{p}_n = 0, \forall n \neq i \wedge n \neq j\},$$

and so forth. Each term in (4.5) is an *unnormalized* Sobol' index for global sensitivity analysis. $\Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i}[\mathcal{F}_h]$ characterize the unnormalized sensitivity indices for each individual stochastic parameter, while the remaining variances quantify the sensitivity of all parameter interactions.

Remark: Recall that Sobol' indices for global sensitivity analysis are computed as ratios of local variances (for example, $\Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i}[\mathcal{F}_h]$, $\Delta \text{Var}_{\mathbf{k}}^{\text{PSP},i,j}[\mathcal{F}_h]$ etc. in (4.5)) and the total variance ($\Delta \text{Var}_{\mathbf{k}}^{\text{PSP}}[\mathcal{F}_h]$ in (4.5)). Since the total variance is a positive constant, it does not affect the ordering of the local variances. In the following, we will refer to the aforementioned local variances as *unnormalized Sobol' indices*.

4.2. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID ALGORITHM

Using decomposition (4.5), we can ascertain *how strongly coupled* the stochastic inputs are (via the unnormalized sensitivities of the parameters interactions) and *how many uncertain inputs and their interaction* are important (via the values of the corresponding unnormalized Sobol' indices). Note, nevertheless, that in general, only a subset of the unnormalized indices in (4.5) are nonzero. Thus identifying which individual inputs or which interactions are stochastically important becomes even more critical. We illustrate this point in the following example.

Example: Assume that $d_{\text{sto}} = 3$. We are given that $\Delta \text{Var}_{\mathbf{k}_1}^{\text{PSP}} = 0.12$ and $\Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP}} = 0.12$ for two multiindices \mathbf{k}_1 and \mathbf{k}_2 . For simplicity, assume that $(\mathbb{E}_{\mathbf{k}_1}^{\text{PSP}}[\mathcal{F}_h])^2 = (\mathbb{E}_{\mathbf{k}_2}^{\text{PSP}}[\mathcal{F}_h])^2 \approx 0$. Dimension-adaptivity based on standard refinement indicators would thus render the two subspaces as equally important (the adaptive algorithm would choose, e.g., the first subspace for refinement). Let us look closer at the two variance surpluses. Assume that the decomposition (4.5) yields

$$\Delta \text{Var}_{\mathbf{k}_1}^{\text{PSP}} = \Delta \text{Var}_{\mathbf{k}}^{\text{PSP},1}[\mathcal{F}_h],$$

where $\Delta \text{Var}_{\mathbf{k}_1}^{\text{PSP},1}[\mathcal{F}_h] = 0.12$, and

$$\Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP}} = \Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP},1}[\mathcal{F}_h] + \Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP},3}[\mathcal{F}_h] + \Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP},1,3}[\mathcal{F}_h],$$

where $\Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP},1}[\mathcal{F}_h] = 0.03$, $\Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP},3}[\mathcal{F}_h] = 0.05$, $\Delta \text{Var}_{\mathbf{k}_2}^{\text{PSP},1,3}[\mathcal{F}_h] = 0.04$. We see that in the subspace associated to \mathbf{k}_1 only the unnormalized Sobol' index associated to the first direction is nonzero, whereas in the second subspace three unnormalized indices have nonnegligible values. Moreover, the inputs are completely decoupled in the first subspaces, whereas in the second we have an interaction between the first and third stochastic input. Therefore, it makes more sense to refine the subspace associated to the *second* multiindex in this case.

We also want to obtain a variance decomposition such as (4.5) for Lagrange interpolation. However, there is no direct link between the Lagrange basis and Sobol' expansions. To this end, we first perform a transformation from the Lagrange basis to the orthonormal basis of the same (multivariate) degree to obtain the equivalent pseudo-spectral coefficients (see, e.g., [60, 66]). Recall that for a multiindex \mathbf{k} , the multivariate Lagrange interpolation operator reads

$$\mathcal{U}_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \mathcal{F}_h(\boldsymbol{\theta}_{\mathbf{p}}) \mathbf{L}_{\mathbf{p}}(\boldsymbol{\theta}),$$

where $\mathcal{P}_{\mathbf{k}} := \{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}} : \mathbf{0} \leq \mathbf{k} \leq \mathbf{P}_{\mathbf{k}}\}$. Since the level-to-nodes mapping for interpolation is $m(k_i) = k_i$ (see Section 2.5.2), then $\mathbf{P}_{\mathbf{k}} = (k_1 - 1, k_2 - 1, \dots, k_{d_{\text{sto}}} - 1)$.

Remark: For the employed level-to-nodes mappings, interpolation and PSP have the same bases, $\mathcal{P}_{\mathbf{k}}$, and therefore the same number of PSP coefficients. However, the level-to-nodes mapping for PSP is $m(k) = 2k - 1$, whereas $m(k) = k$ for interpolation. To achieve the same accuracy, we thus expect interpolation to require fewer grid points.

The change of basis from Lagrange interpolation and PSP is done as follows. We rewrite $\mathcal{U}_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}_h]$ in terms of a PSP basis of the same degree, i.e.,

$$\mathcal{U}_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}_h] = \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \mathcal{F}_h(\boldsymbol{\theta}_{\mathbf{p}}) \mathbf{L}_{\mathbf{p}}(\boldsymbol{\theta}) = \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} c_{\mathbf{p}} \boldsymbol{\Phi}_{\mathbf{p}}(\boldsymbol{\theta}), \quad (4.6)$$

where $\{\boldsymbol{\Phi}_{\mathbf{p}}\}_{\mathbf{p}=\mathbf{0}}^{\mathbf{P}_{\mathbf{k}}}$ is the equivalent PSP basis and $\{c_{\mathbf{p}}\}_{\mathbf{p}=\mathbf{0}}^{\mathbf{P}_{\mathbf{k}}}$ are the corresponding PSP coefficients. To find these coefficients, we simply solve

$$\sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} c_{\mathbf{p}} \boldsymbol{\Phi}_{\mathbf{p}}(\boldsymbol{\theta}_j) = \mathcal{U}_{\mathbf{k}}^{\text{in-g}}[\mathcal{F}_h](\boldsymbol{\theta}_j)$$

for all (L)-Leja points θ_j associated to the multiindex \mathbf{k} . Afterwards, we compute the surpluses (4.2) and perform the variance decomposition (4.5) to obtain the unnormalized Sobol' indices.

In this section, we showed (i) how to obtain variance information for each subspace corresponding to a multiindex \mathbf{k} in the active set, \mathcal{A} and (ii) how to decompose, via Sobol' expansions, the variance surplus in each such subspace to obtain unnormalized Sobol' indices for global sensitivity analysis, for PSP and interpolation operators. This approach remains valid for any other approximation operator provided that the approximation is unique and a mapping from the underlying basis to the PSP basis can be found. In the follow, we introduce the proposed *sensitivity scoring system* on which we build our context-aware, adaptive algorithm.

4.2.3 Sensitivity scores

We propose a *sensitivity scoring system* to perform dimension-adaptivity for sparse grid-based uncertainty propagation in computationally expensive, higher-dimensional problems. Our goal is to design a context-aware refinement strategy that exploits the anisotropic coupling of the stochastic input parameter w.r.t. the given input probability density function. To this end, we make use of the variance decomposition outlined in (4.5); for interpolation, we (i) perform the basis transformation (4.6) and (ii) compute the squared L^2 norm using (4.5) to obtain the unnormalized Sobol' indices. Based on this information, we compute a sensitivity score $s_{\mathbf{k}}^{\text{op}} \in \mathbb{N}$ for each multiindex $\mathbf{k} \in \mathcal{A}$ and use it as refinement indicator in the dimension-adaptive algorithm, i.e.,

$$\epsilon(\mathbf{k}) = s_{\mathbf{k}}^{\text{op}}. \quad (4.7)$$

For a broader overview, we formulate two strategies to assess sensitivity scores.

Let \mathbf{k} be a multiindex in the active set \mathcal{A} . In the first strategy, we employ *all* unnormalized Sobol' indices given by (4.5), i.e., $2^{d_{\text{sto}}} - 1$ indices in total for d_{sto} stochastic input parameters. These indices quantify the importance of all stochastic inputs and all possible combinations thereof. To this end, we employ $2^{d_{\text{sto}}} - 1$ user-defined *tolerances* $\boldsymbol{\tau}^{\text{op}} = (\tau_1^{\text{op}}, \tau_2^{\text{op}}, \dots, \tau_{2^{d_{\text{sto}}}-1}^{\text{op}})$ to compute a sensitivity score $s_{\mathbf{k}}^{\text{op}} \in \mathbb{N}$ associated to the multiindex \mathbf{k} . We initialize $s_{\mathbf{k}}^{\text{op}} = 0$. Afterwards, $s_{\mathbf{k}}^{\text{op}}$ is incremented by one whenever an unnormalized Sobol' index in (4.5) exceeds the associated tolerance. For example, $s_{\mathbf{k}}^{\text{op}}$ increases by one if $\Delta \text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h] \geq \tau_i^{\text{op}}$ for $i = 1, 2, \dots, d_{\text{sto}}$. Thus, the sensitivity score ascertains how many of the $2^{d_{\text{sto}}} - 1$ sensitivities of all stochastic inputs and combinations thereof are important in the subspace associated with \mathbf{k} . With this strategy, $s_{\mathbf{k}}^{\text{op}}$ thus takes integer values between 0 and $2^{d_{\text{sto}}} - 1$.

Example: Let $d_{\text{sto}} = 3$, which means that $2^{d_{\text{sto}}} - 1 = 7$. If $\Delta \text{Var}_{\mathbf{k}}^{\text{op},1}[\mathcal{F}_h]$, $\Delta \text{Var}_{\mathbf{k}}^{\text{op},2}[\mathcal{F}_h]$ and $\Delta \text{Var}_{\mathbf{k}}^{\text{op},1,2}[\mathcal{F}_h]$ exceed the associated tolerances, then $s_{\mathbf{k}}^{\text{op}} = 3$. This would mean that in the subspace associated to \mathbf{k} only the first two directions as well as their interaction are important. The sensitivity score thus reflects (i) how strongly coupled the d_{sto} stochastic inputs are (this is reflected by the number of ones due to interactions between stochastic inputs exceeding their tolerance; in this example, only one interaction was important), and (ii) what the intrinsic dimensionality in the underlying subspace is. In this example, we see that only two parameters are important since only unnormalized Sobol' indices corresponding to their individual contributions as well as their interaction are nonnegligible.

We hereby summarize the first strategy to assess sensitivity scores in Algorithm 4.2. The inputs are the probability density $\boldsymbol{\pi}$, the local tolerances, $\boldsymbol{\tau}^{\text{op}}$ and the hierarchical surplus $\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$. Note that the choice of tolerances $\boldsymbol{\tau}^{\text{op}}$ is problem dependent. For example, if, based on expert opinion or pre-existing knowledge, certain input parameters or their interaction are known to be more important, the corresponding tolerances in $\boldsymbol{\tau}^{\text{op}}$ should be chosen accordingly.

4.2. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID ALGORITHM

Algorithm 4.2 Sensitivity Scores Computation: one sensitivity score for each interaction

```

1: procedure COMPUTESENSITIVITYSCORE( $\boldsymbol{\pi}$ ,  $\boldsymbol{\tau}^{\text{op}}$ ,  $\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$ )
2:    $s_{\mathbf{k}}^{\text{op}} := 0$ 
3:   Compute  $\|\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2$  w.r.t.  $\boldsymbol{\pi}$  to obtain the variance  $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$  via (4.4)
4:   Decompose  $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$  via (4.5) to obtain all  $2^{d_{\text{sto}}} - 1$  unnormalized Sobol' indices


$$\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] = \sum_{i=1}^{d_{\text{sto}}} \Delta\text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h] + \sum_{i,j=1}^{d_{\text{sto}}} \Delta\text{Var}_{\mathbf{k}}^{\text{op},i,j}[\mathcal{F}_h] + \dots + \Delta\text{Var}_{\mathbf{k}}^{\text{op},1,2,\dots,d_{\text{sto}}}[\mathcal{F}_h]$$


5:   for  $n \leftarrow 1, 2, \dots, 2^{d_{\text{sto}}} - 1$  do
6:     if  $\Delta\text{Var}_{\mathbf{k}}^{\text{op},n}[\mathcal{F}_h] \geq \tau_n^{\text{op}}$  then
7:        $s_{\mathbf{k}}^{\text{op}} = s_{\mathbf{k}}^{\text{op}} + 1$ 
8:     end if
9:   end for
10:  return  $s_{\mathbf{k}}^{\text{op}}$ 
11: end procedure

```

Algorithm 4.3 Sensitivity Scores Computation: one sensitivity score for all interactions

```

1: procedure COMPUTESENSITIVITYSCORE( $\boldsymbol{\pi}$ ,  $\boldsymbol{\tau}^{\text{op}}$ ,  $\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$ )
2:    $s_{\mathbf{k}}^{\text{op}} := 0$ 
3:   Compute  $\|\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2$  w.r.t.  $\boldsymbol{\pi}$  to obtain the variance  $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$  via (4.4)
4:   Decompose  $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$  via (4.5) to obtain all unnormalized Sobol' indices


$$\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] = \sum_{i=1}^{d_{\text{sto}}} \Delta\text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h] + \sum_{i,j=1}^{d_{\text{sto}}} \Delta\text{Var}_{\mathbf{k}}^{\text{op},i,j}[\mathcal{F}_h] + \dots + \Delta\text{Var}_{\mathbf{k}}^{\text{op},1,2,\dots,d_{\text{sto}}}[\mathcal{F}_h]$$


5:   for  $i \leftarrow 1, 2, \dots, d_{\text{sto}}$  do
6:     if  $\Delta\text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h] \geq \tau_i^{\text{op}}$  then
7:        $s_{\mathbf{k}}^{\text{op}} = s_{\mathbf{k}}^{\text{op}} + 1$ 
8:     end if
9:   end for
10:  Compute the unnormalized Sobol' index due to all interactions

```

$$\Delta\text{Var}_{\mathbf{k}}^{\text{op,inter}}[\mathcal{F}_h] := \Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] - \sum_{i=1}^{d_{\text{sto}}} \Delta\text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h]$$

```

11:  if  $\Delta\text{Var}_{\mathbf{k}}^{\text{op,inter}}[\mathcal{F}_h] \geq \tau_{d_{\text{sto}}+1}^{\text{op}}$  then
12:     $s_{\mathbf{k}}^{\text{op}} = s_{\mathbf{k}}^{\text{op}} + 1$ 
13:  end if
14:  return  $s_{\mathbf{k}}^{\text{op}}$ 
15: end procedure

```

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

However, when no such knowledge is available, we recommend a conservative choice in which all components of $\boldsymbol{\tau}^{\text{op}}$ are equal.

In the first strategy we employed all $2^{d_{\text{sto}}} - 1$ unnormalized Sobol' indices to compute the sensitivity score. Although assessing the importance of all interactions between the d_{sto} stochastic inputs offers the most detailed stochastic characterization of the subspace associated to a multiindex \mathbf{k} , $2^{d_{\text{sto}}} - 1$ increases exponentially with d_{sto} , and it thus becomes prohibitively large. For example, if $d_{\text{sto}} = 30$, then $2^{d_{\text{sto}}} - 1 \in O(10^9)$. Therefore, in situations when d_{sto} is large or when the input parameters are (relatively) decoupled such that only a few interactions are important, we assess the sensitivity score as follows. Let $\boldsymbol{\tau}^{\text{op}} = (\tau_1^{\text{op}}, \tau_2^{\text{op}}, \dots, \tau_{d_{\text{sto}}}^{\text{op}}, \tau_{d_{\text{sto}}+1}^{\text{op}})$ be $d_{\text{sto}} + 1$ user-defined tolerances. As before, initially $s_{\mathbf{k}}^{\text{op}} = 0$. For $i = 1, 2, \dots, d_{\text{sto}}$, we increase $s_{\mathbf{k}}^{\text{op}}$ by 1 if the *individual* unnormalized Sobol' indices satisfy $\Delta \text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h] \geq \tau_i^{\text{op}}$. Hence, after this step, $s_{\mathbf{k}}^{\text{op}}$ can be at most d_{sto} , which is identical to the sensitivity score computation in the first strategy. However, instead of taking all $2^{d_{\text{sto}}} - d_{\text{sto}} - 1$ interactions separately as before, we compute a score for *all interactions* taken together. That is, we define

$$\Delta \text{Var}_{\mathbf{k}}^{\text{op,inter}}[\mathcal{F}_h] := \Delta \text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] - \sum_{i=1}^{d_{\text{sto}}} \Delta \text{Var}_{\mathbf{k}}^{\text{op},i}[\mathcal{F}_h]$$

and increase $s_{\mathbf{k}}^{\text{op}}$ by one if $\Delta \text{Var}_{\mathbf{k}}^{\text{op,inter}}[\mathcal{F}_h] \geq \tau_{d_{\text{sto}}+1}^{\text{op}}$. Therefore, with this strategy, $s_{\mathbf{k}}^{\text{op}}$ takes integer values between 0 and $d_{\text{sto}} + 1$. We summarize these steps in Algorithm 4.3.

Example: Let us revisit the previous example. Therein, $d_{\text{sto}} = 3$ and only $\Delta \text{Var}_{\mathbf{k}}^{\text{op},1}[\mathcal{F}_h]$, $\Delta \text{Var}_{\mathbf{k}}^{\text{op},2}[\mathcal{F}_h]$ and $\Delta \text{Var}_{\mathbf{k}}^{\text{op},1,2}[\mathcal{F}_h]$ exceeded their associated tolerances. Since only one interaction is important, the sensitivity score will take the value three using both presented strategies. However, if $\Delta \text{Var}_{\mathbf{k}}^{\text{op},1,3}[\mathcal{F}_h]$ exceeds the corresponding tolerance as well, then the first strategy yields a sensitivity score $s_{\mathbf{k}}^{\text{op}} = 4$, whereas with the second strategy we still have $s_{\mathbf{k}}^{\text{op}} = 3$. In our experiments from Sections 4.3 and 4.4, we employ both strategies in real world applications in which only a subset of stochastic parameters and interactions thereof are important.

We depict, in Figures 4.2 and 4.3, a visual comparison between standard refinement indicators (4.12) and our sensitivity score-based strategy employing the variance decomposition (4.5). Moreover, we also depict the two presented strategies to assess sensitivity scores. We show how each strategy makes use of a given hierarchical surplus $\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}(\boldsymbol{\theta})$. The standard refinement indicator uses $\|\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}$ directly, which is the summation between the squared expectation contribution and the total variance surplus associated to multiindex \mathbf{k} (see (4.4)). This is depicted in Figure 4.2. In contrast, the refinement indicator in our proposed approach first decomposes the variance in $\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}(\boldsymbol{\theta})[\mathcal{F}_h]$ into contributions due to each individual parameter and either (i) all possible interactions between all stochastic inputs (Figure 4.3a) or (ii) the summation of all interactions between inputs (Figure 4.3b). The sensitivity score is then increased by one if the associated unnormalized Sobol' index exceeds a user-defined tolerance.

The introduced scoring system employs the Sobol' decomposition of the variance (4.5) to assess the importance of sensitivities due to both individual stochastic parameters and their interactions. In this way, we can take advantage of the anisotropic coupling of the stochastic inputs and of the lower intrinsic dimensionality to preferentially refine only subspaces with a large sensitivity score. Note that our scoring system distinguishes between individual unnormalized Sobol' indices and Sobol' indices due to interactions between stochastic inputs. Thus, if the interactions are unimportant, the sensitivity score will reflect that and will prevent the algorithm to refine these directions extensively. We see therefore that our refinement strategy based on sensitivity is *context-aware*.




Figure 4.2: The refinement indicator in standard dimension-adaptivity employs $\|\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}$ directly. This comprises the squared expectation and the total variance surplus contribution associated to \mathbf{k} (see (4.4)).




Figure 4.3: The refinement indicator in our proposed approach decomposes the variance surplus $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$ associated to \mathbf{k} , using (4.5), into contributions due to each individual parameters and all possible interactions between stochastic inputs (left figure) or all interactions summed together (right figure). The associated score, which is either 0 or 1, is obtained by comparing the components of the variance decomposition with user-defined tolerances.

The sensitivity score is assessed for each multiindex $\mathbf{k} \in \mathcal{A}$. Next, we present a classification procedure for situations in which two are more sensitivity scores are equal.

4.2.4 Maximum sensitivity score

Refinement indicators based on sensitivity scores (4.7) are computed for each multiindex \mathbf{k} in the active set \mathcal{A} . The dimension-adaptive algorithm proceeds by refining the multiindex with the largest refinement indicator. However, it can happen that two or more subspaces have the same sensitivity score. In this situation, we need to distinguish between the associated subspaces such that only one subspace is refined. Two or more subspaces have the same sensitivity score if the same number of unnormalized Sobol' indices are larger than the prescribed local tolerances. A straightforward strategy to distinguish between equal scores would be to compare the unnormalized Sobol' indices that give rise to the equal scores or to compare their sum. However, those indices are not necessarily the same. For example, assume that $d_{\text{sto}} = 3$ and $s_{\mathbf{k}_1}^{\text{op}} = s_{\mathbf{k}_2}^{\text{op}} = 3$ for two multiindices $\mathbf{k}_1, \mathbf{k}_2 \in \mathcal{A}$. The first score equals three because $\Delta\text{Var}_{\mathbf{k}_1}^{\text{op},1}[\mathcal{F}_h], \Delta\text{Var}_{\mathbf{k}_1}^{\text{op},2}[\mathcal{F}_h], \Delta\text{Var}_{\mathbf{k}_1}^{\text{op},3}[\mathcal{F}_h]$ are larger than the corresponding tolerances. On the other hand, $s_{\mathbf{k}_2}^{\text{op}} = 3$ because $\Delta\text{Var}_{\mathbf{k}_2}^{\text{op},1}[\mathcal{F}_h], \Delta\text{Var}_{\mathbf{k}_2}^{\text{op},1,2}[\mathcal{F}_h], \Delta\text{Var}_{\mathbf{k}_2}^{\text{op},1,3}[\mathcal{F}_h]$ are larger than the associated tolerances. We see in this example that only one of the three unnormalized indices are the same. Therefore, instead of comparing or summing together the unnormalized Sobol' indices that lead to the same score, we formulate a different comparison strategy.

We compute the total variance surplus $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$ via (4.4) for all subspaces having equal

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

scores and then we simply select the subspace with the largest variance surplus. In this way, we select the subspace with the largest global contribution. The strategy for finding the subspace with the maximum sensitivity score is summarized in Algorithm 4.4. The first two inputs are \mathcal{S} , a set comprising all current scores, and \mathcal{D} , which contains all current surpluses in the active set \mathcal{A} . In lines 4 – 5, we find the maximum sensitivity scores. If only one maximum score exists, then the algorithm stops (lines 7 – 9). In the case when several maximum scores exist, we select the subspace with the largest variance, $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$ (lines 11 – 16). In the unlikely case that $\Delta\text{Var}_{\mathbf{k}}^{\text{op}}$ is the same for several sensitivity scores, the algorithm returns the multiindex corresponding to the last score. Having introduced the sensitivity scoring system

Algorithm 4.4 Maximum Sensitivity Score Computation

```

1: procedure FINDINDEXMAXIMUMSCORE( $\mathcal{S}, \mathcal{D}$ )
2:    $I := []$ 
3:    $\Delta\text{Var}_{\max}^{\text{op}}[\mathcal{F}_h] := 0, \quad \mathbf{k}_{\max}^{\text{op}} := \mathbf{1}$ 
4:   Find the maximum sensitivity scores  $s_{\mathbf{k}_n}^{\text{op}}$  from  $\mathcal{S}$ 
5:   Append to  $I$  the scalars  $n = 1, 2, \dots, n_{\max}$ , where  $n_{\max}$  is the number of max scores
6:    $n_{\max} := |I|$ 
7:   if  $n_{\max} = 1$  then
8:      $p := I[1]$ 
9:      $\mathbf{k}_{\max}^{\text{op}} := \mathbf{k}_p$ 
10:  else
11:    for  $m \leftarrow 1, 2, \dots, n_{\max}$  do
12:       $q := I[m]$ 
13:      Take  $\Delta_{\mathbf{k}_q}^{\text{op}}[\mathcal{F}_h]$  from  $\mathcal{D}$  and compute  $\Delta\text{Var}_{\mathbf{k}_q}^{\text{op}}[\mathcal{F}_h]$  via (4.4)
14:      if  $\Delta\text{Var}_{\mathbf{k}_q}^{\text{op}}[\mathcal{F}_h] \geq \Delta\text{Var}_{\max}^{\text{op}}[\mathcal{F}_h]$  then
15:         $\Delta\text{Var}_{\max}^{\text{op}}[\mathcal{F}_h] := \Delta\text{Var}_{\mathbf{k}_q}^{\text{op}}[\mathcal{F}_h]$ 
16:         $\mathbf{k}_{\max}^{\text{op}} := \mathbf{k}_q$ 
17:      end if
18:    end for
19:  end if
20:  return  $\mathbf{k}_{\max}^{\text{op}}$ 
21: end procedure

```

and our strategy to assess the maximum sensitivity score, we are hereby ready to formulate our proposed Sensitivity-driven dimension-adaptive sparse grid algorithm, designed for uncertainty propagation in higher-dimensional, computationally expensive problems.

4.2.5 Sensitivity-driven dimension-adaptive sparse grid algorithm

Our strategy is built on the dimension-adaptive algorithm of [71,90], as follows. The refinement indicators are based on sensitivity scores computed either via Algorithm 4.2 or Algorithm 4.3. In case two or more sensitivity scores are equal, we distinguish between them using Algorithm 4.4 such that a maximum score is always found. The employed sensitivity scoring system naturally leads also to the following stopping criterion: if the sensitivity scores of *all* multiindices in the active set, \mathcal{A} , are zero, i.e., all unnormalized Sobol' indices fall below the prescribed tolerances,

4.2. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID ALGORITHM

then we can stop the algorithm. Mathematically, we have

$$\begin{aligned}
 \left\| \sum_{\mathbf{k} \in \mathcal{K}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] - \sum_{\mathbf{k} \in \mathcal{O}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] \right\|_{L^2}^2 &= \left\| \sum_{\mathbf{k} \in \mathcal{O} \cup \mathcal{A}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] - \sum_{\mathbf{k} \in \mathcal{O}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] \right\|_{L^2}^2 \\
 &= \left\| \sum_{\mathbf{k} \in \mathcal{A}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] \right\|_{L^2}^2 \leq \sum_{\mathbf{k} \in \mathcal{A}} \|\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2 \\
 &= \sum_{\mathbf{k} \in \mathcal{A}} \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \Delta c_{\mathbf{p}}^2,
 \end{aligned}$$

where we used the triangle inequality of norms, and in the last equality, we used Parseval's identity (4.3). Therefore, an upper bound of the squared L^2 norm between the current approximation and approximation w.r.t. the old index set, \mathcal{O} , is given by the sum of sums of squared PSP coefficients corresponding to the multiindices in the active set, \mathcal{A} . When all scores in our algorithm fall below the imposed tolerances, we have that

$$\sum_{\mathbf{k} \in \mathcal{A}} \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \Delta c_{\mathbf{p}}^2 \leq \sum_{\mathbf{k} \in \mathcal{A}} \sum_{j=1}^{N_{\mathbf{k}}^{\text{scores}}} \tau_j^{\text{op}} =: \varepsilon_{\text{scores}}^2,$$

where $N_{\mathbf{k}}^{\text{scores}}$ denotes the number of nonzero scores associated to multiindex \mathbf{k} . This ensures an upper bound $\varepsilon_{\text{scores}}^2$ of the squared L^2 error, defined above. In other words, we can assume that the algorithm refined all stochastic direction sufficiently and thus there is no gain in adding further subspaces. This also means that we do not need to employ a surrogate for the global approximation error as in the standard dimension-adaptive algorithm [71, 90], since in our strategy we are not interested in a heuristic for the global approximation error. In general, we assume that the one-dimensional operators on which these sparse grid approximations are defined are convergent, thus the multivariate, tensorized approximation (2.9) is convergent as well (see Section 2.5). Thus, the representation of the high-fidelity model using global sparse grid approximations reads

$$\mathcal{F}_h(\boldsymbol{\theta}) = \sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h](\boldsymbol{\theta}).$$

It follows that

$$\begin{aligned}
 \left\| \sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] - \sum_{\mathbf{k} \in \mathcal{K}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] \right\|_{L^2}^2 &= \left\| \sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}} \setminus \mathcal{K}} \Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h] \right\|_{L^2}^2 \leq \sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}} \setminus \mathcal{K}} \|\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2 \\
 &= \sum_{\mathbf{k} \in \mathbb{N}^{d_{\text{sto}}} \setminus \mathcal{K}} \sum_{\mathbf{p} \in \mathcal{P}_{\mathbf{k}}} \Delta c_{\mathbf{p}}^2,
 \end{aligned}$$

thus upper bound of the sparse grid approximation error is given by the squared surplus PSP coefficients of the multiindices in the complement of the multiindex set, \mathcal{K} . If the underlying sparse grid operators converge in L^2 , these surplus coefficients will decay [199] and we will have a convergent approximation of \mathcal{F}_h . Nevertheless, the other two stopping criteria inherited from the original adaptive algorithm are $\mathcal{A} = \emptyset$ or reaching a user-defined maximum level, $K_{\text{max}}^{\text{op}}$.

Remark: For the condition $\mathcal{A} = \emptyset$ to be satisfied, there are two necessary conditions. First, \mathcal{A} needs to contain only one multiindex in the previous refinement step; this multiindex will have the largest associated refinement indicator. Next, the multiindex is moved to the old index set, \mathcal{O} . The second condition is that *none* of the forward neighbours of this multiindex preserve

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

the admissibility of the global multiindex set, \mathcal{K} . This situation is however rarely encountered in practice, as generally at least one forward neighbour is admissible. The next termination condition, imposing a maximum reachable level K_{\max}^{op} , is typically chosen having a bound on the overall computational cost in mind. This is especially important in applications with an expensive forward model. Ideally, however, we would like the proposed adaptive algorithm to terminate when all sensitivity scores of the multiindices from the active set are zero.

We summarize all steps of the proposed approach in Algorithm 4.5. The inputs are the high-

Algorithm 4.5 Sensitivity-driven Dimension-adaptive Sparse Grid Algorithm

```

1: procedure SENSITIVITYDRIVENADAPTIVESPARSEGRIDAPPROX( $\pi, \mathcal{F}_h, \tau^{\text{op}}, K_{\max}^{\text{op}}$ )
2:    $\mathbf{1} := (1, 1, \dots, 1)$ 
3:    $\mathcal{O} := \emptyset, \mathcal{A} := \{\mathbf{1}\}$ 
4:    $\mathcal{S} := \emptyset, \mathcal{D} := \emptyset$ 
5:    $s_{\mathbf{1}}^{\text{op}} := \text{COMPUTESENSITIVITYSCORE}(\pi, \tau^{\text{op}}, \Delta_{\mathbf{1}}^{\text{op}}[\mathcal{F}_h])$ 
6:    $\mathcal{S} = \mathcal{S} \cup \{s_{\mathbf{1}}^{\text{op}}\}, \mathcal{D} = \mathcal{D} \cup \{\Delta_{\mathbf{1}}^{\text{op}}[\mathcal{F}_h]\}$ 
7:   while all( $\mathcal{S}$ )  $\neq 0$  or  $\mathcal{A} \neq \emptyset$  or  $\max(\mathcal{K}) < K_{\max}^{\text{op}}$  do
8:      $\mathbf{k} := \text{FINDINDEXMAXIMUMSCORE}(\mathcal{S}, \mathcal{D})$ 
9:      $\mathcal{A} = \mathcal{A} \setminus \{\mathbf{k}\}, \mathcal{O} = \mathcal{O} \cup \{\mathbf{k}\}$ 
10:     $\mathcal{S} = \mathcal{S} \setminus \{s_{\mathbf{k}}^{\text{op}}\}, \mathcal{D} = \mathcal{D} \setminus \{\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\}$ 
11:    for  $i \leftarrow 1, 2, \dots, d_{\text{sto}}$  do
12:       $\mathbf{r} \leftarrow \mathbf{k} + \mathbf{e}_i$ 
13:      if  $\mathbf{r} - \mathbf{e}_q \in \mathcal{O}$  for all  $q = 1, 2, \dots, d_{\text{sto}}$  then
14:         $\mathcal{A} = \mathcal{A} \cup \{\mathbf{r}\}$ 
15:         $s_{\mathbf{r}}^{\text{op}} := \text{COMPUTESENSITIVITYSCORE}(\pi, \tau^{\text{op}}, \Delta_{\mathbf{r}}^{\text{op}}[\mathcal{F}_h])$ 
16:         $\mathcal{S} = \mathcal{S} \cup \{s_{\mathbf{r}}^{\text{op}}\}, \mathcal{D} = \mathcal{D} \cup \{\Delta_{\mathbf{r}}^{\text{op}}[\mathcal{F}_h]\}$ 
17:      end if
18:    end for
19:  end while
20:   $\mathcal{K} = \mathcal{O} \cup \mathcal{A}$ 
21:  Determine the PSP coefficients  $\{c_{\mathbf{k}}\}_{\mathbf{k} \in \mathcal{K}}$ 
22:  Compute  $\hat{\mathbb{E}}[\mathcal{F}_h], \hat{\text{Std}}[\mathcal{F}_h], \hat{S}_1^T, \hat{S}_2^T, \dots, \hat{S}_{d_{\text{sto}}}^T$  using the PSP coefficients  $\{c_{\mathbf{k}}\}_{\mathbf{k} \in \mathcal{K}}$ 
23:  return  $\hat{\mathbb{E}}[\mathcal{F}_h], \hat{\text{Std}}[\mathcal{F}_h], \hat{S}_1^T, \hat{S}_2^T, \dots, \hat{S}_{d_{\text{sto}}}^T$ 
24: end procedure

```

fidelity forward model, \mathcal{F}_h , the probability density π , the vector τ^{op} of local tolerances used to assess the sensitivity scores and the maximum grid level that can be reached, K_{\max}^{op} . At lines 2 – 3, we initialize \mathcal{O} and \mathcal{A} as in the standard dimension-adaptive algorithm. We initialize two new data structures, \mathcal{S} and \mathcal{D} , as well to store the scores and surpluses for all indices in the active set \mathcal{A} (line 4). We proceed by computing the sensitivity score of the first multiindex in \mathcal{A} using either Algorithm 4.2 or Algorithm 4.3 and update \mathcal{S} and \mathcal{D} accordingly. In each refinement step, we determine the multiindex with the maximum score (line 8) based on Algorithm 4.4 and then update the active set, \mathcal{A} , and the old index set, \mathcal{O} . Moreover, we also append the largest score to \mathcal{S} and its associated surplus to \mathcal{D} . The algorithm proceeds by adding the forward neighbours of the multiindex with the largest sensitivity score provided that the total multiindex set remains admissible. Further, the sensitivity score is assessed for each neighbour and the sets \mathcal{O} , \mathcal{A} , \mathcal{S} and \mathcal{D} are updated accordingly. When the algorithm terminates, we have the multiindex set \mathcal{K} and thus the dimension-adaptive sparse grid approximation (2.10) is fully specified. Therefore, we have a surrogate of the high-fidelity forward model, \mathcal{F}_h . However, in uncertainty propagation,

we are usually interested in computing quantities of interest to the output of interest. To this end, denoting $\{c_{\mathbf{k}}\}_{\mathbf{k} \in \mathcal{K}}$ the PSP coefficients at the end of the refinement process, we estimate $\hat{\mathbb{E}}[\mathcal{F}_h]$, $\hat{\text{Std}}[\mathcal{F}_h]$ and $\hat{S}_1^T, \hat{S}_2^T, \dots, \hat{S}_{d_{\text{sto}}}^T$ from $\{c_{\mathbf{k}}\}_{\mathbf{k} \in \mathcal{K}}$ via (2.20), (2.21) and (2.47).

4.2.6 Computational cost

Throughout this work we assume that the underlying high-fidelity forward model is computationally expensive. Hence, the cost of the proposed Sensitivity-driven dimension-adaptive sparse grid algorithm is given by the number of high-fidelity evaluations needed by the adaptive refinement process. Assume that one evaluation of the forward model \mathcal{F}_h is performed on P processes ($P = 1$ if the simulation is performed serially) and that the cost of one simulation is $\mathcal{C}_h^{\text{hi-fi}}$. Furthermore, let N^{op} denote the total number of evaluations of \mathcal{F}_h needed by the refinement process; recall that the termination of our adaptive algorithm depends on user-defined tolerances τ^{op} and on a maximum reachable level, $K_{\text{max}}^{\text{op}}$.

Therefore, the computational cost \mathcal{C}^{op} of the proposed approach amounts to

$$\mathcal{C}^{\text{op}} = N^{\text{op}} \mathcal{C}_h^{\text{hi-fi}}.$$

The remaining costs, such as the evaluation of quantities of interest (mean, standard deviation, and total Sobol' indices for sensitivity analysis) are negligible since these quantities of interest are assessed directly from the PSP coefficients. When we perform the interpolation, we need to solve a linear system of equations to find the corresponding PSP coefficients; see (4.6). However, this system has typically a small size and is hence computationally cheap to solve.

Remark: (i) The proposed dimension-adaptive approach can be performed using multiple layers of parallelism. At each refinement step, we need to run simulations for each sparse grid point from the newly added subspaces. These simulations are independent from each other and can be performed in parallel, each on P processes. That is, we can use *two layers* of parallelism. Further layers can be added, e.g., by adding more groups of simulations.

(ii) The solution of the high-fidelity model is sometimes computed using solvers that have significantly varying runtimes across simulations, such as iterative linear algebra solvers. When these simulations are performed in parallel, the processes used by simulations which terminate earlier are kept idle, wasting computational resources. Therefore, when these types of numerical experiments are performed at large-scales, varying runtimes across batches of simulations need to be properly accounted for to reduce computational bottlenecks. Some efforts in this direction were done, e.g., by [108], where a polynomial chaos expansion was used to approximate the *runtime* in UQ simulations, i.e., the runtime itself was the quantity of interest. The polynomial chaos expansion was then used to *predict* the runtime of future simulations and thus to adjust the resource allocation accordingly.

4.2.7 Illustrative examples

Before presenting our results in two plasma microturbulence test cases, we first use the proposed methodology in two problems in which the forward model, \mathcal{F} , is available analytically. We use dimension-adaptive sparse grid interpolation and compare our proposed refinement indicator based on sensitivity scores with the standard indicator (2.25). Since the purpose of these examples is to compare our proposed sensitivity scores refinement with the standard refinement approach, we use here only the strategy in which one score is computed for all interactions, i.e., $d_{\text{sto}} + 1$ scores in total (see Algorithm 4.3).

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

We compare the standard refinement strategy with our approach in terms of the L^2 approximation error

$$\mathcal{E}^2(\mathcal{F} - \mathcal{U}_{\mathcal{L}}[\mathcal{F}]) := \sqrt{\int_{\mathbf{X}} (\mathcal{F}(\boldsymbol{\theta}) - \mathcal{U}_{\mathcal{L}}[\mathcal{F}(\boldsymbol{\theta})])^2 d\boldsymbol{\theta}} \quad (4.8)$$

and the relative error of the expectation approximation

$$\mathcal{E}_{\text{rel}}(\mathbb{E}[\mathcal{F}] - \hat{\mathbb{E}}[\mathcal{F}]) := |1 - \mathbb{E}[\mathcal{F}]/\hat{\mathbb{E}}[\mathcal{F}]|, \quad (4.9)$$

where $\hat{\mathbb{E}}[\mathcal{F}]$ is estimated using the first PSP coefficient. For a fair comparison, we employ small tolerances ($\text{tol}^{\text{in}} = 10^{-8}$ in the standard version and $\boldsymbol{\tau}^{\text{in}} = 10^{-16} \cdot \mathbf{1}_{d_{\text{sto}}+1}$ in our approach). We consider two models, one with five input parameters and another with two inputs.

Five-dimensional example

First, we consider a five-dimensional model available analytically, $\mathcal{F} : [0, 1]^5 \rightarrow \mathbb{R}$,

$$\mathcal{F}(\boldsymbol{\theta}) = 1 + \cos(\pi + 1.5\theta_1 + 0.5\theta_2 + 0.05\theta_3 + 0.1\theta_4 + 0.002\theta_5), \quad (4.10)$$

in which $\boldsymbol{\theta}$ is uniformly distributed in the 5D hypercube, i.e., $\boldsymbol{\pi}(\boldsymbol{\theta}) = U(0, 1)^5$. We estimate the five total Sobol' indices using PSP with 10^5 Gauss-Legendre nodes and obtain

$$\hat{S}_1^T = 0.9034, \quad \hat{S}_2^T = 0.0098, \quad \hat{S}_3^T = 0.0009, \quad \hat{S}_4^T = 0.0039, \quad \hat{S}_5^T = 1.5714 \cdot 10^{-6}.$$

Therefore, the first stochastic parameter is significantly more important than all other four, while θ_5 is the least important parameter. Since the L^2 error (4.8) and relative error (4.9) cannot be estimated analytically, we estimate them numerically using 10000 MC samples.

We depict the results in Figure 4.4. We use *inter* to refer the sensitivity score computation strategy in which one score is assessed for all interactions. In the left plot, we visualize the L^2 approximation error for the two adaptive schemes and we observe that for similar L^2 errors, our approach is cheaper than the standard scheme. For example, for an L^2 error of around $6 \cdot 10^{-8}$, our approach requires about 15% fewer (L)-Leja points than the standard algorithm. In the right figure, we depict the estimate for $\mathcal{E}_{\text{rel}}(\mathbb{E}[\mathcal{F}] - \hat{\mathbb{E}}[\mathcal{F}])$ using the same number of Leja points as for the L^2 error. We observe again that our approach is, for a lower computational cost, similarly accurate as the standard adaptive method. Therefore, for the considered 5D test case




Figure 4.4: MC estimate of the L^2 approximation error (4.8) (left) and the estimate of the relative error of the expectation approximation (4.9) (right) for the example in (4.10).

(4.10), we showed that our approach is computationally cheaper than the dimension-adaptive approach based on the standard refinement indicator (2.25).

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE

Two-dimensional example

In general, we expect the proposed refinement indicator based on sensitivity scores to be more accurate than the standard approach in high(er)-dimensional uncertainty propagation problems with anisotropically coupled inputs and lower intrinsic dimensionality. In contrast, in problems with low stochastic dimensionality or problems in which the uncertain inputs are isotropically coupled, we generally do not obtain a significant benefit from using our approach. We illustrate this point by considering a test case in which we have two uncertain inputs such that (i) one of them is significantly more important than the other and (ii) their interaction is stochastically insignificant. To this end, consider $\mathcal{F} : [0, 1]^2 \rightarrow \mathbb{R}$,

$$\mathcal{F}(\boldsymbol{\theta}) = \sin(\theta_1 + 0.5\theta_2), \quad (4.11)$$

with $\boldsymbol{\pi}(\boldsymbol{\theta}) = U(0, 1)^2$. The associated total Sobol' indices estimated on a Gauss-Legendre grid comprising $20^2 = 400$ nodes are

$$\hat{S}_1^T = 0.9630, \quad \hat{S}_2^T = 0.0384,$$

whose values tell us that θ_1 is significantly more important than θ_2 .

The two errors are estimated as in the previous example and depict the results in Figure 4.5. We observe that in this example, our approach yields results of similar accuracy as the standard approach, for a similar cost both from an approximation or expectation estimation perspectives. Hence, this example, although relatively simplistic, underlines that our proposed sensitivity




Figure 4.5: MC estimate of the L^2 approximation error (4.8) (left) and the estimate of the relative error of the expectation approximation (4.9) (right) for the example in (4.11).

scores-based approach does not always outperform the standard method. The behaviour of our scheme depends on the structure of the underlying problem. Nevertheless, these two examples underline the benefits of dimension-adaptive algorithms in problems in which the uncertain inputs are anisotropically coupled and/or the intrinsic dimensionality is lower than d_{sto} .

4.3 Numerical results: computed dominant amplitude eigenmode

We apply the proposed Sensitivity-driven dimension-adaptive algorithm for the quantification of uncertainty in a real-world problem, the analysis of plasma microinstability (recall Section 3.2). For comparison purposes, we additionally employ the standard dimension-adaptive strategy summarized in Algorithm 2.1. Therein, we employ

$$\mathcal{I}(\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h], C_{\mathbf{k}}^{\text{op}}) = \|\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2} / \Delta L_{\mathbf{k}}^{\text{op}} \quad (4.12)$$

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

as refinement indicator, where $\Delta L_{\mathbf{k}}^{\text{op}}$ represents the number of (surplus) (L)-Leja points necessary to evaluate $\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]$. Note that for interpolation, $\Delta L_{\mathbf{k}}^{\text{in-g}} = 1$ (see Section 2.5.2). We assess $\|\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\|_{L^2}$ using Equation (4.3) from Section 4.2.2.

Initially, we consider two linear local (flux-tube) test cases in which the output of interest is the *computed dominant amplitude eigenmode*. In both test cases, microinstabilities are characterised using the linear eigenvalue solver from the gyrokinetic code GENE (recall Section 3.2.3). In Section 4.3.1, we consider a modified gyrokinetic benchmark scenario with eight uncertain parameters to obtain initial insights into the behaviour of the proposed approach. To test the usefulness of the proposed approach in real-world plasma microinstability analysis problems, we consider, in Section 4.3.2, a particular discharge of the ASDEX Upgrade experiment [63]. For this second test case, the analysis is performed in two steps. First, only three uncertain input parameters are considered in Section 4.3.3. Second, the number of uncertain inputs is extended to 12 in Section 4.3.4. We remark that besides the formulation of a novel context-aware strategy for uncertainty propagation, another novelty of this chapter is the undertaking, to the best of our knowledge, of one of the first UQ studies in plasma microinstability analysis.

In our experiments, we assume that the uncertain inputs are independent and uniformly distributed in an interval $[a_i, b_i]$, i.e., $\pi_i = U(a_i, b_i)$. The specific setup for each test case is presented in the following. Before going further, we remark that the presented methodology is not restricted in any way to uniform measures. Choosing uniform bounds was based on expert opinion from plasma physicists and supported by experimental data.

We performed the GENE simulations on two computing clusters. We used the CoolMUC-2 cluster¹ at the Leibniz Supercomputer Center and the Marconi supercomputer at the CINECA Supercomputing Center². On each machine we employ 32 processes for a single simulation. On the CoolMUC-2 cluster we used two Intel Xeon E5-2697 (Haswell) nodes per simulation, whereas on the Marconi supercomputer we needed 32 cores on an Intel Xeon 8160 (SkyLake) accelerator. The linear gyrokinetic eigenvalue problems were solved using the PETSc/SLEPc high-performance linear algebra libraries, and the parallelization of a single run is done using MPI. In addition, we employed Slurm³ for scheduling jobs on the clusters. All simulations were performed using standard double precision arithmetics. Using these resources, one linear eigenvalue GENE simulation required roughly between one and 120 minutes in total on 32 cores. Since the proposed Sensitivity-driven dimension-adaptive sparse grid algorithm allows performing the UQ simulations using multiple layers of parallelism (see Section 4.2.6), we performed multiple instances of GENE in parallel as well, on groups of 32 processes, whenever possible.

4.3.1 Modified Cyclone Base Case

The first test case is based on the so-called Cyclone Base Case (CBC) proposed in [47]. We chose the CBC because (i) it is a popular benchmark, thus a well understood test case in the gyrokinetic community and (ii) it is known to display a significant sensitivity to changes in the temperature and density gradients, which therefore calls for a quantification of uncertainty in numerical simulations. The original CBC benchmark have been restricted to one gyrokinetic ion species, abbreviated as i , assuming an adiabatic electron response and thus only electrostatic fluctuations. However, we modify the original setting to allow for more choices of stochastic parameters and thus better mimic realistic applications.

In Table 4.1, we outline the extensions and deviations from the parameters from [47] as well as the assumptions, stemming from expert opinion, about the uncertainty of each parameter.

¹<https://www.lrz.de/services/compute/linux-cluster/>

²<http://www.hpc.cineca.it/hardware/marconi>

³<https://slurm.schedmd.com/>

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE

θ	parameter name	symbol	left bound	right bound
θ_1	plasma beta	β	0.598×10^{-3}	0.731×10^{-3}
θ_2	collision frequency	ν_c	0.238×10^{-2}	0.322×10^{-2}
θ_3	i log temperature gradient	$-L_s \partial_x \ln T_i$	7.500	12.500
θ_4	e log temperature gradient	$-L_s \partial_x \ln T_e$	7.500	12.500
θ_5	temperature ratio	T_i/T_e	0.950	1.050
θ_6	i/e log density gradient	$-L_s \partial_x \ln n$	1.665	2.775
θ_7	magnetic shear	$\hat{s} = \frac{r}{q} \frac{\partial q}{\partial r}$	0.716	0.875
θ_8	safety factor	q	1.330	1.470

Table 4.1: The eight stochastic parameters in the modified CBC test case. The first six parameters characterize the two particle species, ions and electrons, whereas the last two inputs characterize the magnetic geometry. The temperature gradient is varied per species while the density gradients of the two particles are always equal to each other due to the quasi-neutrality condition in plasma physics.

We consider a total of eight stochastic parameters, six of which characterize the two considered particle species, ions and electrons, and two parameters which characterize the magnetic geometry. Specifically, the ions temperature and density gradients are considered to be uncertain. The electrons, which we abbreviate as (e), are treated fully gyrokinetically, thus their logarithmic temperature gradient, $-L_s \partial_x \ln T_e$, and density gradient, $-L_s \partial_x \ln n_e$, are considered as well. Moreover, these two gradients are affected by uncertainty as well. The electron temperature gradient is taken in the same range as the ion temperature gradient, $-L_s \partial_x \ln T_i$, but it varies independently. However, to ensure that the quasi-neutrality constraint holds true, the logarithmic density gradient of the electrons is forced to the exact value of its ion counterpart. Therefore, the density gradient of the two species is one stochastic parameter. The logarithmic density gradient also fixes the density ratio to one while the temperature ratio, T_i/T_e , can be treated as an uncertain parameter as well. Adding an electron species furthermore allows us to consider electromagnetic effects in the gyrokinetic simulations. The strength of these effects is determined by the kinetic-to-magnetic pressure ratio, β , which is therefore taken as another stochastic input. Another important parameter which is often avoided in benchmarks due to rather different implementations is the collision operator. Here, we employ a linearized Landau-Boltzmann collision operator and vary the corresponding normalized collision frequency, ν_c , as listed in Table 4.1. The aforementioned six uncertain parameters characterize the particle species. Regarding the magnetic geometry, uncertainties in the circular magnetic ($\hat{s} - \alpha$) equilibrium can be considered by attributing lower and upper limits to the safety factor, q , i.e., the ratio of toroidal turns of a magnetic field line per poloidal turn, as well as to its normalized radial derivative, the magnetic shear, $\hat{s} = r/q \partial q / \partial r$, where r is the radial coordinate labelling flux surfaces. We model the eight stochastic parameters as independent. The corresponding bounds in the last two columns in Table 4.1 are symmetric around the nominal value typically employed in deterministic simulations.

The high-fidelity forward model, \mathcal{F}_h , is the eigenvalue solver from GENE discretized using

$$h = (N_x = 15, \quad N_y = 1, \quad N_z = 24, \quad N_{v_{\parallel}} = 32, \quad N_{\mu_m} = 16),$$

where N_x is the number of (radial) grid points in the x direction, N_y is the number of Fourier modes in the y direction, which is always one for linear simulations and N_z represents the number of grid points in the z direction. Moreover, $N_{v_{\parallel}}$ is the number of grid points in the parallel velocity direction and N_{μ} is the number of points in the magnetic moment direction

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

(see Section 3.2). The output of interest in UQ simulations is the growth rate (amplitude), $\gamma[c_s/L_s]$, of the computed dominant eigenmode, where $c_s = \sqrt{T_e/m_i}$ is the *ion sound speed* and L_s is the *characteristic length*. In our simulations, the output of interest is computed with six digits of precision. Linear eigenvalue problems are characterized by a deterministic parameter, the *normalized perpendicular wavenumber*, $k_y\rho_s$. For a more in-depth overview of the influence of the eight stochastic parameters, we perform uncertainty propagation for multiple $k_y\rho_s$:

$$k_y\rho_s = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9. \quad (4.13)$$

We treat $k_y\rho_s$ as a deterministic free-parameter and for each value in (4.13), we perform uncertainty propagation using Algorithm 4.5.

To obtain the first insights into the underlying microinstabilities, we first perform a *deterministic* simulation using the nominal, i.e., the mean values of the eight parameters from Table 4.1. We compute the linear growth rates and real frequencies, $\omega[c_s/L_s]$, whose sign gives us the microinstability mode (negative real frequency \rightarrow electron-driven mode, positive real frequency \rightarrow ion-driven mode), corresponding to the dominant and first subdominant eigenmodes for all for all $k_y\rho_s$ in (4.13). Around $k_y\rho_s = 0.6$, a change in the frequency sign is clearly visible; the positive frequency indicates a microinstability propagating in the ion-diamagnetic drift direction, whereas negative frequency is associated with a mode with (opposite) electron-diamagnetic drift direction. For the CBC benchmark it is well known that this mode transition is from ITG to a trapped-electron-mode (TEM)/ETG hybrid mode, see, e.g., [79] and the references therein. Given the steep increase of the TEM/ETG mode at large wavenumbers, the transition is marked by a very sharp growth rate gradient at $k_y\rho_s = 0.6$. Since our approach is formulated in terms of global sparse grid approximations with dimension-adaptivity, these approximations are ill-conditioned at $k_y\rho_s = 0.6$. Thus, we discard $k_y\rho_s = 0.6$ and perform uncertainty propagation for the remaining values of the normalized perpendicular wavenumber in (4.13).




Figure 4.6: Growth rate (left) and frequency spectra (right) for the CBC test case of the dominant and first subdominant mode, obtained using the nominal values, i.e. the expectation of the left and right bounds from Table 4.1.

To perform uncertainty propagation, we consider both dimension-adaptive sparse grid interpolation and PSP. The adaptive refinement is performed using our proposed sensitivity scores-based approach detailed in Section 4.2. In addition, we compare our approach with standard dimension-adaptivity based on the refinement indicator (4.12), which is used to obtain reference results. We prescribe $K_{\max}^{\text{op}} = 20$ in both adaptive approaches. In our approach, we prescribe $\tau^{\text{op}} = 10^{-6} \cdot \mathbf{1}_{255}$ (when we employ all unnormalized Sobol' indices to compute the score; we have an eight-dimensional setting, thus $2^8 - 1 = 255$ unnormalized Sobol' indices) or $\tau^{\text{op}} = 10^{-6} \cdot \mathbf{1}_9$, when we use the second strategy to assess the scores, i.e., using all eight individual contributions plus one for all interactions (cf. Section 4.2.3). To make the comparison between our

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE

approach and the standard adaptive algorithm fair, we prescribe for the standard approach a (global) tolerance of $tol^{op} = \sqrt{10^{-6}} = 10^{-3}$. For each considered normalized perpendicular wave number $k_y \rho_s$ we use the aforementioned adaptive sparse grid approximations to assess the expectation, standard deviation and total Sobol' indices of the growth rate, $\gamma[c_s/L_s]$. Recall that all these quantities can be computed from the pseudo-spectral coefficients associated to either PSP or interpolation (see Sections 2.5 and 2.9). We hereby present our results, focusing on the comparison between the proposed sensitivity scores-based adaptivity and the standard dimension-adaptive algorithm using the ratio between the L^2 norm of the surpluses and the cost. Moreover, we also overview the relevance of these results for the plasma physics community.

In Figure 4.7, we depict the expected value of $\gamma[c_s/L_s]$ as well as one standard deviation yielded by all four adaptive sparse grid approaches: adaptive interpolation and PSP plus our proposed approach and interpolation and PSP together with the standard adaptive approach. Moreover, we also illustrate the deterministic growth rates for all considered normalized perpendicular wavenumbers. On the one hand, we observe a good agreement between the results of all sparse grid approaches. This demonstrates that for a similar setup, interpolation and PSP perform comparably well in this test case. In addition, the proposed sensitivity scores approach for adaptivity is as accurate as the standard adaptive approach. On the other hand, the deterministic results overlap almost perfectly with the expectations of the stochastic results; this is not surprising, given the relative simplicity of this test case. The novelty for the plasma physics community is that the computed standard deviations provide a quantitative measure of the uncertainty associated with the expected value. This tells us that the ITG mode appears to be very robust while the onset of the TEM/ETG-branch could be quite different given the much larger standard deviations. These results are already of interest for physicists trying to compare and predict microturbulence in plasmas. The quantification of uncertainty can be however taken even further. To quantitatively understand the total contribution of each stochastic input to the resulting uncertainty, we analyse the total Sobol' indices for global sensitivity analysis.




Figure 4.7: Expected values with one standard deviation as well as the deterministic growth rates of the dominant eigenmode for the CBC test case with eight uncertain parameters.

We show the total Sobol' indices in Figure 4.8. The top two figures depict the results using adaptive interpolation, whereas in the bottom figures we illustrate the total Sobol' indices obtained via adaptive PSP. In addition, the left column corresponds to the standard adaptive approach, whereas the central and right column depict the results for our approach based on

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

one score for all interactions (central column) and one score for each interaction (right column), respectively. The similarity of the total Sobol' indices in all six figures indicates again that (i) interpolation and PSP perform similarly and (ii) that our proposed adaptive approach, with either strategy to assess sensitivity scores, has a similar accuracy as the standard adaptive method. Moreover, we observe that in all plots two stochastic parameters show the largest total Sobol' indices, the logarithmic ion and electron temperature gradients; the other seven stochastic inputs have negligible contributions. This indicates that although a total of eight stochastic parameters is given, only two of them are important for uncertainty propagation in this particular scenario. Such findings are very important for the plasma physics community. The other stochastic parameters are known to have an impact on the modes themselves, but here not within the assumed input uncertainty. This could, therefore, motivate to restrict the much more costly nonlinear studies to the two temperature gradients. Observe that the two temperature gradients have complementary total Sobol' indices: the ion temperature gradient dominates while the electron temperature gradient has a very small total Sobol' index for $k_y\rho_s \leq 0.5$, and the other way around for $k_y\rho_s \geq 0.7$. Thus, except for $k_y\rho_s = 0.5$, where the Sobol' indices for both temperature gradients are non-negligible, the intrinsic stochastic dimensionality is one. The observed behaviour of the total Sobol' indices is consistent with what is known about the CBC benchmark: for $k_y\rho_s \leq 0.5$, the microinstability is driven by an ITG mode, whereas when $k_y\rho_s \geq 0.7$, it is driven by a TEM/ETG mode. Furthermore, at each $k_y\rho_s$, the sum of the total Sobol' indices is close to 1.0: the Sobol' indices due to the eight individual inputs explain almost all output variance, while the interactions of the stochastic inputs are negligible. Therefore, the stochastic model can be well approximated by a linear model. Such information can be very helpful in constructing reduced, e.g., quasi-linear models in the plasma physics community.

Finally, the cost in terms of total number of GENE simulations for all employed adaptive approaches is depicted in Figure 4.9. First, observe that the most expensive problem required only 473 simulations in total, which highlights the benefits of using adaptivity in higher-dimensional UQ problems in which the stochastic inputs are anisotropically couple. To put this number of simulation in perspective, a standard parameter scan, typically used in the plasma physics community, with 10 points in each direction would require a total of 10^8 simulations, which would be computationally prohibitive. As expected, for either standard or sensitivity scores-based adaptivity, PSP is more expensive than interpolation; even the adaptive PSP based on sensitivity scores is generally more expensive than standard adaptive interpolation. By far the cheapest approaches are the adaptive sparse grid interpolation with either of the two methods to assess sensitivity scores. For example, at $k_y\rho_s = 0.7$, our approach depending on all 255 scores is about 3.81 times cheaper than interpolation using standard adaptivity, about 9 times cheaper than adaptive PSP based on the standard method, and more than 5 times cheaper than PSP with adaptivity based on either approach to assess sensitivity scores. Moreover, we observe that the two variants to compute sensitivity scores have comparable costs for both PSP and interpolation, which is due to (i) the weak interaction of the eight uncertain inputs combined with (ii) the very low intrinsic dimensionality, which is at most two.

To summarize, in this section, we considered a modified version of the CBC benchmark scenario [47] in which we had two particle species, ions and electrons, as well as electromagnetic effects and collisions. Linear eigenvalue simulations were performed using the gyrokinetic simulation code GENE to assess the growth rate of the first dominant eigenmode. We used a total of eight stochastic parameters, six of which characterized the two particle species and the remaining two the magnetic geometry of the nuclear reactor. To propagate the uncertainty through the linear gyrokinetic model, we employed the proposed Sensitivity-driven dimension-adaptive algorithm. Moreover, for comparison, we also used the standard dimension-adaptive strategy

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE




Figure 4.8: Total Sobol' indices for global sensitivity analysis via adaptive sparse grid interpolation (top figures) and PSP (bottom figures) based on the standard adaptive approach (left figures), our proposed approach using one sensitivity score for all interactions (central figures) and our approach using all unnormalized Sobol' indices (right figures) for the CBC test case using $k_y \rho_s \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.8, 0.9\}$.




Figure 4.9: Total number of evaluations needed to construct the sparse interpolation and PSP surrogates using our proposed approach based on sensitivity scores (either using only one for all interactions or one score for each interaction) vs. the standard adaptive approach.

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

of [71, 90]. For a broader overview, we considered two strategies for approximation, PSP and interpolation, and two strategies to assess the sensitivity scores. Our results show that the proposed approach is considerably cheaper than the standard strategy for an equivalent setup. Moreover, the sensitivity-driven adaptive PSP and interpolation have similar accuracies, but the significantly lower cost of interpolation makes it our method of choice in the following realistic and computationally more expensive test case.

4.3.2 Realistic test case: overview

The CBC benchmark test case offered the first insights into the behaviour of the proposed sparse grid approach for uncertainty propagation. However, the CBC test case is not necessarily a realistic scenario since corresponding nonlinear simulations were found to dramatically overestimate the transport levels obtained in experiments. This can to some degree be explained by idealizations in (i) the choice of parameters and (ii) by an incomplete physical description, such as missing external shear flows.

To test the power and usefulness of the proposed Sensitivity-driven dimension-adaptive sparse grid algorithm in a realistic plasma microturbulence test case, we consider a particular GENE validation study from [63], performed for the ASDEX Upgrade experiment. Therein, nonlinear simulation results have been compared to experimentally determined ion and electron heat fluxes as well as various turbulence observables such as electron temperature fluctuation levels, radial correlation functions and cross phases with electron density fluctuations. Linear simulations furthermore revealed that the parameters associated with the discharge of interest are very close to the dominant mode transitions and therefore represent a challenging parameter set which could (i) be influenced by further stochastic parameters and (ii) turn out to be subject to cross interactions among these parameters. We took these concerns as motivation to perform a two-step UQ analysis. First, we consider three uncertain parameters characterizing the particle species in Section 4.3.3. Afterwards, in Section 4.3.4, we consider 12 uncertain inputs characterizing the particle species (seven parameters) and the magnetic geometry (five parameter). As for the CBC test case, the output of interest is the first dominant amplitude eigenmode. For a more comprehensive overview of this test case, in Section 4.4, we also perform uncertainty analysis for the specific microinstability amplitude eigenmode (TEM/ETG vs. ITG). As for the CBC test case, the high-fidelity forward model is the GENE eigenvalue solver using

$$h = (N_x = 15, \quad N_y = 1, \quad N_z = 24, \quad N_{v_{||}} = 48, \quad N_{\mu_m} = 16).$$

The modelling of the uncertain parameters is based on expert opinion supported by experimental measurements: we were provided with experimentally determined expected temperature and density gradients, as well as expectations of further input parameters affected by uncertainty. The bounds of the uncertainty intervals stem from expert opinion. As for the modified CBC benchmark, the left and right bounds are symmetric about the aforementioned expected values.

4.3.3 Realistic test case with three stochastic parameters

First, we perform uncertainty propagation using three stochastic parameters as listed in Table 4.2. These parameters are the logarithmic density and temperature gradients for ions and electrons. Their left and right bounds are obtained by respectively varying by 25% the nominal value of the gradients. We perform simulations for multiple values of $k_y \rho_s$ as in the CBC scenario from Section 4.3.1. However, since this test case is more representative for real-world

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE

θ	parameter name	symbol	left bound	right bound
θ_1	i/e log density gradient	$-L_s \partial_x \ln n$	1.156	1.927
θ_2	i log temperature gradient	$-L_s \partial_x \ln T_i$	2.096	3.494
θ_3	e log temperature gradient	$-L_s \partial_x \ln T_e$	4.040	6.733

Table 4.2: Summary of the three stochastic parameters considered for the ASDEX Upgrade test case with ranges estimated from experimental measurements.

problems, wave numbers up to the hyper-fine electron gyroradius scales are considered:

$$k_y \rho_s = 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.9, 1.0, 2.0, 3.0, 4.0, 5.0, 7.5, 10.0, 12.5, 15.0, 17.5, 20.0, 22.5, 25.0, 27.5, 30.0. \quad (4.14)$$

To obtain some first insights into the underlying microinstabilities, deterministic runs are performed using the expected values of all input parameters in Table 4.2. We assess the growth rate, $\gamma[c_s/L_s]$, and frequency spectra, $\omega[c_s/L_s]$, of the dominant eigenmode. Given the large wave number and amplitude range, the results are depicted in a logarithmic scale in Figure 4.10 (note the negative sign of the frequency). Both curves are smooth and monotonous, suggesting that the dominant mode changes neither its (electron diamagnetic) drift direction nor its character in a drastic way. The slightly different slopes in frequency, i.e., dispersion relations, and the prominent curvature change around $k_y \rho_s = 1.0$ in growth rate imply that nevertheless two or more different microinstabilities such as pure TEM and TEM/ETG-hybrids are excited.




Figure 4.10: Growth rate (left) and frequency spectra (right) of the dominant eigenmode for the ASDEX Upgrade test case obtained using the nominal values, i.e., the expectation of the input parameters.

For uncertainty propagation we employ adaptive sparse grid interpolation based on our sensitivity-driven approach in which the sensitivity scores are computed using the two strategies summarized in Section 4.2.3. For comparison, we also employ interpolation in which adaptivity is performed using the standard dimension-adaptive algorithm. The same setup as in the CBC test case is used for adaptivity: $K_{\max}^{\text{in-g}} = 20$ is set in both adaptive approaches, whereas $\tau^{\text{in-g}} = 10^{-6} \cdot \mathbf{1}_7$ when we employ all unnormalized Sobol' indices to compute the score or $\tau^{\text{in-g}} = 10^{-6} \cdot \mathbf{1}_4$ when we use the second strategy to assess the scores, i.e., using all individual contributions plus one for all interactions. In addition, in standard adaptivity, $\text{tol}^{\text{in-g}} = 10^{-3}$. The output of interest, the growth rate $\gamma[c_s/L_s]$ of the dominant eigenmode, is computed with six digits of precision. In postprocessing, we evaluate its expectation, standard deviation and total Sobol' indices for each $k_y \rho_s$. Due to the larger growth rates and simpler mode structure,

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

the runtime significantly decreased with growing $k_y\rho_s$ values: at $k_y\rho_s = 0.2$ the runtime of one simulation is about 30 minutes on 32, whereas at $k_y\rho_s = 30$, one simulation required around 65 seconds on 32 cores. We depict the expectation and one standard deviation as well as the deterministic growth rate for each $k_y\rho_s$ in Figure 4.11. As in the previous test case, we observe an overlap between the results obtained with our proposed approach (with the two variants to assess the sensitivity scores) and the standard dimension-adaptive strategy. Hence, our approach is similarly accurate as the standard adaptive method for this test case as well. In addition, the deterministic result is similar to the expectations; the (absolute) difference varies uniformly, roughly between 0.01 and 0.1. However, these differences are more significant for $k_y\rho_s \leq 1.0$, a wave number range that is usually considered to be the most important for a correct assessment of the *ion heat flux*. Therefore, results from corresponding nonlinear simulations (with nominal values) should be carefully analysed.




Figure 4.11: Expected values and one standard deviation as well as the deterministic growth rates for the ASDEX Upgrade test case with three uncertain parameters and $k_y\rho_s$ as in (4.14).

The total Sobol' indices are depicted in Figure 4.12 (top two figures: standard adaptive approach, center figures: our proposed approach in which one sensitivity score is computed for all interactions, bottom figures: our approach in which a sensitivity score is computed for each interaction). For a clearer illustration, we split the results in two figures corresponding to $k_y\rho_s \leq 1$ (left) and $k_y\rho_s \geq 2$ (right). Again a good agreement between the two variants of our proposed approach and the standard adaptive strategy is observed. On the one hand, when $k_y\rho_s \leq 0.8$, all three stochastic parameters have non-negligible total Sobol' indices, i.e., the intrinsic stochastic dimensionality is full. Moreover, for $k_y\rho_s = 0.9$ and $k_y\rho_s = 1$, the two logarithmic temperature gradients are the most important, that is, the ambient stochastic dimensionality is two. Even more, for $k_y\rho_s \leq 1$, the interactions between the stochastic inputs are nonnegligible (the sum of the three total Sobol' indices exceeds 100%, thus the contributions due to parameters interactions are different from zero), thus the corresponding stochastic model exhibits a nonlinear behaviour. When $k_y\rho_s \geq 2$, on the other hand, the uncertainty in the logarithmic electron temperature gradient is the most important stochastic parameter. Moreover, the sum of the three total Sobol' indices is close to one, which shows that the stochastic model can be well approximated by a linear model depending only on the contribution due to the logarithmic electron temperature gradient. Therefore, the intrinsic stochastic dimensionality is one when $k_y\rho_s \geq 2$. In addition to information about the total contribution of the three

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE

stochastic inputs to the output uncertainty, Figure 4.12 provides information on the underlying microinstabilities as well. For example, for $k_y \rho_s \geq 2$, we clearly have an ETG mode: the logarithmic electron temperature gradient is most important parameter. On the other hand, for $k_y \rho_s \leq 1$, the electron temperature gradient is more important than the ion temperature gradient and the contribution of the density gradient decreases as $k_y \rho_s$ increases. Hence, we have a mixture of TEM/ETG modes for $k_y \rho_s \leq 1$ which is, however, also affected by the logarithmic ion temperature gradient and may be in competition with subdominant ITG modes; we analyse this further in Section 4.4.

We see that in terms of accuracy of quantities of interest, the proposed Sensitivity-driven dimension-adaptive sparse grid algorithm performs similarly to the standard dimension-adaptive strategy employing refinement indicators based on the L^2 norm of the surpluses and their associated number of grid points. Next, we show the computational costs of the two approaches. We visualize, in Figure 4.13, the cost in terms of total number of grid points, i.e., GENE evaluations (left: $k_y \rho_s \leq 1$, right: $k_y \rho_s \geq 2$). First, we see that our approach requires fewer GENE runs than the standard adaptive approach for all normalized perpendicular wavenumbers $k_y \rho_s$. For example, at $k_y \rho_s = 3$, our approach requires around 3.3 times fewer GENE evaluations than the standard approach. For $k_y \rho_s \leq 0.9$, we see that our sensitivity scores-based approach leads to significant cost savings. However, for $k_y \rho_s \geq 1$, the savings yielded by our approach are not very significant because both approaches are computationally cheap, requiring a small, roughly the same number of GENE evaluations for all these $k_y \rho_s$; this is mainly because the intrinsic stochastic dimensionality is one. Additionally, the behaviour observed in the total Sobol' indices is reflected in the cost: it is computationally more expensive to perform uncertainty propagation for $k_y \rho_s \leq 1$ than for $k_y \rho_s \geq 2$ due to the higher intrinsic stochastic dimension.

To further illustrate the computational savings due to our algorithm, we next look at the multiindices and the associated sparse grids corresponding to a specific perpendicular wavenumber. Moreover, we also assess the local Sobol' indices. As an example, we consider $k_y \rho_s = 0.9$ and depict our results in Figure 4.14 (left: standard adaptive approach, center: sensitivity scores with all interactions, right: sensitivity scores). With the employed stochastic setup, the local Sobol' indices listed in Table 4.3 indicate that the logarithmic electron temperature gradient is the most important individual direction ($S_3 = 0.7960$), the logarithmic density gradient is the second most important ($S_1 = 0.0292$), while the third important individual direction is the ion logarithmic temperature gradient ($S_2 = 0.0902$). Moreover, the interaction between the electron temperature gradient and the density gradient is the most significant ($S_{23} = 0.0742$), while the Sobol' indices associated to the remaining interactions are less significant, with values below 10^{-1} . The multiindex sets, depicted in the top figures, show that both algorithms detect the third direction as the most important. Moreover, most effort in the mixed directions is invested in the electron logarithmic temperature gradient - logarithmic density gradient direction. However, the standard adaptive algorithm adds many (unnecessary) interaction multiindices in the other mixed directions as well, while our approach better exploits the fact that (i) one direction is significantly more important than the other two (ii) the interactions between the stochastic inputs are nonnegligible in only one direction. Furthermore, between the two variants of our algorithm, the approach that computes a separate sensitivity score for *each* interaction is more effective in taking advantage of the anisotropic coupling between the three stochastic terms. Regarding costs, the standard approach needs 264 grid points in total, whereas our approach needs a total of only 99 grid points, when a sensitivity score is computed for the summation of all interactions, and of only 70 points, when a sensitivity score is assessed for each interaction, to yield results similar up to four digits.

In this section, we considered a realistic plasma microturbulence analysis scenario based on ASDEX Upgrade experiment in which we used three uncertain parameters: the logarithmic

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS




Figure 4.12: Total Sobol' indices for global sensitivity analysis via adaptive sparse grid interpolation for the ASDEX Upgrade test case with three uncertain parameters and $k_y \rho_s$ as in (4.14) (top figures: standard dimension-adaptivity is employed; middle figures: we use our approach in which we compute one score for all interactions; in the bottom figures, we compute one score for each interaction in our proposed approach).

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE




Figure 4.13: Total number of GENE evaluations needed to construct the sparse interpolation surrogate using our proposed approach based on sensitivity scores (using either one sensitivity scores for all interactions or one score for each interactions between the stochastic inputs) vs. the standard adaptive approach for the ASDEX Upgrade test case with three uncertain parameters.

\hat{S}_1	\hat{S}_2	\hat{S}_3	\hat{S}_{12}	\hat{S}_{13}	\hat{S}_{23}	\hat{S}_{123}
0.0292	0.0902	0.7960	0.0079	0.0014	0.0742	0.0007

Table 4.3: Individual Sobol' indices as well Sobol' indices due to all interactions between the three stochastic parameters from Table 4.2, at $k_y \rho_s = 1.0$.




Figure 4.14: Multiindices and their associated sparse grids at $k_y \rho_s = 1.0$ for adaptive interpolation for the 3D ASDEX Upgrade scenario (left: standard adaptive approach, center: out proposed approach with one score for all interactions; right: our approach with one score for each interaction).

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

temperature gradients of both ions and electrons as well as the logarithmic density gradient. Next, we extend our analysis and consider a scenario with a total of 12 uncertain parameters of the same ASDEX Upgrade test case.

4.3.4 Realistic test case with 12 stochastic inputs

In a second uncertainty analysis step, we enhance the previous scenario with nine additional stochastic terms shown in Table 4.4, resulting in a total of 12 stochastic inputs for the ASDEX Upgrade test case, i.e., $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_{12})$. The first seven parameters characterize the two particle species, ions and electrons. In addition to the logarithmic temperature and density gradients, we consider five other stochastic parameters associated to the two species: the kinetic-to-magnetic pressure ratio, β , the normalized collision frequency, ν_c , of the employed (linearized) Landau collision operator, the ratio of the temperature of the two species, and the effective ion charge, Z_{eff} . Moreover, the latter five parameters are associated to the magnetic geometry.

θ	parameter name	symbol	left bound	right bound
θ_1	plasma beta	β	0.488×10^{-3}	0.597×10^{-3}
θ_2	collision frequency	ν_c	0.641×10^{-2}	0.867×10^{-2}
θ_3	i/e log density gradient	$-L_s \partial_x \ln n$	1.156	1.927
θ_4	i log temperature gradient	$-L_s \partial_x \ln T_i$	2.096	3.494
θ_5	temperature ratio	T_i/T_e	0.610	0.670
θ_6	e log temperature gradient	$-L_s \partial_x \ln T_e$	4.040	6.733
θ_7	effective ion charge	$Z_{\text{eff}} = \sum_i n_i q_i^2 / \sum_i n_i$	1.280	1.920
θ_8	safety factor	q	2.170	2.399
θ_9	magnetic shear	$\hat{s} = \frac{r}{q} \frac{dq}{dr}$	1.992	2.435
θ_{10}	elongation	k	0.128	0.141
θ_{11}	elongation gradient	$s_k = \frac{r}{k} \frac{\partial k}{\partial r}$	0.200	0.250
θ_{12}	triangularity	δ	0.710	0.870

Table 4.4: Summary of the 12 stochastic inputs considered for the ASDEX Upgrade test case.

We employ an equivalent setup in our proposed approach and in the standard adaptive approach to the ones in the 3D scenario: we employ the same maximum level $K_{\text{max}}^{\text{in-g}} = 20$ and the same tolerances for both algorithms (the tolerances for sensitivity scores are extended from three to 12 stochastic dimensions).

We visualize the expectation, one standard deviation and the deterministic growth rates in Figure 4.15. Moreover, for comparison, we depict the expectation and the one standard deviation from the 3D test case as well. On the one hand, we see that for this scenario too, our proposed approach with either of the two ways of assessing sensitivity scores yields results very similar to the standard adaptive strategy. Moreover, the difference between the deterministic and UQ results is quantitatively very similar to the 3D case. On the other hand, we observe that the expectations and standard deviations for the 12D case overlap almost perfectly with the 3D results. Therefore, it looks like the extra nine uncertain parameters contribute insignificantly to the overall results. This would be an important piece of information for on-going gyrokinetic validation studies: our uncertainty propagation results indicate that in a nonlinear turbulence simulation, the focus should be on only the sensitivities of the logarithmic gradients.

To ascertain the assumption that the extra nine stochastic parameters are stochastically insignificant, we depict the total Sobol' indices for sensitivity analysis in Figure 4.16 (top: standard adaptive interpolation, center: adaptive interpolation with sensitivities scores in which

4.3. NUMERICAL RESULTS: COMPUTED DOMINANT AMPLITUDE EIGENMODE




Figure 4.15: Expected values with one standard deviation as well as the deterministic growth rates for the ASDEX Upgrade test case with 12 uncertain parameters and $k_y\rho_s$, as in (4.14).

we compute on score for all interaction, bottom: adaptive interpolation in which we compute a score for each interaction). Note, first of all, that as for the 3D scenario, our proposed approach produces results very similar to the standard dimension-adaptive approach. We see that besides some small contributions due to the magnetic shear, \hat{s} , at $k_y\rho_s \leq 0.5$, the total Sobol' indices associated to the logarithmic temperature and density gradients have the largest values for $k_y\rho_s \leq 1.0$, while the remaining nine parameters have negligible Sobol' indices. Moreover, when $k_y\rho_s \geq 2.0$, the logarithmic electron temperature gradient is the most important parameter and the other 11 are negligible. Throughout the considered $k_y\rho_s$ domain, the total Sobol' indices of the aforementioned important parameters are quantitatively very similar to the indices from the 3D scenario. Hence, we conclude that the nine extra stochastic input parameters contribute insignificantly to the output uncertainty. In addition, the underlying microinstabilities are the same as for the 3D scenario.

We visualize the cost comparison between our proposed approach with the two strategies to assess sensitivity scores and the standard adaptive approach in Figure 4.17 (left: $k_y\rho_s \leq 1$, right: $k_y\rho_s \geq 2$). We observe that for all $k_y\rho_s$, dimension-adaptivity based on sensitivity scores requires significantly fewer GENE evaluations than the standard adaptive approach. We discuss first the cost for $k_y\rho_s \leq 1$. On the one hand, our approach in which we compute one sensitivity score for all interactions requires between around 1.8 and 9.4 times fewer grid points than the standard dimension-adaptive approach. On the other hand, when we assess a sensitivity score for all interactions, we reduce the number of (L)-Leja points from a factor of 2.3 to a factor of 13 compared to the standard adaptive method. For $k_y\rho_s \geq 2$, our approach with either strategy of computing the sensitivity scores exploits the fact that the ambient stochastic dimensionality is only one and significantly reduces the number of grid points as compared to the standard dimension-adaptive variant. We observe that the cost of our approach is similar for all $k_y\rho_s$ and it is kept between 105 and 167 number of (L)-Leja points. For example, at $k_y\rho_s = 15$, our approach in which we assess a sensitivity score for each interaction requires 14.5 times fewer evaluations than the standard adaptive approach. To put these cost savings into perspective, a full-tensor grid used in typical parameter scans in the plasma physics community, constructed using 10 points in each direction would require 10^{12} GENE evaluations in total, which is computationally prohibitive. In contrast, the standard adaptive strategy required at most

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS




Figure 4.16: Total Sobol' indices for global sensitivity analysis via adaptive sparse grid interpolation for the ASDEX Upgrade test case with 12 uncertain parameters and $k_y \rho_s$ as in (4.14) (top figures: standard dimension-adaptivity is employed; middle figures: we use our approach in which we compute one score for all interactions; in the bottom figures, we compute one score for each interaction in our proposed approach).

4.4. NUMERICAL RESULTS: TRACING PARTICULAR MICROINSTABILITIES

2283 GENE evaluations, whereas the largest number of simulations in the proposed approach was 546 for the stochastic scenario with 12 stochastic inputs.




Figure 4.17: Total number of GENE evaluations needed to construct the sparse interpolation surrogate using our proposed approach based on sensitivity scores vs. the standard adaptive approach.

To conclude these two stochastic scenarios for the ASDEX Upgrade test case, the proposed Sensitivity-driven dimension-adaptive sparse grid algorithm proved to be comparatively accurate as the standard adaptive approach at significantly reduced computational cost. In addition, we saw that dimension-adaptivity enables the uncertainty propagation and sensitivity analysis in problems in which standard approaches would be almost impossible to use. Moreover, when the intrinsic stochastic dimensionality is smaller than the given dimensionality, which is usually the case in practice, our algorithm explores and exploits this structure to significantly reduce the cost of uncertainty propagation. Finally, the results yielded by our UQ analysis aided the physical interpretation and reveal insights about the underlying microturbulence.

4.4 Numerical results: tracing particular microinstabilities

In Section 4.3, we employed our proposed context-aware sensitivity-driven approach in two (stochastic) plasma microturbulence test cases in which the output of interest was the computed dominant eigenmode of the amplitude. This UQ analysis is already valuable for the plasma physics community, as it provides (i) a quantitative measure of uncertainty (in our numerical experiments, the standard deviation) associated to the expected value of the dominant amplitude eigenmode, which could help physicists to compare and predict microturbulence in plasmas and (ii) a quantitative measure of importance of stochastic parameters, via total Sobol' indices, which can help finding improved reduced physics models and can simplify the parameter setup in the considerably more expensive nonlinear simulations.

For a more comprehensive overview, we employ our proposed method for uncertainty propagation and compare it with the standard dimension-adaptive strategy (4.12) when we *trace* a particular microinstability mode, that is, when the output of interest is a specific eigenmode of the amplitude, such as TEM/ETG, that is, an electron-driven mode, or ITG, an ion-drive mode (see Section 3.2). For simplicity, we restrict our attention to the ASDEX Upgrade test case [63] with three stochastic parameters considered in Section 4.3.3.

We performed these GENE simulations on 32 cores on a Intel Xeon 8160 (SkyLake) accelerator from the Marconi supercomputer at the CINECA Supercomputing Center. Since, as we will

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

see next, they required computing the *first dominant* and the *first subdominant* eigenmodes, the runtime was larger than before and varied between roughly 10 and 70 minutes on 32 cores.

4.4.1 Realistic test case with three stochastic parameters

The three uncertain inputs are the logarithmic density and temperature gradients for ions and electrons with values listed in Table 4.2. As before, we perform simulations for multiple values of $k_y \rho_s$, but, for simplicity, we restrict the wave number range to

$$k_y \rho_s = 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.9, 1.0. \quad (4.15)$$

The rest of the setup is identical to what we considered in Section 4.3.3. Therefore, to assess the quantities of interest, we employ dimension-adaptive sparse grid interpolation in which refinement indicator is given by the two proposed methods to assess sensitivity scores. Moreover, we also compare our approach with the standard indicator (4.12).

Recall the left and right bounds of the three logarithmic gradients, listed in columns four and five in Table 4.2, which were obtained by respectively varying by 25% the nominal value of the gradients, stemming from experiments. Let us compute *both* the dominant and first subdominant eigenmode spectra using these values. We visualize the results in Figure 4.18. On the one hand, as $k_y \rho_s$ increases, the transition to an ETG mode is clearly visible, since the




Figure 4.18: Growth rate (left) and frequency spectra (right) of the dominant eigenmode (black) and first subdominant eigenmode (red) for the ASDEX Upgrade test case obtained using the nominal values, i.e., the expectation of the input parameters.

associated frequency is negative, while the first subdominant frequency is positive. On the other hand, especially for $k_y \rho_s \leq 0.5$, where the two deterministic modes are closer to each other, it can happen that in UQ simulations, electron-driven eigenmodes, such as TEM/ETG, and ion-driven modes, e.g., ITG, are *interchanged*. For example, if the underlying mode is electron-driven (corresponding to negative frequency), the computed dominant amplitude eigenmode might have complementary sign, i.e., it is an ion-driven mode. We illustrate this point in Figures 4.19–4.20. In Figure 4.19, we depict the computed dominant (in black) and the first subdominant (in red) modes for the frequency spectra for all UQ simulations at $k_y \rho_s = 0.5$. We observe that the two computed modes are not clearly separated, i.e., for some simulations, the electron- and ion-driven modes are interchanged. Furthermore, in Figure 4.20 we depict the *sensitivity* of the two computed modes w.r.t. each individual input parameter. That is, one of the logarithmic gradients (density, ion temperature and electron temperature) varies globally in the associated uniform domain (see Table 4.2), while the other two parameters are *fixed* to their nominal value. In the top figures, we depict the results for the amplitude and the bottom figures, the results for the frequency. We indeed observe interchanged frequency modes w.r.t. all

4.4. NUMERICAL RESULTS: TRACING PARTICULAR MICROINSTABILITIES




Figure 4.19: Computed dominant (black) and first subdominant (red) eigenmodes for the frequency spectra at $k_y \rho_s = 0.5$ for the ASDEX Upgrade test case obtained using the nominal values, i.e., the expectation of the input parameters.




Figure 4.20: Sensitivity of the two computed (black: first, red:second) eigenmodes w.r.t. to each individual input logarithmic gradients (left: density gradient, center: ion temperature gradient and right: electron temperature gradient). In the top figures, we depict the results for the amplitude and the bottom figures, the results for the frequency. In each figure, one logarithmic gradient varies globally in the uniform domain listed in Table 4.2, while the other two parameters are *fixed* to their nominal value.

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

three input gradients. Therefore, in some UQ simulations, the second subdominant eigenmode is the actual electron- or ion-driven mode. This motivates to quantify the uncertainty when the output of interest is a *specific* microinstability mode as well.




Figure 4.21: Underlying (black) and subdominant (red) eigenmodes for the frequency spectra at $k_y \rho_s = 0.5$ for the ASDEX Upgrade test case obtained using the expectations of the input parameters. To determine the two modes, we employed (4.16).




Figure 4.22: Sensitivity of the underlying (black) and subdominant (red) amplitude (top figures) and frequency eigenmodes (bottom figures) w.r.t. to each individual input logarithmic gradients (left: density gradient, center: ion temperature gradient and right: electron temperature gradient). In each figure, one logarithmic gradient varies globally in the uniform domain listed in Table 4.2, while the other two parameters are *fixed* to their nominal value.

Sparse grid approximations are defined in terms of an admissible multiindex set, \mathcal{K} (see Section 2.5). Moreover, recall from Section 2.5.2 that for interpolation, the number of surplus (L)-Leja points for any subspace corresponding to a multiindex \mathbf{k} is always one. Let therefore $\boldsymbol{\theta}_{\mathbf{k}}$ denote the surplus (L)-Leja point associated to \mathbf{k} . Moreover, let $a_1(\boldsymbol{\theta}_{\mathbf{k}}), a_2(\boldsymbol{\theta}_{\mathbf{k}})$ denote the first two computed *amplitude* eigenmodes and $f_1(\boldsymbol{\theta}_{\mathbf{k}}), f_2(\boldsymbol{\theta}_{\mathbf{k}})$ the first two computed *frequency* eigenmodes using the (L)-Leja point $\boldsymbol{\theta}_{\mathbf{k}}$. We denote by $a_p(\boldsymbol{\theta}_{\mathbf{k}})$ the *underlying amplitude eigenmode*

4.4. NUMERICAL RESULTS: TRACING PARTICULAR MICROINSTABILITIES

associated to $\boldsymbol{\theta}_{\mathbf{k}}$. For a multiindex $\mathbf{k} \in \mathcal{K}$, we define its *backward neighbour set* as

$$\text{backneigh}(\mathbf{k}) := \{\mathbf{j} \in \mathcal{K} : |\mathbf{k} - \mathbf{j}|_1 = 1\}.$$

For example, if $\mathcal{K} = \{(1, 1); (1, 2); (2, 1); (2, 2); (3, 1)\}$, then $\text{backneigh}((2, 1)) = \{(1, 1)\}$, whereas $\text{backneigh}((2, 2)) = \{(1, 2), (2, 1)\}$.

Since dimension-adaptivity increases the cardinality of \mathcal{K} sequentially, we determine $a_p(\boldsymbol{\theta}_{\mathbf{k}})$ for each $\mathbf{k} \in \mathcal{K}$ sequentially, too, as follows. If $\mathbf{k} = \mathbf{1}$, then $a_p(\boldsymbol{\theta}_{\mathbf{k}}) = a_1(\boldsymbol{\theta}_{\mathbf{1}})$, i.e., for the first (L)-Leja point (which is nothing else but the expected value of the stochastic inputs), the underlying mode is the first computed mode. For the remaining multiindices in \mathcal{K} , we define:

$$a_p(\boldsymbol{\theta}_{\mathbf{k}}) = \begin{cases} a_1(\boldsymbol{\theta}_{\mathbf{k}}), & \text{if } \sum_{\mathbf{j} \in \text{backneigh}(\mathbf{k})} |f_1(\boldsymbol{\theta}_{\mathbf{k}}) - f_1(\boldsymbol{\theta}_{\mathbf{j}})| \leq \sum_{\mathbf{j} \in \text{backneigh}(\mathbf{k})} |f_2(\boldsymbol{\theta}_{\mathbf{k}}) - f_2(\boldsymbol{\theta}_{\mathbf{j}})| \\ a_2(\boldsymbol{\theta}_{\mathbf{k}}), & \text{else.} \end{cases} \quad (4.16)$$

In other words, we compute the sum of *absolute frequency eigenmode distances* between the current dominant and subdominant modes, and the dominant and subdominant modes of the Leja points in the backward neighbour set. The underlying amplitude mode is hence given by the computed amplitude mode with the *minimum* associated frequency distance. We illustrate the results of the aforementioned algorithm to determine the underlying mode in Figures 4.21 and 4.22, for the same data as in Figure 4.19.

We visualize, in Figure 4.23, right, the expected value and one standard deviation yielded by our context-aware approach with either strategy to assess sensitivity scores and by the dimension-adaptive interpolation with the standard refinement indicator. Moreover, in the left figure we depict, for comparison, the equivalent results from Section 4.3.3, i.e., when the output of interest was the computed dominant amplitude eigenmode. For this scenario, too, our approach is comparatively accurate as the standard method. We observe that when the output of interest is the underlying amplitude mode, the expectation overlaps almost perfectly with the deterministic spectra. Compared again with the results from the left figure, we observe that the standard deviations obtained here are larger for most values of $k_y \rho_s$. We see thus that performing uncertainty propagation for two different outputs of interest for the same simulation setup offers a broad perspective on the impact and influence of the input uncertainty.

In Figure 4.12, we depict the total Sobol' indices of the three uncertain inputs yielded by each refinement strategy. We see once again that our approach has similar accuracy to the standard adaptive method. Moreover, if we compare these results to the total Sobol' indices obtained in Section 4.3.3, we observe that the indices in Figure 4.24 are qualitatively similar: for $k_y \rho_s \leq 0.7$, the logarithmic electron temperature gradient is the most important parameter, the logarithmic density gradient the second and the logarithmic ion temperature gradient the least important. Moreover, in this wave number range, all three parameters are important. For $k_y \rho_s \geq 0.8$, the logarithmic density gradient becomes unimportant and thus the intrinsic stochastic dimensionality decreases to two. Finally, the values of total Sobol' indices in Figure 4.24 indicate indeed a stronger coupling between the input parameters than what was observed in Figure 4.12 (this was indicated by the larger standard deviations in Figure 4.23 as well). For example, at $k_y \rho_s = 0.2$, the total Sobol' indices from here sum up to 1.4392, whereas in Figure 4.12, their sum was 1.2270. Moreover, at $k_y \rho_s = 0.7$, the three Sobol' indices in Figure 4.24 give a total of 1.2526, compared to a summation value of 1.0235 in Figure 4.12.

Finally, we plot the cost in terms of total number of GENE evaluations for our proposed context-aware approach and the standard refinement method. We see that at most perpendicular wave numbers, our approach is cheaper than the standard refinement strategy. For example, at $k_y \rho_s = 0.5$, we reduce the cost by a factor of around 2.5. Moreover, from the two variants of

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS




Figure 4.23: Expected values with one standard deviation as well as the deterministic growth rates for the ASDEX Upgrade test case with three uncertain parameters and $k_y \rho_s$, as in (4.15). In the left figure, we depict the results when the output of interest is the *computed dominant amplitude eigenmode* (see Section 4.3.3). On the right, we visualize the results when the output of interest is the *underlying amplitude eigenmode*.




Figure 4.24: Total Sobol' indices for global sensitivity analysis via adaptive sparse grid interpolation for the ASDEX Upgrade test case with three uncertain parameters and $k_y \rho_s$ as in (4.15) (upper-left figure: standard dimension-adaptivity is employed; upper-right figure: our approach in which we compute one score for all interactions; in the lower-left figure, we compute one score for each interaction in our proposed approach).

our sensitivity scores approach, the strategy in which we assess a sensitivity score for each interaction is slightly cheaper than the variant which computes one score for all interactions. Due to the stronger nonlinear coupling between the stochastic inputs, the cost savings are not as large as in Section 4.3.3. However, our approaches required at most 236 GENE evaluations, compared to 363, the largest number required by the standard approach, both of which are smaller than the number of simulations required by standard techniques in plasma physics community.




Figure 4.25: Total number of GENE evaluations needed to construct the sparse interpolation surrogate using our proposed approach based on sensitivity scores vs. the standard adaptive approach.

In this test case, we saw that considering different outputs of interest for the same simulation scenario provides a more comprehensive understanding of the impact and influence of the input uncertainty. Moreover, we again showed that for a similar setup, our approach is as accurate as the standard approach, but at a smaller computational cost.

Conclusion and remarks

In this chapter, we presented a Sensitivity-driven dimension-adaptive sparse grid algorithm for efficient uncertainty propagation in higher-dimensional, computationally expensive real-world problems. As a representative example, we considered the analysis of plasma microinstability, a problem of high practical relevance for efforts such as the ITER project, which aims to demonstrate a net gain from nuclear fusion. The proposed method employs model hierarchies appearing in sparse grid approximations formulated in terms of *global* operators, such as interpolation in terms of Lagrange polynomials or PSP. To construct these approximations, we used (L)-Leja sequences, which were proved to be accurate at low cardinality. Moreover, (L)-Leja sequences are nested, have very fine (linear) granularity and can be constructed for arbitrary probability measures, which makes them an excellent choice for constructing sparse grid approximations. The starting point in formulating our approach was the dimension-adaptive sparse grid algorithm [71,90]. The essential ingredient in this algorithm is a refinement indicator which is used to guide the adaptive process. However, standard refinement indicators are not computed using stochastic information, but rather heuristics for the deterministic approximation error. Additionally, the employed heuristics in standard refinement indicators are generally based on global information in each subspace, usually a norm.

CHAPTER 4. SENSITIVITY-DRIVEN DIMENSION-ADAPTIVE SPARSE GRID APPROXIMATIONS

To this end, we designed refinement indicators that employ stochastic information about each uncertain parameter and interactions thereof. We first showed how the (squared) L^2 norm of the surpluses associated to each multiindex can be decomposed into expectation and variance surpluses, for both interpolation and PSP. Furthermore, exploiting the connection between PSP coefficients and Sobol' indices, we also showed how the variance surplus can be further decomposed into a summation of unnormalized Sobol' indices characterizing both the individual input parameters and their interaction. From this variance decomposition, we formulated two sensitivity scoring strategies, one in which all interactions are considered separately, which is feasible in low to medium dimensionality, and another in which all interactions are considered together, which is feasible in arbitrary settings. The obtained sensitivity score was then used as refinement indicator in the dimension-adaptive sparse grid algorithm. Our goal with this strategy was to explore and exploit the rich structure typically found in real-world problems, such as the anisotropic coupling of the input parameters and the lower intrinsic stochastic dimensionality. To test the proposed approach, we did a comprehensive numerical study involving two simple examples with algebraic solutions and two plasma microturbulence test cases. For the latter, we first considered a modified gyrokinetic benchmark with eight uncertain parameters to obtain initial insights into the behaviour of the proposed approach. Therein, we compared the proposed sensitivity score-based refinement with the standard refinement indicator involving the L^2 norm of the surpluses in terms of accuracy and cost of assessing the quantities of interest, for both interpolation and PSP. We concluded that for the same setup, sparse grid interpolation is computationally cheaper than PSP and we therefore only considered interpolation in the second plasma microturbulence test case. Next, we considered a real-world plasma microinstability analysis problem, a particular discharge [63] of the ASDEX Upgrade experiment. For this second test case, the uncertainty analysis was done for two stochastic scenarios, one with three and the other with 12 stochastic inputs, and for two outputs of interest. Furthermore, we compared our sensitivity scores approach with the standard dimension-adaptive refinement. Our results showed that to obtain a similar accuracy for both approximation and computation of the quantities of interest, our proposed approach is computationally cheaper than the standard approach up to factors of 14.5. In these two test cases, the output of interest was the computed dominant amplitude eigenmode. We showed, however that in certain situations, the computed mode can be either electron- or ion-driven. This suggested a complementary UQ analysis, in which we trace the particular amplitude mode associated to a wavenumber. We performed this analysis in the ASDEX Upgrade scenario with three uncertain inputs and again showed that our approach is more efficient than the standard approach. Our entire UQ analysis is valuable for the plasma physics community, as it provides: (i) a quantitative measure of uncertainty (in our numerical experiments, via the standard deviation) associated to the expected value of the output of interest, which could help physicists to compare and predict microturbulence in plasmas, (ii) a quantitative measure of importance of stochastic parameters, via total Sobol' indices, which can help finding improved reduced physics models and can simplify the parameter setup in the considerably more expensive nonlinear simulations and (iii) a flexible framework for performing uncertainty propagation and sensitivity analysis for arbitrary outputs of interest. In addition, for stochastic problems with up to 12 stochastic parameters, our proposed approach required at most 546 runs of the forward model. Even when such simulations are computationally more expensive, on today's supercomputers, 546 runs are feasible. We see therefore that exploring and exploiting structure in real-world problems is key to enable the quantification of uncertainty and sensitivity analysis in such problems.

The presented methodology is based on the assumptions that the uncertain inputs are stochastically independent and that the forward map parametrized in terms of these inputs is sufficiently smooth for interpolation or PSP to be employed. When the first assumption

is not fulfilled, we could employ nonlinear transformations, such as transport maps [121], to write the dependent inputs as functions of independent parameters. Nevertheless, any type of mapping in this case needs to be used with care, as there is no guarantee that in general, the known properties or convergence rates will still hold true after the mapping is employed. When the forward parametrized map is not smooth or has sharp gradients, as it was the case in one of our plasma physics test cases, the proposed approach can be reformulated, for example, in terms of basis functions with local support (see also Chapter 5), such as hat functions, piecewise polynomials or wavelets. In addition, the proposed sensitivity scores approach can be extended in several ways. For example, instead of assessing one sensitivity score for all input interactions or one sensitivity score for each interaction, one could compute a score for all interactions involving each input. That is, for a problem with d_{sto} stochastic inputs, d_{sto} scores for each individual parameters and d_{sto} additional scores for all interactions involving each input would be computed. Furthermore, the sensitivity scoring refinement can be enhanced, for example, with filtering step that prevents the algorithm to add multiindices in the directions that were rendered unimportant (for a similar filtering idea, see the adaptive refinement strategy based on directional variances presented in Section 7.3.2). We end this chapter by pointing out again that the behaviour of the proposed approach depends on how much structure the underlying stochastic problem has. When the stochastic inputs are strongly coupled, or when the intrinsic dimensionality is large, then alternative techniques, e.g., based on MC sampling, would be a more appropriate choice.

This concludes the first contribution chapter. In the next chapter, we present our second contribution for uncertainty propagation, a novel algorithm employed multilevel model hierarchies in which we perform context-aware, online stochastic dimensionality reduction.

“Evolution has ensured that our brains just aren’t equipped to visualise 11 dimensions directly. However, from a purely mathematical point of view it’s just as easy to think in 11 dimensions, as it is to think in three or four.”

Stephen Hawking

5

Multilevel adaptive spectral projections with online dimensionality reduction

In the previous chapter, we presented a context-aware dimension-adaptive sparse grid algorithm for uncertainty propagation, in which sensitivity information stemming from unnormalized Sobol’ indices was employed to drive the adaptive process. In the fifth chapter of this work, we propose a multilevel adaptive spectral projection approach in which we employ sensitivity information to perform *online stochastic dimensionality reduction* in uncertainty propagation settings. As the multilevel decomposition is performed, we check, online, how many uncertain inputs are stochastically important and we reduce the dimensionality accordingly. This makes our approach *context-aware* since the outcomes of the dimensionality reduction procedure depend on the properties of the problem under consideration. Furthermore, at each level in the multilevel hierarchy, we formulate a new approach to assess the pseudo-spectral coefficients based on spatially adaptive sparse grids [147] and one-dimensional quadrature. Since this approach of assessing pseudo-spectral coefficients is based on sparse grids with hierarchical basis functions summarized in Section 2.6, it introduces an additional model hierarchy. Therefore, the proposed multilevel approach with online dimensionality reduction employs two layers of model hierarchies. Our algorithm is designed for black-box uncertainty propagation in higher-dimensional, computationally expensive problems. As a representative example of such problems, we consider and FSI test case with five stochastic inputs (see Section 5.4.2).

The algorithm ideas and results presented in this chapter are mainly based on two of our research works, [56] and [57]. In [57], we formulated the new approach to assess PSP coefficients, the Spatially adaptive sparse pseudo-spectral approximation methods, which we present here in Section 5.3.1. We do not explicitly show our numerical results from that work, but we reference our results and conclusions from [57] throughout this chapter. We extended the spatially adaptive pseudo-spectral approach to the multilevel version and online stochastic dimensionality reduction in [56], which is the main focus of this chapter. The two test cases considered in [56] are also presented in Section 5.4, but we here we do more numerical experiments in Section 5.4.1 for a more comprehensive overview. Note that in [56], we employed two approaches to assess PSP coefficients, one based on spatially adaptive sparse grids and another on dimension-adaptive Leja interpolation, but our focus in this chapter is on spatially adaptive sparse grids.

This chapter is organized as follows. In Section 5.1, we discuss how UQ in FSI simulations, including how we incorporate uncertainty in the considered FSI test case in Section 5.4.2. In Section 5.2, we overview existing techniques for reducing the stochastic dimensionality in uncertainty propagation problems. Examples include local sensitivity analysis, Morris screening and active subspaces. We present in detail our proposed *context-aware multilevel spectral projection with dimensionality reduction* approach in Section 5.3. First, we introduce a novel technique for assessing PSP coefficients based on *spatially adaptive sparse grid approximations* and *one-dimensional quadrature*, which we call Spatially adaptive pseudo-spectral approach (see

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

Section 5.3.1). We then extend this approach to a multilevel version in which we additionally check, online, if the stochastic dimensionality can be reduced at *no additional computational cost*. Specifically, we exploit the fact that a multilevel decomposition entails (i) solving the same problem with different resolutions and (ii) that this set of problems can be solved independently. To this end, we first compute a *subset* of solutions using the given stochastic dimensionality and assess the total Sobol' indices for global sensitivity analysis. In the next step, based on user-defined thresholds, we ascertain the importance of each stochastic input. The inputs that are rendered unimportant are *replaced* by a corresponding *deterministic value*, usually their *expectation*, and the remaining subset of problems is solved using samples from the space of reduced dimensionality. Additionally, the results from the previous step computed using the original dimensionality are mapped onto the space of reduced dimensionality. In Section 5.4, we present our numerical results. We consider two test cases, a damped oscillator and an FSI simulation. For a more comprehensive overview of the proposed UQ approach, we consider two stochastic scenarios in each test case. For the oscillator test case, we employ multilevel decompositions with or without adaptivity or dimensionality reduction and perform a comprehensive numerical study to test the impact of stochastic dimensionality reduction on the accuracy and cost of computing quantities of interest. The FSI scenario is more complex and computationally more expensive, therefore we employ only the most promising approach from the first test case.

Highlights and novelties

- We employ multilevel model hierarchies and formulate a novel context-aware multilevel approach for uncertainty propagation in which we perform online stochastic dimensionality reduction at no additional cost;
- At each level in the multilevel hierarchy, we formulate a new method to assess PSP coefficients, based on spatially adaptive sparse grids and one-dimensional quadrature;
- We employ the proposed approach, which we call Multilevel adaptive spectral projection with online stochastic dimensional reduction, in two test cases: an oscillator test case, first considering five and then six uncertain inputs, and in a more complex FSI scenario with five uncertain inputs;
- We show, numerically, that the proposed approach leads to dimensionality reduction in the two test cases while (i) maintaining the accuracy of standard approaches and (ii) drastically reducing the computational cost;
- In the FSI test case, the standard multilevel approach based on telescoping sums required about 70.2 CPU hours, for one stochastic scenario, and 129.4 for another; our approach reduced the runtime with at least 54 CPU hours in the first scenario and at least 22 hours in the second.

5.1 Quantifying uncertainty in fluid-structure interaction simulations

FSI problems play a prominent role in many domains and pose many challenges due to their non-linearity and multidisciplinary nature. Representative examples include the study of blood flow, calcification, and aneurysms, see e.g. [113] and aeronautics. For the latter, flow-induced vibrations have an important influence on the stability and durability of aircrafts (see e.g. [58]). Parachutes [177] or noise prediction [158] are other examples, to only name a few. We study

FSI not only due its practical relevance, but also as a representative of a challenging and computationally expensive multi-physics application. In general, FSI features several physical parameters which are affected by uncertainty as they typically stem from measurements: both fluid and structure density, the viscosity of the fluid, the structural elasticity and plasticity parameters, possibly different in every spatial dimension.

A non-exhaustive list of previous research concerning UQ in FSI includes [195], where the focus has been on uncertainty propagation in unsteady oscillatory problems, using a stochastic collocation approach with Newton-Cotes quadrature in simplex elements. In [118], the authors employed a two-step collocation approach in a linear piston problem. In addition, in [196], a generalized spectral projection based on Gram-Schmidt orthogonalization was employed to analyse uncertainties in a single-degree of freedom stall flutter model. In [151], pressure induced oscillations of a disc-like structure have been analysed via a model reduction approach and sparse grid spectral projections.

In Section 5.4.2, we consider a 2D fluid flow scenario and a Saint-Venant-Kirchhoff model for the structure. This leaves us with a five-dimensional stochastic space: fluid and structure density, the dynamic fluid viscosity, the structural elasticity module and Poisson’s ratio (recall Section 3.1). We do not consider uncertain quantities with spatial variation, such as uncertainty in boundary conditions or non-uniform materials. Possible outputs of interest include the frequency and amplitude of vibrations or oscillations, the maximum stress and its location, or the mean flow rate. In the FIS test case in Section 5.4.2, we look at the structure’s x -deflection. Next, we discuss existing techniques for dimensionality reduction in uncertainty propagation.

5.2 Stochastic dimensionality reduction

Uncertainty propagation in higher-dimensional, computationally expensive problems is very challenging for a number of reasons. Since uncertainty propagation is an *outer-loop* around the underlying high-fidelity forward model, repeated evaluations of a computationally expensive model can be prohibitive. Even on a supercomputer, if one run requires most of the resources, performing even a few simulations is already challenging. The number of total high-fidelity evaluations that we need to perform in uncertainty propagation is directly related to the number of uncertain inputs: the more inputs, the more simulations we need to perform. However, as we already discussed in Chapter 4, usually only a *subset* of stochastic inputs is important in a UQ problem. That is equivalent to saying that in general, the intrinsic stochastic dimensionality is smaller than the given dimensionality, d_{sto} .

There are several approaches to ascertain and reduce stochastic dimensionality in UQ simulations. For example, *local sensitivity analysis* [156] assesses the local impact of uncertain inputs by varying them in a neighbourhood around a reference value; the sensitivity measure is usually the value of the gradient of the underlying model’s response. Another established technique for stochastic dimensionality reduction is *Morris screening* [131]. In Morris screening, one uncertain input is varied at a time (it is a so-called *one-step-at-a-time method*) while the others are kept to their nominal values. This results in an *elementary effect* for each stochastic inputs. The magnitude of these effects ascertains the importance of each input parameter. Besides sensitivity analysis, there are also other approaches to identifying the subset of important stochastic inputs. For example, in [157], the dimensionality of forward UQ problems was reduced by using sparse Bayesian learning to retain only the most relevant basis terms of a spectral projection approximation. In addition, another prominent approach for dimensionality reduction is based on *active subspaces* (see [36] for a comprehensive introduction into this topic). Loosely speaking, in its standard form, this approach seeks to find the subspace of

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

“active” directions by using the decay of the spectrum of the matrix formed as the outer product of the model’s response gradient. Active subspaces have the advantage that they are also applicable to UQ problems characterized by *dependent* inputs. Some recent developments in this direction include [109], where multifidelity methods were blended with active subspaces for dimensionality reduction in engineering models. We also mention the recent work [93] in which a new machinery for dimensionality reduction was proposed. Specifically, since the challenges associated with higher-dimensionality stem from assuming that the underlying model solution is *Lipschitz continuous* (see, e.g., [49]), [93] proposed using a different, more general approach to measure distances between points, called *Lipschitz Matrix*, to delay or even overcome the curse of dimensionality. In most of these approaches, reducing the stochastic dimensionality either requires an additional preprocessing step and with it, additional computational work, or it is required to compute or approximate gradients of the system’s response. In the following, we present in detail our proposed approach in which we employ *multilevel spectral decompositions* and *spatially adaptive sparse grid approximations* to perform *online stochastic dimensionality reduction* at no additional computational cost.

5.3 Multilevel adaptive spectral projection with online dimensionality reduction

Motivated by the challenges of quantifying uncertainty in higher-dimensional, computationally expensive problems such as FSI simulations, our goal is to formulate an algorithm that:

1. *exploits* the *local structure* of the underlying problem;
2. *handles arbitrary* input probability *densities*;
3. employs *model hierarchies* to *decrease* the overall *computational cost*;
4. *exploits* the *lower intrinsic dimensionality* of UQ problems.

To achieve our first goal, we employ sparse grid interpolation with *hierarchical, locally supported* bases and *spatial adaptivity*, summarized in Section 2.6. Besides the capability to perform spatially adaptive refinement, we employ these sparse grid approximations because they allow, as we will show in Section 5.3.1, to efficiently¹ approximate multivariate integrals with products between hierarchical surplus and products of 1D integrals. Spatially adaptive sparse grid approximations are designed for functions supported in the unit hypercube, $[0, 1]^{d_{\text{sto}}}$. Thus, to achieve the second goal, we employ a *mapping* in terms of the *inverse cumulative distribution function* of the stochastic inputs to map $[0, 1]^{d_{\text{sto}}}$ to arbitrary probability spaces, \mathbf{X} . For the third goal, we make use of *multilevel model hierarchies* (see Section 2.7). Finally, to achieve our fourth and last goal, we take advantage of the *linearity* and *independence* of quantities entering standard multilevel approximations to perform *online stochastic dimensionality reduction*.

To this end, we formulate the *Multilevel adaptive spectral projection with online dimensionality reduction*, a novel approach for performing *online stochastic dimensionality reduction* in uncertainty propagation settings. Our approach is based on standard multilevel decompositions and PSP in which the quantities of interest estimated with the multilevel approach are the PSP coefficients. At each level in the multilevel hierarchy, we employ a new method to compute PSP coefficients based on *spatially adaptive sparse grid interpolation* and one-dimensional integration: with spatial adaptivity, we can exploit the *local structure* of the model’s solution

¹The efficiency of these approximations does not come for free, of course, but depends on, e.g., the smoothness and local structure of the underlying forward model

5.3. MULTILEVEL ADAPTIVE SPECTRAL PROJECTION WITH ONLINE DIMENSIONALITY REDUCTION

to create a low cost approximation, while the one-dimensional integrals can either be exactly evaluated or very accurately approximated via, e.g., Gauss-Legendre quadrature. We call the new approach to compute PSP coefficients *Spatially adaptive sparse pseudo-spectral approximation*. To perform online dimensionality reduction, we exploit the fact that standard multilevel decompositions employ a linear combination of quantities than can be assessed independently. Therefore, we compute a subset of the multilevel PSP coefficients via Spatially adaptive sparse pseudo-spectral approximations on $[0, 1]^{d_{\text{sto}}}$. Based on these PSP coefficients, we compute the d_{sto} total Sobol' indices (recall Section 2.9) and use them to *ascertain* the number of *important* input parameters.

Remark: We employ total Sobol' indices to ascertain the importance of a stochastic input because they comprise its local as well as contributions due to interactions with the other parameters. Therefore, a small total Sobol' index means that an uncertain input is unimportant both individually and relative to the other input parameters.

If a subset of the initial d_{sto} stochastic input parameters is rendered unimportant, we *replace* the unimportant inputs with a deterministic value, e.g., their expectation, and continue the multilevel approximation using a stochastic space of reduced dimensionality. For consistency, we also *project* the previous solution obtained on the d_{sto} -dimensional space onto the stochastic space with reduced dimensionality. Therefore, our proposed approach based on multilevel model hierarchies is *context-aware*, since the structure of the underlying problem is exploited to perform the stochastic dimensionality reduction. Moreover, the local structure of the solution at each level is exploited via spatially adaptive refinement in the Spatially adaptive sparse pseudo-spectral approximation algorithm. We note that in the worst-case scenario, if no dimensionality reduction is performed, we do a multilevel approximation of the PSP coefficients in d_{sto} dimensions, thus still reducing the computational cost compared to a single level approach. Before going further, we illustrate the usage of total Sobol' indices for stochastic dimensionality and the effects of replacing unimportant stochastic parameters with different deterministic values.

Example: We consider an uncertainty propagation problem in which the forward model is available analytically and depends on a six-dimensional input, uniformly distributed in $[0, 1]^6$:

$$\mathcal{F}(\boldsymbol{\theta}) = \cos(\theta_1 + 0.8\theta_2 + 0.2\theta_3 + 0.1\theta_4 + 0.05\theta_5 + 0.01\theta_6). \quad (5.1)$$

We estimate the six total Sobol' indices for global sensitivity analysis using a full Gauss-Legendre grid with $8^6 = 262144$ points. We obtain

$$\hat{S}_1 = 0.5987, \quad \hat{S}_2 = 0.3815, \quad \hat{S}_3 = 0.0237, \quad \hat{S}_4 = 0.0059, \quad \hat{S}_5 = 0.0014, \quad \hat{S}_6 = 5.9311 \cdot 10^{-5}.$$

We observe that the first two total indices have non-negligible values, whereas the remaining four are unimportant, with values smaller than 0.03. These values suggest that the stochastic dimensionality can be reduced from six to two, i.e., the intrinsic stochastic dimensionality of the considered uncertainty propagation problem is *two*.

We therefore reduce the stochastic dimensionality as indicated above. To obtain an overview of the effects of the reduction, we proceed as follows. First, we perform an MC simulation using all six parameters and we use the resulting ensemble of forward model evaluations to estimate the probability density function of the output, which we denote by $\pi(\mathcal{F})$, via kernel density estimation with a Gaussian kernel. Afterwards, we gradually reduce the number of uncertain inputs from six to five (we neglect θ_6), then to four, three and lastly, to two, i.e., only θ_1 and θ_2 are considered for uncertainty propagation. Moreover, we estimate $\pi(\mathcal{F})$ for each case. A neglected stochastic input is replaced by its uniform *expected value*, i.e., 0.5. In each simulation,

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

we use 10^5 MC samples. In Figure 5.1, we compare the output densities when using all six parameters (black dotted line) with the density obtained with the reduced number of stochastic inputs (red dotted line): from the upper left to the lower right figure, the dimensionality is reduced to five, four, three and lastly, to two, respectively. We observe that in all four cases, the two densities overlap, therefore the neglected parameters are indeed insignificant.




Figure 5.1: Comparison of the probability density functions of the output, $\pi(\mathcal{F})$, estimated using 10^5 MC samples and kernel density estimation, when using all six uniform uncertain inputs and a reduced number of parameters. From top left to bottom right, the dimensionality is reduced to five, four, three and two; we denote the reduced dimensionality by d_{red} . Moreover, the unimportant inputs are replaced with their expected value, 0.5.

Next, we want to assess the impact of replacing an unimportant stochastic input with a deterministic value different than its expectation. We consider the same forward model (5.1), but now the uncertain inputs are standard normal random variables, $N(0, 1)$. The total Sobol' indices obtained using a Gauss-Hermite grid with $8^6 = 262144$ points are:

$$\hat{S}_1 = 0.8594, \quad \hat{S}_2 = 0.6589, \quad \hat{S}_3 = 0.0558, \quad \hat{S}_4 = 0.0140, \quad \hat{S}_5 = 0.0035, \quad \hat{S}_6 = 0.0001.$$

Qualitatively, the Sobol' indices are similar to the case when the inputs were uniform, which is not surprising since sensitivity analysis reflects properties of the *model*. Therefore, we reduce the stochastic dimensionality from six to two. In this setting, we replace the four unimportant inputs, θ_3 to θ_6 , with either zero, i.e., their expected value or with a different value; in our experiments, we used $\theta_3 = \theta_4 = \theta_5 = \theta_6 := 1$. We visualize, in Figure 5.2, the output density functions when using either all six inputs or only the two important parameters, estimated with kernel density estimation using 10^5 MC evaluations. In the left figure, we depict the case in which we replace the unimportant inputs with their expected value, zero; on the right, we replaced them with the different value, i.e., one. We observe that the value with which we replace the unimportant stochastic parameters has no influence on the output density function.

5.3. MULTILEVEL ADAPTIVE SPECTRAL PROJECTION WITH ONLINE DIMENSIONALITY REDUCTION




Figure 5.2: Comparison of $\pi(\mathcal{F})$, estimated using 10^5 MC samples and kernel density estimation, when using all six standard normal uncertain inputs ($d_{sto} = 6$) and only the first two parameters ($d_{red} = 2$). In the left plot, the four unimportant inputs are replaced with their expectations, whereas in the right plot, with the value 1.

In summary, we saw that total Sobol' indices are an useful tool to assess the importance of stochastic inputs and to hence reduce the stochastic dimensionality. Moreover, we also showed that replacing stochastic rendered unimportant with a deterministic value different than their expectation has insignificant influence on the output results.

Remark: Another approach to verify whether dimensionality reduction is appropriate is to (i) perform uncertainty propagation with all uncertain inputs to obtain synthetic observations, (ii) employ Bayesian inversion using these observations to obtain the posterior density of the parameters (iii) verify whether the marginal posterior corresponding to the important inputs is similar, in a suitable metric, to the posterior density obtained using observational data w.r.t. to the forward propagation of the important parameters only. In other words, if the ignored parameters are indeed unimportant, this implies that their posterior and the aforementioned posterior marginal should be similar to each other. For more details on Bayesian inversion, see Section 2.4.2 and Chapter 7.

We note, however, that in uncertainty propagation we are generally interested in (scalar) quantities of interest such as expectation or standard deviation of outputs of interest, rather than in the entire output density. Thus, it is sufficient to ascertain the effects of dimensionality reduction directly in the quantities of interest. In other words, it is sufficient to have a statistical moments matching rather than a full density matching between the model using all stochastic parameters and the model using only the parameters rendered important.

In the following, we present the algorithm steps of the proposed multilevel approach in detail. First, we present the proposed method to assess PSP via spatially adaptive sparse grid interpolation and one-dimensional integrals in Section 5.3.1. Afterwards, in Section 5.3.2, we introduce some of our notation and summarize the setup used in our multilevel approach. Section 5.3.3 is focused on our approach for online stochastic dimensionality reduction. We end with discussing the computational cost, in Section 5.3.4, and the termination of the proposed multilevel approach, in Section 5.3.5.

5.3.1 Spatially adaptive sparse pseudo-spectral approximations

In this section, we present a novel approach to assess PSP coefficients, the Spatially adaptive sparse pseudo-spectral approximation algorithm, which can be found in our work [57]. Therein, we did a comprehensive numerical study using two FSI test case, the one considered in Section 5.4.2, and the FSI-3 benchmark [182], in which we showed that the proposed approach is more accurate than classical, Gaussian-based sparse grid quadrature approaches.

Consider a multivariate PSP expansion

$$\mathcal{U}^{\text{PSP}}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{P}^{\text{PSP}}} c_{\mathbf{p}} \Phi_{\mathbf{p}}(\boldsymbol{\theta}), \quad (5.2)$$

where \mathcal{P}^{PSP} is a set containing the multivariate PSP degrees (recall Sections 2.5.1 and 2.9). For simplicity, denoting $P^{\text{PSP}} = |\mathcal{P}^{\text{PSP}}|$, we use instead a scalar index $p = 1, 2, \dots, P^{\text{PSP}}$ in (5.2), i.e.,

$$\mathcal{U}^{\text{PSP}}[\mathcal{F}_h] := \sum_{p=1}^{P^{\text{PSP}}} c_p \Phi_p(\boldsymbol{\theta}). \quad (5.3)$$

Moreover, recall from Section 2.5.1 that the PSP coefficients c_p are obtained by imposing the residual of (5.3) to be orthonormal to the space spanned by the orthonormal basis, that is,

$$c_p = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \Phi_p(\boldsymbol{\theta}) \boldsymbol{\pi}(\boldsymbol{\theta}) d\boldsymbol{\theta} =: \mathbb{E}[\mathcal{F}_h \Phi_p]. \quad (5.4)$$

The computational cost of PSP is driven by evaluating (5.9) since this involves evaluations of the high-fidelity model, \mathcal{F}_h . Generally, the PSP coefficients are evaluated via quadrature w.r.t. the density $\boldsymbol{\pi}$, such as sparse grid quadrature [199]. In the following, we formulate an alternative methodology to assess these coefficients in which we employ spatially adaptive sparse grids, summarized in Section 2.6 and one-dimensional quadrature.

Let us take a closer look at the PSP coefficients c_p in (5.4). By Assumption 2.2, the stochastic inputs are independent and thus the input density, $\boldsymbol{\pi}$, has a product structure, i.e., $\boldsymbol{\pi} = \prod_{i=1}^{d_{\text{sto}}} \pi_i$. In addition, we know that the multivariate PSP basis polynomials are obtained via *tensorising* 1D orthonormal polynomials, i.e., $\Phi_p = \prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}$. Therefore, (5.4) can be written as

$$c_p = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(\theta_i) \prod_{i=1}^{d_{\text{sto}}} \pi_i(\theta_i) \right) d\boldsymbol{\theta} = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(\theta_i) \pi_i(\theta_i) \right) d\boldsymbol{\theta}. \quad (5.5)$$

Remark: When \mathcal{F}_h is nonlinear and computationally expensive, (5.5), standard non-adaptive quadrature techniques will quickly become prohibitive as d_{sto} increases, due to the curse of dimensionality. However, as we saw in Chapter 4, \mathcal{F}_h generally depends anisotropically on $\boldsymbol{\theta}$ and only a subset of the d_{sto} input parameters in $\boldsymbol{\theta}$ are stochastically important. This makes the numerical approximation (5.5) computationally tractable even in high-dimensional settings.

If \mathcal{F}_h had a product structure, e.g., $\mathcal{F}_h = \prod_{i=1}^{d_{\text{sto}}} \mathcal{F}_{h_i}$, (5.5) would become

$$\begin{aligned} c_p &= \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(\theta_i) \pi_i(\theta_i) \right) d\boldsymbol{\theta} = \int_{\mathbf{X}} \left(\prod_{i=1}^{d_{\text{sto}}} \mathcal{F}_{h_i}(\theta_i) \prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(\theta_i) \prod_{i=1}^{d_{\text{sto}}} \pi_i(\theta_i) \right) d\boldsymbol{\theta} \\ &= \int_{\mathbf{X}} \left(\prod_{i=1}^{d_{\text{sto}}} \mathcal{F}_{h_i}(\theta_i) \Phi_{p_i}(\theta_i) \pi_i(\theta_i) \right) d\theta_1 d\theta_2 \dots d\theta_{d_{\text{sto}}} = \prod_{i=1}^{d_{\text{sto}}} \int_{X_i} \mathcal{F}_{h_i}(\theta_i) \Phi_{p_i}(\theta_i) \pi_i(\theta_i) d\theta_i, \end{aligned} \quad (5.6)$$

5.3. MULTILEVEL ADAPTIVE SPECTRAL PROJECTION WITH ONLINE DIMENSIONALITY REDUCTION

where we used Assumption 2.1, i.e., $\mathbf{X} = \bigotimes_{i=1}^{d_{\text{sto}}} X_i$ to obtain the last equality. Therefore, a tensorized decomposition of the high-fidelity model, \mathcal{F}_h , transforms the multivariate integral (5.4) which defines the PSP coefficients into a product of one-dimensional integrals (5.6) which can be efficiently approximated via, e.g., Gaussian or Leja quadrature. Unfortunately, in problems characterized by complex models, \mathcal{F}_h does not usually have a product structure.

To arrive at an expression similar to (5.6), we approximate \mathcal{F}_h via sparse grid interpolation with hierarchical basis functions, summarized in Section 2.6. Recall that these approximations are defined for multivariate functions defined on $[0, 1]^{d_{\text{sto}}}$. To be able to employ them on arbitrary domains, \mathbf{X} , we need a (possibly nonlinear) transformation $\mathbf{T} : [0, 1]^{d_{\text{sto}}} \rightarrow \mathbf{X}$. The hierarchical sparse grid approximation depending on \mathbf{T} (recall (2.30), (2.31)) reads

$$\mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h \circ \mathbf{T}] = \sum_{\mathbf{k} \in \mathcal{K}} \alpha_{\mathbf{k}} \varphi_{\mathbf{k}}(\mathbf{u}), \quad (5.7)$$

where the transformation \mathbf{T} enters the computation of the hierarchical surpluses, $\alpha_{\mathbf{k}}$, i.e.,

$$\alpha_{\mathbf{k}} = \mathcal{F}_h(\mathbf{T}(\mathbf{u}_{\mathbf{k}})) - \mathcal{U}_{\mathcal{K} \setminus \{\mathbf{k}\}}^{\text{in},1}[\mathcal{F}_h(\mathbf{T}(\mathbf{u}_{\mathbf{k},v}))].$$

By Assumptions 2.1 and 2.2, $\boldsymbol{\theta}$ has *independent* components, hence we define the transformation component-wisely, i.e., $\mathbf{T}(\mathbf{u}) = (T_1(u_1), T_2(u_2), \dots, T_{d_{\text{sto}}}(u_{d_{\text{sto}}}))$. Therefore, in the following we make use of \mathbf{T} to map $[0, 1]^{d_{\text{sto}}}$, the domain of spatially adaptive sparse grid approximations, to the underlying stochastic domain, \mathbf{X} . The stochastic inputs θ_i are *continuous* random variables, being characterized by a density function, π_i , and a *cumulative distribution function*, G_i , which satisfies $G_i(X_i) = [0, 1]$ for $i = 1, 2, \dots, d_{\text{sto}}$. Since G_i is invertible, the natural way to define the mapping \mathbf{T} is

$$\mathbf{T}(\mathbf{u}) = (G_1^{-1}(u_1), G_2^{-1}(u_2), \dots, G_{d_{\text{sto}}}^{-1}(u_{d_{\text{sto}}}))). \quad (5.8)$$

Next, we show that (5.8) leads to integrating (5.4) w.r.t. the uniform density on $[0, 1]^{d_{\text{sto}}}$.

Lemma: 5.3.1 *Let $\mathbf{T} : [0, 1]^{d_{\text{sto}}} \rightarrow \mathbf{X}$ be defined as in (5.8). Then, (5.6) becomes*

$$c_p = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \Phi_p(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{[0,1]^{d_{\text{sto}}}} \mathcal{F}_h(\mathbf{T}(\mathbf{u})) \Phi_p(\mathbf{T}(\mathbf{u})) d\mathbf{u}, \quad p = 1, 2, \dots, P^{\text{PSP}}.$$

Proof: We have that $\mathbf{T}(\mathbf{u}) = (G_1^{-1}(u_1), G_2^{-1}(u_2), \dots, G_{d_{\text{sto}}}^{-1}(u_{d_{\text{sto}}}))$. By definition, $u_i = G_i(\theta_i)$ is uniformly distributed in $[0, 1]$ and $\theta_i = G_i^{-1}(u_i)$ is distributed according to π_i in X_i . We make the substitution $\boldsymbol{\theta} = \mathbf{T}(\mathbf{u})$ in

$$c_p = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \Phi_p(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}.$$

Standard probability theory ensures that \mathbf{T} defined in (5.8) is bijective, Lebesgue measurable and differentiable [170]. We have that

$$d\boldsymbol{\theta} = |\det(D\mathbf{T}(\mathbf{u}))| d\mathbf{u},$$

where $D\mathbf{T}(\mathbf{u})$ is the Jacobian of \mathbf{T} . Therefore,

$$c_p = \int_{\mathbf{X}} \mathcal{F}_h(\boldsymbol{\theta}) \Phi_p(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{[0,1]^{d_{\text{sto}}}} \mathcal{F}_h(\mathbf{T}(\mathbf{u})) \Phi_p(\mathbf{T}(\mathbf{u})) \pi(\mathbf{T}(\mathbf{u})) |\det(D\mathbf{T}(\mathbf{u}))| d\mathbf{u}.$$

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

Moreover, the transformation theorem for probability densities ensures that

$$\pi(T(\mathbf{u}))|\det(DT(\mathbf{u}))| = 1,$$

i.e., the uniform density on $[0, 1]^{d_{\text{sto}}}$ [170]. We hence obtain:

$$c_p = \int_{[0,1]^{d_{\text{sto}}}} \mathcal{F}_h(T(\mathbf{u})) \Phi_p(T(\mathbf{u})) d\mathbf{u}. \quad \blacksquare$$

Note that by definition, the mapping $T(\mathbf{u})$ has an inverse $T^{-1}(\boldsymbol{\theta}) := (G_1, G_2, \dots, G_{d_{\text{sto}}})$. Before going further, we show an example of mapping $T(\mathbf{u})$.

Example: If π is the uniform density supported in $\mathbf{X} = \bigotimes_{i=1}^{d_{\text{sto}}} [a_i, b_i]$, the i th cumulative distribution is $G_i(\theta_i) = \frac{\theta_i - a_i}{b_i - a_i}$ and $G_i^{-1}(u_i) = a_i + (b_i - a_i)u_i$ for $i = 1, 2, \dots, d_{\text{sto}}$. Therefore, the mapping $\mathbf{T}(\mathbf{u}) = (a_1 + (b_1 - a_1)u_1, a_2 + (b_2 - a_2)u_2, \dots, a_{d_{\text{sto}}} + (b_{d_{\text{sto}}} - a_{d_{\text{sto}}})u_{d_{\text{sto}}})$ is linear.

Remark: When π is a multivariate Gaussian with a mean $\boldsymbol{\mu}$ and covariance matrix C , then $\boldsymbol{\theta} = \boldsymbol{\mu} + C^{1/2} \cdot \boldsymbol{\zeta}$ where $\boldsymbol{\zeta} \sim N(\mathbf{0}, I)$ is a standard multivariate normal density. For the standard normal, $G_i(\zeta_i) = 1/2\text{erf}(\zeta_i/\sqrt{2})$, where erf is the error function. Therefore, the inverse mapping, \mathbf{T}^{-1} , and with that, \mathbf{T} , are nonlinear. In general, we expect a nonlinear mapping \mathbf{T} to deteriorate the accuracy of spatially adaptive sparse grid approximations. Nevertheless, in [57], we showed numerically that a nonlinear mapping leads to more accurate results than linear mappings and standard sparse grid approaches based on Gaussian points.

Applying Lemma 5.3.1 in (5.5), we obtain

$$c_p = \int_{[0,1]^{d_{\text{sto}}}} \mathcal{F}_h(T(\mathbf{u})) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(T_i(u_i)) \right) d\mathbf{u}. \quad (5.9)$$

To obtain our proposed approach for assessing PSP coefficients, we plug (5.7) into (5.9), i.e.,

$$\begin{aligned} c_p &= \int_{[0,1]^{d_{\text{sto}}}} \mathcal{F}_h(T(\mathbf{u})) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(T_i(u_i)) \right) d\mathbf{u} = \int_{[0,1]^{d_{\text{sto}}}} \mathcal{U}_{\mathcal{K}}^{\text{in},1}[\mathcal{F}_h \circ \mathbf{T}] \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(T_i(u_i)) \right) d\mathbf{u} \\ &= \int_{[0,1]^{d_{\text{sto}}}} \sum_{\mathbf{k} \in \mathcal{K}} \alpha_{\mathbf{k}} \varphi_{\mathbf{k}}(\mathbf{u}) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(T_i(u_i)) \right) du_1 du_2 \dots du_{d_{\text{sto}}} \\ &= \sum_{\mathbf{k} \in \mathcal{K}} \alpha_{\mathbf{k}} \int_{[0,1]^{d_{\text{sto}}}} \prod_{i=1}^{d_{\text{sto}}} \varphi_{k_i}(u_i) \left(\prod_{i=1}^{d_{\text{sto}}} \Phi_{p_i}(T_i(u_i)) \right) du_1 du_2 \dots du_{d_{\text{sto}}} \\ &= \sum_{\mathbf{k} \in \mathcal{K}} \alpha_{\mathbf{k}} \left(\prod_{i=1}^{d_{\text{sto}}} \int_{[0,1]} \varphi_{k_i}(u_i) \Phi_{p_i}(T_i(u_i)) du_i \right) = \sum_{\mathbf{k} \in \mathcal{K}} \alpha_{\mathbf{k}} B_{p,\mathbf{k}} = \langle \boldsymbol{\alpha}, \mathbf{B}_p \rangle, \end{aligned} \quad (5.10)$$

for $p = 1, 2, \dots, P^{\text{PSP}}$, where $B_{p,\mathbf{k}} := \prod_{i=1}^{d_{\text{sto}}} \int_{[0,1]} \varphi_{k_i}(u_i) \Phi_{p_i}(T_i(u_i)) du_i$, and $\boldsymbol{\alpha}$ and \mathbf{B}_p contain all hierarchical surpluses and integrals $B_{p,\mathbf{k}}$ associated with the multiindex set, \mathcal{K} .

Therefore, (5.10) transforms the computation of PSP coefficients in terms of integrals into an inner product between *surpluses* stemming from interpolation with hierarchical basis with local support and a *product of one-dimensional integrals* involving the interpolation and PSP bases. Moreover, from Lemma 5.3.1, we perform all computations w.r.t. the uniform measure on the unit hypercube $[0, 1]^{d_{\text{sto}}}$. If the mapping \mathbf{T} is *linear*, e.g., if the input density π is uniform

5.3. MULTILEVEL ADAPTIVE SPECTRAL PROJECTION WITH ONLINE DIMENSIONALITY REDUCTION




Figure 5.3: Visual summary of the proposed *Spatially adaptive sparse pseudo-spectral approximation* for uncertainty propagation. This approach is used to assess PSP coefficients via spatial adaptive interpolation formulated in terms of hierarchical bases and one-dimensional quadrature. The high-fidelity model is treated like a black-box. At each adaptive step, this model is evaluated for each sample θ added by the spatially adaptive algorithm. Since spatially adaptive sparse grids are defined on $[0, 1]^{d_{\text{sto}}}$, we additionally need a mapping \mathbf{T} with range \mathbf{X} (see Lemma 5.3.1). Note that the size of this outer-loop process is not known a priori, since we employ adaptivity. This approach is extended in Section 5.3 to multilevel hierarchy in which we additionally test, online, whether the stochastic dimensionality can be reduced.

in $\mathbf{X} = \bigotimes_{i=1}^{d_{\text{sto}}} [a_i, b_i]$, then \mathbf{B}_p can be evaluated *exactly*, up to the employed numerical precision since all integrands in $B_{\mathbf{k}} := \prod_{i=1}^{d_{\text{sto}}} \int_{[0,1]} \varphi_{k_i}(u_i) \Phi_{p_i}(T_i(u_i)) du_i$ are polynomials. In contrast, if \mathbf{T} is *nonlinear*, then $\Phi_{p_i}(T_i(u_i))$ is no longer a polynomial, thus $B_{p,\mathbf{k}}$ cannot be integrated exactly. However, since the integrals entering in $B_{p,\mathbf{k}}$ are 1D, they can be evaluated very accurately via, e.g., Gauss-Legendre quadrature. Therefore, regardless of the linearity of the integrands in $B_{p,\mathbf{k}}$, $B_{p,\mathbf{k}}$ can be evaluated very accurately and at a low computational cost. The computationally expensive part in (5.10) is evaluating the hierarchical coefficients in α . To reduce this cost and to exploit the *local* structure of the underlying high-fidelity model, \mathcal{F}_h , we employ spatial adaptivity, as summarized in Section 2.6.2. A intuitive illustration of this algorithm is provided in Figure 5.3. The proposed Spatially adaptive sparse pseudo-spectral approximation is an *outer loop* scenario of a priori unknown size, since we employ (spatial) adaptivity. The forward model $\mathcal{F}_h \circ \mathbf{T}$ is evaluated at each sparse grid point to find the hierarchical surpluses (5.14) which are then employed in (5.10) to assess PSP coefficients.

We summarize our proposed computational procedure to assess PSP coefficients in Algorithm 5.6. The inputs are the high-fidelity forward model, \mathcal{F}_h , and the transformation \mathbf{T} defined in (5.8), used to mapping the hypercube to the underlying stochastic domain, \mathbf{X} . The next input, K_{init} , is the initial grid level, before adaptivity is performed. As we discussed in Section 2.6, usually $K_{\text{init}} = 3$. The four input $r := (N_{\text{adapt}}; \text{ref_perc}; \text{ref_pol})$, comprises the parameters used for spatial adaptivity: N_{adapt} is the number of refinement steps to be performed and ref_perc represents the percentage of the grid points to be refined locally using the refinement policy ref_pol . The last input is the stochastic dimensionality w.r.t. which we perform the computations, d_{sto} . The two refinement policies considered in this work, summarized in Section 2.6.2 are the *maximum absolute value* of the surplus and *maximum volume* under the basis

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

functions. The algorithm begins by computing the hierarchical surplus, $\alpha_{\mathbf{k}}$, on the initial sparse grid of level K_{init} (step 2). The computationally expensive part is between steps 3 – 6, where we perform spatial adaptivity. After this step, we have the vector of hierarchical surpluses, $\boldsymbol{\alpha}$. Afterwards, we compute the products of 1D integrals $B_{p,\mathbf{k}} = \prod_{i=1}^{d_{\text{sto}}} \int_{[0,1]} \varphi_{k_i}(u_i) \Phi_{p_i}(T_i(u_i)) du_i$ at step 10. We save all products $B_{p,\mathbf{k}}$ into a vector \mathbf{B}_p and compute the PSP coefficients c_p as in (5.10) for $p = 1, 2, \dots, P^{\text{PSP}}$. At the end, the algorithm returns all PSP coefficients.

Algorithm 5.6 Spatially adaptive sparse pseudo-spectral approximation

```

1: procedure SPATIALLYADAPTIVESPARSEPSP( $\mathcal{F}_h, \mathbf{T}, K_{\text{init}}, r, d_{\text{sto}}$ )
2:   Compute  $\alpha_{\mathbf{k}}$  on the initial grid of level  $K_{\text{init}}$  using  $\mathcal{F}_h \circ \mathbf{T}$ 
3:   for  $n \leftarrow 1, 2, \dots, N_{\text{adapt}}$  do
4:     Refine, via spatial adaptivity, ref_perc of the current number of points using ref_pol
5:     Compute  $\alpha_{\mathbf{k}}$  using  $\mathcal{F}_h \circ \mathbf{T}$ 
6:   end for
7:   Append all hierarchical surpluses  $\alpha_{\mathbf{k}}$  to a vector  $\boldsymbol{\alpha}$ 
8:   for  $p \leftarrow 1, 2, \dots, P^{\text{PSP}}$  do
9:     for  $\mathbf{k} \in \mathcal{K}$  do
10:      Compute, via Gauss-Legendre quadrature,

```

$$B_{p,\mathbf{k}} = \prod_{i=1}^{d_{\text{sto}}} \int_{[0,1]} \varphi_{k_i}(u_i) \Phi_{p_i}(T_i(u_i)) du_i$$

```

11:    end for
12:    Append all  $B_{p,\mathbf{k}}$  to a vector  $\mathbf{B}_p$ 
13:    Compute the PSP coefficients  $c_p$  via (5.10)
14:  end for
15:  Append all  $c_p$  to a vector  $\mathbf{c}$ 
16:  return  $\mathbf{c}$ 
17: end procedure

```

The presented Spatially adaptive sparse pseudo-spectral approximation algorithm is based on hierarchies given by hierarchical sparse grid approximations. Next, we extend it to a multilevel version in which we additionally employ a grid-based model hierarchy (see Section 2.7). Additionally, we exploit the structure of the resulting multilevel decomposition to formulate a novel context-aware approach in which we perform *online* stochastic dimensionality reduction.

5.3.2 Multilevel spectral projections

First, we introduce the notation used in our proposed multilevel approach with online stochastic dimensionality reduction. The multilevel decomposition is standard and performed as summarized in Section 2.7. The novelty of our algorithm is the stochastic dimensionality reduction. Let again $2 \geq J \in \mathbb{N}$ denote the number of levels in multilevel decompositions and let $j = 1, 2, \dots, J$. Note that we use levels to characterize both sparse grid and multilevel decompositions. To avoid confusion, we will explicitly specify what is meant by level in each context.

Recall that standard multilevel decompositions are used to estimate linear quantities of interest. In our approach, multilevel decompositions are used to assess PSP coefficients, which we compute, at each level, via the proposed Spatially adaptive sparse pseudo-spectral approach. Note that either in the standard formula (5.3), based on integration, or in our proposed approach (5.10), the PSP coefficients are *linear quantities*, thus standard multilevel approaches based on

telescoping sums can be employed to assess them.

Remark: Standard multilevel methods are used to estimate linear quantities of interest, such as the expectation of the output of interest. In our multilevel approach, we estimate a linear quantity as well, the PSP coefficients. However, as discussed in Sections 2.5.1 and 2.9, PSP coefficients can be used to additionally estimate the standard deviation/variance of the output of interest and the Sobol' indices for sensitivity analysis. Therefore, we can easily employ the computed multilevel *linear* quantities to assess *nonlinear* quantities of interest in the underlying uncertainty propagation problem.

In multilevel decompositions, we need two discretization hierarchies. First, we employ a hierarchy $h_1 \leq h_2 \leq \dots \leq h_J$ for the discretization of the forward model. The second hierarchy is used for the sampling of the stochastic space, \mathbf{X} . For this task, we use spatially adaptive sparse grids. Let $r_1 \leq r_2 \leq \dots \leq r_J$ denote a sequence of parameters associated to spatial adaptivity; we specify what r_j means in our context in the next section. Since spatially adaptive sparse grids are nested, we employ the multilevel decomposition (2.35), in which, for the same discretization of the forward model, estimations from the previous sparse grid level can be used at the current level, i.e.,

$$c_{p,J} \approx \sum_{j=1}^J (c_{p,(J-j+1,j)} - c_{p,(J-j+1,j-1)}), \quad (5.11)$$

for $p = 1, 2, \dots, P^{\text{PSP}}$. To simplify the notation, we use the subscript $\ell(j^c, j)$ to refer to the combination $(J-j+1, j)$, where by j^c we refer to the *complement of j w.r.t. J* , i.e., $j^c = J-j+1$, and the subscript $\ell(j^c, j-1)$ to refer to $(J-j+1, j-1)$. Since the results computed using $\ell(j^c, j-1)$ are a subset of the results computed via $\ell(j^c, j)$, we use $\ell(j^c, j)$ to refer to “level j ” in our multilevel approach. Moreover, we employ the notation $\ell(\delta j)$ to refer to differences between results obtained with $\ell(j^c, j)$ and $\ell(j^c, j-1)$. With this notation, the multilevel decomposition of the PSP coefficients (5.11) becomes

$$c_{p,J} \approx \sum_{j=1}^J (c_{p,(J-j+1,j)} - c_{p,(J-j+1,j-1)}) = \sum_{j=1}^J c_{p,\ell(j^c,j)} - c_{p,\ell(j^c,j-1)} = \sum_{j=1}^J c_{p,\ell(\delta j)}, \quad (5.12)$$

for $p = 1, 2, \dots, P^{\text{PSP}}$. The nestedness of spatially adaptive sparse grids imply that to evaluate $c_{p,\ell(\delta j)}$ in (5.12), we need only evaluations of \mathcal{F}_{J-j+1} w.r.t. r_j , that is, the computationally expensive part is the evaluation of $c_{j^c,j}$, since $c_{j^c,j-1}$ is obtained from a subset of evaluations associated with $c_{j^c,j}$. Having presented the multilevel decomposition used to assess PSP coefficients, we hereby present our proposed approach for online stochastic dimensionality reduction.

5.3.3 Online stochastic dimensionality reduction

To the best of our knowledge, online stochastic dimensionality reduction has not been considered in previous multilevel formulations. Let $\mathcal{J} = \{c_{p,\ell(1^c,1)}, c_{p,\ell(2^c,2)}, \dots, c_{p,\ell(j^c,j)}\}$ denote the set of J PSP coefficients in (5.12); the remaining $J-1$ coefficients $\{c_{p,\ell(2^c,1)}, c_{p,\ell(3^c,2)}, \dots, c_{p,\ell(j^c,j-1)}\}$ are computed from subsets of model evaluations associated to the coefficients in \mathcal{J} . The PSP coefficients in \mathcal{J} can be assessed independently from each other and thus a subset of these coefficients can be used to learn properties of the underlying stochastic model, such as intrinsic stochastic dimensionality. Thus, to formulate our online stochastic dimensionality reduction approach, for a given $j \geq 1$ and an integer $I \in \mathbb{N}$, $I < J$, we consider a subset $\mathcal{I} = \{c_{p,\ell(j^c,j)}, c_{p,\ell((j+1)^c,j+1)}, \dots, c_{p,\ell((j+I-1)^c,j+I-1)}\} \subset \mathcal{J}$ containing I PSP coefficients from \mathcal{J} . Note that the PSP coefficients in \mathcal{I} can be computed *independently* from each other. Finally, let

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

$\boldsymbol{\tau}^{\text{in},1} := (\tau_1^{\text{in},1}, \tau_2^{\text{in},1}, \dots, \tau_{d_{\text{sto}}}^{\text{in},1}) \in (0, 1]^{d_{\text{sto}}}$ denote user-defined *thresholds for stochastic importance* associated to each of the d_{sto} stochastic inputs.

Initially, the stochastic dimension is the original one, i.e., d_{sto} . We first compute the coefficients from \mathcal{I} depending on d_{sto} -dimensional spatially adaptive sparse grids. Afterwards, we obtain the multilevel approximation of $c_{p,I}$,

$$c_{p,I} = \sum_{n=j}^{j+I-1} (c_{p,\ell(n^c,n)} - c_{p,\ell(n^c,n-1)}) = \sum_{n=j}^{j+I-1} c_{p,\ell(\delta n)}, \quad p = 1, 2, \dots, P^{\text{PSP}}. \quad (5.13)$$

We use the PSP coefficients $c_{p,I}$ in (2.47) to estimate the total Sobol' indices corresponding to each stochastic input and obtain a set

$$\mathcal{S} = \{\hat{S}_1^T, \hat{S}_2^T, \dots, \hat{S}_{d_{\text{sto}}}^T\}.$$

Note that up to this point, the stochastic dimensionality is still d_{sto} . Next, we analyse whether we can reduce the stochastic dimensionality. That is, we compare each $\hat{S}_i^T \in \mathcal{S}$ to the user defined threshold $\tau_i^{\text{in},1}$; the choice of $\boldsymbol{\tau}^{\text{in},1}$ is heuristic and can depend on the number of uncertain inputs and desired results accuracy. If $\hat{S}_i^T < \tau_i^{\text{in},1}$, we consider the i th input as stochastically unimportant and we ignore it. By ignoring an uncertain input, we mean that we no longer consider it uncertain, but rather *deterministic* or *known with certainty*. In this work, the deterministic value associated to an uncertain input θ_i is its expectation, $\hat{\theta}_i := \mathbb{E}[\theta_i]$. However, as discussed in Section 5.2, another deterministic value can be considered as well. We denote the number of unimportant stochastic inputs by $d_{\text{sto}}^{\text{unimp}}$ with $0 \leq d_{\text{sto}}^{\text{unimp}} \leq d_{\text{sto}} - 1$. In contrast, when $\hat{S}_i^T \geq \tau_i^{\text{in},1}$, the i th uncertain input remains stochastic.

In the remainder of the algorithm, we assess the PSP coefficients from $\mathcal{J} \setminus \mathcal{I}$. On the one hand, if $d_{\text{sto}}^{\text{unimp}} < d_{\text{sto}}$ uncertain inputs are ignored, we compute the coefficients from $\mathcal{J} \setminus \mathcal{I}$ using a stochastic grid of dimensionality $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$. Since the coefficients from \mathcal{J} and $\mathcal{J} \setminus \mathcal{I}$ are computed using different stochastic dimensionality, we need to transplant the results from the original grid onto the reduced grid. We compute the multilevel PSP coefficients via (5.12) using the grid of dimensionality $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$. On the other hand, if no uncertain input is ignored, we calculate the coefficients from $\mathcal{J} \setminus \mathcal{I}$ using the original stochastic grid and, afterwards, compute the multilevel PS coefficients using (5.12). We provide a visual illustration of the proposed approach for online stochastic dimensionality reduction in the following example.

Example: Recall that once the stochastic dimensionality is reduced, the results at level $\ell(2, 2)$ are transplanted to a 2D stochastic space; this step is not depicted in Figure 5.4. Let $J = 3$ and assume that $d_{\text{sto}} = 3$. The complementary levels are $1^c = 3, 2^c = 2$ and $3^c = 1$ and the set of PSP coefficients at the three levels is $\mathcal{J} := \{c_{p,\ell(3,1)}, c_{p,\ell(2,2)}, c_{p,\ell(1,3)}\}$ for $p = 1, 2, \dots, P^{\text{PSP}}$. We take $\mathcal{I} := \{c_{p,\ell(2,2)}\}$, i.e., we employ the PSP coefficients at level $\ell(2, 2)$ to compute the total Sobol' indices, which are then used to ascertain the importance of the three stochastic inputs. Let us assume that one uncertain input is rendered unimportant. Thus, the dimensionality is reduced from three to two and the PSP coefficients at levels $\ell(1, 3)$ and $\ell(3, 1)$ are computed using 2D stochastic domains. In Figure 5.4, we illustrate the hierarchy of model discretizations and sparse grids used to sample the stochastic domain, based on the aforementioned setup. Note that we depict the coefficients $\{c_{p,\ell(2,1)}, c_{p,\ell(1,2)}\}$ as well, which are computed from a subset of the model evaluations used to assess $\{c_{p,\ell(3,1)}, c_{p,\ell(1,3)}\}$, due to the nestedness of sparse grids.

To transplant the results from the stochastic grid with dimensionality d_{sto} onto a grid with dimensionality $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$, we simply evaluate the d_{sto} -dimensional spatially adaptive interpolant at the $d_{\text{sto}}^{\text{unimp}}$ deterministic inputs. Let ν be a permutation of indices $i = 1, 2, \dots, d_{\text{sto}}$




Figure 5.4: Visual illustration of the proposed context-aware multilevel approach with online stochastic dimensionality reduction. We used $J = 3$, $d_{\text{sto}} = 3$ and $\mathcal{I} := \{c_{p, \ell(2, 2)}\}$. The PSP coefficients at level $\ell(2, 2)$ are computed using the initial stochastic dimensionality, three. The dimensionality is then reduced to two, thus the PSP coefficients at levels $\ell(1, 3)$ and $\ell(3, 1)$ are computed using 2D spatially adaptive sparse grids.

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

such that the inputs indexed from $\nu(d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}} + 1)$ to $\nu(d_{\text{sto}})$ are rendered stochastically unimportant; the remaining stochastic parameters $\theta_{\nu(1)}, \theta_{\nu(2)}, \dots, \theta_{\nu(d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}})}$ are therefore important. The PSP coefficients (5.13) transplanted on the $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$ dimensional space are

$$c_{p,\ell(n^c,n)} = \sum_{\mathbf{k} \in \mathcal{K}} \bar{\alpha}_{\mathbf{k}} \left(\prod_{i=1}^{d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}} \int_{[0,1]} \varphi_{k_{\nu(i)}}(u_{\nu(i)}) \Phi_{p_{\nu(i)}}(T_{\nu(i)}(u_{\nu(i)})) du_{\nu(i)} \right),$$

where

$$\bar{\alpha}_{\mathbf{k}} = \alpha_{\mathbf{k}} \cdot \prod_{i=d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}} + 1}^{d_{\text{sto}}} \varphi_{k_{\nu(i)}}(\hat{u}_{\nu(i)}), \quad \hat{u}_{\nu(i)} = T_{\nu(i)}^{-1}(\hat{\theta}_i). \quad (5.14)$$

Therefore, to evaluate $\bar{\alpha}_{\mathbf{k}}$ we simply multiply the surpluses $\alpha_{\mathbf{k}}$ computed on the original sparse grid of dimension d_{sto} with the product of hierarchical basis functions evaluated at the expectation of the unimportant stochastic inputs. Since these expectations need to reside in $[0, 1]$, we use the inverse mapping to transport them from X_i to the unit interval. Hence, denoting

$$\bar{B}_{p,\mathbf{k}} := \prod_{i=1}^{d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}} \int_{[0,1]} \varphi_{k_{\nu(i)}}(u_{\nu(i)}) \Phi_{p_{\nu(i)}}(T_{\nu(i)}(u_{\nu(i)})) du_{\nu(i)},$$

the mapped PSP coefficients read

$$c_{p,\ell(n^c,n)} = \sum_{\mathbf{k} \in \mathcal{K}} \bar{\alpha}_{\mathbf{k}} \bar{B}_{p,\mathbf{k}} = \langle \bar{\alpha}, \bar{B}_p \rangle. \quad (5.15)$$

To avoid unnecessarily complicating the notation, we denote the transplanted PSP coefficients as the ones computed using the original dimensionality, i.e., $c_{p,\ell(n^c,n)}$. In essence, if the dimensionality is reduced, the mapped coefficients are used in the multilevel decomposition (5.3).

We summarize the proposed Multilevel adaptive spectral projection with online dimensionality reduction procedure in Algorithm 5.7. The inputs are the number of levels J used in multilevel approximations, the hierarchy of forward model discretizations $\mathbf{h} := (h_1, h_2, \dots, h_J)$, K_{init} , the hierarchy of parameters used in spatial adaptivity, $\mathbf{r} := (r_1, r_2, \dots, r_J)$, and the forward operator, \mathcal{F}_h . Moreover, we also have the original stochastic dimensionality, d_{sto} , the two integers j and I with which we determine size of the subset of PSP coefficients to be computed w.r.t. the original stochastic dimensionality, and the last input parameter, the thresholds for stochastic importance, $\boldsymbol{\tau}^{\text{in},1} := (\tau_1^{\text{in},1}, \tau_2^{\text{in},1}, \dots, \tau_{d_{\text{sto}}}^{\text{in},1})$. Between steps 2 – 8, we assess the multilevel PSP coefficients $c_{p,I}$ for $p = 1, 2, \dots, P^{\text{PSP}}$ using (5.10), (5.13) w.r.t. the uniform measure on $[0, 1]^{d_{\text{sto}}}$. Next, at steps 10 – 15, we assess the d_{sto} total Sobol' indices using (2.47) and ascertain the number of unimportant stochastic inputs via the thresholds $\boldsymbol{\tau}^{\text{in},1}$. We count the number of unimportant inputs with the integer $d_{\text{sto}}^{\text{unimp}}$. If at least one stochastic input is rendered unimportant, we transplant the I previously computed PSP coefficients on a sparse grid of dimensionality $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$ via (5.15) (step 18). Next, we compute the remaining $J - I$ PSP coefficients at step 23. Having now the PSP coefficients at all levels, we finally compute their multilevel approximations $c_{p,J}$ via (5.12) for $p = 1, 2, \dots, P^{\text{PSP}}$. We save all multilevel PSP coefficients into a vector \mathbf{c}_J , which is returned before the algorithm terminates. Once the coefficients are assessed, we can compute quantities of interest such as expectation, standard deviation or Sobol' indices of the output of interest (see Sections 2.5.1 and 2.9).

Before ending this section, we note that although in our formulation we analyse only once whether the stochastic dimensionality can be reduced, this can be easily generalized. For instance, we could consider a sequence $\mathcal{I}_1, \mathcal{I}_2, \dots \subset \mathcal{J}$, possibly such that $\mathcal{I}_v \subseteq \mathcal{I}_{v+1}, v = 1, 2, \dots$

5.3. MULTILEVEL ADAPTIVE SPECTRAL PROJECTION WITH ONLINE DIMENSIONALITY REDUCTION

Algorithm 5.7 Multilevel adaptive spectral projection with online dimensionality reduction

```

1: procedure MLSPONLINEDIMRED( $J, \mathbf{h}, K_{\text{init}}, \mathbf{r}, \mathcal{F}_h, \mathbf{T}, d_{\text{sto}}, j, I, \tau^{\text{in},1}$ )
2:   for  $n \leftarrow j + I - 1, j + I, \dots, I$  do
3:     for  $p \leftarrow 1, 2, \dots, P^{\text{PSP}}$  do
4:        $c_{p,\ell(n^c,n)} := \text{SpatiallyAdaptiveSparsePSP}(\mathcal{F}_{h_{J-n+1}}, \mathbf{T}, K_{\text{init}}, r_n, d_{\text{sto}})$ 
5:     end for
6:     Compute  $c_{p,\ell(n^c,n-1)}$  w.r.t.  $r_{n-1}$  from the evaluations of  $\mathcal{F}_{h_{J-n+1}}$  w.r.t.  $r_n$ 
7:   end for
8:   Compute the multilevel coefficients  $c_{p,I}$  via (5.13)  $\triangleright$  the underlying dimensionality is  $d_{\text{sto}}$ 
9:    $d_{\text{sto}}^{\text{unimp}} := 0$ 
10:  for  $i \leftarrow 1, 2, \dots, d_{\text{sto}}$  do
11:    Compute the total Sobol' indices  $\hat{S}_i^T$  from the PSP coefficients  $c_{p,I}$  via (2.47)
12:    if  $\hat{S}_i^T < \tau_i^{\text{in},1}$  then
13:       $d_{\text{sto}}^{\text{unimp}} = d_{\text{sto}}^{\text{unimp}} + 1$ 
14:    end if
15:  end for
16:  if  $d_{\text{sto}}^{\text{unimp}} \geq 1$  then  $\triangleright$  the underlying dimensionality is  $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$ 
17:    Find  $\nu$  such that the unimportant inputs are from  $\nu(d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}} + 1)$  to  $\nu(d_{\text{sto}})$ 
18:    Map the PSP coefficients to the grid of dimensionality  $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$ 

```

$$\bar{\alpha}_{\mathbf{k}} = \alpha_{\mathbf{k}} \cdot \prod_{i=d_{\text{sto}}-d_{\text{sto}}^{\text{unimp}}+1}^{d_{\text{sto}}} \varphi_{k_{\nu(i)}}(\hat{u}_{\nu(i)}), \quad \hat{u}_{\nu(i)} = T_{\nu(i)}^{-1}(\hat{\theta}_i)$$

$$c_{p,\ell(n^c,n)} = \sum_{\mathbf{k} \in \mathcal{K}} \bar{\alpha}_{\mathbf{k}} \bar{B}_{p,\mathbf{k}}$$

```

19:     Analogously, map  $c_{p,\ell(n^c,n-1)}$  to the grid of dimensionality  $d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}}$ 
20:   end if
21:  for  $n \leftarrow \{1, 2, \dots, J\} \setminus \{j + I - 1, j + I, \dots, I\}$  do
22:    for  $p \leftarrow 1, 2, \dots, P^{\text{PSP}}$  do
23:       $c_{p,\ell(n^c,n)} := \text{SpatiallyAdaptiveSparsePSP}(\mathcal{F}_{h_{J-n+1}}, \mathbf{T}, K_{\text{init}}, r_n, d_{\text{sto}} - d_{\text{sto}}^{\text{unimp}})$ 
24:    end for
25:    Compute  $c_{p,\ell(n^c,n-1)}$  w.r.t.  $r_{n-1}$  from the evaluations of  $\mathcal{F}_{h_{J-n+1}}$  w.r.t.  $r_n$ 
26:  end for
27:  for  $p \leftarrow 1, 2, \dots, P^{\text{PSP}}$  do
28:    Compute, using (5.12), the multilevel PSP coefficients

```

$$c_{p,J} \approx \sum_{j=1}^J c_{p,\ell(j^c,j)} - c_{p,\ell(j^c,j-1)} = \sum_{j=1}^J c_{p,\ell(\delta_j)}$$

```

29:  end for
30:  Append the PSP coefficients  $c_{p,J}$  to a vector  $\mathbf{c}_J$ 
31:  return  $\mathbf{c}_J$ 
32: end procedure

```

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

After we compute the PSP coefficients from \mathcal{J}_v , we analyse whether the stochastic dimensionality could be reduced. This approach with an arbitrary number of checks for stochastic dimensionality reduction is suitable, e.g., when the initial number of uncertain parameters is large or when some total Sobol' indices are close to the associated thresholds.

5.3.4 Computational cost

Throughout this work, the forward model is assumed to be computationally expensive. Thus, the most cost intensive operation in our proposed approach is the computation of the hierarchical surpluses (5.14) at each level $\ell(1^c, 1), \ell(2^c, 2), \dots, \ell(J^c, J)$; recall that due to the nestedness of hierarchical sparse grids, the surpluses at levels $\ell(2^c, 1), \ell(3^c, 2), \dots, \ell(J^c, J-1)$ are found from subsets of model evaluations used to find the surpluses at the aforementioned J levels. Assume that one evaluation of the forward operator is performed on P processes, where $P = 1$ if the simulation is performed serially. To this end, for $j = 1, 2, \dots, J$, we denote by \mathcal{C}_{h_j} the cost of one evaluation of $\mathcal{F}_{h_j}(\boldsymbol{\theta})$ on P processes, where $\mathcal{F}_{h_j}(\boldsymbol{\theta})$ denotes the discretization of the forward operator depending on h_j . Moreover, since in our proposed approach we check whether the stochastic dimensionality can be reduced, let d_j denote the dimensionality at level $\ell(j^c, j)$.

Therefore, we denote by $N_{J-j+1}^{\text{in}, d_j}$ the total number of evaluations of $\mathcal{F}_{h_j}(\boldsymbol{\theta})$ needed to find the surpluses at level $\ell(j^c, j)$ using a spatially adaptive sparse grid of dimensionality d_j depending on the refinement parameter r_{J-j+1} . The computational cost $\mathcal{C}_J^{\text{in}, 1}$ of the proposed approach when using J levels and P processes amounts to

$$\mathcal{C}_J^{\text{in}, 1} = \sum_{j=1}^J N_{J-j+1}^{\text{in}, d_j} \mathcal{C}_{h_j}. \quad (5.16)$$

In the worst case scenario in which the stochastic dimensionality cannot be reduced, that is, when $d_1 = d_2 = \dots = d_J = d_{\text{sto}}$, $\mathcal{C}_J^{\text{in}, 1}$ represents the cost of the standard multilevel approach. In contrast, when stochastic dimensionality reduction is possible, the cost of our proposed approach will be smaller than the cost of an equivalent standard multilevel approach. The remaining costs, such as computing the products of integrals in (5.10) or transplanting the PSP coefficients to the lower dimensional sparse grid (5.13) when the stochastic dimensionality is reduced, are negligible.

5.3.5 Termination of the multilevel algorithm

In general, standard multilevel decompositions terminate when a prescribed accuracy, ε , is attained by the MSE of the estimator. This accuracy, in general, is split equally between the bias term, due to approximating the given continuous model, \mathcal{F}_h , with \mathcal{F}_{h_j} , and an error term due to using sampling to approximate the quantity of interest. However, this framework requires a detailed theoretical knowledge of the underlying problem and a subsequent mathematical derivation of the bias and sampling terms, which is not readily available for complex problems such as FSI simulations. In our proposed approach, we have more flexibility since we employ spatial adaptivity to compute the quantity of interest. In addition, since the main focus in our proposed approach is on stochastic dimensionality reduction in the quantification of uncertainty in computationally expensive problems, we terminate our proposed approach by imposing a maximum number of levels, J . J should be chosen having in mind a worst-case scenario computational cost. The worst-case scenario in our approach happens when no dimensionality reduction is performed at any level, i.e., with the notation from the previous section, $d_1 = d_2 = \dots = d_J = d_{\text{sto}}$. Note, however, that the cost $\mathcal{C}_J^{\text{in}, 1}$ in (5.16) cannot be, in general,

estimated a priori since we do not know the values of $N_{J-j+1}^{\text{in},d_j}$, which are known only at the end of the adaptive process. Thus, a worst-case scenario cost would be a theoretically derived *upper bound* of $C_J^{\text{in},1}$ in which a static sparse grid is chosen at each level. In general, we expect both the spatial adaptivity and the dimensionality reduction to yield significant cost reductions as compared to the worst-case scenario. In our numerical results, presented next, choosing $J = 3$ was sufficient to estimate the quantities of interest accurately and to obtain a detailed overview of the proposed approach.

5.4 Numerical results

We apply the proposed context-aware multilevel spectral projection with dimensionality reduction in two test cases, in which, at each level, we employ the Spatially adaptive sparse pseudo-spectral approach to assess the PSP coefficients. In addition, for a more comprehensive overview, we consider two stochastic scenarios in each test case. We compare our approach with an equivalent multilevel decomposition (2.35) with no dimensionality reduction in terms of *accuracy* (relative error w.r.t. a reference solution) and *cost* (total number of grid points and runtime). In both test cases, the uncertain parameters are modelled as independent uniform random variables in an interval $[a_i, b_i]$ for $i = 1, 2, \dots, d_{\text{sto}}$. Therefore, the mapping \mathbf{T} is linear:

$$\mathbf{T}(\mathbf{u}) = (G_1^{-1}(u_1), G_2^{-1}(u_2), \dots, G_{d_{\text{sto}}}^{-1}(u_{d_{\text{sto}}})) , \quad G_i^{-1}(u_i) = a_i + (b_i - a_i)u_i.$$

Moreover, we employ hierarchical polynomial bases of second degree in all sparse grid computations. The functionality for spatially adaptive sparse grid approximations was provided by the open-source library **SG++** [147]. All numerical experiments were performed using standard double precision arithmetics.

In Section 5.4.1, we consider a damped oscillator, modelled as a second order ODE system. In the first stochastic scenario, we have five stochastic inputs and in the second, we consider an extra parameter to end up with a total of six uncertain inputs. The uniform bounds of the six uncertain parameters are the same as in [134, Section 5.1] and they ensure an underdamped regime. In the Spatially adaptive sparse pseudo-spectral approach, we employ the two adaptive refinement criteria summarized in Section 2.6: the maximum absolute value of the surplus and the maximum volume contained by the hierarchical basis functions. We consider three levels in the proposed multilevel approach. For a better understanding of the effects of stochastic dimensionality reduction, we first test whether the dimensionality can be reduced using the Sobol' indices computed at one level. Afterwards, we test whether dimensionality reduction is possible by employing PSP coefficients computed at two levels. These numerical experiments were performed on a desktop computer with a four core Intel i7-4790 CPU and 24 GB of RAM.

In Section 5.4.2, we employ our context-aware multilevel approach in a multi-physics problem, a 2D flow over an elastic vertical structure, in which we consider five uncertain inputs: two characterize the fluid flow and the remaining three parameters characterize the elastic structure. At each level, we employ spatially adaptive refinement based on the maximum absolute value of the surpluses. For a broader overview, we employ the proposed multilevel approach with online dimensionality reduction to quantify uncertainty at two instances in the time domain. These simulations were performed on a Intel SandyBridge-EP Xeon E5-2670 processor and 128 GB of RAM on the CoolMAC cluster².

²http://www.mac.tum.de/wiki/index.php/MAC_Cluster

5.4.1 Damped oscillator

The first considered test case is a damped oscillator subject to external forces, modelled as a second-order ODE system

$$\begin{cases} \frac{d^2y}{dt^2}(t) + c\frac{dy}{dt}(t) + ky(t) = f \cos(\omega t) \\ y(0) = y_0 \\ \frac{dy}{dt}(0) = y_1, \end{cases} \quad (5.17)$$

where $c[N \cdot \text{sec}/m]$ is the damping coefficient, $k[N/m]$ the spring constant, $f[N]$ the forcing amplitude, and $\omega[\text{rad}/\text{sec}]$ is the frequency. Furthermore, $y_0[m]$ represents the initial position, whereas $y_1[m/s]$ is the initial velocity. Throughout our simulations, $t \in [0, 20]$. The solution of interest in (5.17) is the displacement, $y(t)$. To obtain it numerically, i.e., $y_{\delta t}(t) \approx y(t)$, we discretize (5.17) via Adams predictor-corrector methods from the `scipy.integrate`³ package. In our experiments, the output of interest is $y_{\delta t}(t = t_{\text{fixed}})$ where

$$t_{\text{fixed}} = 10.$$

To gain a broader overview of our approach, we consider two stochastic scenarios. In the first one, we consider five stochastic inputs $\theta_1 = (c, k, f, y_0, y_1)$ and fix the frequency $\omega = 1.05$. In the second scenario, we consider the frequency to be uncertain as well and thus have a total of six stochastic parameters $\theta_2 = (c, k, f, \omega, y_0, y_1)$. In Table 5.1, we list the left and right bounds of the uncertain parameters (columns 4 and 5) as well as their deterministic values, i.e., their expectation, in column 6. The uniform bounds are taken from [134, Section 5.1] and they ensure an underdamped regime. Note that all deterministic values reside in $[0, 1]$, therefore it is not needed to map them via \mathbf{T}^{-1} (cf. Section 5.3.3).

$\theta_{1,2}$	parameter name	symbol	left bound	right bound	$\hat{\theta} := \mathbb{E}[\theta]$
θ_1	damping coefficient	$c[N \cdot \text{sec}/m]$	0.080	0.120	0.100
θ_2	spring constant	$k[N/m]$	0.003	0.004	0.035
θ_3	forcing amplitude	$f[N]$	0.008	0.120	0.100
θ_4	frequency	$\omega[\text{rad}/\text{sec}]$	0.800	1.200	1.000
θ_5	initial position	$y_0[m]$	0.450	0.550	0.500
θ_6	initial velocity	$y_1[m/s]$	-0.050	0.050	0.000

Table 5.1: Second-order oscillator test case: uniform bounds for θ_1 (comprising (c, k, f, y_0, y_1) ; the frequency which is fixed to $w := 1.05$) and θ_2 (all six parameters). In the last column, we list the deterministic values of all parameters, i.e., their expectations.

We employ three levels in our proposed context-aware multilevel spectral projection with dimensionality reduction. Therefore, $J = 3$, which means that the complementary levels are $1^c = 3, 2^c = 2$ and $3^c = 1$. Recall that to obtain the displacement at each level, we employ Adams predictor-corrector time integrators, which yield the semidiscrete displacement parametrized in terms of the stochastic parameters. Afterwards, we employ sparse grid approximations with hierarchical bases to approximate the stochastically parametrized displacement to obtain the quantities of interest. We list the multilevel setup in Table 5.2. The employed time steps are shown in the second column. The spatially adaptive sparse grid approximations are characterized by a list of three parameters $r_j = (N_{\text{adapt},j}; \text{ref_perc}; \text{ref_pol})$, in which we additionally need the initial level, K_{init} . We employ

$$K_{\text{init}} = 3, \quad N_{\text{adapt},j} := j - 1, \quad \text{ref_perc} := 20\%,$$

³<https://docs.scipy.org/doc/scipy-0.18.1/reference/integrate.html>

Table 5.2: Multilevel setup for the damped oscillator model (5.17).

Level	h	r
$\ell(1, 3)$	$h_1 = \delta t_1 = 2.0 \cdot 10^{-3}$	$r_3 = (N_{\text{adapt},3}, \text{ref_perc}, \text{ref_pol})$
$\ell(2, 2)$	$h_2 = \delta t_2 = 5.0 \cdot 10^{-4}$	$r_2 = (N_{\text{adapt},2}, \text{ref_perc}, \text{ref_pol})$
$\ell(3, 1)$	$h_3 = \delta t_3 = 2.5 \cdot 10^{-5}$	$r_1 = (N_{\text{adapt},1}, \text{ref_perc}, \text{ref_pol})$

that is, we use the same initial sparse grid level at all levels, which also the sparse grid level at $\ell(3, 1)$ we refine once at level $\ell(2, 2)$ and twice at level $\ell(1, 3)$, each time 20% of the grid points having the largest refinement indicators. The refinement policy, `ref_pol`, is either based on the maximum absolute value of the surpluses (*Max*) or on the maximum volume (*Vol*); see Section 2.6.2. For comparison, we also employ non-adaptive hierarchical sparse grids with

$$K_{\max,1}^{\text{in},l} := K_{\text{init}}, \quad K_{\max,2}^{\text{in},l} := K_{\max,1}^{\text{in},l} + 1, \quad K_{\max,3}^{\text{in},l} := K_{\max,2}^{\text{in},l} + 1.$$

At each level, we employ a PSP expansion with total degree multiindex set

$$\mathcal{P} = \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}^j} : |\mathbf{p}|_1 \leq 4\},$$

where d_{sto}^j is the underlying stochastic dimensionality. Finally, the quantities of interest are the expectation and standard deviation of the output of interest. To ascertain the accuracy of the estimated expectations and standard deviations with all considered multilevel approaches, denoted generically by \hat{Q} , we assess the relative error of \hat{Q} w.r.t. a reference value Q_{ref} , i.e.,

$$\mathcal{E}_{\text{rel}}(Q_{\text{ref}} - \hat{Q}) := |1 - \hat{Q}/Q_{\text{ref}}|.$$

Since $J = 3$, the set of PSP coefficients that require forward model evaluations is

$$\mathcal{J} := \{c_{p,\ell(3,1)}, c_{p,\ell(2,2)}, c_{p,\ell(1,3)}\}.$$

Recall that hierarchical sparse grids are nested, thus the remaining coefficients, $\{c_{p,\ell(2,1)}, c_{p,\ell(1,2)}\}$, are assessed from subsets of forward model evaluations associated to $\{c_{p,\ell(2,2)}, c_{p,\ell(1,3)}\}$. To analyse the possibility for dimensionality reduction, we consider two choices for $\mathcal{I} \subset \mathcal{J}$. In the first choice, $I = 1$ and $\mathcal{I}_1 = \{c_{p,\ell(2,2)}\}$. We choose the “middle” level because it comprises approximations of intermediate accuracy, thus avoiding potentially inaccurate results. In the second choice, we additionally consider level $\ell(1, 3)$, i.e., $I = 2$ and $\mathcal{I}_2 = \{c_{p,\ell(2,2)}, c_{p,\ell(1,3)}\}$. Moreover, the threshold for stochastic importance is

$$\tau^{\text{in},l} := 0.05 \cdot \mathbf{1}. \tag{5.18}$$

We first present our results for the scenario with five stochastic parameters. Initially, we compute reference quantities of interest using the finest time discretization and a full Gauss-Legendre grid comprising $8^5 = 32768$ points, which gives

$$\mathbb{E}_{5D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.0195, \quad \text{Std}_{5D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.1014.$$

Before analyzing the possibility of dimensionality reduction, we employ the multilevel algorithm using $d_{\text{sto}} = 5$ at all levels, in which to assess the PSP coefficients at each level, we employ either spatially adaptive or static sparse grids, with the setup presented above. In Table 5.3, we present the results and the associated relative errors. To simplify the notation, we use *SML* to refer to

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

Table 5.3: Expectation and standard deviation estimation of the displacement at $t_{\text{fixed}} = 10$ when the stochastic input is θ_1 . Here, we do not perform dimensionality reduction.

Strategy in 5D	$\mathbb{E}_{5D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\mathbb{E}_{5D}^{\text{ref}}[\mathcal{F}_{h_3}] - \mathbb{E}_{5D}^{\text{ML}}[\mathcal{F}_{h_3}])$	$\text{Std}_{5D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\text{Std}_{5D}^{\text{ref}}[\mathcal{F}_{h_3}] - \text{Std}_{5D}^{\text{ML}}[\mathcal{F}_{h_3}])$
SML	0.0195	$1.4000 \cdot 10^{-5}$	0.1014	$1.0177 \cdot 10^{-5}$
AML Max	0.0195	$1.3443 \cdot 10^{-4}$	0.1014	$1.3166 \cdot 10^{-5}$
AML Vol	0.0195	$1.6322 \cdot 10^{-4}$	0.1014	$1.5541 \cdot 10^{-5}$

Table 5.4: Expectation estimation of the displacement at $t_{\text{fixed}} = 10$ when the stochastic input is θ_1 . We analyse the possibility of stochastic dimensionality reduction using the PSP coefficients at level $\ell(2, 2)$, i.e., from \mathcal{I}_1 (first three rows) and the PSP coefficient from \mathcal{I}_2 , that is, at levels $\ell(2, 2)$ and $\ell(1, 3)$ (last three rows).

Strategy with dim. red.	$\mathbb{E}_{2D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\mathbb{E}_{5D}^{\text{ref}}[\mathcal{F}_{h_3}] - \mathbb{E}_{2D}^{\text{ML}}[\mathcal{F}_{h_3}])$	$\mathcal{E}_{\text{rel}}(\mathbb{E}_{2D}^{\text{ref}}[\mathcal{F}_{h_3}] - \mathbb{E}_{2D}^{\text{ML}}[\mathcal{F}_{h_3}])$
SML + \mathcal{I}_1	0.0198	$1.6921 \cdot 10^{-2}$	$1.9472 \cdot 10^{-5}$
AML Max + \mathcal{I}_1	0.0198	$1.6422 \cdot 10^{-2}$	$5.0981 \cdot 10^{-4}$
AML Vol + \mathcal{I}_1	0.0198	$1.6422 \cdot 10^{-2}$	$5.0981 \cdot 10^{-4}$
SML + \mathcal{I}_2	0.0198	$1.7153 \cdot 10^{-2}$	$2.0453 \cdot 10^{-4}$
AML Max + \mathcal{I}_2	0.0198	$1.5966 \cdot 10^{-2}$	$1.0693 \cdot 10^{-3}$
AML Vol + \mathcal{I}_2	0.0198	$1.5833 \cdot 10^{-2}$	$1.0875 \cdot 10^{-3}$

Table 5.5: Standard deviation estimation of the displacement at $t_{\text{fixed}} = 10$ when the stochastic input is θ_1 . We analyse the possibility of stochastic dimensionality reduction using the PSP coefficients at level $\ell(2, 2)$, i.e., from \mathcal{I}_1 (first three rows) and the PSP coefficient from \mathcal{I}_2 , that is, at levels $\ell(2, 2)$ and $\ell(1, 3)$ (last three rows).

Strategy with dim. red.	$\text{Std}_{2D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\text{Std}_{5D}^{\text{ref}}[\mathcal{F}_{h_3}] - \text{Std}_{2D}^{\text{ML}}[\mathcal{F}_{h_3}])$	$\mathcal{E}_{\text{rel}}(\text{Std}_{2D}^{\text{ref}}[\mathcal{F}_{h_3}] - \text{Std}_{2D}^{\text{ML}}[\mathcal{F}_{h_3}])$
SML + \mathcal{I}_1	0.1001	$1.2672 \cdot 10^{-2}$	$3.2151 \cdot 10^{-6}$
AML Max + \mathcal{I}_1	0.1002	$1.1177 \cdot 10^{-2}$	$1.3801 \cdot 10^{-3}$
AML Vol + \mathcal{I}_1	0.1002	$1.1177 \cdot 10^{-2}$	$1.3801 \cdot 10^{-3}$
SML + \mathcal{I}_2	0.1000	$1.3686 \cdot 10^{-2}$	$1.0953 \cdot 10^{-3}$
AML Max + \mathcal{I}_2	0.1002	$1.2552 \cdot 10^{-2}$	$1.9033 \cdot 10^{-5}$
AML Vol + \mathcal{I}_2	0.1001	$1.2550 \cdot 10^{-2}$	$1.6201 \cdot 10^{-5}$

static multilevel approaches and *AML* to refer to adaptive multilevel decompositions. Moreover, we also use the abbreviation *CA* to refer to the context-aware versions in which we verify, via the proposed approach, whether stochastic dimensionality reduction is possible. In the first row, we show the expectation and standard deviation yielded by the non-adaptive approach, while in the second and third rows, we list the results of the spatially adaptive version using the *Max* and *Vol* refinement policy, respectively. We observe that all estimations are accurate, with relative errors in $O(10^{-4})$ and $O(10^{-5})$. The relative error of the expectation yielded by the adaptive approach is about one order of magnitude smaller than the relative error of the static approach. However, the cost benefit of using adaptivity is more significant: the adaptive approaches lead to about 3.6 times fewer points at level $\ell(3, 1)$ and about 1.9 times fewer points at level $\ell(2, 2)$ (see Figure 5.5). Therefore, we can conclude that in standard multilevel decompositions with no dimensionality reduction, the adaptive approaches are more efficient.

Next, we employ Algorithm 5.7 and analyse the possibility of stochastic dimensionality reduction using (i) the PSP coefficients from \mathcal{I}_1 and (ii) the PSP coefficient from \mathcal{I}_2 . The estimations of the five total Sobol' indices computed using the PSP coefficients from \mathcal{I}_1 read

$$\mathcal{I}_1 : \hat{S}_1 = 0.0212, \quad \hat{S}_2 = 0.1060, \quad \hat{S}_3 = 7.3012 \cdot 10^{-5}, \quad \hat{S}_4 = 3.9172 \cdot 10^{-4}, \quad \hat{S}_5 = 0.8791,$$

whereas the estimations due to the coefficients from \mathcal{I}_2 are

$$\mathcal{I}_2 : \hat{S}_1 = 0.0213, \quad \hat{S}_2 = 0.1060, \quad \hat{S}_3 = 7.3032 \cdot 10^{-5}, \quad \hat{S}_4 = 3.9143 \cdot 10^{-4}, \quad \hat{S}_5 = 0.8791.$$

In both cases, considering $\tau^{\text{in},1}$ in (5.18), we have that \hat{S}_1, \hat{S}_3 and \hat{S}_4 are unimportant, i.e., the *intrinsic stochastic dimensionality is two*. Our proposed context-aware multilevel approach makes thus use of this information and reduces the stochastic dimensionality from five to two at levels $\ell(1, 3)$ and $\ell(3, 1)$, when starting from the PSP coefficients from \mathcal{I}_1 , and at level $\ell(3, 1)$, when starting from the PSP coefficients from \mathcal{I}_2 . To better understand the impact of dimensionality reduction, we also compute the quantities of interest directly on a reference 2D grid with $15^2 = 225$ Gauss-Legendre nodes (the three parameters rendered unimportant were replaced with their expectation values), which read

$$\mathbb{E}_{2D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.0198, \quad \text{Std}_{2D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.1001.$$

In Tables 5.4 and 5.5, we list the results.

We observe that relative to the reference 5D results, the relative error of the quantities of interest computed using online dimensionality reduction is below 2%. If we further look at the relative errors w.r.t. the reference results computed directly on the reference 2D (listed in the last column in Tables 5.4 and 5.5), we see that the relative error of our results is below 0.2%. Therefore, reducing the stochastic dimensionality from five to two has an insignificant effect on the accuracy of the quantities of interest.

In Figure 5.5, we visualize the cost comparison in terms of number of sparse grid points corresponding to the static or adaptive multilevel approaches, and with or without stochastic dimensionality reduction, at all three levels. In the upper two figures, we depict the number of grid points associated to our multilevel approach with stochastic dimensionality reduction when using \mathcal{I}_1 . We observe that our proposed approach leads to a significant reduction of the number of grid points at levels $\ell(1, 3)$ and $\ell(3, 1)$, which is where the stochastic dimensionality is decreased from five to two. For example, at level $\ell(1, 3)$, the adaptive approaches require 35.9 times fewer grid points as compared to the static 5D approach and around 10 times fewer points as compared to the adaptive 5D version. Moreover, at level $\ell(3, 1)$, the reduction factor is around 4.2 for all variants, since at this level we do not perform adaptivity. We get the same reduction factor of around 4.2 at level $\ell(3, 1)$ when using \mathcal{I}_2 as well.




Figure 5.5: Number of sparse grid points at levels $\ell(1, 3)$, $\ell(2, 2)$ and $\ell(2, 2)$ for the oscillator test case (5.17) with five uncertain inputs. In all figures, the dark red and black bars are the number of grid points corresponding to the static and adaptive approach without dimensionality reduction. The red and grey bars represent the number of points of our context-aware approach with or without spatial adaptivity. In the top two figures, the PSP coefficients from \mathcal{I}_1 are used to check if the stochastic dimensionality can be reduced. In the bottom figures, we depict the results for when \mathcal{I}_2 is used.

To summarize, we showed that stochastic dimensionality reduction can lead to at least a factor of 4.2 fewer grid points compared to an equivalent approach that employs the given stochastic dimensionality at all levels, while keeping the relative error of the quantities of interest below 2%. Moreover, in the tested scenario, we observe that both \mathcal{I}_1 and \mathcal{I}_2 lead to similarly accurate results. Hence, for the five-dimensional scenario, it suffices to test whether the stochastic dimensionality can be reduced using only the PSP coefficients from \mathcal{I}_1 .

We now consider the scenario in which the input is six-dimensional, i.e., θ_2 . We make the same numerical experiments as in the 5D case. Reference quantities of interest are computed using the finest time discretization and a full Gauss-Legendre grid comprising $8^6 = 262144$ points, which read

$$\mathbb{E}_{6D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.0285, \quad \text{Std}_{6D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.1162.$$

As previously, we first employ the standard and adaptive approaches without stochastic dimensionality reduction. In the Spatially adaptive sparse pseudo-spectral approach, we consider $K_{\text{init}} = 4$. In the static version, $K_{\text{max},1} = 4$, $K_{\text{init},2} = 5$ and $K_{\text{max},3} = 6$. We list the results in Table 5.6. We observe that all relative errors are in $O(10^{-5})$, therefore the static and the adaptive approaches are similarly accurate. However, as in the 5D case, the adaptive approaches lead to significant cost reduction, making spatial adaptivity more efficient in this scenario as well. For example, at level $\ell(1, 3)$, spatial adaptivity leads to around 2.9 times fewer grid points and to 1.8 times fewer points at level $\ell(2, 2)$ (see Figure 5.6).

We analyse the possibility of stochastic dimensionality reduction using \mathcal{I}_1 and \mathcal{I}_2 . The estimations of the six total Sobol' indices corresponding to \mathcal{I}_1 are

$$\mathcal{I}_1 : \hat{S}_1 = 0.0181, \quad \hat{S}_2 = 0.0824, \quad \hat{S}_3 = 0.0034, \quad \hat{S}_4 = 0.2359, \quad \hat{S}_5 = 0.0003, \quad \hat{S}_6 = 0.6675.$$

Table 5.6: Expectation and standard deviation estimation of the displacement at $t_{\text{fixed}} = 10$ when the stochastic is θ_2 . Here, we do not perform dimensionality reduction.

Strategy in 6D	$\mathbb{E}_{6D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\mathbb{E}_{6D}^{\text{ref}}[\mathcal{F}_{h_3}] - \mathbb{E}_{6D}^{\text{ML}}[\mathcal{F}_{h_3}])$	$\text{Std}_{6D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\text{Std}_{6D}^{\text{ref}}[\mathcal{F}_{h_3}] - \text{Std}_{6D}^{\text{ML}}[\mathcal{F}_{h_3}])$
SML	0.0285	$2.3662 \cdot 10^{-5}$	0.1162	$1.9051 \cdot 10^{-5}$
AML Max	0.0285	$1.8163 \cdot 10^{-5}$	0.1162	$1.8727 \cdot 10^{-5}$
AML Vol	0.0285	$2.2072 \cdot 10^{-5}$	0.1162	$1.8778 \cdot 10^{-5}$

Table 5.7: Expectation estimation of the displacement at $t_{\text{fixed}} = 10$ when the stochastic input is θ_2 . We analyse the possibility of stochastic dimensionality reduction using the PSP coefficients at level $\ell(2, 2)$, i.e., from \mathcal{I}_1 (first three rows) and the PSP coefficient from \mathcal{I}_2 , that is, at levels $\ell(2, 2)$ and $\ell(1, 3)$ (last three rows).

Strategy with dim. red.	$\mathbb{E}_{3D}^{\text{ref}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\mathbb{E}_{6D}^{\text{ref}}[\mathcal{F}_{h_3}] - \mathbb{E}_{3D}^{\text{ML}}[\mathcal{F}_{h_3}])$	$\mathcal{E}_{\text{rel}}(\mathbb{E}_{3D}^{\text{ref}}[\mathcal{F}_{h_3}] - \mathbb{E}_{3D}^{\text{ML}}[\mathcal{F}_{h_3}])$
SML + \mathcal{I}_1	0.0288	$1.1572 \cdot 10^{-2}$	$2.4312 \cdot 10^{-5}$
AML Max + \mathcal{I}_1	0.0289	$1.6622 \cdot 10^{-2}$	$5.0163 \cdot 10^{-3}$
AML Vol + \mathcal{I}_1	0.0289	$1.6623 \cdot 10^{-2}$	$5.0165 \cdot 10^{-3}$
SML + \mathcal{I}_2	0.0288	$1.2133 \cdot 10^{-2}$	$5.7491 \cdot 10^{-4}$
AML Max + \mathcal{I}_2	0.0290	$1.7184 \cdot 10^{-2}$	$5.5656 \cdot 10^{-3}$
AML Vol + \mathcal{I}_2	0.0290	$1.7187 \cdot 10^{-2}$	$5.5668 \cdot 10^{-3}$

Table 5.8: Standard deviation estimation of the displacement at $t_{\text{fixed}} = 10$ when the stochastic input is θ_2 . We analyse the possibility of stochastic dimensionality reduction using the PSP coefficients at level $\ell(2, 2)$, i.e., from \mathcal{I}_1 (first three rows) and the PSP coefficient from \mathcal{I}_2 , that is, at levels $\ell(2, 2)$ and $\ell(1, 3)$ (last three rows).

Strategy with dim. red.	$\text{Std}_{3D}^{\text{ML}}[\mathcal{F}_{h_3}]$	$\mathcal{E}_{\text{rel}}(\text{Std}_{6D}^{\text{ref}}[\mathcal{F}_{h_3}] - \text{Std}_{3D}^{\text{ML}}[\mathcal{F}_{h_3}])$	$\mathcal{E}_{\text{rel}}(\text{Std}_{3D}^{\text{ref}}[\mathcal{F}_{h_3}] - \text{Std}_{3D}^{\text{ML}}[\mathcal{F}_{h_3}])$
SML + \mathcal{I}_1	0.1148	$1.2322 \cdot 10^{-2}$	$1.8941 \cdot 10^{-5}$
AML Max + \mathcal{I}_1	0.1149	$1.1183 \cdot 10^{-2}$	$4.3763 \cdot 10^{-4}$
AML Vol + \mathcal{I}_1	0.1149	$1.1183 \cdot 10^{-2}$	$4.3763 \cdot 10^{-4}$
SML + \mathcal{I}_2	0.1137	$2.2501 \cdot 10^{-2}$	$9.9752 \cdot 10^{-3}$
AML Max + \mathcal{I}_2	0.1137	$2.1981 \cdot 10^{-2}$	$9.4712 \cdot 10^{-3}$
AML Vol + \mathcal{I}_2	0.1137	$2.1981 \cdot 10^{-2}$	$9.4712 \cdot 10^{-3}$

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

In addition, when using the PSP coefficients from \mathcal{I}_2 , the estimated total Sobol' indices are

$$\mathcal{I}_2 : \hat{S}_1 = 0.0181, \quad \hat{S}_2 = 0.0825, \quad \hat{S}_3 = 0.0033, \quad \hat{S}_4 = 0.2357, \quad \hat{S}_5 = 0.0003, \quad \hat{S}_6 = 0.6677.$$

In both cases, \hat{S}_1, \hat{S}_3 and \hat{S}_5 fall below the prescribed threshold of importance, 0.05, thus the intrinsic stochastic dimensionality is three. Therefore, the proposed context-aware multilevel approach reduces the stochastic dimensionality from six to three when using either \mathcal{I}_1 or \mathcal{I}_2 .

Remark: In the five-dimensional scenario, the most important stochastic parameters were the spring constant, k (total Sobol index of about 0.1060), and the initial velocity, y_1 , with a total Sobol' index of approximately 0.8791. When adding the frequency as a sixth uncertain input, the list of important parameters is extended to three, with the frequency being an important parameter as well. Moreover, in this case, the overall importance of both k and y_1 is decreased. Finally, in both five- and six-dimensional scenarios, the sum of the total Sobol' indices is approximately one, which tells us that the interactions between the stochastic inputs are negligible.

As for the 5D scenario, we also compute the quantities of interest on a 3D grid comprising $15^3 = 3375$ Gauss-Legendre nodes, which gives

$$\mathbb{E}_{3D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.0288, \quad \text{Std}_{3D}^{\text{ref}}[\mathcal{F}_{h_3}] = 0.1148.$$

We present our results in Tables 5.7 and 5.8. We observe that similar to the 5D scenario, the relative error w.r.t. the reference results computed using all six stochastic parameters is below 3%; the errors w.r.t. the reference 3D results are below 1%. In addition, we observe that we obtain very similar results regardless of whether we use the PSP coefficients from \mathcal{I}_1 or \mathcal{I}_2 . We can hence conclude again that (i) reducing the dimensionality from six to three has little impact on the accuracy of the obtained quantities of interest and (ii) it suffices to check whether the stochastic dimensionality can be reduced using the PSP coefficients from \mathcal{I}_1 .

Figure 5.6 depicts the costs of all employed multilevel approaches in terms of number of grid points at all three levels. Our context-aware multilevel method leads to a significant cost reduction for this scenario as well. At levels $\ell(1,3)$ and $\ell(3,1)$, the number of grid points is decreased by a factor of approximately 10, when using static sparse grids, by 25.5 when employing spatial adaptivity at level $\ell(1,3)$, and by a factor of about 4.9 at level $\ell(3,1)$, when using a static grid in all approaches.

Remark: Observe the effect of using exponentially-growing point sets in static sparse grids in Figures 5.5 and 5.6. For example, in 6D, for a grid of level four, we have 545 Newton-Cotes points, we have 2561 points at level five and 10625 grid points at level six. The reduction in the number of grid points due to spatial adaptivity is therefore very significant in this case. For a comparison, we refer the reader to the other three contribution chapters in this work, Chapters 4, 6 and 7, in which we employed linearly increasing weighted (L)-Leja sequences.

To summarize the first test case, we considered a damped oscillator model and performed a comprehensive numerical study in which we compared multilevel decompositions with or without (i) spatial adaptivity and (ii) stochastic dimensionality reduction in two stochastic scenarios. The comparison was done in terms of the relative error of the quantities of interest w.r.t. reference results computed both using the original and the reduced dimensionality, and in terms of computational cost measured in number of samples, i.e., sparse grid points at each level. Our results show, first of all, that adaptivity leads to results of comparable accuracy to using a static sparse grid, but a considerably decreased cost. For example, in the 6D scenario,




Figure 5.6: Number of sparse grid points at levels $\ell(1,3)$, $\ell(2,2)$ and $\ell(2,2)$ for the oscillator test case (5.17) with six uncertain inputs. In all figures, the dark red and black bars are the number of grid points corresponding to the static and adaptive approach without dimensionality reduction. The red and grey bars represent the number of points of our context-aware approach with or without spatial adaptivity. In the top two figures, the PSP coefficients from \mathcal{I}_1 are used to check if the stochastic dimensionality can be reduced. In the bottom figures, we depict the results for when \mathcal{I}_2 is used.

we obtained similar accuracy at a cost decreased by a factor of 2.9 at level $\ell(1,3)$ and 1.8 at level $\ell(2,2)$. In addition, as expected, the reduction of the stochastic dimensionality has little effect on the accuracy of the quantities of interest, but reduces the cost quite significantly. For example, in the 5D scenario, we obtained reductions of up to 35.9 orders of magnitude. Next, we consider a more complex test case, a multi-physics simulation in which, due to the significant computational cost, we employ only our context-aware multilevel approach with dimensionality reduction, in which we use the *Max* refinement policy for spatial adaptivity.

5.4.2 Fluid flow over an elastic vertical flap

In the second test case, we consider a more complex and computationally more expensive test case in which we apply the proposed Multilevel adaptive spectral projection with dimensionality reduction. We consider an FSI simulation, a channel flow in a two-dimensional physical space with an elastic vertical flap, depicted in Figure 5.7. A parabolic inflow at the left boundary excites a periodic bending movement of the vertical structure. That is, we prescribe $v_{f,x} = 1 - \cos(20\pi t)$ and $v_{f,y} = 0$ at the left boundary. At the right boundary, we have an outflow given as a homogeneous Neumann condition $\frac{\partial v_f}{\partial \mathbf{n}} = 0$, where \mathbf{n} is the unit normal vector. Finally, no-slip conditions are prescribed at the top and bottom boundaries. The fluid flow is governed by the incompressible Navier-Stokes equations, formulated in the arbitrary Lagrangian-Eulerian approach [48] to allow for moving geometries (recall Section 3.1). A non-linear Saint-Venant-Kirchhoff model governs the elastic structure. The flow and structure domains are discretized via FE. The FE meshes match at the boundary, where a balancing between stresses and displacements is enforced. To couple the fluid flow and structure solvers,




Figure 5.7: Channel flow with an elastic vertical flap. We have a parabolic inflow at the left and an outflow at the right boundary. Furthermore, at the top and bottom boundaries, we prescribe no-slip conditions. The inflow excites a periodic bending movement of the vertical structure.

we use a Dirichlet-Neumann implicit partitioned approach and sub-iterate until convergence in every time step. A quasi-Newton scheme [129] is used to stabilize and accelerate the sub-iterations. In this way, both solvers can be simultaneously executed, which allows for two layers of parallelism in a single simulation: each solver is executed in parallel (first layer) and the two solvers are simultaneously executed as well (second layer).

To simulate the fluid flow and structure solvers, we employ the high-performance multi-physics code *Alya* [188] as a black box. For coupling, we use the open-source library *preCICE*⁴ [21]; see Section 3.1. For more details about the *Alya*–*preCICE* coupling, see [184]. Note that since the proposed multilevel approach for uncertainty propagation is nonintrusive, we can add a third layer of parallelism by performing FSI simulations corresponding to different sparse grid points in parallel. As an example, we visualize, in Figure 5.8, the three layers of parallelism used in our work [57], in which we employed the single level Spatially adaptive sparse pseudo-spectral approach for uncertainty propagation in two FSI test cases: the test case considered in this section and the FSI-3 benchmark [182]. For those simulations, we employed a 28 core Intel Xeon E5-2697 v3 processor from the SuperMUC Phase 2 supercomputer⁵: 20 cores were used for the fluid flow and eight for the structure solver. With this setup, a benchmark simulation with a time domain $[0, 3]$ seconds required between 288 and 336 CPU hours. Moreover, multiple FSI simulations corresponding to different sparse grid points were performed in parallel on different nodes. In this test case, we use three layers of parallelism as well, as described above. An FSI simulation is performed on a 16 core Intel Sandy Bridge processor on the CoolMAC cluster, with 12 cores for the fluid flow and four cores for the structure solver.

To discretize the problem domain, we employ 2D triangular FE for both the fluid flow (f) and the structure (s) domains. Moreover, the time domain is $[0, 0.5]$, which we discretize using a time step $\delta t = 10^{-3}$ at all three levels. Therefore, we simulate a total of 500 time steps. For both flow and structure solvers, we use implicit time integrators (implicit Euler for the flow solver and a Newmark scheme [136] for the structure solver) such that we can use the same timestep size for all meshes. We therefore do not expect stability issues and have also not observed any in our numerical experiments. Each FSI simulation is performed as follows. In an offline phase, the initial data in the flow domain is obtained via precomputing 100 time step) with a rigid structure to ensure a stabilized initial flow field before starting the coupled simulation; the structure is initially at rest. Afterwards, we use this precomputed fluid field to start the coupled simulations.

⁴<http://www.precice.org/>

⁵<https://www.lrz.de/services/compute/supermuc/>




Figure 5.8: The three parallel layers of a UQ-FSI simulation: each solver runs on parallel cores (layer one), the two solvers run in parallel to each other (layer two), and simulations corresponding to different sparse grid points are computed in parallel (layer three). The simulations are performed on 28 core Intel Xeon E5-2697 v3 processors from the SuperMUC Phase 2 supercomputer. In each simulation, we use 20 cores for the fluid solver (NASTIN in Alya) and eight for the structure solver (SOLIDZ in Alya).

In this test case, we consider five stochastic parameters $\theta = (\rho_f, \mu_f, \rho_s, E, \nu_s)$ characterizing both the flow and structure domains: $\theta = (\rho_f, \mu_f, \rho_s, E, \nu_s)$, where $\rho_f[10^3 kg/m^3]$ is the density and $\mu_f[10^6 kg/(m \cdot sec)]$ the dynamic viscosity of the fluid, whereas $\rho_s[10^3 kg/m^3]$ is the density, $E[10^6 kg/(m \cdot sec^2)]$ the Young's modulus, and ν_s the Poisson's ratio, a dimensionless parameter, corresponding to the elastic structure. In Table 5.9, we list the left and right bounds of the five uncertain inputs (columns four and five) as well as the associated deterministic value, in column six.

θ	parameter name	symbol	left bound	right bound	$\hat{\theta} := \mathbb{E}[\theta]$
θ_1	fluid density	$\rho_f[10^3 kg/m^3]$	0.740	1.260	1.000
θ_2	dynamic viscosity	$\mu_f[10^6 kg/(m \cdot sec)]$	0.074	0.126	0.100
θ_3	structure density	$\rho_s[10^3 kg/m^3]$	0.740	1.260	1.000
θ_4	Young's modulus	$E[10^6 kg/(m \cdot sec^2)]$	0.370	0.630	0.500
θ_5	Poisson's ratio	ν_s	0.222	0.378	0.300

Table 5.9: FSI test case: uniform bounds for θ (columns four and five) and associated deterministic values (column six).

As in the first test case, we consider $J = 3$ levels in our proposed context-ware multilevel spectral projection with dimensionality reduction. We list the setup at each level in Table 5.10. The values of h_1, h_2 and h_3 are in the second column and they comprise the number of elements for the flow domain, NE_f , and the number of elements for the structure domain, NE_s ; the time step is $\delta t = 10^{-3}$ at all levels.

Remark: Considering hierarchies of both spatial discretizations and time steps in a problem governed by a time-dependent PDE is not possible in a multilevel scheme, but in a so-called multiindex scheme, such as multiindex MC [86] or multiindex stochastic collocation [85].

On the employed 16 core Intel Sandy Bridge processors, one FSI run using h_1 requires about 10 minutes (2.6 CPU hours), around 15 minutes (4 CPU hours) for h_2 and around 60 minutes (16

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

Table 5.10: Multilevel setup for FSI test case with five stochastic parameters

Level	h	r
$\ell(1, 3)$	$h_1 = (\text{NE}_f = 1568; \text{NE}_s = 40)$	$r_3 = (n_{\text{ref},3} = 2; n_{\text{perc},3} = 20\%)$
$\ell(2, 2)$	$h_2 = (\text{NE}_f = 6272; \text{NE}_s = 160)$	$r_2 = (n_{\text{ref},2} = 1; n_{\text{perc},2} = 20\%)$
$\ell(3, 1)$	$h_3 = (\text{NE}_f = 25088; \text{NE}_s = 640)$	$r_1 = K_{\text{init}}$

CPU hours) when using h_3 . We see therefore that these simulations are computationally rather expensive, making the quantification of uncertainty challenging even though the stochastic dimensionality is only five. At each level, we employ the Spatially adaptive pseudo-spectral approach with

$$K_{\text{init}} = 3, \quad N_{\text{adapt},j} := j - 1, \quad \text{ref_perc} := 20\%,$$

using the *Max* refinement policy. We use PSP expansions with total degree multiindex set

$$\mathcal{P} = \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}} : |\mathbf{p}|_1 \leq 4\}$$

The output of interest is the *x-axis displacement of the upper right corner of the elastic structure*. Moreover, as in Section 5.4.1, the quantities of interest are the expectation and the standard deviation of the output of interest. To analyse the possibility for dimensionality reduction, we consider $\mathcal{I} = \{c_{p,\ell(2,2)}\}$ and a threshold for stochastic importance

$$\tau^{\text{in},1} := 0.05 \cdot \mathbf{1}. \quad (5.19)$$

For a more comprehensive overview, we perform uncertainty propagation with our proposed multilevel approach at two fixed time steps,

$$t_{\text{fixed},1} = 0.235 \quad \text{and} \quad t_{\text{fixed},2} = 0.500.$$

In Table 5.11, we show the estimated values of the five total Sobol' indices at both fixed time steps, using the PSP coefficients from \mathcal{I} . On the one hand, we observe that at $t_{\text{fixed},1} = 0.235$, the Sobol' indices associated to the fluid and structure density, and to Poisson's ratio are below 0.05, which indicates that the stochastic dimensionality can be reduced from five to two. Moreover, the five indices sum up to only a bit more than 100%, i.e., the interactions between the uncertain inputs are very weak. On the other hand, the total Sobol' indices at $t_{\text{fixed},2} = 0.500$ add up to approximately 140%, i.e., the interactions between the uncertain inputs are stronger than at $t_{\text{fixed},2} = 0.235$. Here, only \hat{S}_5 , the Sobol' index associated to the Poisson's ratio falls below the prescribed threshold of stochastic importance. Therefore, at $t_{\text{fixed},2} = 0.235$ the stochastic dimensionality can be reduced from five to four.

Fixed time step	\hat{S}_1	\hat{S}_2	\hat{S}_3	\hat{S}_4	\hat{S}_5
$t_{\text{fixed},1} = 0.235$	0.4231	0.0073	0.0091	0.5983	0.0021
$t_{\text{fixed},2} = 0.500$	0.2972	0.3545	0.1636	0.5722	0.0072

Table 5.11: Estimated values of the total Sobol' indices for at $t_{\text{fixed},1} = 0.235$ and $t_{\text{fixed},2} = 0.500$ for the FSI test case.

We proceed as follows. We first compute the PSP coefficients from \mathcal{I} using the Spatially adaptive sparse pseudo-spectral approach in 5D. These coefficients are then mapped onto the grid of the corresponding reduced dimensionality (two for $t_{\text{fixed},1}$ and four for $t_{\text{fixed},2}$), on which

the PSP coefficients at $\ell(1,3)$ and $\ell(3,1)$ are computed as well. At the end, the multilevel PSP coefficients and the quantities of interest are assessed. In Table 5.12, we list the estimated values of the expectation (column three) and standard deviation (column four) of the x -axis displacement, using the upper script *CA-AML*. Moreover, for a consistency check, we also list, in columns three and five, the expectation and standard deviation obtained with our proposed context-aware multilevel approach in which, at each level, we used global, dimension-adaptive sparse grid interpolation with (L)-Leja points (listed with upper script *CA-Leja*; see our work [56] for a more detailed discussion). The stochastic dimensionality was reduced in the same way as when using spatially adaptive interpolation. We observe that all results agree with at least two digits of precision, thus they are consistent. Moreover, we see that at $t_{\text{fixed},1}$, both expectation and standard deviation are larger than at $t_{\text{fixed},2}$, which is the upper limit of the considered time domain. The larger value of the expectation indicates that at $t_{\text{fixed},1}$, the fluid flow bends more the elastic structure than at the latter time step. In addition, considering the estimated Sobol' indices in Table 5.11, as the value of the expectation increases, the fluid density and Young's modulus become the most important stochastic parameters, while the other three are insignificant w.r.t. the value of $\tau^{\text{in},1}$ in (5.19).

Fixed time step	$\mathbb{E}^{\text{CA-AML}}[\mathcal{F}_{h_3}]$	$\mathbb{E}^{\text{CA-Leja}}[\mathcal{F}_{h_3}]$	$\text{Std}^{\text{CA-AML}}[\mathcal{F}_{h_3}]$	$\text{Std}^{\text{CA-Leja}}[\mathcal{F}_{h_3}]$
$t_{\text{fixed},1} = 0.235$	0.3231	0.3233	0.0285	0.0284
$t_{\text{fixed},2} = 0.500$	0.2884	0.2880	0.0214	0.0228

Table 5.12: Estimated values of the expectation and variance of the x - axis displacement at $t_{\text{fixed},1} = 0.235$ and $t_{\text{fixed},2} = 0.500$ using the proposed multilevel approach for the FSI test case.

The cost in terms of number of grid points at all three levels is listed in Table 5.13. In terms of runtime, at $t_{\text{fixed},1}$, the total compute time is around 70.2 hours, whereas at $t_{\text{fixed},2}$, it is about 129.4 hours. Reducing the stochastic dimensionality at $\ell(1,3)$ and $\ell(3,1)$ led to considerable computational savings. For example, at $t_{\text{fixed},1}$, decreasing the number of grid points from 71 to 17 at level $\ell(3,1)$ means that we save approximately 54 CPU hours of compute time (the total compute time with our approach at this fixed time step was around 70.2 hours), whereas at $t_{\text{fixed},2}$, the reduction from 71 to 49 grid points means saving around 22 compute hours (the total compute time here was 129.4 hours). Moreover, at level $\ell(1,3)$, we save $6 \cdot 10 = 60$ minutes of compute time for each six fewer points than on an equivalent 5D grid.

Fixed time step	No. SG points $\ell(3,1)$	No. SG points $\ell(2,2)$	No. SG points $\ell(1,3)$
$t_{\text{fixed},1} = 0.235$	17	184	43
$t_{\text{fixed},2} = 0.500$	49	179	214

Table 5.13: Number of grid points at levels $\ell(3,1)$, $\ell(2,2)$, $\ell(1,3)$ for the FSI test case. At the level $\ell(2,2)$, the stochastic dimensionality is the original one, five, and the remaining two levels the stochastic dimensionality is reduced.

To summarize the FSI test case, the proposed Multilevel adaptive spectral projection with dimensionality reduction led to considerable reduction of the computational cost for uncertainty propagation. The reduction of the stochastic dimensionality from five to two at $t_{\text{fixed},1}$ led to saving around 54 CPU hours at level $\ell(3,1)$, while the reduction from five to four stochastic dimensions at $t_{\text{fixed},2}$ meant saving about 22 hours of compute time. Moreover, we saw that even when the stochastic inputs are strongly coupled, as it was the case at $t_{\text{fixed},2}$, stochastic dimensionality reduction may still be possible. For complex, computationally expensive prob-

CHAPTER 5. MULTILEVEL ADAPTIVE SPECTRAL PROJECTIONS WITH ONLINE DIMENSIONALITY REDUCTION

lems, reducing the number of necessary simulations even by a few can lead to significant cost reductions. For example, if a single run requires 24 hours, decreasing the number of simulations from 20 to 10 translates into saving the equivalent of 10 days of computations, which constitutes a considerable reduction.

Conclusion and remarks

In this chapter, we presented a Multilevel adaptive spectral projection approach with online stochastic dimensionality reduction. Our algorithm is designed for the quantification of uncertainty in higher-dimensional, computationally expensive problems. As a representative example of such problems, we considered a channel flow in a two-dimensional physical domain with an elastic vertical flap, which required approximately 16 CPU hours for one simulation using the finest discretization resolution. Our context-aware approach exploits the fact that (i) in a multilevel decomposition we need to solve the same problem but with different resolutions and (ii) the solutions to these problems are independent. To this end, the major novelty of our algorithm is the *online stochastic dimensionality reduction* performed at *no additional computational cost*. Specifically, we first compute a *subset* of solutions from the total set associated to the multilevel decomposition using the given stochastic dimensionality and assess the total Sobol' indices for global sensitivity analysis. We use total Sobol' indices because they comprises both the local contribution of an uncertain input and contributions due its interactions with all other parameters to the total resulted uncertainty. In the next step, based on user-defined *thresholds for stochastic importance*, we quantify the importance of each stochastic input. If a subset of stochastic inputs are rendered unimportant, they are replaced by a suitable chosen deterministic value, which, in our experiments, was their expected value, and the remaining subset of problems are solved using samples from the space of reduced dimensionality. Moreover, for consistency, the results from the previous step are mapped onto the space of reduced dimensionality as well. At the end, we assess the multilevel PSP coefficients using telescoping sums, from which we estimate the quantities of interest. In the worst case scenario in which no dimensionality reduction is performed, we compute the PSP coefficients via standard multilevel decompositions based on samples from the original stochastic space. Note, nevertheless, that we still reduce the computational cost compared to a single level approach. In our formulation we analysed only once whether the stochastic dimensionality can be reduced. However, the proposed approach can be easily generalized. For instance, one could consider a sequence of nested subsets of multilevel solutions and, starting from the first subset, sequentially analyse whether the stochastic dimensionality can be reduced. This generalized approach is suitable, for example, when the initial number of uncertain parameters is large or when some total Sobol' indices in the current step are close to the associated thresholds.

To find the PSP coefficients at each level, we formulated a novel Spatially adaptive sparse pseudo-spectral approach. In this method, the multivariate integrals to find the PSP coefficients are transformed into an *inner product* between the *hierarchical surpluses* of a spatially adaptive sparse grid interpolant and products of one-dimensional integrals involving the univariate PSP polynomials and the hierarchical bases used in spatially adaptive interpolation. With spatial adaptivity, we exploit the local properties of the underlying forward model, thus reducing the cost of approximation, while the one-dimensional integrals can be either exactly or arbitrarily accurately computed. We tested the proposed context-aware multilevel spectral projection with dimensionality reduction in two test cases: (i) a damped oscillator and (ii) a multi-physics problem, i.e., a channel flow in a two-dimensional physical domain with an elastic vertical flap. For a more comprehensive overview, we considered two stochastic scenarios in each test case. In both test cases, the uncertain parameters were modelled as independent uniform random

variables and the quantities were the expectation and standard deviation of the output of interest. Moreover, we employed hierarchical polynomial bases of second degree in all sparse grid computations. In the oscillator test case, we performed a comprehensive numerical study in which we compared multilevel decompositions with or without (i) spatial adaptivity and (ii) stochastic dimensionality reduction in two stochastic in a damped oscillator model. The comparison was done in terms of the relative error of the quantities of interest w.r.t. reference results computed using both the original and the reduced dimensionality, and the computational cost measured in terms of number of samples (sparse grid points) at each level. We considered stochastic setups with five and six uncertain parameters, respectively. Our results showed that adaptivity leads to results of comparable accuracy to using a static sparse grid, but a considerably decreased cost. For example, in the 6D scenario, we obtained similar accuracy at a cost decreased by a factor of 2.9, at the first level, and 1.8, at the second level. In addition, in the 5D scenario, we could reduce the dimensionality to two, while in the second considered scenario we reduced the number of stochastic inputs from six to three. The relative errors of the quantities of interest w.r.t. either reference results were below 1%, showing that the neglected stochastic parameters were indeed unimportant. At the same time, the cost was significantly reduced. For example, we obtained reductions of up to 35.9 orders of magnitude in the 5D scenario and up to 25.5 orders of magnitude in the 6D scenario. Since the multi-physics test case was computationally more expensive, we considered only one variant of our approach for stochastic dimensionality reduction. In this problem, we performed uncertainty propagation at two fixed time steps, one close to the middle and another at the right boundary of the considered time domain. The proposed Multilevel adaptive spectral projection approach with dimensionality reduction led to considerable reduction of the computational cost for uncertainty propagation. The reduction of the stochastic dimensionality from five to two at the first fixed time step led to saving around 54 CPU hours at the first level alone, while the reduction from five to four stochastic dimensions at the second fixed time step translated into saving at least 22 hours of compute time. Moreover, we saw that even when the stochastic inputs are strongly coupled, as it was the case at the second fixed time step, stochastic dimensionality reduction may still be possible. Our algorithm assumes that the uncertain inputs are independent. We end by point out once again that when the inputs are dependent, or when their coupling is strong relative to the underlying model, stochastic dimensionality reduction has to be handled with care as it might lead to wrongly neglecting important inputs.

This concludes the second contribution chapter of this work. Next, we present our third and last algorithm for uncertainty propagation, in which we switch focus from model hierarchies based on deterministic sampling to model hierarchies in which we employ pseudo-random samples to estimate quantities of interest of high-fidelity model outputs.

PART III

CONTEXT-AWARE MULTIFIDELITY MONTE CARLO SAMPLING FOR UNCERTAINTY PROPAGATION

"As the statistician George E. P. Box wrote, "All models are wrong, but some models are useful." What he meant by that is that all models are simplifications of the universe, as they must necessarily be. As another mathematician said, "The best model of a cat is a cat." ... The key is in remembering that a model is a tool to help us understand the complexities of the universe, and never a substitute for the universe itself."

Nate Silver



Context-aware model reduction for multifidelity Monte Carlo sampling

In Chapters 4 and 5, we presented two context-aware uncertainty propagation algorithms based on deterministic methods, i.e., sparse grid and multilevel model hierarchies. In this chapter, we switch focus from deterministic approaches to sampling-based methods and present our third and last contribution to uncertainty propagation. We exploit *multifidelity model hierarchies* and formulate a *context-aware multifidelity sampling* algorithm for the estimation of statistics of complex and computationally expensive high-fidelity model outputs. In particular, we split the given computational budget between creating and improving the low-fidelity models (to reduce their deterministic approximation error) and sampling the high- and low-fidelity models (to reduce the statistical error). Moreover, we show that there is no need to improve the low-fidelity models indefinitely, but rather that there is a trade-off between exploring and exploiting the underlying models. Our proposed multifidelity approach is designed for quantifying uncertainty in higher-dimensional, computationally expensive problems. As a practically relevant application, we consider again a plasma microturbulence analysis test case similar to what we analysed in Section 4.3.4. Moreover, to demonstrate the generality of the proposed algorithms in this work, we employ the Sensitivity-driven dimension-adaptive sparse grid algorithm from Chapter 4 as a low-fidelity model in multifidelity approach formulated in this chapter. We have in preparation an article in which we will present the proposed algorithm and numerical results.

This chapter is organized as follows. In Section 6.1, we discuss some aspects of model reduction in UQ. We briefly summarize the main goals of traditional model reduction and remind the reader how multifidelity or multilevel methods make use of reduced models. In Section 6.2, we present in details our proposed context-aware multifidelity Monte Carlo algorithm for uncertainty propagation. Starting from the accuracy and cost rate functions of the low-fidelity models, we find an upper bound of the MSE of standard MFMC, summarized in Section 2.8.3, and minimize this upper bound to find the number of high-fidelity evaluations to construct and improve the surrogates. In an offline stage, we thus use part of the given computational budget to construct context-aware reduced models using the aforementioned number of high-fidelity model evaluations. The remaining budget is used in the online, multifidelity sampling stage to estimate the expected value of the high-fidelity model. We present our numerical results in two test cases in Section 6.3. First, we consider a heat conduction problem in a two-dimensional domain, in which we employ a scenario with one and another with two low-fidelity models with heterogeneous accuracy rate functions. In the second test case, we revisit the ASDEX Upgrade scenario with 12 uncertain inputs from Section 4.3.4 and apply the proposed context-aware multifidelity sampling approach to a setting in which the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm from Chapter 4 is used as a surrogate.

Highlights and novelties

- We formulate a novel Context-aware multifidelity Monte Carlo sampling algorithm for uncertainty propagation;
- In our approach, low-fidelity models are explicitly constructed for being used together with the high-fidelity model to estimate statistics of the high-fidelity model outputs;
- We show that there is quasi-optimal trade off between exploring and exploiting the surrogates;
- Low-fidelity models with heterogeneous accuracy and cost rates can be employed in the new approach;
- We show, in a thermal block problem, in which we consider a projection-based surrogate and a data-fit, regression low-fidelity model, and in a plasma microturbulence scenario in which the Sensitivity-driven dimension-adaptive algorithm from Chapter 4 is used as surrogate, that our context-aware approach is more efficient than standard MC and MFMC sampling;
- We show that in contrast with the traditional model reduction perspective, low-fidelity models can be too accurate for multifidelity sampling methods;
- To the best of our knowledge, another novelty of this chapter is that our numerical experiments using the plasma microturbulence test case are among the first in which multifidelity methods have been used to quantify uncertainty.

6.1 Reduced order models for uncertainty propagation

Model reduction traditionally aims to construct low-cost, low-fidelity models to replace complex and expensive high-fidelity models for speeding up computations. The goal of model reduction is thus to construct low-fidelity surrogates that (i) satisfy the accuracy requirements of the underlying problem and (ii) replace the high-fidelity model with the low-fidelity approximation. In contrast, in multifidelity methods (see Section 2.8.3), high- and low-fidelity models are used *together* and so the primary purpose of the surrogates is to support computations with the high-fidelity models rather than to approximate and replacing them [146]. In particular, the low-fidelity models are leveraged for speedup, whereas the high-fidelity model is kept in the loop for accuracy guarantees. Therefore, in multifidelity methods, it is unnecessary that the surrogates achieve the accuracy required by the problem at hand: the occasional recourse to the high-fidelity model corrects the overall result and hence the global accuracy is ensured even if the surrogates are not very accurate from the perspective of traditional model reduction.

In standard multifidelity sampling methods, it is assumed that the underlying low-fidelity models are *fixed* and cannot be further *improved* or *refined* such that the *entire* computational budget is spent on multifidelity sampling. Additionally, the correlation coefficients and the costs of the underlying models are typically estimated from sampling, in a preprocessing, offline step (see Section 2.8.3). Indeed, there are situations in which low-fidelity surrogates are fixed or cannot be improved. For example, consider a reduced physics model in which the simplified equations are fixed and thus any improvement would necessitate changing the model altogether. However, when the low-fidelity models can either be improved, refined, or even created *ab initio*

(this is especially the case for data-fit or projection-based reduced models), this represents an additional degree-of-freedom which we want to incorporate in the multifidelity sampling procedure. In this situation, we want to split the given computational budget between creating and improving the surrogates and sampling the high- and low-fidelity models. Nevertheless, this splitting needs to be handled with care such that the budget for finding and improving the surrogates, and the remaining budget used for multifidelity sampling are properly balanced.

Extensions of the standard multilevel Monte Carlo sampling algorithm [74] include adaptive hierarchies of non-uniform time discretization [92], goal-oriented a posteriori error estimators [52, 102], time and dimension-adaptivity [72] and Bayesian calibration of parameters related to cost and (weak) error across levels [34]. Nevertheless, in these approaches, the costs of constructing low-fidelity models are typically considered to be negligible and are therefore ignored. In [87], a bifidelity model hierarchy that relies on the low-rank structure of the forward map was proposed. Specifically, an estimate for the error due to the bifidelity model is derived, which is used to determine if a given pair of low- and high-fidelity models lead to an accurate bifidelity approximation. Additionally, in [133, 208, 209], collocation-based low-fidelity reduced models were created or improved using information from theoretically derived error bounds. However, these works consider the construction or improvement of the models only and ignore trading off the improvement and sampling costs of models in multifidelity hierarchies. To this end, in [141], the cost of creating adaptive low-fidelity models was balanced with the sampling cost in a bifidelity setting in which the low-fidelity surrogate has algebraic accuracy and cost rates. Next, we present in details our proposed context-aware multifidelity sampling approach.

6.2 Context-aware multifidelity Monte Carlo sampling

In this section, we present our proposed *Context-aware multifidelity Monte Carlo sampling algorithm* for uncertainty propagation. Our goal is to formulate a sampling method based on multifidelity model hierarchies such that:

1. reduced models are *explicitly* constructed for being used *together* with high-fidelity models;
2. a *fraction* of the computational budget is spent on creating the surrogates;
3. *heterogeneous* low-fidelity models can be employed;
4. there is *no need* for *improving* low-fidelity models *indefinitely*;
5. there is a *trade-off* between *exploring* and *exploiting* the surrogates.

To achieve the first goal, starting from the *accuracy rate*, measured in terms of *Pearson correlation coefficients* between the high- and low-fidelity models, and the *evaluation cost rate* of the surrogates, we find an upper bound of the MSE of the multifidelity estimator. By minimizing this upper bound, we find the number of high-fidelity evaluations needed to *construct* and *improve* the low-fidelity models. The associated cost is subtracted from the given computational budget, thus the second goal is achieved. The remaining budget is spent on the multifidelity sampling. To achieve the third goal, we design our proposed context-aware sampling approach to hand handle *heterogeneous* low-fidelity models, which, in our context means handling surrogates with arbitrary accuracy and cost rates, such as algebraic, geometric or exponential, characteristic to standard approximation theory [181] and model reduction [10]. Finally, the fourth and fifth goals refer to trading off the exploration with the exploitation of the high- and

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

low-fidelity models. Specifically, we want to show that there is (i) no need to improve the surrogates indefinitely, which is in stark contrast with traditional model reduction and (ii) a cost balance between creating and improving the low-fidelity models, and multifidelity sampling.

To this end, we formulate a *Context-aware multifidelity Monte Carlo sampling algorithm* for uncertainty propagation in which an arbitrary number of heterogenous low-fidelity models can be employed. Based on the accuracy and cost rate functions of the surrogates, we find an upper bound of the MSE of the multifidelity sampling estimator. This upper bound is minimized to determine *the number of high-fidelity model evaluations* necessary to construct and improve the surrogates. Hence, the low-fidelity models are created with the outer-loop result in mind, which makes them *context-aware*. A fraction of the computational budget is invested into creating and improving the surrogates; this step can be seen as the *exploration* of the low-fidelity surrogates. In the next stage, we *exploit* the context-aware surrogates and *perform* the multifidelity *sampling* according to the model management strategy in standard MFMC [145]. Note that exploration versus exploitation is also central to reinforcement learning or Bayesian optimization. This division of the computational budget is in contrast with the static MFMC approach in which the entire budget is used for sampling. The exploration and the exploitation of the low-fidelity models need to be handled with care: if the budget invested in improving the surrogates is too large, then the remaining budget for sampling is too small, and vice versa. We, nevertheless, stress out that the cost of finding reduced models cannot be neglected in problems in which the surrogates are constructed from expensive high-fidelity model evaluations. In this chapter, we will show that (i) the quasi-optimal number of high-fidelity evaluations to construct and improve the surrogates is unique and (ii) it is bounded even if the computational budget is infinite. Therefore, in our approach there is no need to improve the surrogates indefinitely, or, in other words, low-fidelity models can be *too accurate* for our purposes. Note that throughout this chapter, we use the term *quasi-optimal* and not *optimal* because, as we will see later, we are minimizing an upper bound of the MSE, not the MSE itself. This implies that there is a quasi-optimal trade-off between exploring the surrogates and exploiting the high- and low-fidelity models in MFMC sampling.

Next, we present in detail the algorithmic steps of our proposed context-aware multifidelity approach. In Section 6.2.1, we recall some of the notation and formulas introduced in the theoretical overview in Section 2.8.3. We introduce our assumptions and some preliminary lemmata in Section 6.2.2, which we need in our mathematical proofs. The proposed algorithm is detailed in Section 6.2.3. For simplicity, we first concentrate on the bifidelity setting in which we have a low-fidelity model with arbitrary accuracy and cost rates. We afterwards consider the generic setting in which we have $m - 1$ heterogenous low-fidelity models with arbitrary rates.

6.2.1 Setup

We assume we have available m models $f^{(1)}, f^{(2)}, \dots, f^{(m)}$, where $f^{(1)} \equiv \mathcal{F}_h$ is the high-fidelity model and the rest of $m - 1$ models are low-fidelity surrogates. We consider low-fidelity models that are explicitly constructed from a set of high-fidelity model evaluations, such as data-fit or projection-based reduced models. Therefore, letting N_j denote the number of high-fidelity evaluations from which $f^{(j)}$ is determined, we explicitly denote the dependency of $f^{(j)}$ on N_j by $f_{N_j}^{(j)}$, for $j = 2, 3, \dots, m$. For simplicity, in the rest of this section, unless specified otherwise, we use $f_{N_j}^{(j)}$ to refer to the low-fidelity models for $j = 2, 3, \dots, m$ without explicitly specifying that j takes values from 2 to m .

The variances of the m models are denoted by $\sigma_1^2 := \text{Var}[f^{(1)}]$ and $\sigma_j^2 := \text{Var}[f_{N_j}^{(j)}]$. Moreover, recall that the Pearson correlation coefficient between the high-fidelity model, $f^{(1)}$, and the

surrogate, $f_{N_j}^{(j)}$, is $\rho_j = \text{Cov}[f^{(1)}, f_{N_j}^{(j)}]/(\sigma_1 \sigma_j)$. In addition, \mathcal{C}_j denotes the cost of evaluating $f_{N_j}^{(j)}$. Without loss of generality, we normalize the costs with the cost of the the high-fidelity model, i.e., $\bar{\mathcal{C}}_1 = 1$ and $\bar{\mathcal{C}}_j = \mathcal{C}_j/\mathcal{C}_1$. Next, we state the assumptions we make in this chapter and some preliminary lemmata which we will use in our proofs.

6.2.2 Preliminaries

Throughout this work, we assume that the underlying stochastic space, \mathbf{X} , and its associated probability density function, $\boldsymbol{\pi}$, have a product structure (recall Assumptions 2.1 and 2.2). To formulate our proposed context-aware sampling algorithm, we need additional assumptions, which we state next.

Remark: MC estimators can be defined for dependent input as well, i.e., for multivariate density functions not having a product structure. Therefore, Assumptions 2.1 and 2.2 are not very strict in this chapter. Nevertheless, the formulas for the MSE in Section 2.8, which we will also use to formulate our context-aware algorithm, were derived assuming i.i.d. input samples. Should the stochastic inputs be dependent, we would need to rederive these error formulas. Interestingly, a possible research direction in multifidelity sampling methods is to investigate whether employing dependent MC estimators in MFMC or its variants can lead to further variance reduction.

The first assumption is that σ_j is bounded from above and below.

- Assumption 6.1: There exist $\underline{\sigma}_j, \bar{\sigma}_j \in \mathbb{R}, 0 < \underline{\sigma}_j \leq \bar{\sigma}_j$ such that

$$\underline{\sigma}_j \leq \sigma_j \leq \bar{\sigma}_j.$$

Next, we assume that the accuracy of the low-fidelity surrogate $f_{N_j}^{(j)}$, measured in terms of its Pearson correlation w.r.t. the high-fidelity model is bounded from above by a decaying rate function.

- Assumption 6.2: Let $p > 0$ and $r_a : [1, p] \rightarrow \mathbb{R}_+$ be a convex function in $C^3([1, p])$. There exist a constant $c_1^{(j)} > 0$ and a rate $\alpha^{(j)} > 0$ such that

$$1 - \rho_j^2 \leq c_1^{(j)} r_a^{(j)}(\alpha^{(j)}, N_j).$$

In addition, we assume that the normalized costs $\bar{\mathcal{C}}_j$ are bounded from above as well.

- Assumption 6.3: Let $p > 0$ and $r_c : [1, p] \rightarrow \mathbb{R}_+$ be a function in $C^3([1, p])$. There exist a constant $c_2^{(j)} > 0$ and a rate $\beta^{(j)} > 0$ such that

$$\bar{\mathcal{C}}_j \leq c_2^{(j)} r_c^{(j)}(\beta^{(j)}, N_j).$$

Note that the accuracy and cost rates defined in Assumptions 6.2 and 6.3 are defined on *compact* domains $[1, p]$.

Remark: For some problems, the accuracy and cost rates in Assumptions 6.2 and 6.3 are available analytically. As an example, when the high-fidelity is an elliptic PDE, the rates for low-fidelity surrogates such as interpolation [139] or reduced basis approximation [153] are known (see Section 6.3.1). If the rates are not analytically known but (at least two) evaluations of the underlying models are available, e.g., from previous experiments, the rates can be estimated via

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

importance sampling [3]. Otherwise, the accuracy and cost rates need to be inferred numerically. First, the correlation coefficient and the runtime of the low-fidelity models are evaluated for an increasing number of samples, and then the rates are estimated via, e.g., (possibly nonlinear) regression. Note that estimating the accuracy and cost rates numerically inquires an extra computational cost, which can be explicitly taken into account in our proposed algorithm.

Before going any further, we provide some examples of accuracy and cost rates.

Example: Examples of accuracy rate functions are *algebraic*, i.e.,

$$r_a^{(j)}(\alpha^{(j)}, N_j) := N_j^{-\alpha^{(j)}}$$

and *exponential*, i.e.,

$$r_a^{(j)}(\alpha^{(j)}, N_j) := \exp(-aN_j^{\alpha^{(j)}}),$$

for an $a \in \mathbb{R}_+$. Moreover, an example of cost rates which we consider in this work are *algebraic*:

$$r_c^{(j)}(\beta^{(j)}, N_j) := N_j^{\beta^{(j)}}.$$

In Section 6.3, we consider low-fidelity models with exponential accuracy rate and algebraic cost rate as well as algebraic accuracy and cost rates.

Since $r_a^{(j)}(\alpha^{(j)}, N_j)$ and $r_c^{(j)}(\beta^{(j)}, N_j)$ describe the accuracy and evaluation cost of the surrogates, it is straightforward to make the following assumption.

- Assumption 6.4: If $N_j \leq M_j$, then

$$\begin{aligned} r_a^{(j)}(\alpha^{(j)}, N_j) &\geq r_a^{(j)}(\alpha^{(j)}, M_j), & \text{i.e., } (r_a^{(j)}(\alpha^{(j)}, N_j))' &\leq 0 \\ r_c^{(j)}(\beta^{(j)}, N_j) &\leq r_c^{(j)}(\beta^{(j)}, M_j), & \text{i.e., } (r_c^{(j)}(\beta^{(j)}, N_j))' &\geq 0. \end{aligned}$$

That is, the accuracy rates are monotonically *decreasing*, whereas the cost rates are monotonically *increasing*: as we invest more effort in determining the low-fidelity models, we expect to have more accurate surrogates, which are also more expensive to evaluate. In general, a reduced model is constructed globally and that it is characterized by a single accuracy and thus a single cost function as well. We additionally assume that the rate functions have a global behaviour.

- Assumption 6.5: The rate functions $r_a^{(j)}(\alpha^{(j)}, N_j)$ and $r_c^{(j)}(\beta^{(j)}, N_j)$ have the same monotonicity and convexity on their entire (compact) domain, $[1, p]$.

Remark: In Assumptions 6.2 and 6.3, the rate functions are assumed to be in $C^3([1, p])$, which provides sufficient smoothness for our purposes. Nevertheless, these functions are generally elementary (polynomial, exponential, logarithmic etc.) and thus *analytic*, i.e. in $C^\infty([1, p])$.

In addition, the accuracy rates are assumed to be *monotonically decreasing* and *convex* on $[1, p]$. In contrast, the cost rate function can be either concave (for example, an algebraic cost rate with $\beta^{(j)} \leq 1$) or convex (algebraic cost rate with $\beta^{(j)} \geq 1$).

Next, we prove a generalization of Lemma 1 from [141], which shows that Assumption 6.2 is implied if the same decay rate is assumed for $\text{Var}[f^{(1)} - f_{N_j}^{(j)}]$. This result is related to an analogous assumption made in multilevel Monte Carlo sampling [74], in which $\text{Var}[f^{(1)} - f_{N_j}^{(j)}]$ is bounded from above. We employ the following notation: $a \lesssim b$ if there exists $c > 0$ not depending on a or b such that $a \leq cb$.

Lemma: 6.2.1 [141, Lemma 1] $\text{Var}[f^{(1)} - f_{N_j}^{(j)}] \lesssim r_a^{(j)}(\alpha^{(j)}, N_j) \Rightarrow 1 - \rho_j^2 \lesssim r_a^{(j)}(\alpha^{(j)}, N_j)$.

Proof: We have

$$\begin{aligned} \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &= \text{Var}[f^{(1)}] + \text{Var}[f_{N_j}^{(j)}] - 2\text{Cov}[f^{(1)}, f_{N_j}^{(j)}] \Rightarrow \\ \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &\geq 2\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]} - 2\text{Cov}[f^{(1)}, f_{N_j}^{(j)}] \Leftrightarrow \\ \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &\geq 2\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]} \left(1 - \frac{\text{Cov}[f^{(1)}, f_{N_j}^{(j)}]}{\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]}} \right) \Leftrightarrow \\ \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &\geq 2\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]}(1 - \rho_j). \end{aligned}$$

Using $2 \geq 1 + \rho_j$, we get

$$\begin{aligned} \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &\geq 2\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]}(1 - \rho_j) \Leftrightarrow \\ \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &\geq (1 + \rho_j)\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]}(1 - \rho_j) \Leftrightarrow \\ \text{Var}[f^{(1)} - f_{N_j}^{(j)}] &\geq \sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]}(1 - \rho_j^2). \end{aligned}$$

Since $\sqrt{\text{Var}[f^{(1)}]\text{Var}[f_{N_j}^{(j)}]}$ is bounded from above by a constant (cf. Assumption 6.1), the conclusion follows. ■

We showcase Lemma 6.2.1 in the following example.

Example: Consider a problem with two uniformly distributed stochastic inputs in $[0, 1]^2$. The high-fidelity model is available analytically and it reads:

$$f^{(1)}(\boldsymbol{\theta}) = \cos(\pi + 2\theta_1 + 1.5\theta_2).$$

We consider a single low-fidelity model, $f_{N_2}^{(2)}$, obtained via adaptive sparse grid interpolation using the Sensitivity-driven dimension-adaptive sparse grid algorithm from Chapter 4. Since the stochastic setting is two-dimensional, either strategy to compute sensitivity scores requires assessing three scores (see Section 4.2.3). In this context, N_2 represents the number of high-fidelity evaluations to obtain the adaptive sparse grid interpolation surrogate. We consider tolerances $\boldsymbol{\tau}^{in-g}$ between $(10^{-6}, 10^{-6}, 10^{-6})$ and $(10^{-9}, 10^{-9}, 10^{-9})$. Moreover, we estimate $\text{Var}[f^{(1)} - f_{N_2}^{(2)}]$ and $1 - \rho_2^2$ numerically using 1000 MC samples. Lemma 6.2.1 says that as N_2 increases, $\text{Var}[f^{(1)} - f_{N_2}^{(2)}]$ and $1 - \rho_2^2$ should decrease with the *same rate* and the two decay curves should be *parallel* to each other. We depict our results in Figure 6.1. We observe that on a logarithmic scale, for the considered setup, the two curves decrease linearly, i.e., the decay rates are *algebraic*. Moreover, the two curves are parallel, i.e., they differ by a constant independent on the decay rate and N_2 . Therefore, Lemma 6.2.1 is verified.

Next, we introduce some preliminary lemmata which we need in proofs later on.

The first lemma states a necessary and sufficient condition for a single variable function $g \in C^2([a, b])$ to be convex.

Lemma: 6.2.2 Let $g : [a, b] \subset \mathbb{R} \rightarrow \mathbb{R}, g \in C^2([a, b])$. g is convex if $\forall x \in [a, b] \Rightarrow g''(x) \geq 0$.

Moreover, the summation of two convex functions is also convex.




Figure 6.1: $\text{Var}[f^{(1)} - f_{N_2}^{(2)}]$ and $1 - \rho_2^2$, calculated estimated using 1000 MC samples, for increasing values of N_2 , i.e., decreasing sparse grid interpolation tolerances.

Lemma: 6.2.3 *Let $g, h : [a, b] \subset \mathbb{R} \rightarrow \mathbb{R}, g, h \in C^2([a, b])$. If g and h are convex, $g + h$ is convex as well.*

We next remind the reader of Weierstrass' Extreme Value theorem for univariate functions defined on compact domains.

Theorem 6.2.4 *(Weierstrass' Extreme Value theorem). A monotonically increasing continuous function g defined on a compact domain $[a, b] \subset \mathbb{R}$ satisfies $g(a) \leq g(x) \leq g(b)$. Conversely, if f is monotonically decreasing, then $g(b) \leq g(x) \leq g(a)$.*

Finally, we recall how the sign of a second degree polynomial depends on the sign of its roots and the sign of the second degree monomial.

Lemma: 6.2.5 *Let $Q : \mathbb{R} \rightarrow \mathbb{R}$ be a second degree polynomial, i.e., $Q(x) = Ax^2 + Bx + C$. Let $\mathcal{D} := B^2 - 4AC$ be the discriminant and x_1^* and x_2^* the two roots of $Q(x) = 0$. If $\mathcal{D} < 0$, then $x_1^*, x_2^* \in \mathbb{C}$ are complex conjugate. When $\mathcal{D} \geq 0$, then $x_1^*, x_2^* \in \mathbb{R}$ with $x_1^* = x_2^*$ if $\mathcal{D} = 0$. We have the following conditions regarding the sign of $Q(x)$:*

- (i) *if $A \geq 0$ and $\mathcal{D} \leq 0$, then $Q(x) \geq 0$ for all $x \in \mathbb{R}$;*
- (ii) *if $A \leq 0$ and $\mathcal{D} \leq 0$, then $Q(x) \leq 0$ for all $x \in \mathbb{R}$;*
- (iii) *if $A \geq 0$, $\mathcal{D} \geq 0$, then $Q(x) \geq 0$ for all $x \in \mathbb{R} \setminus [x_1^*, x_2^*]$ and $Q(x) \leq 0$ for all $x \in [x_1^*, x_2^*]$;*
- (iv) *if $A \leq 0$, $\mathcal{D} \geq 0$, then $Q(x) \leq 0$ for all $x \in \mathbb{R} \setminus [x_1^*, x_2^*]$ and $Q(x) \geq 0$ for all $x \in [x_1^*, x_2^*]$.*

In the following, we present in details the algorithmic steps of our proposed context-aware multifidelity sampling approach. First, we show how we obtain an upper bound of the MSE (2.42) of the MFMC estimator (2.40) introduced in Section 2.8.3 depending on the accuracy and cost rates of the low-fidelity models. Afterwards, we show how to construct context-aware low-fidelity models and how to perform MFMC sampling in a bifidelity setting in which we have a low-fidelity model with arbitrary accuracy and cost rates. Finally, we generalize this setting to the situation in which we have $m - 1$ heterogenous low-fidelity surrogates.

6.2.3 Context-aware multifidelity Monte Carlo sampling algorithm

Let us recall the MSE (2.42) of the MFMC estimator (2.40) from Section 2.8.3,

$$\text{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) = \text{Var}[\hat{f}_{\text{MFMC}}^{(1)}] = \underbrace{\frac{\text{Var}[f^{(1)}]}{p}}_{\text{MSE single model}} \underbrace{\left(\sum_{j=1}^m \sqrt{c_j (\rho_j^2 - \rho_{j+1}^2)} \right)^2}_{\text{variance reduction}}.$$

This MSE depends on the variance of the high-fidelity model, $\text{Var}[f^{(1)}]$, the computational budget, p , and the correlation coefficients and costs of the m employed models. For a given computational budget p , we can perform $N_1 := N_{\text{hi-fi}}$ high-fidelity model evaluations, i.e.,

$$p = N_1 c_1 = N_{\text{hi-fi}} c_1.$$

Therefore, the above MSE transforms into

$$\text{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) = \frac{\text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(\sum_{j=1}^m \sqrt{\bar{c}_j (\rho_j^2 - \rho_{j+1}^2)} \right)^2. \quad (6.1)$$

In case that p/c_1 is not an integer, we simply round it up. In this work we focus on complex applications, hence the associated computational budgets for uncertainty propagation are usually large, which means that rounding p/c_1 up has a negligible impact. Our target is to obtain an upper bound of the above MSE in (6.1) that depends on the accuracy and cost rates of the low-fidelity surrogates, stated in Assumptions 6.2 and 6.3.

Remark: Observe that we have no bias term w.r.t. the high-fidelity model in the MSE (6.1), even though $m - 1$ low-fidelity surrogates are employed. This offers us great flexibility to construct heterogeneous surrogates, such as projection-based or data-fit, with heterogeneous cost and accuracy rates, for example, exponential and algebraic. We note that when the high-fidelity model is an approximation of a continuous model, as it is the case in this work, this introduces a bias associated to the high-fidelity model, which is controlled by the employed approximation. However, as the high-fidelity model is assumed to be sufficiently accurate to yield results with any desired accuracy, the associated bias is small.

In the first step, we derive an upper bound of (6.1) in which we eliminate the square roots and the squared summation. To this end, we make use of the multinomial expansion

$$\left(\sum_{n=1}^N a_n \right)^2 = \sum_{n=1}^N a_n^2 + \sum_{1 \leq n_1 \neq n_2 \leq N} 2a_{n_1} a_{n_2}, \quad a_1, a_2, \dots, a_N \in \mathbb{R},$$

and of the inequality between the geometric and arithmetic means of $x, y \in \mathbb{R}_+$,

$$2\sqrt{xy} \leq x + y.$$

We therefore obtain

$$\begin{aligned}
\text{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) &= \frac{\text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(\sum_{j=1}^m \sqrt{\bar{C}_j (\rho_j^2 - \rho_{j+1}^2)} \right)^2 \\
&= \frac{\text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(\sum_{j=1}^m \bar{C}_j (\rho_j^2 - \rho_{j+1}^2) + \sum_{1 \leq j_1 \neq j_2 \leq m} 2\sqrt{\bar{C}_{j_1} (\rho_{j_1}^2 - \rho_{j_1+1}^2)} \sqrt{\bar{C}_{j_2} (\rho_{j_2}^2 - \rho_{j_2+1}^2)} \right) \\
&\leq \frac{\text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(\sum_{j=1}^m \bar{C}_j (\rho_j^2 - \rho_{j+1}^2) + \sum_{1 \leq j_1 \neq j_2 \leq m} (\bar{C}_{j_1} (\rho_{j_1}^2 - \rho_{j_1+1}^2) + \bar{C}_{j_2} (\rho_{j_2}^2 - \rho_{j_2+1}^2)) \right) \\
&= \frac{m \text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(\sum_{j=1}^m \bar{C}_j (\rho_j^2 - \rho_{j+1}^2) \right).
\end{aligned}$$

To make use of the accuracy and cost rates of the low-fidelity models stated in Assumptions 6.2 and 6.3, we need to do some further arithmetic manipulations in the term $\sum_{j=1}^m \bar{C}_j (\rho_j^2 - \rho_{j+1}^2)$. To simplify the notation, we use $c_1^{(j)} r_a^{(j)} := c_1^{(j)} r_a^{(j)}(\alpha^{(j)}, N_j)$ and $c_2^{(j)} r_c^{(j)} := c_2^{(j)} r_c^{(j)}(\beta^{(j)}, N_j)$ in the following. We have

$$\begin{aligned}
\sum_{j=1}^m \bar{C}_j (\rho_j^2 - \rho_{j+1}^2) &= 1 - \rho_2^2 + \sum_{j=2}^{m-1} \bar{C}_j (\rho_j^2 - \rho_{j+1}^2) + \bar{C}_m \rho_m^2 \\
&= 1 - \rho_2^2 + \sum_{j=2}^{m-1} \bar{C}_j ((1 - \rho_{j+1}^2) - (1 - \rho_j^2)) + \bar{C}_m \rho_m^2 \\
&\leq c_1^{(2)} r_a^{(2)} + \sum_{j=2}^{m-1} c_1^{(j)} r_c^{(j)} (c_1^{(j+1)} r_a^{(j+1)} - c_1^{(j)} r_a^{(j)}) + \bar{C}_m \rho_m^2 \\
&\leq c_1^{(2)} r_a^{(2)} + \sum_{j=2}^{m-1} c_2^{(j)} r_c^{(j)} (c_1^{(j+1)} r_a^{(j+1)} - c_1^{(j)} r_a^{(j)}) + c_2^{(m)} r_c^{(m)},
\end{aligned}$$

where we used the accuracy and cost rates defined in Assumptions 6.2 and 6.3, and, in the last step we used the fact that $\rho_m^2 \leq 1$.

We hereby obtain the following upper bound of the MSE of the MFMC estimator:

$$\text{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) \leq \frac{m \text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(c_1^{(2)} r_a^{(2)} + \sum_{j=2}^{m-1} c_2^{(j)} r_c^{(j)} (c_1^{(j+1)} r_a^{(j+1)} - c_1^{(j)} r_a^{(j)}) + c_2^{(m)} r_c^{(m)} \right), \tag{6.2}$$

which depends *only* on the accuracy and cost rates of all low-fidelity models. We want to use the upper bound in (6.2) in our proposed algorithm to find *context-aware* low-fidelity models. In Figure 6.2, we provide a visual summary of the proposed Context-aware multifidelity Monte Carlo sampling algorithm.

For simplicity, we first consider the bifidelity setting, i.e., the high-fidelity model and one surrogate, which can have arbitrary accuracy and cost rates satisfying Assumptions 6.2 – 6.5. We find the conditions for which the objective used to find the number of high-fidelity evaluations, N_2^* , to construct the surrogate is *convex* (Theorem 6.2.6) and N_2^* is *bounded* w.r.t. the budget

p (Theorem 6.2.7). Afterwards, we generalize the multifidelity model hierarchy to the situation in which we have $m - 1$ heterogenous low-fidelity models with $m \geq 3$. We show how the results from the bifidelity setting can be sequentially extended to the generic situation with $m \geq 3$ models in total and leave the convexity and boundness proofs for future research. In the numerical results in Section 6.3, we consider $m = 2$, i.e., a bifidelity setting and $m = 3$.




Figure 6.2: Visual summary of the proposed Context-aware multifidelity Monte Carlo sampling algorithm for uncertainty propagation. We denote the total computational budget by p . First, $m - 1$ context-aware surrogates $f_{N_2}^{(2)}, f_{N_3}^{(3)}, \dots, f_{N_m}^{(m)}$ are created using N_2, N_3, \dots, N_m high-fidelity model evaluations. N_2, N_3, \dots, N_m are quasi-optimally found from minimizing an upper bound of the MSE of the MFMC estimator. We denote the cost corresponding to the $\sum_{j=2}^m N_j$ high-fidelity evaluations to find the $m - 1$ surrogates by $p_{\text{lo-fi}}$. The remaining budget, $p - p_{\text{lo-fi}}$, is used for multifidelity sampling to estimate the expectation of the high-fidelity model, $\hat{f}_{\text{CA-MFMC}}^{(1)}$, where CA stands for context-aware.

One context-aware low-fidelity model

We consider first the setup in which we have the high-fidelity model, $f^{(1)}$ and one low-fidelity model, $f_{N_2}^{(2)}$, i.e., $m = 2$ in (6.2). Therefore, (6.2) reads

$$\text{MSE}(\hat{f}_{\text{MFMC}}^{(1)}) \leq \frac{2\text{Var}[f^{(1)}]}{N_{\text{hi-fi}}} \left(c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2) \right). \quad (6.3)$$

The surrogate $f_{N_2}^{(2)}$ is constructed and improved from N_2 high-fidelity evaluations. Our goal is to find N_2 such that the low-fidelity model is *context-aware*. To this end, we proceed as follows. The upper bound in (6.3) depends explicitly on N_2 , which enters in the accuracy and cost rate functions. Thus, to find the (quasi-optimal) number of high-fidelity evaluations, N_2^* ,

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

to construct and improve the surrogate, we minimize the upper bound in (6.3). Importantly, after N_2^* was determined, we subtract it from the total budget which allows $N_{\text{hi-fi}}$ high-fidelity evaluations and we are hence left with the budget $N_{\text{hi-fi}} - N_2^*$ for the multifidelity sampling. We define the following objective:

$$u(N_2) = \frac{1}{N_{\text{hi-fi}} - N_2} \left(c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2) \right),$$

Finding N_2^* simply means solving the following minimization problem:

$$\begin{aligned} \min_{N_2 \in \mathbb{N}} \quad & u(N_2) \\ \text{subject to} \quad & 2 \leq N_2 \leq N_{\text{hi-fi}} - 1. \end{aligned} \tag{6.4}$$

We constrain the optimum in (6.4) to be at least 2 and at most $N_{\text{hi-fi}} - 1$, to allow for at least one remaining high-fidelity evaluation in the multifidelity sampling part. We relax the program (6.4) and solve it for $N_2 \in (1, N_{\text{hi-fi}}) \subset \mathbb{R}$. Thus, in computations, we round it up to the nearest integer \bar{N}_2^* , i.e.,

$$\bar{N}_2^* := \lceil N_2^* \rceil \in \mathbb{N}.$$

Remark: Observe that (6.4) gives the quasi-optimal number of high-fidelity model evaluations to construct and improve the low-fidelity model, but does not prescribe *how* to choose the samples at which to evaluate the high-fidelity model. These samples can be selected *anywhere* in the stochastic input domain, \mathbf{X} , therefore they can be pseudo-random samples, deterministic collocation or quadrature points, etc. The quality of the samples is reflected by the accuracy rate, $\alpha^{(2)}$. Hence, our algorithm does not restrict the type of reduced model to be employed. In our numerical results, we will construct reduced models using pseudo-random samples and weighted (L)-Leja interpolation points.

We want to find the conditions for $u(N_2)$ to be *convex*, i.e., for the minimum of (6.4) to *exist* in $(1, N_{\text{hi-fi}})$ and to be *unique*. To this end, we prove the following theorem.

Theorem 6.2.6 *Consider the function $u : (1, N_{\text{hi-fi}}) \rightarrow \mathbb{R}$ defined in (6.4). Additionally, let $r(N_2) := c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2)$ and define $Q : [1, \infty) \rightarrow \mathbb{R}$, $Q(N_{\text{hi-fi}}) := r(1) + r'(1)(N_{\text{hi-fi}} - 1) + \frac{1}{2}r''(1)(N_{\text{hi-fi}} - 1)^2$. If any of the following three conditions hold true:*

- (i) $r'''(N_2) \leq 0$
- (ii) $r'''(N_2) \geq 0$ and $(r'(1))^2 - 2r(1)r''(1) \leq 0$
- (iii) $r'''(N_2) \geq 0$, $(r'(1))^2 - 2r(1)r''(1) \geq 0$, $r''(1) \geq 0$ and $q_1^* \leq q_2^* < 1$,

where q_1^*, q_2^* are the real roots of Q , then the function u is convex in $(1, N_{\text{hi-fi}})$. Therefore, a unique global minimum N_2^* exists in $(1, N_{\text{hi-fi}})$.

Proof: We first compute the second derivative of u :

$$u''(N_2) = \frac{2}{(N_{\text{hi-fi}} - N_2)^3} \left(r(N_2) + r'(N_2)(N_{\text{hi-fi}} - N_2) + r''(N_2) \frac{(N_{\text{hi-fi}} - N_2)^2}{2} \right).$$

We want to find the conditions when $u''(N_2) \geq 0$ (see Lemma 6.2.2). Let $\bar{u} : [1, N_{\text{hi-fi}}] \rightarrow \mathbb{R}$,

$$\bar{u}(N_2) := r(N_2) + r'(N_2)(N_{\text{hi-fi}} - N_2) + r''(N_2) \frac{(N_{\text{hi-fi}} - N_2)^2}{2}. \tag{6.5}$$

The condition $u''(N_2) \geq 0$ is equivalent to $\bar{u}''(N_2) \geq 0$ since $N_{\text{hi-fi}} - N_2 > 0$ and with that, $2/(N_{\text{hi-fi}} - N_2)^3 > 0$.

We study the sign of \bar{u} in (6.5) by making use of Theorem 6.2.4. We first compute $\bar{u}'(N_2)$:

$$\bar{u}'(N_2) = r'''(N_2) \frac{(N_{\text{hi-fi}} - N_2)^2}{2}.$$

If, on the one hand, $r'''(N_2) \leq 0$, which is equivalent to $\bar{u}'(N_2) \leq 0$, then \bar{u} is monotonically decreasing on $[1, N_{\text{hi-fi}}]$ (see Assumption 6.5). Theorem 6.2.4 thus implies that

$$\bar{u}(N_2) \geq \bar{u}(N_{\text{hi-fi}}) = c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_{\text{hi-fi}}) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_{\text{hi-fi}}) \geq 0,$$

where the last inequality is due to the fact that both rate functions are positive. Hence, if $r'''(N_2) \leq 0$ then $\bar{u}(N_2) \geq 0$ and thus $u(N_2)$ defined in (6.4) is convex. This corresponds to condition (i) of the theorem.

On the other hand, if $r'''(N_2) \geq 0$, then $\bar{u}'(N_2) \geq 0$ and thus \bar{u} is monotonically increasing in $[1, N_{\text{hi-fi}}]$. From Theorem 6.2.4, this means that

$$\bar{u}(N_2) \geq \bar{u}(1) = r(1) + r'(1)(N_{\text{hi-fi}} - 1) + \frac{1}{2}r''(1)(N_{\text{hi-fi}} - 1)^2 =: Q(N_{\text{hi-fi}}).$$

To ascertain the sign of $\bar{u}(1)$, we study the sign of the quadratic polynomial $Q(N_{\text{hi-fi}})$ defined on $[1, \infty)$. First, let us expand $Q(N_{\text{hi-fi}})$:

$$Q(N_{\text{hi-fi}}) = \frac{1}{2}r''(1)N_{\text{hi-fi}}^2 + (r'(1) - r''(1))N_{\text{hi-fi}} + \left(r(1) - r'(1) + \frac{1}{2}r''(1) \right).$$

We study the sign of $Q(N_{\text{hi-fi}})$ using Lemma 6.2.5. We have

$$A_Q := \frac{1}{2}r''(1)$$

and

$$\mathcal{D}_Q := (r'(1) - r''(1))^2 - 2r''(1) \left(r(1) - r'(1) + \frac{1}{2}r''(1) \right) = (r'(1))^2 - 2r(1)r''(1).$$

The discriminant \mathcal{D}_Q is negative if

$$\mathcal{D}_Q \leq 0 \Leftrightarrow (r'(1))^2 - 2r(1)r''(1) \leq 0 \Leftrightarrow (r'(1))^2 \leq 2r(1)r''(1).$$

Since $(r'(1))^2 \geq 0$ and $r(1) \geq 0$ because both rate functions are positive, it follows that $\mathcal{D}_Q \leq 0$ can hold true only if $r''(1) \geq 0$, which further implies that $A_Q \geq 0$. When these are satisfied, condition (ii) in Lemma 6.2.5 ensures that $Q(N_{\text{hi-fi}}) \geq 0$ for all $N_{\text{hi-fi}} \in \mathbb{R}$. In summary, when $r'''(N_2) \geq 0$ and $\mathcal{D}_Q = (r'(1))^2 - 2r(1)r''(1) \leq 0$, condition (ii) from our theorem is verified.

If \mathcal{D}_Q is positive, then

$$\mathcal{D}_Q \geq 0 \Leftrightarrow (r'(1))^2 - 2r(1)r''(1) \geq 0 \Leftrightarrow (r'(1))^2 \geq 2r(1)r''(1)$$

which always holds true if $r''(1) \leq 0$, that is, if $A_Q \leq 0$ or can hold true if $r''(1) \geq 0$, i.e., $A_Q \geq 0$. Assuming that $\mathcal{D}_Q \geq 0$, let $q_1^* \leq q_2^* \in \mathbb{R}$ be the two real roots of $Q(N_{\text{hi-fi}}) = 0$ such that $q_1^* = q_2^*$ if $\mathcal{D}_Q = 0$. If $r''(1) \geq 0$ and $q_1^* \leq q_2^* < 1$ then from condition (iii) in Lemma 6.2.5, $Q(N_{\text{hi-fi}}) \geq 0$ on $\mathbb{R} \setminus [q_1^*, q_2^*] = [1, \infty)$, i.e., $Q(N_{\text{hi-fi}}) \geq 0$ on $[1, \infty)$, which is the domain on which $N_{\text{hi-fi}}$ is defined. This corresponds to condition (iii) of the theorem.

If any other of the remaining conditions hold true, that is, $r''(1) \geq 0$ and at least one of q_1^*, q_2^* is greater than one, or if $r''(1) \leq 0$, then there are $N_{\text{hi-fi}} \in [1, \infty)$ for which $Q(N_{\text{hi-fi}}) \leq 0$ which imply that there are exist budgets p for which $u(N_2)$ is concave. ■

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

Before moving forward, we want to employ the theorem in two specific examples, in which we consider the rate functions of interest in this work.

Example: In [141, Proposition 3], it was shown that $u(N_2)$ defined in (6.4) is *always* convex when the accuracy and cost rate functions are *algebraic*, i.e.,

$$r_a^{(2)}(\alpha^{(2)}, N_2) := N_2^{-\alpha^{(2)}}$$

and

$$r_c^{(2)}(\beta^{(2)}, N_2) := N_2^{\beta^{(2)}}. \quad (6.6)$$

Note that the proof of [141, Proposition 3] was tailored to the specific case in which the rates are algebraic, whereas Theorem 6.2.6 focuses on the setting with generic rate functions.

We prove by contradiction that the conditions from Theorem 6.2.6 for which $u(N_2)$ is *not* convex cannot hold true for algebraic accuracy and cost rates. To this end, it is sufficient to show that the discriminant \mathcal{D}_Q of $Q(N_{\text{hi-fi}})$ cannot be positive. We have

$$\begin{aligned} r(1) &= c_1^{(2)} + c_2^{(2)} \\ r'(1) &= -c_1^{(2)}\alpha^{(2)} + c_2^{(2)}\beta^{(2)} \\ r''(1) &= c_1^{(2)}\alpha^{(2)}(\alpha^{(2)} + 1) + c_2^{(2)}\beta^{(2)}(\beta^{(2)} - 1). \end{aligned}$$

Therefore, for algebraic rates, \mathcal{D}_Q is

$$\begin{aligned} \mathcal{D}_Q &= \left(-c_1^{(2)}\alpha^{(2)} + c_2^{(2)}\beta^{(2)}\right)^2 - 2\left(c_1^{(2)} + c_2^{(2)}\right)\left(c_1^{(2)}\alpha^{(2)}(\alpha^{(2)} + 1) + c_2^{(2)}\beta^{(2)}(\beta^{(2)} + 1)\right) \\ &= -\left(\left(c_1^{(2)}\right)^2 + 2c_1^{(2)}c_2^{(2)}\right)\left(\alpha^{(2)}\right)^2 - \left(\left(c_2^{(2)}\right)^2 + 2c_1^{(2)}c_2^{(2)}\right)\left(\beta^{(2)}\right)^2 - 2c_1^{(2)}c_2^{(2)}\alpha^{(2)}\beta^{(2)} - \\ &\quad \left(\left(c_1^{(2)}\right)^2 + 2c_1^{(2)}c_2^{(2)}\right)\alpha^{(2)} + \left(\left(c_2^{(2)}\right)^2 + 2c_1^{(2)}c_2^{(2)}\right)\beta^{(2)}. \end{aligned}$$

Note that \mathcal{D}_Q is *symmetric* w.r.t. $c_1^{(2)}, c_2^{(2)}, \alpha^{(2)}$ and $\beta^{(2)}$: if $c_1^{(2)}, c_2^{(2)}$ and $\alpha^{(2)}, \beta^{(2)}$ are swapped, \mathcal{D}_Q is unchanged.

Let us assume that $\mathcal{D}_Q \geq 0$ for all $c_1^{(2)}, c_2^{(2)}, \alpha^{(2)}, \beta^{(2)} > 0$. We denote $\underline{c} := \min\{c_1^{(2)}, c_2^{(2)}\}$, $\bar{c} := \max\{c_1^{(2)}, c_2^{(2)}\}$, $\underline{\nu} := \min\{\alpha^{(2)}, \beta^{(2)}\}$ and $\bar{\nu} := \max\{\alpha^{(2)}, \beta^{(2)}\}$. For $\underline{c}, \underline{\nu}$, we have

$$\mathcal{D}_Q = -(\underline{c}^2 + 2\underline{c}^2)\underline{\nu}^2 - (\underline{c}^2 + 2\underline{c}^2)\underline{\nu}^2 - 2\underline{c}^2\underline{\nu}^2 - (2\underline{c}^2 + 2\underline{c}^2)\underline{\nu} + (2\underline{c}^2 + 2\underline{c}^2)\underline{\nu} = -8\underline{c}^2\underline{\nu}^2 < 0, \quad (6.7)$$

that is, \mathcal{D}_Q is negative. By symmetry, \mathcal{D}_Q remains negative also when employing $\bar{c}, \bar{\nu}$, i.e.,

$$\mathcal{D}_Q = -(\bar{c}^2 + 2\bar{c}^2)\bar{\nu}^2 - (\bar{c}^2 + 2\bar{c}^2)\bar{\nu}^2 - 2\bar{c}^2\bar{\nu}^2 - (2\bar{c}^2 + 2\bar{c}^2)\bar{\nu} + (2\bar{c}^2 + 2\bar{c}^2)\bar{\nu} = -8\bar{c}^2\bar{\nu}^2 < 0. \quad (6.8)$$

Thus, (6.7) and (6.8) contradict the assumption that $\mathcal{D}_Q \geq 0$ for all $c_1^{(2)}, c_2^{(2)}, \alpha^{(2)}, \beta^{(2)} > 0$.

To summarize, assuming that $r'''(N_2) \geq 0$ leads to $\mathcal{D}_Q \leq 0$ which satisfies condition (ii) from Theorem 6.2.6. Alternatively, we have that $r'''(N_2) \leq 0$, which is condition (i) from our theorem. In either case, $u(N_2)$ is convex. We hence showed that for any $c_1^{(2)}, c_2^{(2)}, \alpha^{(2)}, \beta^{(2)} > 0$, the objective defined in (6.4) is always convex when the two rate functions are algebraic, which verifies [141, Proposition 3].

In Figure 6.3, we depict the objective $u(N_2)$ constructed using budgets $p = 1, 10, 100$ seconds and the estimated algebraic accuracy and cost rates of the Gaussian process surrogate from Section 6.3.1. We observe that for all budgets, $u(N_2)$ is indeed convex.




Figure 6.3: Objective $u(N_2)$ constructed using budgets $p = 1$ (left), $p = 10$ (center) and $p = 100$ (right), and the estimated algebraic accuracy and cost rates of the Gaussian process surrogate from Section 6.3.1.

Example: Let us now consider the setting in which the accuracy rate is *exponential*, i.e.,

$$r_a^{(2)}(\alpha^{(2)}, N_2) := \exp(-aN_2^{\alpha^{(2)}}),$$

where $a \geq 0$, and the cost rate function is algebraic (6.6), as in the previous example. Working with the exponential accuracy rate $\exp(-aN_2^{\alpha^{(2)}})$ directly in this form is cumbersome due to having the additional parameter a . To simplify things, we use the injectivity of the exponential function N_2^x for $x > 0$ to write

$$\forall N_2 > 0, \forall a \geq 0, \exists \alpha_0 \in \mathbb{R} \text{ s.t. } a = N_2^{\alpha_0}.$$

Thus

$$\exp(-aN_2^{\alpha^{(2)}}) = \exp(-N_2^{\alpha_0} N_2^{\alpha^{(2)}}) = \exp(-N_2^{\alpha_0 + \alpha^{(2)}}) = \exp(-N_2^{\bar{\alpha}^{(2)}}),$$

where $\bar{\alpha}^{(2)} := \alpha_0 + \alpha^{(2)}$ for an $\alpha_0 \in \mathbb{R}$.

In this example, we want to show that $u(N_2)$ in (6.4) is always convex when the accuracy rate is exponential and the cost rate is algebraic. To this end, we employ the same strategy as in the previous example. We have

$$\begin{aligned} r(1) &= c_1^{(2)} e^{-1} + c_2^{(2)} \\ r'(1) &= -c_1^{(2)} \bar{\alpha}^{(2)} e^{-1} + c_2^{(2)} \beta^{(2)} \\ r''(1) &= -c_1^{(2)} \bar{\alpha}^{(2)} (\bar{\alpha}^{(2)} - 1) e^{-1} + c_1^{(2)} (\bar{\alpha}^{(2)})^2 e^{-1} + c_2^{(2)} \beta^{(2)} (\beta^{(2)} - 1). \end{aligned}$$

To simplify the notation further, let $\bar{c}_1^{(2)} := c_1^{(2)} e^{-1}$, which leads to

$$\begin{aligned} r(1) &= \bar{c}_1^{(2)} + c_2^{(2)} \\ r'(1) &= -\bar{c}_1^{(2)} \bar{\alpha}^{(2)} + c_2^{(2)} \beta^{(2)} \\ r''(1) &= -\bar{c}_1^{(2)} \bar{\alpha}^{(2)} (\bar{\alpha}^{(2)} - 1) + \bar{c}_1^{(2)} (\bar{\alpha}^{(2)})^2 + c_2^{(2)} \beta^{(2)} (\beta^{(2)} - 1). \end{aligned}$$

Therefore, the discriminant \mathcal{D}_Q becomes

$$\begin{aligned} \mathcal{D}_Q &= (-\bar{c}_1^{(2)} \bar{\alpha}^{(2)} + c_2^{(2)} \beta^{(2)})^2 - 2(\bar{c}_1^{(2)} + c_2^{(2)}) (-\bar{c}_1^{(2)} \bar{\alpha}^{(2)} (\bar{\alpha}^{(2)} - 1) + \bar{c}_1^{(2)} (\bar{\alpha}^{(2)})^2 \\ &\quad + c_2^{(2)} \beta^{(2)} (\beta^{(2)} - 1)) = (\bar{c}_1^{(2)})^2 (\bar{\alpha}^{(2)})^2 - ((c_2^{(2)})^2 + 2\bar{c}_1^{(2)} c_2^{(2)}) (\beta^{(2)})^2 - \\ &\quad 2\bar{c}_1^{(2)} c_2^{(2)} \bar{\alpha}^{(2)} \beta^{(2)} - ((\bar{c}_1^{(2)})^2 + 2\bar{c}_1^{(2)} c_2^{(2)}) \bar{\alpha}^{(2)} + ((c_2^{(2)})^2 + 2\bar{c}_1^{(2)} c_2^{(2)}) \beta^{(2)}. \end{aligned}$$

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

In this setting, \mathcal{D}_Q is nonsymmetric w.r.t. $c_1^{(2)}, c_2^{(2)}, \alpha^{(2)}$ and $\beta^{(2)}$.

As in the previous example, we assume that $\mathcal{D}_Q \geq 0$ for all $\bar{c}_1^{(2)}, c_2^{(2)}, \bar{\alpha}^{(2)}, \beta^{(2)} > 0$. Let again $\underline{c} := \min\{\bar{c}_1^{(2)}, c_2^{(2)}\}$, $\bar{c} := \max\{\bar{c}_1^{(2)}, c_2^{(2)}\}$, $\underline{\nu} := \min\{\bar{\alpha}^{(2)}, \beta^{(2)}\}$ and $\bar{\nu} := \max\{\bar{\alpha}^{(2)}, \beta^{(2)}\}$. For $\underline{c}, \underline{\nu}$, we have that \mathcal{D}_Q is negative, i.e.,

$$\mathcal{D}_Q = \underline{c}^2 \underline{\nu}^2 - (\underline{c}^2 + 2\underline{c}^2) \underline{\nu}^2 - 2\underline{c}^2 \underline{\nu}^2 - (2\underline{c}^2 + 2\underline{c}^2) \underline{\nu} + (2\underline{c}^2 + 2\underline{c}^2) \underline{\nu} = -4\underline{c}^2 \underline{\nu}^2 < 0. \quad (6.9)$$

\mathcal{D}_Q remains negative even when employing $\bar{c}, \bar{\nu}$, i.e,

$$\mathcal{D}_Q = \bar{c}^2 \bar{\nu}^2 - (\bar{c}^2 + 2\bar{c}^2) \bar{\nu}^2 - 2\bar{c}^2 \bar{\nu}^2 - (2\bar{c}^2 + 2\bar{c}^2) \bar{\nu} + (2\bar{c}^2 + 2\bar{c}^2) \bar{\nu} = -4\bar{c}^2 \bar{\nu}^2 < 0. \quad (6.10)$$

Thus, (6.9) and (6.10) contradict the assumption that $\mathcal{D}_Q \geq 0$ for for all $\bar{c}_1^{(2)}, c_2^{(2)}, \bar{\alpha}^{(2)}, \beta^{(2)} > 0$, which means that the objective defined in (6.4) is always convex when the accuracy rate is exponential and the cost rate is algebraic.

We visualize, in Figure 6.4, the objective $u(N_2)$ constructed using budgets $p = 1, 10, 100$ seconds and the estimated exponential accuracy rate and algebraic cost rate of the reduced basis surrogate used in Section 6.3.1. Observe that $u(N_2)$ is convex for all budgets.




Figure 6.4: Objective $u(N_2)$ constructed using budgets $p = 1$ (left), $p = 10$ (center) and $p = 100$ (right), and the estimated exponential accuracy rate and the algebraic cost rate of the reduced basis surrogate from Section 6.3.1.

In Theorem 6.2.6, we showed that the objective $u(N_2)$ defined in (6.4) is *convex* if the accuracy and cost rates satisfy certain conditions. When $u(N_2)$ is convex, (6.4) has a unique solution N_2^* in $(1, N_{\text{hi-fi}} - 1)$. Next, we want to show that assuming that additionally N_2^* remains bounded even if the total budget p is infinite.

Theorem 6.2.7 *Assume that $u(N_2)$ in (6.4) is convex. As the budget $p \rightarrow \infty$, that is, as $N_{\text{hi-fi}} \rightarrow \infty$, N_2^* stays bounded, i.e., $N_2^* < \infty$.*

Proof: We prove the theorem by contradiction. We thus assume that N_2^* grows indefinitely with p . Let us denote the dependency of p explicitly by $N_2(p)$. Recall that the accuracy rate, $r_a^{(2)}(\alpha^{(2)}, N_2(p))$, is a monotonically decreasing function, whereas the cost rate function, $r_c^{(2)}(\beta^{(2)}, N_2(p))$, is monotonically increasing. We have

$$\exists p_0 > 1, \exists \delta > 0 \text{ s.t. } \forall p > p_0 : c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2(p)) \geq c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2(p_0)) - \delta.$$

Moreover,

$$\exists N_0 \geq 1 \text{ s.t. } \forall p > p_0 : c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p)) = c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p_0) + N_0) = c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p_0)) + \delta.$$

Thus,

$$\begin{aligned}
u(N_2(p)) &= \frac{1}{N_{\text{hi-fi}} - N_2(p)} (c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2(p)) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p))) \\
&\geq \frac{1}{N_{\text{hi-fi}} - N_2(p_0)} (c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2(p_0)) - \delta + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p_0) + N_0)) \\
&= \frac{1}{N_{\text{hi-fi}} - N_2(p_0)} (c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2(p_0)) - \delta + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p_0)) + \delta) \\
&= \frac{1}{N_{\text{hi-fi}} - N_2(p_0)} (c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2(p_0)) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2(p_0))) = u(N_2(p_0)),
\end{aligned}$$

that is, $u(N_2(p)) \geq u(N_2(p_0))$, $\forall p > p_0$, which therefore contradicts the assumption that the minimum N_2^* of $u(N_2(p))$ grows indefinitely with p . \blacksquare

In Theorem 6.2.7, we showed that the optimum N_2^* to construct and improve the low-fidelity model $f_{N_2}^{(2)}$ is bounded w.r.t. the given computational budget, p . This result can be interpreted in several ways. First, it tells us that there is *quasi-optimal trade-off* between constructing and improving the surrogate and sampling the high- and low-fidelity models. That is, the given budget is quasi-optimally balanced between exploring and exploiting the models. Second, the fact that N_2^* remains bounded even for infinite budgets means that low-fidelity models can be *too accurate* for MFMC sampling. This may seem rather counter-intuitive, especially since in classical traditional model reduction the goal is to obtain accurate surrogates. However, in our context, it is unnecessary to improve the low-fidelity model beyond a certain point, since that extra cost exceeds the benefits of having a more accurate surrogate. We will show what a “too accurate” surrogate means in our numerical results in Section 6.3.

We summarize the steps to perform Context-aware Multifidelity Monte Carlo sampling with one low-fidelity model in Algorithm 6.8. The inputs are the computational budget, p , the high-fidelity model, $f^{(1)}$, and its cost \mathcal{C}_1 , the constants and the accuracy and cost rates of the low-fidelity model, and the input density function, π . Between lines 3 – 5, we explore the low-fidelity model by solving the minimization (6.4) and finding the number of high-fidelity evaluations to construct the surrogate, \bar{N}_2^* . Afterwards, we exploit the surrogate in MFMC sampling (lines 6 – 11). We first compute the squared correlation coefficient and the normalized cost of the low-fidelity model using the corresponding rate functions. We then compute the number of samples and control variate coefficient used in multifidelity sampling as presented in Section 2.8.3 (see [141, 145]). The algorithm ends by computing the context-aware MFMC estimator, $\hat{f}_{\text{CA-MFMC}}^{(1)}$, for the mean of the high-fidelity model at line 11. For the particular case in which the accuracy and cost rates of the low-fidelity model are algebraic, we refer the reader to [141, Algorithm 1].

In [141], it was shown that in the limit case in which $\bar{\mathcal{C}}_2 \approx 0$, that is, when the cost of evaluating the low-fidelity model is insignificant compared to that of evaluating the high-fidelity model, then the resulting MFMC sampling achieves a faster convergence rate than standard MC. Specifically, assuming an algebraic accuracy rate $c_1^{(2)} N_2^{-\alpha^{(2)}}$, $\bar{\mathcal{C}}_2 \approx 0$ implies that the MSE of the MFMC estimator is in $O(N_2^{-1-\alpha^{(2)}})$ (see [141, Corollary 6]). This usually happens in the *pre-asymptotic* regime, for small computational budgets. We expect a similar behaviour for the case in which we have an arbitrary accuracy rate, that is, we expect the MSE of our context-aware MFMC algorithm to decay faster than $O(N_2^{-1})$ for small computational budgets.

Algorithm 6.8 Context-aware Multifidelity Monte Carlo sampling with one surrogate

- 1: **procedure** CONTEXTAWAREMFMC1D($p, f^{(1)}, \mathcal{C}_1, c_1^{(2)}, c_2^{(2)}, \alpha^{(2)}, \beta^{(2)}, \boldsymbol{\pi}$)
- 2: $N_{\text{hi-fi}} := \lceil p/\mathcal{C}_1 \rceil$
- 3: Minimize

$$u(N_2) = \frac{1}{N_{\text{hi-fi}} - N_2} \left(c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2) \right)$$

to find N_2^* in $(1, N_{\text{hi-fi}} - 1)$

- 4: $\bar{N}_2^* := \lceil N_2^* \rceil$
- 5: Construct low-fidelity model $f_{\bar{N}_2^*}^{(2)}$ from \bar{N}_2^* evaluations of $f^{(1)}$ ▷ Exploration of $f_{\bar{N}_2^*}^{(2)}$
- 6: Estimate the squared correlation coefficient $\bar{\rho}_2^2$ from Assumption 6.2

$$\bar{\rho}_2^2 = 1 - c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, \bar{N}_2^*)$$

- 7: Estimate the normalized cost of evaluating $f_{\bar{N}_2^*}^{(2)}$ from Assumption 6.3

$$\bar{\mathcal{C}}_2 = c_2^{(2)} r_c^{(2)}(\beta^{(2)}, \bar{N}_2^*)$$

- 8: Determine the number of samples \bar{M}_1, \bar{M}_2 and CV coefficient $\bar{\gamma}_2$ as in MFMC [145]

$$r = \sqrt{\frac{\bar{\rho}_2^2}{\bar{\mathcal{C}}_2(1 - \bar{\rho}_2^2)}}, \quad \bar{M}_1 = \left\lceil \frac{N_{\text{hi-fi}} - \bar{N}_2^*}{1 + \bar{\mathcal{C}}_2 r} \right\rceil, \quad \bar{M}_2 = \lceil r \bar{N}_1 \rceil, \quad \bar{\gamma}_2 = \bar{\rho}_2$$

- 9: Draw \bar{M}_2 realizations $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_{\bar{M}_2}$ of the stochastic input $\boldsymbol{\theta} \sim \boldsymbol{\pi}$
- 10: Compute the Monte Carlo estimators ▷ Exploitation of $f_{\bar{N}_2^*}^{(2)}$

$$\hat{f}_{\bar{M}_1}^{(1)} = \frac{1}{\bar{M}_1} \sum_{n=1}^{\bar{M}_1} f^{(1)}(\boldsymbol{\theta}_n), \quad \hat{f}_{\bar{M}_2}^{(2)} = \frac{1}{\bar{M}_2} \sum_{n=1}^{\bar{M}_2} f_{\bar{N}_2^*}^{(2)}(\boldsymbol{\theta}_n), \quad \hat{f}_{\bar{M}_1}^{(2)} = \frac{1}{\bar{M}_1} \sum_{n=1}^{\bar{M}_1} f_{\bar{N}_2^*}^{(2)}(\boldsymbol{\theta}_n)$$

- 11: Compute the context-aware MFMC estimator

$$\hat{f}_{\text{CA-MFMC}}^{(1)} = \hat{f}_{\bar{M}_1}^{(1)} + \hat{\gamma}_2(\hat{f}_{\bar{M}_2}^{(2)} - \hat{f}_{\bar{M}_1}^{(2)})$$

- 12: **return** $\hat{f}_{\text{CA-MFMC}}^{(1)}$
 - 13: **end procedure**
-

Multiple context-aware low-fidelity models

We want to generalize the results from the previous section to the case in which we have $m \geq 3$ models in total, i.e., at least two-fidelity models. To this end, we again split the given budget p , which allows for $N_{\text{hi-fi}}$ high-fidelity evaluations, between *constructing* and *improving* the surrogates (offline stage) and multifidelity sampling (online phase). The multivariate objective function used to find the number of high-fidelity evaluations to construct the surrogates reads

$$u^{m-1}(N_2, \dots, N_m) = \frac{1}{N_{\text{hi-fi}} - N_{\text{lo-fi}}^{m-1}} \left(c_1^{(2)} r_a^{(2)} + \sum_{j=2}^{m-1} c_2^{(j)} r_c^{(j)} \left(c_1^{(j+1)} r_a^{(j+1)} - c_1^{(j)} r_a^{(j)} \right) + c_2^{(m)} r_c^{(m)} \right),$$

where $N_{\text{lo-fi}}^{m-1} := \sum_{j=2}^m N_j$ denotes the total number of high-fidelity evaluations to construct and improve the $m-1$ surrogates. We used the superscript $m-1$ to explicitly show the dependency on $m-1$ surrogates.

To find $(N_2^*, N_3^*, \dots, N_m^*)$, we need therefore to solve the following minimization problem:

$$\begin{aligned} \min_{N_2, N_3, \dots, N_m \in \mathbb{N}^{m-1}} \quad & u^{m-1}(N_2, N_3, \dots, N_m) \\ \text{subject to} \quad & 2 \leq N_j \leq \frac{N_{\text{hi-fi}} - 1}{m-1}, \quad j = 2, 3, \dots, m. \end{aligned} \tag{6.11}$$

The right bound of the constraints in (6.11) was chosen to ensure that

$$N_{\text{hi-fi}} - \sum_{j=2}^m N_j^* \geq 1,$$

that is, at least one high-fidelity evaluation is possible in the online, multifidelity sampling stage. As for the case with only one surrogate, we relax the minimization (6.11), we solve it for $(N_2, N_3, \dots, N_m) \in (1, N_{\text{hi-fi}}/(m-1))^{m-1} \subset \mathbb{R}^{m-1}$, and then we round the solution up, i.e.,

$$(\bar{N}_2^*, \bar{N}_3^*, \dots, \bar{N}_m^*) := ([N_2^*], [N_3^*], \dots, [N_m^*]) \in \mathbb{N}^{m-1}.$$

Our goal is to extend Theorems 6.2.6 and 6.2.7 to the general case in which we have $m \geq 3$ models in total, i.e., we want to find the conditions for which the objective u^{m-1} in (6.11) is *convex* and the associated minimum is *bounded* for an prescribed budget, p . However, finding these conditions, and especially those for convexity, is not easy mainly because of the nonlinear combination, $c_2^{(j)} r_c^{(j)} \left(c_1^{(j+1)} r_a^{(j+1)} - c_1^{(j)} r_a^{(j)} \right)$, appearing in (6.11).

In a first step, we can nevertheless relax the multidimensional minimization (6.11) by solving a *sequence* of one-dimensional minimizations, for which we can make use of Theorems 6.2.6 and 6.2.6 for the case with only one surrogate. Recall that in the MSE of the MFMC estimator, the low-fidelity models are ordered w.r.t. their Pearson correlation coefficients (see Section 2.8.3). Hence, the surrogates in (6.11) are ordered by their accuracy rate function as well. This ordering naturally allows a *sequential* treatment of the low-fidelity models, as follows.

For $m = 2$, that is, when the model hierarchy comprises the high-fidelity model and the most accurate surrogate, $f_{N_2}^{(2)}$, we use Theorems 6.2.6 and 6.2.7 to respectively ascertain the convexity of the objective

$$u^1(N_2) = \frac{1}{N_{\text{hi-fi}} - N_2} \left(c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) + c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2) \right), \tag{6.12}$$

and the boundness of its minimum, N_2^* . Assuming the conditions in the two theorems are verified, we compute N_2^* and $\bar{N}_2^* := [N_2^*]$. Let us denote the budget in seconds associated to

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

$[N_2^*]$ by $p_{\bar{N}_2^*}$. Once \bar{N}_2^* is computed, the remaining budget is $\bar{p} := p - p_{\bar{N}_2^*}$, with which we can perform $\bar{N}_{\text{hi-fi}} := N_{\text{hi-fi}} - \bar{N}_2^*$ high-fidelity evaluations. When the second low-fidelity model, $f_{N_3}^{(3)}$, is added in the multifidelity model hierarchy, i.e., when m is increased to three, we make use of the previously computed \bar{N}_2^* and define the *sequentially updated* objective as

$$u^1(N_3) := u^2(\bar{N}_2^*, N_3) = \frac{1}{\bar{N}_{\text{hi-fi}} - N_3} \left(\bar{c}_1^{(3)} r_a^{(3)}(\alpha^{(3)}, N_3) + c_2^{(3)} r_c^{(3)}(\beta^{(3)}, N_3) + \bar{c}_{\bar{N}_2^*} \right),$$

where

$$0 < \bar{c}_1^{(3)} := c_1^{(3)} c_2^{(2)} r_c^{(2)}(\beta^{(2)}, \bar{N}_2^*), \quad \bar{c}_{\bar{N}_2^*} := c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, \bar{N}_2^*) \left(1 - c_2^{(2)} r_c^{(2)}(\beta^{(2)}, \bar{N}_2^*) \right).$$

In other words, we *fix* $N_2 := \bar{N}_2^*$ in the bivariate objective, u^2 , which results in a univariate objective depending on N_3 . The resulting objective is analogous to (6.12), but with an additional constant, $\bar{c}_{\bar{N}_2^*}$. Let us have a closer look at $\bar{c}_{\bar{N}_2^*}$. Recall that $r_c^{(2)}$ is the function which gives the evaluation cost of the second low-fidelity model *relative* to the cost of the high-fidelity model. Therefore, in general, $c_2^{(2)} r_c^{(2)}(\beta^{(2)}, \bar{N}_2^*) \leq 1$. Although the alternative is also theoretically possible, we are not considering it here, since in that case, the surrogate is computationally more expensive than the high-fidelity model; in such cases, we can simply discard those low-fidelity models. In addition, if Theorem 6.2.7 holds for \bar{N}_2^* , then \bar{N}_2^* is bounded w.r.t. p which makes $c_2^{(2)} r_c^{(2)}(\beta^{(2)}, \bar{N}_2^*)$ bounded as well. Thus, $c_2^{(2)} r_c^{(2)}(\beta^{(2)}, \bar{N}_2^*) \leq 1$ which implies that $\bar{c}_{\bar{N}_2^*} \geq 0$. We rewrite $u^1(N_3)$ as

$$u^1(N_3) = \frac{1}{\bar{N}_{\text{hi-fi}} - N_3} \left(\bar{c}_1^{(3)} r_a^{(3)}(\alpha^{(3)}, N_3) + c_2^{(3)} r_c^{(3)}(\beta^{(3)}, N_3) \right) + \frac{\bar{c}_{\bar{N}_2^*}}{\bar{N}_{\text{hi-fi}} - N_3}. \quad (6.13)$$

We have that

$$\bar{c}_{\bar{N}_2^*} \geq 0 \Rightarrow \left(\frac{\bar{c}_{\bar{N}_2^*}}{\bar{N}_{\text{hi-fi}} - N_3} \right)'' = \frac{\bar{c}_{\bar{N}_2^*}}{(\bar{N}_{\text{hi-fi}} - N_3)^3} \geq 0,$$

hence the second term in (6.13) is convex. Moreover, when Theorem 6.2.6 is satisfied by the first term in (6.13), the first term in (6.13) is convex as well. Therefore, from Lemma 6.2.3, the objective (6.13) is convex as the summation of two convex functions. Furthermore, since $\bar{c}_{\bar{N}_2^*}$ is a positive constant, bounded w.r.t. p , we can use an analogous argument as in the proof of Theorem 6.2.7 to show that when (6.13) is convex, its minimum is *bounded* w.r.t. \bar{p} . Thus, we compute \bar{N}_3^* , which, provided the above conditions are satisfied, exists, is unique and is bounded w.r.t. \bar{p} .

In general, assuming we sequentially computed $\bar{N}_2^*, \bar{N}_3^*, \dots, \bar{N}_{m-1}^*$ for the first $m-2$ surrogates, to find the minimum for the m th surrogate, \bar{N}_m^* , we minimize

$$\begin{aligned} u^1(N_m) &:= u^{m-1}(\bar{N}_2^*, \bar{N}_3^*, \dots, \bar{N}_{m-1}^*, N_m) \\ &= \frac{1}{\bar{N}_{\text{hi-fi}} - N_m} \left(\bar{c}_1^{(2)} r_a^{(m)}(\alpha^{(m)}, N_m) + c_2^{(m)} r_c^{(m)}(\beta^{(m)}, N_m) + \bar{c}_{\bar{N}_2^*, \bar{N}_3^*, \dots, \bar{N}_{m-1}^*} \right), \end{aligned} \quad (6.14)$$

where $\bar{N}_{\text{hi-fi}} := N_{\text{hi-fi}} - \sum_{j=2}^{m-1} \bar{N}_j^*$, $0 < \bar{c}_1^{(m)} := c_1^{(m)} c_2^{(m-1)} r_c^{(m-1)}(\beta^{(m-1)}, \bar{N}_{m-1}^*)$ and

$$\bar{c}_{\bar{N}_2^*, \bar{N}_3^*, \dots, \bar{N}_{m-1}^*} = \sum_{j=2}^{m-1} c_1^{(j)} r_a^{(j)}(\alpha^{(j)}, \bar{N}_j^*) \left(1 - c_2^{(j)} r_c^{(j)}(\beta^{(j)}, \bar{N}_j^*) \right).$$

Using a similar line of argument as above, that is, assuming that $c_2^{(j)} r_c^{(j)}(\beta^{(j)}, \bar{N}_j^*) \leq 1$ for all $j = 2, 3, \dots, m-1$, then if

$$\frac{1}{\bar{N}_{\text{hi-fi}} - N_m} \left(\bar{c}_1^{(2)} r_a^{(m)}(\alpha^{(m)}, N_m) + c_2^{(m)} r_c^{(m)}(\beta^{(m)}, N_m) \right)$$

satisfies Theorem 6.2.6, then the objective (6.14) is convex as well. Furthermore, we can analogously establish the boundness of N_m^* .

In summary, the sequential treatment of the low-fidelity models allows us to compute the number of high-fidelity evaluations to construct and improve the surrogates by minimizing *one-dimensional* objectives. In this way, we can make use of the theorems from the case with only one surrogate to ascertain the convexity of the objective and the boundness of their minima. However, we are interested in computing *directly* the multivariate minimum of the objective (6.11). Finding the conditions for which (i) the objective u^{m-1} in (6.11) is convex and (ii) the minimum of u^{m-1} is bounded w.r.t. to the budget p for an arbitrary $m \geq 3$ is subject to our future research. The model hierarchies used in our numerical results comprise one low-fidelity model, for which we can make use of Theorems 6.2.6 and 6.2.7, or two-fidelity models. In the latter case, we establish the convexity of the objective (6.11) and the boundness of its minimum numerically for the considered budgets. In the follow remark, we show nevertheless that the previously presented sequential construction of the number of high-fidelity evaluations to construct and improve the surrogates can be used, under certain conditions, to establish the boundness of the multivariate minimum of u^{m-1} .

Remark: Assume that the multivariate objective u^{m-1} defined in (6.11) is *convex* and let $(N_2^{*,\text{glob}}, N_3^{*,\text{glob}}, \dots, N_m^{*,\text{glob}})$ denote its minimum. Assume that (i) $c_2^{(j)} r_c^{(j)}(\beta^{(j)}, \bar{N}_j^*) \leq 1$ and (ii) the conditions in Theorems 6.2.6 and 6.2.7 are verified by (6.14) for all $j = 2, 3, \dots, m-1$. We denote by $(N_2^{*,\text{seq}}, N_3^{*,\text{seq}}, \dots, N_m^{*,\text{seq}})$ the minimum obtained via the sequential process outlined above. From the convexity of u^{m-1} , we have that

$$u^{m-1}(N_2^{*,\text{glob}}, N_3^{*,\text{glob}}, \dots, N_m^{*,\text{glob}}) \leq u^{m-1}(N_2^{*,\text{seq}}, N_3^{*,\text{seq}}, \dots, N_m^{*,\text{seq}}). \quad (6.15)$$

Assuming the conditions in Theorems 6.2.6 and 6.2.7 are verified by (6.14), then all components of $(N_2^{*,\text{seq}}, N_3^{*,\text{seq}}, \dots, N_m^{*,\text{seq}})$ are *unique* and *bounded* for any prescribed budget. If

$$(N_2^{*,\text{glob}}, N_3^{*,\text{glob}}, \dots, N_m^{*,\text{glob}}) \leq (N_2^{*,\text{seq}}, N_3^{*,\text{seq}}, \dots, N_m^{*,\text{seq}})$$

then (6.15) implies that $(N_2^{*,\text{glob}}, N_3^{*,\text{glob}}, \dots, N_m^{*,\text{glob}})$ is bounded as well. When

$$(N_2^{*,\text{glob}}, N_3^{*,\text{glob}}, \dots, N_m^{*,\text{glob}}) \geq (N_2^{*,\text{seq}}, N_3^{*,\text{seq}}, \dots, N_m^{*,\text{seq}})$$

we cannot directly use $(N_2^{*,\text{seq}}, N_3^{*,\text{seq}}, \dots, N_m^{*,\text{seq}})$ to establish the boundness of the global minimum. Nevertheless, if we can, for example, find a finite $\kappa > 0$ such that

$$(N_2^{*,\text{glob}}, N_3^{*,\text{glob}}, \dots, N_m^{*,\text{glob}}) \leq (\kappa N_2^{*,\text{seq}}, \kappa N_3^{*,\text{seq}}, \dots, \kappa N_m^{*,\text{seq}})$$

then we show again that the global minimum of (6.11) is bounded.

Next, we present our numerical results.

6.3 Numerical results

We apply the proposed Context-aware multifidelity Monte Carlo algorithm in two test cases. First, we consider, in Section 6.3.1, a heat conduction problem in a 2D spatial domain similar to [154], governed by an elliptic PDE. The heat conductivity field is parametrized by 9 uniformly distributed stochastic parameters. We chose this test case because it allows the construction of low-fidelity models with heterogeneous rates, such as reduced basis, with exponential accuracy and algebraic cost rates, or Gaussian process regression, with algebraic accuracy and cost rates.

In this test case, we compare the proposed approach with equivalent standard MC and static MFMC sampling summarized in Section 2.8.3. In MFMC, the static low-fidelity models were constructed for three different numbers of high-fidelity evaluations. The comparison is made in terms of the estimators' MSE w.r.t. a reference solution. These numerical experiments have been performed on a desktop computer with a four core Intel i7-4790 CPU and 24 GB of RAM.

In Section 6.3.2, we consider a real world problem, the simulation of plasma microturbulence (see Section 3.2). Specifically, we consider the ASDEX Upgrade test case with 12 uncertain similar to Section 4.3.3. However, for a broader overview of this test case, we consider a different output of interest. The low-fidelity model is an interpolation surrogate obtained via the Sensitivity-driven dimension-adaptive sparse grid algorithm in Chapter 4. We compare again, in terms of MSE, our approach with static MFMC in which the sparse grid surrogate was constructed for two different numbers of high-fidelity evaluations. We performed these simulations on Intel Xeon 8160 (SkyLake) nodes on the Marconi supercomputer at the CINECA Supercomputing Center and we employed Slurm for scheduling jobs on the clusters. All numerical experiments in this section have been performed using standard double precision arithmetics.

6.3.1 Heat conduction in a two-dimensional domain

In the first test case, we consider the heat conduction in a solid two-dimensional domain, a so-called *thermal block* problem, governed by the following elliptic PDE:

$$\begin{aligned} -\nabla \cdot k(x, y, \boldsymbol{\theta}) \nabla u(x, y) &= 0 \text{ in } \Omega \\ u(x, y) &= 0 \text{ on } \Gamma_D \\ k(x, y, \boldsymbol{\theta}) \nabla u(x, y) \cdot \mathbf{n} &= q \text{ on } \Gamma_{N,q}, \quad q = 0, 1, \end{aligned} \tag{6.16}$$

where $\Omega = (0, 1)^2 := \bigcup_{i=1}^{d_{\text{sto}}} \Omega_i$ is the spatial domain divided into $d_{\text{sto}} := B_1 \times B_2$ non-overlapping vertical and horizontal blocks, Ω_i with $i = 1, 2, \dots, d_{\text{sto}}$, Γ_D is a Dirichlet boundary, $\Gamma_{N,0}$ and $\Gamma_{N,1}$ are Neumann boundaries, and

$$k(x, y, \boldsymbol{\theta}) = \sum_{i=1}^{d_{\text{sto}}} \theta_i \chi_{\Omega_i}(x, y)$$

is the heat conductivity field parametrized by $\boldsymbol{\theta} := (\theta_1, \theta_2, \dots, \theta_{d_{\text{sto}}})$, where $\chi_{\Omega_i}(x, y)$ is the indicator function on $\Omega_i(x, y)$. This setup is a slight modification of the thermal block problem considered by Patera and Rozza in [154].

In our numerical experiments, $B_1 = B_2 = 3$, i.e., we divide the 2D problem domain into $3 \times 3 = 9 := d_{\text{sto}}$ blocks. Moreover, the heat conductivity field, $k(x, y, \boldsymbol{\theta})$, is parametrized by 9 uniformly distributed parameters in $[1, 10]^9$. We depict, in Figure 6.5, left, the division of Ω into 9 tiles, where we have the Dirichlet boundary, Γ_D , at the top and the Neumann boundaries $\Gamma_{N,0}$ and $\Gamma_{N,1}$ at the bottom and right, respectively. Furthermore, in the right figure, we visualize an example solution of (6.16) for an instance of $\boldsymbol{\theta} \sim U([1, 10]^9)$. We observe that the solution is non-symmetric and has a rather complex structure; this is due to using (i) more than one horizontal blocks and (ii) non-equal values for θ_i (see [10, Chapter 2]). A typical output of interest, which we also consider here, is the mean heat flow at the Neumann boundary $\Gamma_{N,1}$,

$$o(\boldsymbol{\theta}) = \int_{\Gamma_{N,1}} u(x, y, \boldsymbol{\theta}) dx dy. \tag{6.17}$$

We are hereby interested in the following uncertainty propagation problem: given the uncertain input $\boldsymbol{\theta} \sim U([1, 10]^9)$ that parametrizes the heat conductivity field, $k(x, y, \boldsymbol{\theta})$, in (6.16),




Figure 6.5: Left: Two-dimensional heat conductivity domain, Ω , divided into $3 \times 3 = 9$ tiles. Right: example solution $u(x, y, \boldsymbol{\theta})$ for an instance of $\boldsymbol{\theta} \sim U([1, 10]^9)$. Observe the rather complex, nonsymmetric solution structure which is due to using a $B_1 > 1$ and non-equal values of the uncertain parameters θ_i .

we want compute the expected value of the output of interest, $o(\boldsymbol{\theta})$, i.e., $\mathbb{E}[o(\boldsymbol{\theta})]$. To estimate $\mathbb{E}[o(\boldsymbol{\theta})]$, we employ our context-aware multifidelity sampling approach.

The high-fidelity model is a FE discretization of (6.16) using 7200 elements and thus a mesh width $h = \sqrt{2}/60$. The functionality for finding the numerical solution of the high-fidelity model is provided by the MATLAB library `RBmatlab`¹. On one core from an Intel i7-4790 CPU, the average runtime of the high-fidelity model is

$$\mathcal{C}_1 = 0.1150 \text{ sec.}$$

We chose the thermal block problem because it allows the construction of low-fidelity models with *heterogeneous* accuracy and cost rates. It was shown, for example in [154], that intrusive reduced models, such as reduced basis, have *exponential* accuracy rates and *algebraic* cost rates for the thermal block problem. In addition, most non-intrusive low-fidelity models have *algebraic* accuracy and cost rates; as an example, we consider regression with Gaussian processes [150].

Reduced basis low-fidelity model

In the following, we briefly summarize the reduced basis (RB) surrogate. For a more detailed description of this method, we refer the reader to [154] and [10, Chapter 2]. As a side remark, we note that the RB surrogate is the only intrusive approach considered in this work.

RB is based on offline/online decompositions. In the offline stage, a reduced space and an associated basis, the *reduced basis*, are constructed *directly* from N_2 high-fidelity model evaluations $\{f^{(1)}(\boldsymbol{\theta}_1), f^{(1)}(\boldsymbol{\theta}_2), \dots, f^{(1)}(\boldsymbol{\theta}_{N_2})\}$ at N_2 realizations $\{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_{N_2}\}$ of the random input. To construct the reduced basis, we employ a greedy procedure similar to [189]. Initially, a *training set* $\Theta_{\text{train}} = \{\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}, \dots, \boldsymbol{\theta}^{(N_{\text{max}})}\} \subset \mathbf{X}$ comprising N_{max} realizations of the random variable $\boldsymbol{\theta}$ is prescribed. Θ_{train} is chosen such that it covers the stochastic input space, \mathbf{X} , sufficiently well. In our experiments, we set $N_{\text{max}} = 2000$. The greedy procedure used to construct the reduced basis employs a suitably chosen error indicator, $\mathcal{E}(f^{(1)} - f_n^{(2)})$, usually an *upper bound* of the approximation error

$$\|f^{(1)} - f_n^{(2)}\|$$

¹<https://www.morepas.org/software/rbmatlab/0.11.04/doc/index.html>

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

measured in a certain norm $\|\cdot\|$, where $f_n^{(2)}$ is the RB surrogate constructed using n bases. For $n = 1$, the first sample, $\boldsymbol{\theta}_1$, is chosen at random from Θ_{train} . In addition, the corresponding high-fidelity solution, $f^{(1)}(\boldsymbol{\theta}_1)$, and the first reduced basis are computed. For $n \geq 2$, the reduced basis is enriched using the sample from Θ_{train} with the largest $\mathcal{E}(f^{(1)} - f_{n-1}^{(2)})$. The procedure continues until either $n = N_2$ reduced bases are found, N_{max} is reached or $\mathcal{E}(f^{(1)} - f_{N_2}^{(2)})$ falls below an imposed tolerance, ε_{RB} . We prescribe $\varepsilon_{\text{RB}} := 10^{-16}$. The large value of N_{max} and the small $\varepsilon_{\text{RB}} := 10^{-16}$ ensured finding N_2 reduced basis for any value of N_2 in our numerical experiments. For improved numerical stability, the reduced basis is made *orthogonal* by applying the Gram-Schmidt procedure to the set of high-fidelity solutions $\{f^{(1)}(\boldsymbol{\theta}_1), f^{(1)}(\boldsymbol{\theta}_2), \dots, f^{(1)}(\boldsymbol{\theta}_{N_2})\}$. In this work, $\mathcal{E}(f^{(1)} - f_{N_2}^{(2)})$ is an *a posteriori error indicator*, based on the norm of the residual of the weak form of (6.16) depending on $f_n^{(2)}$ (see [10, Chapter 2]). After the reduced basis is found, in the online stage, an evaluation of the RB surrogate entails assembling and solving a dense linear system of equations. Therefore, the RB method pays-off when this system is small-sized, i.e., the dimensionality of the reduced space is small. The functionality for constructing and evaluating the RB surrogate was provided by the MATLAB library `RBMatlab`.

In works such as [154] or [10, Chapter 2], it was showed that for the thermal block problem, the RB surrogate has an exponential accuracy rate,

$$c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) = c_1^{(2)} \exp(-aN_2^{\alpha^{(2)}}).$$

Moreover, since evaluating the RB surrogate entails linear algebra operations such as assembling and solving the reduced, dense system, the cost rate function is algebraic, i.e.,

$$c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2) = c_2^{(2)} N_2^{\beta^{(2)}}.$$

We estimate the constants $c_1^{(2)}, c_2^{(2)}$ and a , and the two rates, $\alpha^{(2)}$ and $\beta^{(2)}$ numerically via regression, as follows. We create RB surrogates with increasing accuracy and cost for $N_2 \in [2, 25]$. For each considered N_2 , we assess the correlation coefficients, ρ_2 , using 1000 MC high- and low-fidelity evaluations. In addition, the evaluation costs $\bar{\mathcal{C}}_2$ of the low-fidelity models relative to \mathcal{C}_1 are determined by averaging the runtime of evaluating the surrogates at 10^5 MC samples. We obtain the estimates

$$c_1^{(2)} = 0.2124, \quad a = 0.1435, \quad \alpha^{(2)} = 1.4170, \quad c_2^{(2)} = 8.8124 \cdot 10^{-5}, \quad \beta^{(2)} = 0.5380.$$

We depict these results in Figure 6.6.

Gaussian process low-fidelity model

The second reduced model considered in the thermal block test case is regression with Gaussian processes (GP) [150]. We provide only a brief overview of the implementation details of this surrogate and refer the reader to, e.g., [149, 150] for a more detailed description of GP.

We use the implementation `libsvm` [33] to construct the GP surrogate model via ε -support-vector machine (ε -SVM) regression with Gaussian kernels. Similar to the RB model, in the offline stage, the GP surrogate is trained using a set of N_2 high-fidelity model evaluations $\{f^{(1)}(\boldsymbol{\theta}_1), f^{(1)}(\boldsymbol{\theta}_2), \dots, f^{(1)}(\boldsymbol{\theta}_{N_2})\}$ at N_2 realizations $\{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_{N_2}\}$ of the random input. The flags for `libsvm` are `-s 3` for the ε -SVM model and `-t 2` for using Gaussian kernels. Moreover, ε is set to $\varepsilon = \varepsilon_{\text{GP}} := 10^{-2}$. The kernel bandwidth and the costs parameters are selected with five-fold cross validation. The stochastic input is first transformed into the unit hypercube $[0, 1]^9$ before the low-fidelity model is trained. We note that GP surrogates with a similar setup have been also considered in previous multifidelity sampling research, see [141, 145].




Figure 6.6: Estimated values for $c_1^{(2)}$, the constant a and the exponential accuracy rate, $\alpha^{(2)}$, (left figure), and $c_2^{(2)}$ and $\beta^{(2)}$, the algebraic cost rate (right figure) for the RB surrogate. The low fidelity model was constructed using N_2 between 2 and 25. Moreover, the correlation coefficients, ρ_2 , were estimated numerically from 1000 MC high- and low-fidelity evaluations. The evaluation costs of the RB surrogates, relative to the cost \mathcal{C}_1 of the high-fidelity model, was determined by averaging its runtime at 10^5 MC samples.

It is known that the accuracy and cost rates of GP surrogates are algebraic, see [150], i.e.,

$$c_1^{(3)} r_a^{(3)}(\alpha^{(3)}, N_3) := c_1^{(3)} N_3^{-\alpha^{(3)}}, \quad c_3^{(3)} r_c^{(3)}(\beta^{(3)}, N_3) := c_2^{(3)} N_3^{\beta^{(3)}}.$$

We estimate $c_1^{(3)}$ and $c_2^{(3)}$, and the two rates $\alpha^{(3)}$ and $\beta^{(3)}$ numerically, via regression, for GP surrogates constructed using $N_3 \in [10^2, 10^5]$. To assess the correlation coefficients, ρ_2 , for all N_3 , we use 1000 MC high- and low-fidelity evaluations. In addition, we average the runtimes of the 1000 surrogate evaluations to assess their costs relative to cost of the high-fidelity model. We obtain the following estimates:

$$c_1^{(3)} = 0.8534, \quad \alpha^{(3)} = 0.4542, \quad c_2^{(3)} = 9.3245 \cdot 10^{-7}, \quad \beta^{(3)} = 0.5695,$$

which are depicted in Figure 6.7.

Next, we apply the proposed context-aware sampling approach and compare it, in terms of MSE, with standard MC and static MFMC. For a more comprehensive overview, we consider two multifidelity scenarios. First, we employ a hierarchy consisting of the high-fidelity model and the RB surrogate. In the second scenario, we augment this hierarchy with the GP surrogate.

Scenario 1: One low-fidelity model

In the first scenario, we consider a multifidelity hierarchy consisting of the high-fidelity model and the RB surrogate. Recall that the accuracy rate of this low-fidelity model is exponential, whereas its cost rate is algebraic. In Section 6.2.3, we showed that for this setup, the objective (6.4) is convex, thus the conditions in Theorems 6.2.6 and 6.2.7 are verified. In Figure 6.8, we depict the minima of (6.4) for all budgets $p \in [10^0, 10^3]$ seconds. We observe that as the theory has predicted, \bar{N}_2^* is bounded with p . Moreover, due to the high accuracy of the surrogate, \bar{N}_2^* is at most 19. In other words, our context-aware algorithm needs at most 19 high-fidelity model evaluations to construct and improve the RB surrogate.

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING




Figure 6.7: Estimated values for $c_1^{(3)}$ and the algebraic accuracy rate, $\alpha^{(3)}$, (left figure), and $c_2^{(3)}$ and $\beta^{(3)}$, the algebraic cost rate (right figure) for the GP surrogate constructed using $N_3 \in [10^2, 10^5]$ high-fidelity model evaluations. Moreover, the correlation coefficients, ρ_3 , and the evaluation costs of the GP surrogates relative to \mathcal{C}_1 , were estimated numerically using 1000 MC samples.

We apply the proposed context-aware multifidelity sampling algorithm (Algorithm 6.8) to estimate the expectation of the output of interest (6.17). We consider budgets p between 1 and 1000 seconds. The values of \bar{N}_2^* used in the offline stage are depicted in Figure 6.8. In Figure 6.9, we depict the split of the total budget between the offline and the online stages in our approach. Observe that relative to the total budget, the boundness of \bar{N}_2^* w.r.t. p implies that the effort invested in the offline stage decreases with p . Moreover, in the online, multifidelity sampling stage of the algorithm, more than 99.5% of the total number of samples are used to compute the MC estimators of the low-fidelity model, i.e., the high-fidelity model is evaluated for only 0.5% of the total number of samples. In contrast, in a standard MC estimator, the high-fidelity model is employed for *all* samples. We compare our approach with (i) standard MC in which




Figure 6.8: The minimum \bar{N}_2^* of (6.4) for budgets $p \in [10^0, 10^3]$ seconds for the RB model.




Figure 6.9: Budget split between constructing and improving the surrogate (offline stage) and sampling the high- and low-fidelity models (online stage) in the thermal block problem in which we employ the high-fidelity model and the RB surrogate.

we use only the high-fidelity model and the number of samples, $N_{\text{hi-fi}}$, allowed by each p , and (ii) static MFMC. For the static MFMC, we assume we are given the RB surrogate for three values of N_2 . First, we consider two values that are *smaller* than the *largest* \bar{N}_2^* in Figure 6.8, i.e., $N_2 = 2$ and $N_2 = 8$. In addition, we also consider an MFMC setup in which a larger number of bases, $N_2 = 50$, was used to construct the surrogate. For a fair comparison between our approach and MFMC, the sampling budget in MFMC is $p - p_{N_2}$, that is, we subtract the budget p_{N_2} corresponding to N_2 high-fidelity evaluations that were necessary to construct the RB surrogate. The comparison between all approaches is done in terms of the MSE of their estimated expectations w.r.t. a reference solution. The reference solution is computed using standard MC sampling with a budget $p_{\text{ref}} = 10^4$ seconds, i.e., 86957 high-fidelity evaluations. In addition, all results are averaged over 50 runs. The reference expectation reads

$$\hat{f}_{\text{ref}}^{(1)} = 0.2206.$$

The MSE of all employed sampling approaches are depicted in Figure 6.10. First, let us look at the results of standard MC and static MFMC. We see that, as expected, the corresponding MSE decrease linearly with p . In addition, the variance reduction of the MFMC estimator leads to a smaller MSE by at least one order of magnitude: even a RB surrogate with two bases is sufficient to decrease the variance of the standard MC estimator. Let us now look closer at the MSE of our context-aware multifidelity sampling algorithm. We see that our approach is the most accurate: for larger budgets, we decrease the variance of standard MC by about four orders of magnitude, and our results are at least half an order of magnitude more accurate than the results yielded by all three MFMC setups. In addition, as we discussed in Section 6.2.3 and similar to what Peherstorfer observed for algebraic accuracy and cost rates in [141], we also see in Figure 6.10 that for small budgets ($p \leq 10$), the MSE of our approach decreases *exponentially* fast. Finally, our context-aware approach, which requires at most 19 reduced basis to construct and improve the surrogate, is more accurate than the static MFMC in which the RB surrogate was constructed using $N_2 = 50$ reduced bases. This underlines two of the main properties of the proposed algorithm: (i) there is a quasi-optimal trade-off between improving the surrogates and sampling them, that is to say, there is no need to improve reduced models indefinitely, and (ii) low-fidelity models can be too accurate for MFMC sampling. In a classical model reduction




Figure 6.10: Comparison between the MSE of our context-aware approach (CA-MFMC) in which we employ the RB surrogate, and standard MC sampling and static MFMC. In static MFMC, we assume that the RB surrogate was constructed using $N_2 = 2, 8$ or 50 reduced bases. All results were averaged over 50 runs.

setting, $N_2 = 50$ reduced bases would certainly lead to a more accurate reduced model than using only $N_2 = 19$ bases. In contrast, in multifidelity sampling methods, we seek low-fidelity models with an accuracy/cost ratio that leads to the highest variance reduction. In most cases, these surrogates do not actually have to be very accurate in the classical model reduction sense.

Scenario 2: Two low-fidelity models

In the second scenario, we augment the previously considered bifidelity hierarchy with the GP low-fidelity model. Therefore, we have two low-fidelity models. The first surrogate, the RB model, has a fast decaying exponential accuracy rate function, whereas the GP model has a rather slowly decreasing algebraic rate function. However, the GP surrogate is about 100 times cheaper to evaluate than the RB model. Therefore, we have a scenario with a very accurate first surrogate and a second, less accurate, but significantly cheaper low-fidelity model. We note that in a classical model reduction setting, the GP surrogate in this form would not be sufficiently accurate to replace the high-fidelity model.

We consider budgets $p \in [10^1, 10^3]$ seconds. For each budget, we ascertain numerically the convexity of the objective (6.11) for $m = 3$, in which we plug the rates of the RB and GP surrogates. We conclude that the bivariate objective is indeed convex. For example, we visualize, in Figure 6.11, the bivariate objectives corresponding to budgets $p = 10$ (top, left) and $p = 50$ (bottom, left), which we see that are convex. Moreover, we also depict the minima corresponding to the two budgets in the right figures. Therefore, the bivariate minimum of (6.11) exists and is unique. Additionally, we illustrate, in Figure 6.12, the sequentially updated one-dimensional objectives discussed in Section 6.2.3, for the same budgets as in Figure 6.11. In the left figures, we depict the 1D objective (6.4) corresponding to the RB surrogate, whose




Figure 6.11: Bivariate objective (6.11) in which we use the rates of the RB and GP surrogates for budgets $p = 10$ seconds (top figures) and $p = 50$ seconds (bottom figures).




Figure 6.12: On the left, we illustrate the 1D objective (6.4) corresponding to the RB surrogate, whose minimum is $\bar{N}_2^{*,seq}$. In the right figures, we depict the sequentially updated objective (6.13) depending on $\bar{N}_2^{*,seq}$, with minimum is $\bar{N}_3^{*,seq}$. In the top figures, the budget is $p = 10$ seconds and in the bottom figures, $p = 50$ seconds.

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

minimum is $\bar{N}_2^{*,seq}$; from Theorems 6.2.6 and 6.2.7, $\bar{N}_2^{*,seq}$ is bounded w.r.t. p . In the right figures, we depict the sequentially updated objective (6.13) corresponding to the RB surrogate with minimum is $\bar{N}_3^{*,seq}$; recall that we use the superscript *seq* to refer to the fact that the minimum corresponds to sequentially updated objectives. This objective is convex since the cost rate function of the RB model is bounded above by one. Therefore, its minimum, $\bar{N}_3^{*,seq}$, is unique and bounded w.r.t. any budget.

We next compute the global minimum, $(\bar{N}_2^{*,glob}, \bar{N}_3^{*,glob})$, and compare it with $(\bar{N}_2^{*,seq}, \bar{N}_3^{*,seq})$. We depict the obtained results in Figure 6.13 for budgets $p \in [10^1, 10^5]$; we considered budgets larger than 10^3 seconds to better illustrate the trend of the minima. We observe that (i) the




Figure 6.13: Global and sequentially updated minima of the bivariate objective (6.11) depending on the rates of the RB and GP surrogates. We see that although $(\bar{N}_2^{*,seq}, \bar{N}_3^{*,seq})$ is a lower bound of $(\bar{N}_2^{*,glob}, \bar{N}_3^{*,glob})$, multiplying the sequentially computed optima by only 1.2 provides an upper bound of the global minimum, thus showing that $(\bar{N}_2^{*,glob}, \bar{N}_3^{*,glob})$ w.r.t. p . To better illustrate the trend of the minima, we considered budgets $p \in [10^1, 10^5]$ seconds.

global minimum is bounded w.r.t. the budget, p , and (ii) the sequentially updated minima are slightly *smaller* than the global optima. As discussed in Section 6.2.3, when the minimum corresponding to the sequentially updated objectives is smaller than the global minimum, we can, for example, multiply the sequential minimum by a finite number, κ , to obtain an upper bound of $(\bar{N}_2^{*,glob}, \bar{N}_3^{*,glob})$. Here, we choose $\kappa = 1.2$ and see that indeed $(1.2 \cdot \bar{N}_2^{*,seq}, 1.2 \cdot \bar{N}_3^{*,seq})$ is an upper bound of $(\bar{N}_2^{*,glob}, \bar{N}_3^{*,glob})$, i.e., $(\bar{N}_2^{*,glob}, \bar{N}_3^{*,glob})$ is bounded w.r.t. p .

In the offline stage of the proposed Context-aware multifidelity Monte Carlo algorithm, we use the global minima illustrated in Figure 6.11 restricted to budgets $p \in [10^1, 10^3]$ seconds used in the offline stage. We visualize, in Figure 6.14, the split of the total budget between the offline and online stages in our approach. We observe again that the boundness of the bivariate global minimum w.r.t. p leads to a decreasing offline percentage as the budget increases. Moreover, in the online, multifidelity sampling phase, more than 99.9% of the total number of samples are used to compute the MC estimators of the two surrogates.

Finally, in Figure 6.15, we depict the MSE of our context-aware approach employing both surrogates compared to all MSE from Figure 6.10. We see that although the GP surrogate is not very accurate from a classical model reduction perspective, its low evaluation cost leads to further variance reduction when used in our context-aware sampling approach. Thus, the most accurate estimate of the expectation of the output of interest (6.17) is our proposed multifidelity method in which we use *both* surrogates.




Figure 6.14: Budget split between constructing and improving the surrogates (offline stage) and sampling the high- and low-fidelity models (online stage) in the thermal block problem in which we employ the high-fidelity model, and the RB and GP surrogates.




Figure 6.15: Comparison between the MSE of our context-aware approach (CA-MFMC) in which we employ either the RB surrogate or both the RB and GP low-fidelity models, and standard MC sampling and static MFMC. In static MFMC, we assume that the RB surrogate was constructed using $N_2 = 2, 8$ or 50 reduced bases. All results were averaged over 50 runs.

To summarize, we considered a thermal block problem in a 2D spatial domain with a heat conductivity field parametrized by nine uniformly distributed stochastic parameters in $[1, 10]^9$. In the proposed context-aware multifidelity algorithm, we employed two low-fidelity models, a RB surrogate having exponential accuracy rate and algebraic cost rate, and a GP low-fidelity model with both rates algebraic. For a more comprehensive overview, we considered two scenarios, one in which the multifidelity hierarchy comprised the high-fidelity model and the RB surrogate, and another in which we employ the high-fidelity model and both the RB and GP low-fidelity models. We showed that compared to either standard MC or static MFMC sampling in which the low-fidelity model has been constructed using either fewer or more reduced basis than in the context-aware approach, our algorithm is the most accurate. In this way, we underlined that (i) there is a trade-off between exploring and exploiting low-fidelity models and (ii) surrogates can be too accurate for multifidelity sampling algorithms.

6.3.2 ASDEX Upgrade experiment

Next, we employ the proposed approach in a real-world application, the simulation of plasma microturbulence. To the best of our knowledge, this is one of the first times when multifidelity methods are employed to quantify uncertainty in such problems. We also mention the recent Bachelor thesis [105] developed within our group in which the static MFMC algorithm was applied to the CBC benchmark considered in Section 4.3.1. Therein, the Sensitivity-driven dimension-adaptive sparse grid approximation and a neuronal network have been used as low-fidelity models and the results showed that static MFMC is orders of magnitude more efficient than standard MC sampling.

In Chapter 4, we performed a comprehensive uncertainty propagation study in two plasma microturbulence test cases, the Modified Cyclon Base Case (see Section 4.3.1) and the ASDEX Upgrade test case; recall Sections 4.3.2 and 4.4.1. Therein, the output of interest was either the *computed* or the *underlying dominant amplitude* eigenmode. For a broader overview of the ASDEX Upgrade test case, we employ the proposed context-aware sampling approach in the scenario with 12 uncertain inputs using the same setup as in Section 4.3.4. Here, the output of interest is the *underlying dominant frequency eigenmode* determined similarly to the underlying amplitude mode in Section 4.4.1. Therefore, the high-fidelity model, $f^{(1)}$, is the eigenvalue solver from GENE discretized using (recall Section 3.2)

$$N_x = 15, \quad N_y = 1, \quad N_z = 24, \quad N_{v_{\parallel}} = 48, \quad N_{\mu_m} = 16.$$

Here, we restrict our attention to one value of the perpendicular wave number, i.e., $k_y \rho_s = 0.8$.

The context-aware multifidelity algorithm depends on evaluating the high- and low-fidelity models at samples from the stochastic input space, \mathbf{X} . To determine the dominant underlying frequency eigenmode of the high-fidelity model evaluated at samples $\boldsymbol{\theta}_n \in \mathbf{X}$, we proceed as follows. As in Section 4.4.1, we employ the GENE eigenvalue solver to compute the dominant and the first subdominant frequency eigenmodes, which we denote by $f_1(\boldsymbol{\theta}_n)$ and $f_2(\boldsymbol{\theta}_n)$. We visualize, in Figure 6.16, left, the dominant (black dots) and the first subdominant (red squares) frequency eigenmodes at 1000 MC samples; we later use these evaluations to assess the correlation coefficient between the high- and low-fidelity model and the runtime of the high-fidelity model. Moreover, we also depict the horizontal lines that pass through the two computed modes of the *first* sample. We observe that we have both negative frequencies, corresponding to electron-driven modes, such as TEM/ETG, and positive frequencies, associated to ion-driven modes, e.g., ITG. In the numerical experiments in Sections 4.3.2 and 4.4.1, we saw that at $k_y \rho_s = 0.8$, the underlying mode is *electron-driven*, hence the frequency is *negative* (see, for example, Figure 4.18). We thus determine the underlying frequency mode, which we denote




Figure 6.16: ASDEX Upgrade test case: dominant and first subdominant frequency eigenmodes corresponding to $k_y \rho_s = 0.8$ at 1000 MC samples. Left: computed dominant and first subdominant mode. We observe an overlap between the two modes, i.e., between electron-driven (negative frequency) and ion-driven (positive frequency) modes. Right: tracing the dominant electron mode using (6.18). With this procedure, the electron-driven mode (black disks) is separated from the ion-driven subdominant mode (red squares) for 97.9% of the samples.

by $f_p(\boldsymbol{\theta}_n)$, corresponding to sample $\boldsymbol{\theta}_n$ as follows. When $n = 1$, the underlying mode is the computed *negative* frequency. For example, in Figure 6.16, left, this is the *second* computed eigenmode. For $n = 2, 3, \dots$, we employ the following rule:

$$f_p(\boldsymbol{\theta}_n) = \begin{cases} f_1(\boldsymbol{\theta}_n), & \text{if } |f_1(\boldsymbol{\theta}_n) - f_p(\boldsymbol{\theta}_1)| \leq |f_2(\boldsymbol{\theta}_n) - f_p(\boldsymbol{\theta}_1)| \\ f_2(\boldsymbol{\theta}_n), & \text{else.} \end{cases} \quad (6.18)$$

That is, the underlying dominant mode associated to sample n , $n \geq 2$ is the computed mode with the *smallest* absolute distance to the underlying mode of the first sample. We depict the results of using this procedure in Figure 6.16, right. The underlying mode is correctly determined, except at 21 samples, i.e., 2.1% of the total number of samples. These samples have small associated negative frequency modes which are further away from $f_p(\boldsymbol{\theta}_1)$ than the positive frequency modes. Nevertheless, this small percentage of misclassified modes have no impact on the computation of the correlation coefficient.

Remark: In Section 4.4.1, we determined the underlying dominant mode based on the smallest absolute distance w.r.t. the underlying modes of the backward neighbours. This was possible because the input samples were deterministic, i.e., (L)-Leja points, hence we could define what a backward neighbour is. In contrast, when the input samples are pseudo-random we cannot precisely define what are the backward neighbours of a sample. We point out that for scenarios in which the microinstability modes are more tightly intermixed, a different procedure to find the underlying dominant mode might be preferred.

Sensitivity-driven dimension-adaptive sparse grid low-fidelity model

In Chapter 4, we formulated the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm to assess quantities of interest such as expectation, standard deviation or total Sobol' indices for global sensitivity analysis. Since this algorithm is based on sparse grid interpolation, it can also be used as a reduced model. To further test the capabilities of this

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

algorithm, we thus employ it here as a low-fidelity model in the proposed context-aware sampling algorithm. For simplicity, we compute the sensitivity scores via the strategy in which one score is assessed for all stochastic input interactions (see Section 4.2.3). That is, we assess a total of $12 + 1 = 13$ scores.

In uncertainty propagation, parametrized high-fidelity models generally have a *finite* number of bounded derivatives, which means that interpolation surrogates have *algebraic* accuracy rates, see [181]. In addition, evaluating the sparse grid interpolation surrogate entails assessing the multivariate surpluses in (2.9) for all multiindices in \mathcal{K} , obtained via dimension-adaptivity. Moreover, for interpolation the number of (L)-Leja points is the same as the cardinality of \mathcal{K} (see Section 2.5.2). Thus, the evaluation cost of the surrogate is also algebraic (in fact, we expect it to be linear in N_2). Therefore,

$$c_1^{(2)} r_a^{(2)}(\alpha^{(2)}, N_2) := c_1^{(2)} N_2^{-\alpha^{(2)}}, \quad c_2^{(2)} r_c^{(2)}(\beta^{(2)}, N_2) := c_2^{(2)} N_2^{\beta^{(2)}}.$$

We estimate $c_1^{(2)}$, $c_2^{(2)}$, $\alpha^{(2)}$ and $\beta^{(2)}$ numerically for budgets $p \in [3 \cdot 10^5, 2 \cdot 10^7]$ seconds, as follows. We construct the sensitivity-driven low-fidelity model using a small tolerance, i.e., $\tau^{in-g} = 10^{-12} \cdot \mathbf{1}$, which required a total of only 623 high-fidelity evaluations. As a side remark, this small number of high-fidelity evaluations underlines once again the benefits of using sensitivity-driven adaptivity in problems in which the stochastic inputs are anisotropically coupled and the intrinsic dimensionality is small. We exploit the nestedness of the (L)-Leja points and reuse these high-fidelity evaluations to construct surrogates for tolerances down to $\tau^{in-g} = 10^{-2} \cdot \mathbf{1}$. Afterwards, we assess the correlation coefficients, ρ_2 , for all obtained values of N_2 using 1000 MC samples. To determine the evaluation costs \mathcal{C}_2 of the low-fidelity model relative to \mathcal{C}_1 , we average the runtime of evaluating the surrogates at the same MC 1000 samples; the average runtime of the high-fidelity model is

$$\mathcal{C}_1 = 850.8476 \text{ sec.}$$

We obtain, via regression, the following estimates of the four constants:

$$c_1^{(2)} = 1.0665, \quad \alpha^{(2)} = 1.0584, \quad c_2^{(2)} = 2.3676 \cdot 10^{-5}, \quad \beta^{(2)} = 1.0137,$$

which we also depict in Figure 6.17. We see that the sparse grid surrogate has good approximation properties: for an accuracy in $O(N_2^{-1})$, it requires a cost in $O(N_2)$.

For algebraic accuracy and cost rates, the conditions in Theorems 6.2.6 and 6.2.7 are verified (recall that this was showed in [141] as well), that is, the number \bar{N}_2^* of high-fidelity model evaluations to construct the low-fidelity model is *unique* and *bounded* for any budget, p . We visualize, in Figure 6.18, the values of \bar{N}_2^* for budgets p between $3 \cdot 10^5$ and $2 \cdot 10^7$ seconds. For a total budget of up to 20 million seconds, our context-aware algorithm requires at most only 179 high-fidelity model evaluations to construct and improve the low-fidelity model.

In the offline stage of the context-aware sampling algorithm, the low-fidelity models are constructed and improved using \bar{N}_2^* high-fidelity model evaluations. To ensure that any value of \bar{N}_2^* is feasible, low-fidelity models are generally constructed using point sets with arbitrarily fine granularity. For example, in Section 6.3.1, we employed RB and GP regression surrogates constructed from high-fidelity snapshots at pseudo-random samples, which are arbitrarily granular. Here, in contrast, we employ deterministic (L)-Leja sequences to construct the sparse grid surrogate. Nevertheless, as discussed in Section 2.5.2, (L)-Leja points have fine granularity: for each multiindex $\mathbf{k} \in \mathcal{K}$, only one surplus Leja point is added. In addition, since the sensitivity-driven algorithm exploits the structure of the underlying problem, as more refinement steps are performed, fewer subspaces and thus fewer (L)-Leja points are added per refinement step.




Figure 6.17: Estimated values for $c_1^{(2)}$ and the algebraic accuracy rate, $\alpha^{(2)}$, (left figure), and $c_2^{(2)}$ and $\beta^{(2)}$, the algebraic cost rate (right figure) for the sensitivity-drive dimension-adaptive sparse grid interpolation surrogate. The low-fidelity model was determined using tolerances up to $\tau^{in-g} = 10^{-12} \cdot \mathbf{1}$ and the correlation coefficient, ρ_2 , and the evaluation costs of the surrogate were estimated evaluations at 1000 samples from the 12D stochastic input domain.




Figure 6.18: Estimated value of the quasi-optimal number of high-fidelity model evaluations, \bar{N}_2^* , to construct the sensitivity-driven surrogate in the *offline* phase of the underlying context-aware sampling algorithm. On the x -axis, we have the number of high-fidelity evaluations, $N_{\text{hi-fi}}$, corresponding to budgets p between $3 \cdot 10^5$ and $2 \cdot 10^7$ seconds. \bar{N}_2^* is bounded by 179.

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

In our numerical experiments, we were able to construct the low-fidelity models using exactly \bar{N}_2^* high-fidelity evaluations for any value of \bar{N}_2^* .

Remark: When low-fidelity models are constructed using high-fidelity evaluations at nested, deterministic samples, as it is the case for (L)-Leja points, the high-fidelity evaluations from previous steps can be reused as \bar{N}_2^* increases.

We employ Algorithm 6.8 to estimate the expectation of the electron-driven dominant frequency mode. The considered budgets are between $3 \cdot 10^5$ and $2 \cdot 10^7$ seconds. We compare our context-aware algorithm with two static MFMC setups, one in which the sparse grid surrogate was constructed using $N_2 = 49$, i.e., the tolerance $\tau^{in-g} = 10^{-6} \cdot \mathbf{1}$, and another in which $N_2 = 623$, yielded by $\tau^{in-g} = 10^{-12} \cdot \mathbf{1}$. Note that the first N_2 is smaller while the second N_2 is larger than the largest \bar{N}_2^* in our context-aware algorithm (see Figure 6.18). The reference value for the expectation of the electron-driven dominant frequency mode is

$$\hat{f}_{\text{ref}}^{(1)} = -0.5648,$$

computed using static sparse grid interpolation of level $K_{\text{max}}^{\text{in-g}} = 6$ comprising 6686 (L)-Leja points. For a fair comparison between the context-aware algorithm and the static MFMC, the sampling budget for MFMC is $p - p_{N_2}$ for any considered p , where p_{N_2} is the budget corresponding to the N_2 high-fidelity evaluations used to construct the surrogate.

The values of \bar{N}_2^* used in the offline stage of the context-aware algorithms are illustrated in Figure 6.18. We visualize, in Figure 6.19, the split of the total budget, p , between the offline and the online stages in our algorithm. We see that as the budget increases, the percentage for the offline stage decreases significantly, which is again due to \bar{N}_2^* being bounded w.r.t. p (Theorem 6.2.7). Moreover, in the online stage, more than 99.5% of the total number of samples are used in the MC estimators of the low-fidelity model. We note that the low-fidelity model is up to 3131 times cheaper to evaluate than the high-fidelity model.




Figure 6.19: Budget split between constructing and improving the sensitivity-driven dimension-adaptive surrogate (offline stage) and sampling the high- and low-fidelity models (online stage) in the ASDEX Upgrade scenario with 12 uncertain inputs.

In Figure 6.20, we visualize the estimated MSE of the proposed context-aware approach compared with static MFMC the two static MFMC setups, one in which the sparse grid surrogate was constructed using $N_2 = 49$, and another in which $N_2 = 623$ high-fidelity evaluations were used to construct the low-fidelity model. We observe that our approach is the most accurate. That is, the context-aware multifidelity algorithm in which at most $\bar{N}_2^* = 179$ high-fidelity

model evaluations were used to construct and improve the surrogate is more accurate than static MFMC in which the surrogate was constructed using $N_2 = 49$ high-fidelity evaluations, which is not surprising, but also than static MFMC in which $N_2 = 623$. Therefore, we see once again that surrogates can be too accurate for multifidelity sampling purposes and that there is a trade-off between exploring and exploiting low-fidelity models in multifidelity sampling. In addition, we also see that, as it was shown in [141], for budgets $p \leq 10^5$ seconds, the MSE of our approach decreases exponentially fast.




Figure 6.20: Comparison between the MSE of our context-aware approach (CA-MFMC) in which we employ the dimension-adaptive sparse grid surrogate, and standard MC sampling and static MFMC. In static MFMC, we assume that the sparse grid surrogate was constructed using $N_2 = 49$ or 623 high-fidelity evaluations. The sampling results were averaged over 10 runs.

In summary, we applied the proposed context-aware sampling approach to a real-world problem, the simulation of plasma microturbulence. Specifically, we considered the ASDEX Upgrade scenario with 12 uncertain inputs, also considered in Section 4.3.4. We considered a scenario with one low-fidelity model, which was the Sensitivity-driven dimension-adaptive sparse grid algorithm from Chapter 4, having algebraic accuracy and cost rates. The estimates of the rates show that this surrogate has good approximation properties. Our proposed context-aware multifidelity approach was more accurate than a static MFMC estimate in which both fewer and more high-fidelity evaluations have been used to construct the static surrogate. We therefore showed again that low-fidelity models can be too accurate for multifidelity sampling and that there is a trade-off between exploring and exploiting low-fidelity models in multifidelity sampling approaches.

Conclusion and remarks

We presented a Context-aware multifidelity Monte Carlo sampling algorithm for quantifying uncertainty in higher-dimensional, computationally expensive real-world applications. The major novelty of the proposed approach was that low-fidelity models were explicitly constructed

CHAPTER 6. CONTEXT-AWARE MODEL REDUCTION FOR MULTIFIDELITY MONTE CARLO SAMPLING

for being together with the high-fidelity model to compute statistics of high-fidelity model outputs. This is especially the case for data-fit or projection-based reduced models, which are constructed from sets of high-fidelity model evaluations. To this end, we determined an upper bound of the MSE of the standard multifidelity sampling estimator depending on the accuracy and cost rates of the surrogates. By minimizing this upper bound, we found the quasi-optimal number of high-fidelity evaluations to construct and improve the surrogates. In this way, the given computational budget was split between exploring the low-fidelity models (offline stage) and exploiting them, together with the high-fidelity model in the online, multifidelity sampling stage. In the proposed approach, low-fidelity models with heterogeneous accuracy and cost rates, such as exponential, algebraic or geometric can be considered; in our numerical examples, we had surrogates with exponential and algebraic accuracy rates, and algebraic cost rates.

For simplicity, we first looked at the bifidelity case, that is, the model hierarchy comprising the high- and one low-fidelity model. We found the conditions under which the objective used to find the number of high-fidelity evaluations to construct and improve the surrogate is convex. When these conditions hold true, the minimizer thus exists and is unique. We showed that these conditions are always satisfied for algebraic rates and for exponential accuracy rate and algebraic cost rate. We also proved that this minimizer is bounded w.r.t. the computational budget, which means that there is a trade off between exploring and exploiting the surrogates, i.e., low-fidelity models can be too accurate for multifidelity sampling methods, which is in contrast with standard model reduction. We afterwards considered the case with an arbitrary number of low-fidelity models. We showed how a sequential treatment of the multivariate objective used to construct and improve the surrogates allows us to reuse the results from the bifidelity case, which means that the sequentially obtained minimizer (i) exists, (ii) is unique and (iii) is bounded w.r.t. any given budget. Finding the conditions under which the multivariate objective is convex and its minimum bounded w.r.t. the computational budget is subject to our future research.

To test the capabilities of the new approach, we considered two test cases, a thermal block problem with a 2D spatial domain, with nice uncertain inputs, and the plasma microturbulence test case with 12 uncertain parameters considered in Section 4.3.4 as well. In the first test case, two surrogates were used, a reduced basis low-fidelity model, which has an exponential accuracy rate and an algebraic cost rate, and a Gaussian process regression surrogate with both rates algebraic. The first low-fidelity model was very accurate, but about 100 times more expensive to evaluate than the less accurate Gaussian process surrogate. We considered initially the bifidelity setting, in which the reduced basis surrogate was employed, and then we looked at the trifidelity setting. For the latter, we ascertained numerically the convexity of the objective used to find the number of high-fidelity model evaluations from which the surrogates are constructed and the boundness of its global minimum. In both settings, we compared our approach, in terms of MSE, with standard MC and static MFMC with the RB surrogate which was constructed using either fewer or more high-fidelity evaluations than were used in our context-aware approach. Our results clearly showed that our approach is the most accurate. That is, our method was more accurate even than static MFMC in which the surrogate was constructed using more high-fidelity evaluations than in our approach, thus showing that low-fidelity models can indeed be too accurate for multifidelity methods. In the plasma physics test case, we employed the Sensitivity-driven dimension-adaptive sparse grid algorithm from Chapter 4 as surrogate, which has algebraic rates. We compared our context-aware approach with static MFMC in which the surrogate was again constructed using either fewer or more high-fidelity evaluations than in our method, and we showed that the proposed algorithm is the most accurate.

This concludes this chapter. Next, we present our fourth and last algorithm proposed in this work, where we switch focus from uncertainty propagation to Bayesian inference.

PART IV

CONTEXT-AWARE MULTILEVEL DECOMPOSITIONS FOR BAYESIAN INFERENCE

“The most important maxim for data analysis to heed, and one which many statisticians seem to have shunned, is this: Far better an approximate answer to the right question, which is often vague, than an exact answer to the wrong question, which can always be made precise.”

John Tukey

7

Multilevel adaptive sparse grid Leja approximations for Bayesian inference

In the previous chapters, the efficacy of context-aware sparse grid approximations (Chapters 4 and 5) and of context-aware multifidelity Monte Carlo sampling (Chapter 6) was demonstrated in uncertainty propagation settings. In the fourth and last contribution of this work, we switch focus from uncertainty propagation to parameter inference, and demonstrate the effectiveness of context-aware model hierarchies in Bayesian inverse problems. Specifically, we focus on Bayesian parameter estimation, another example of an outer-loop scenario. We employ multilevel approximations to reduce the computational cost of finding efficient posterior-focused surrogates of the parameter-to-observable map. Moreover, we also want to reduce the cost of integration w.r.t. the posterior density via posterior-focused dimension-adaptive sparse grids.

The main reference for the ideas and results presented in this chapter is our work [55], which has been recently accepted for publication. Therein, we formulated the Multilevel adaptive sparse Leja algorithm and applied it in the four test cases considered in Section 7.4. Nevertheless, we have several novelties in this chapter. First, we discuss two termination criteria of the proposed multilevel approach. Moreover, in Section 7.4.3, in which we consider the test case with bimodal observation data to highlight the main limitation of the proposed approach, we perform a more detailed numerical study than in [55].

This chapter is organized as follows. In Section 7.1, we discuss posterior-focused surrogate models for Bayesian inference. Specifically, we first summarize the standard techniques for constructing approximations in Bayesian parameter estimations, focusing on the main challenges of finding such approximations efficiently in higher-dimensional, computationally expensive problems. Addressing some of these challenges constitutes the motivation for our proposed methodology, the *Multilevel adaptive sparse Leja algorithm*. The proposed multilevel approach is detailed in Section 7.2. It is a fully deterministic, sampling-free, multilevel context-aware strategy which relies on sparse grid approximations at each level. Our algorithm is formulated as a methodology, independent thus of the specific implementation of the sparse grid approximations. To exploit the anisotropic coupling of the input parameters and thus the intrinsic lower-dimensional structure of the problem (context-awareness), we employ again dimension-adaptive sparse grid approximations, which we present in Section 7.3. For a better overview of the behaviour of such algorithms in our context, we consider two dimension-adaptive strategies, a standard approach and an enhanced strategy based on directional variances. Finally, in Section 7.4 we present our results. First, we employ a simple test case with analytically known forward model. Afterwards, we consider two source inversion problems defined in two-dimensional spatial domains. In both problems, we infer two stochastic parameters. The second source inversion is used to underline the major limitation of the proposed approach, which will be addressed in our future research. Finally, as a representative of a more complex application, we consider an inverse problem with a forward model defined on a 3D spatial domain, in which we infer eight parameters.

Highlights and novelties

- We formulate a novel context-aware algorithm for efficient Bayesian inference based on model hierarchies appearing multilevel decompositions and global sparse grid approximations;
- The main novelty of our approach, which we call the Multilevel adaptive sparse Leja algorithm, is the sequential update of the reference information;
- The sequential update of the reference information leads to the placement of sparse grid points in the region of high posterior probability and thus to the construction of posterior-focused surrogates, at no additional cost;
- The multilevel decomposition in the proposed approach is performed implicitly, not via telescoping sums;
- Besides assessing quantities of interest, another outcome of the new algorithm is an approximation of the posterior density, which can be used, for example, in uncertainty propagation settings;
- The proposed approach is agnostic regarding the implementations of the sparse grid approximations;
- In our numerical results, we consider standard dimension-adaptive sparse approximations and an enhanced adaptive interpolation approach in which the stochastic directions rendered as unimportant are not further refined by the algorithm;
- We employ our approach in four test cases with elliptic forward models, including a more complex test case with eight stochastic inputs and a forward model defined on a 3D domain, and a scenario in which we highlight the limitations of the method.

7.1 Posterior-focused surrogate models in Bayesian inversion

The posterior solution of Bayesian inference is often intractable in the sense that it does not admit closed form analytic expressions. Closed form analytic expressions are only possible in very few cases which usually involve linear forward models and specifically chosen prior densities such as Gaussians. In the following example, we consider an inverse problem with a linear forward model and Gaussian prior.

Example: Let us assume that the forward model, \mathcal{F} , is analytically available and it is linear w.r.t. the unknown parameters $\boldsymbol{\theta}$. For example, consider $\mathcal{F}(x, y, \boldsymbol{\theta}) = \theta_1 \cos(x) + \theta_2 \sin(y)$ is linear w.r.t. θ_1 and θ_2 . In addition, assume that the prior density is Gaussian, that is $\pi_0 = N(\boldsymbol{\mu}_{\text{pr}}, C_{\text{pr}})$ with mean $\boldsymbol{\mu}_{\text{pr}} \in \mathbb{R}^{d_{\text{sto}}}$ and a positive definite symmetric covariance matrix $C_{\text{pr}} \in \mathbb{R}^{d_{\text{sto}} \times d_{\text{sto}}}$. In [171], it was shown that with this setup, the posterior is also Gaussian, i.e., $\pi^{\mathbf{y}} = N(\boldsymbol{\mu}_{\text{post}}, C_{\text{post}})$, where $\boldsymbol{\mu}_{\text{post}} = \boldsymbol{\mu}_{\text{pr}} - C_{\text{post}} \Delta Y(\boldsymbol{\mu}_{\text{pr}})$ and $C_{\text{post}}^{-1} = \Delta^2 Y + C_{\text{pr}}^{-1}$, with $Y(\boldsymbol{\theta}) = \frac{1}{2} (\mathbf{y} - \mathcal{F}_h(\boldsymbol{\theta}))^T \Gamma^{-1} (\mathbf{y} - \mathcal{F}_h(\boldsymbol{\theta}))$ and Γ is the covariance matrix of the assumed additive Gaussian noise model. We employ the following setup:

$$d_{\text{sto}} = 2, \quad \boldsymbol{\mu}_{\text{pr}} = (0, 0), \quad C_{\text{pr}} = \begin{pmatrix} 0.8 & 0.1 \\ 0.1 & 0.4 \end{pmatrix}, \quad \Gamma = I.$$

7.1. POSTERIOR-FOCUSED SURROGATE MODELS IN BAYESIAN INVERSION

Moreover, the measurement locations are $(x, y) = (0.2i, 0.2j)$ for $i, j = 1, 2, 3, 4$.

We depict, in Figure 7.1, the Gaussian prior and the Gaussian posterior, obtained as described above. Since the posterior is available in close form, it is straightforward to sample from it or to assess quantities of interest depending on the posterior.




Figure 7.1: Left: Gaussian prior $\pi_0 = N(\boldsymbol{\mu}_{\text{pr}}, C_{\text{pr}})$. Right: Gaussian posterior $\pi^{\mathbf{y}} = N(\boldsymbol{\mu}_{\text{post}}, C_{\text{post}})$ with mean and covariance matrix obtained as described above.

Our assumption throughout this work is that the forward operator is nonlinear and computationally expensive, usually given in terms of PDEs. In this case, the posterior is not available in close form. We illustrate the case in which the forward model is nonlinear in the next example.

Example: Let us revisit the previous example. We consider the same prior, but we make the forward model nonlinear by swapping x, y with θ_1, θ_2 , i.e., $\mathcal{F}(x, y, \boldsymbol{\theta}) = x \cos(\theta_1) + y \sin(\theta_1)$. In addition, we assume a lower noise level

$$\Gamma = 0.04I.$$

To find the posterior density, we apply Bayes' theorem (2.3), in which we assess the evidence using $10^2 = 100$ Gauss-Legendre nodes. In Figure 7.2, visualize the prior and posterior densities.




Figure 7.2: Left: Gaussian prior $\pi_0 = N(\boldsymbol{\mu}_{\text{pr}}, C_{\text{pr}})$. Right: Posterior density obtained via Bayes' theorem (2.3) for a nonlinear forward model.

The nonlinear forward model leads to a more complex density, not available in closed form, from

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

which we cannot sample directly. Moreover, any computations w.r.t. this posterior involve one evaluation of the forward, which enters the likelihood, thus these computations can be very expensive. Addressing the challenges of dealing with nonlinear forward models in Bayesian inference is one of the main goals of the proposed approach in this chapter.

To this end, approximations need to be used in practice. Finding accurate approximations efficiently is however challenging when the forward operator is computationally expensive or when the number of parameters to be inferred is large.

Remark: Surrogates or reduced models of the forward operator are typically constructed w.r.t. the prior density. This means that the forward model is evaluated at samples from the support of the prior density, and the resulting ensemble of forward model evaluations is then used to construct the approximation. Thus, surrogate model construction based on the prior density is equivalent to an *uncertainty propagation* problem. Note, however, that in uncertainty propagation the input density is assumed to be known. That is, the input density stems from a preprocessing step in which the uncertainty in the input parameters has been quantified. In contrast, in Bayesian inversion, the prior density incorporates any prior knowledge, e.g., expert opinion, available data, or simply the “best available guess” on the input parameters before measurements are performed. Thus, the prior density typically overestimates the uncertainty in the input parameters. To construct surrogates in Bayesian inversion based on a density that correctly quantifies the uncertainty in the input parameters, we would need the posterior density. Constructing surrogates from (sequentially updated) posterior densities is supported by the proposed algorithm in this chapter.

Theoretical results from [122,171] show that if the prior-based surrogate of the forward model converges at a specific rate w.r.t. the prior weighted L^2 -norm, then the resulting approximate posterior converges to the exact posterior with at least the same rate. This result has been improved in [203] where it was showed that the convergence rate of the posterior approximation is at least twice as large as the convergence rate of the surrogate, for general priors. Nevertheless, constructing a surrogate over the entire support of the prior might not be feasible and is in fact often unnecessary. Indeed, in inference problems in which the measurement data are informative, the posterior differs significantly from the prior distribution, i.e., its support is a (small) subset of the support of the prior. This suggests to adapt and localize the surrogate construction to the support of the posterior, which is, however, not an easy task since the posterior is available only after the Bayesian inference is performed.

The idea of posterior-focused surrogates is not new, but it has received little attention to date in the literature. One of the first works in this direction was [115], in which Li and Marzouk employed ideas from statistics to construct an efficient polynomial chaos surrogate associated with a density that minimizes the cross entropy between the posterior and a family of multivariate normal distributions. Recently, Jiang and Ou [98] proposed a two-stage surrogate based on generalized multiscale FE and least-squares stochastic collocation. In addition, Yan and Zhou [204] suggested a multifidelity polynomial chaos surrogate which combines a large number of inexpensive low-fidelity model evaluations with a small number of expensive high-fidelity model evaluations. Moreover, Chen, Villa and Ghattas [30] proposed using a Laplace approximation of the posterior at the maximum a posteriori point to efficiently perform adaptive sparse grid quadrature in Bayesian inverse problems with Gaussian prior. Next, we present our proposed multilevel approach in which we construct posterior-focused surrogates via multilevel decompositions and context-aware sparse grid approximations.

7.2 Multilevel adaptive sparse Leja algorithm

In this section, we formulate a computational methodology for addressing higher-dimensional, computationally expensive Bayesian inverse problems. To this end, we aim to:

1. perform the inversion problem from *informative reference information*;
2. *exploit* the *anisotropic coupling* of the input parameters w.r.t. the posterior;
3. *handle arbitrary* probability density functions;
4. *perform integration* w.r.t. the posterior density solution *efficiently*.

To achieve the first goal, we employ multilevel model hierarchies and sequentially update the reference information such that starting with the second level, the previous level posterior is used as prior density at the current level; at the first level, the reference information is given by the prior density. In this way, the corresponding Bayesian inverse problem starts from more informative prior information. The sequential update of the reference information provides the mechanism to achieve the second goal as well. Specifically, reusing the previous level posterior as the prior at the current level enables us to exploit the sparsity, i.e., anisotropic coupling of the input parameters w.r.t. the updated reference density and thus to construct efficient, posterior-focused surrogate models. One challenge of posterior-focused surrogates is the need to handle arbitrary densities which can deviate significantly from the prior which is usually a classical density such as uniform or Gaussian. To handle arbitrary density functions, we construct our approach using weighted (L)-Leja sequences (see, e.g., [82, 134]). Lastly, to achieve the fourth goal we again exploit the sequential update of the reference density and employ deterministic quadrature performed w.r.t. the updated density.

To this end, we formulate a *multilevel, deterministic, adaptive, sampling-free* methodology which we call *Multilevel adaptive sparse Leja algorithm*. The *multilevel* decomposition enables us to sequentially update the reference information such that starting with the second level, the previous level posterior is used as the current level's prior density.

Remark: Recall that sparse grids are defined on spaces with weight functions having a product structure (cf. Assumption 2.2). Since posterior densities in Bayesian inversion do not generally have a product structure, the updated reference density in our approach is thus not in a product form. To address this issue, we employ a biasing density that leads to the desired product structure, which is explained in Section 7.2.3.

At each level we perform *deterministic*, dimension-adaptivity sparse grid interpolation of the potential or data misfit function, $\Phi(\boldsymbol{\theta}; \mathbf{y})$, w.r.t. the updated reference information, i.e, the interpolation knots reside within the support of the updated density; recall that in our implementation, the interpolation points are weighted (L)-Leja sequences (recall Section 2.5.2). With adaptivity, we explore and exploit the lower-dimensional structure of the underlying problem due to the intrinsic sparsity of the forward operator w.r.t. the updated reference density to construct cost-efficient and accurate sparse grid interpolation surrogates. This makes our approach *context-aware*. Our motivation for approximating the potential function is two-fold. First, evaluating the potential function entails evaluating the forward operator, $\mathcal{F}(\boldsymbol{\theta})$ (recall (2.2)), which is the computationally expensive part. Second, even if the forward operator is vector-valued, the potential function is a norm and thus a scalar-valued function. Approximating vector-valued functions with adaptive sparse grids implies that a separate approximation is typically needed for each output component, which can become computationally prohibitive when the number of components is large. Once the sparse grid interpolant of the potential function is obtained,




Figure 7.3: Intuitive depiction of the proposed approach for Bayesian inference, consisting of two stages. In the first stage (left figure), the proposed *Multilevel adaptive sparse Leja algorithm* is used to find surrogates of the potential or data misfit function, $\Phi(\boldsymbol{\theta}; \mathbf{y})$, which we denote $\Phi_{\text{SG}}(\boldsymbol{\theta}; \mathbf{y})$. In the second step, $\Phi_{\text{SG}}(\boldsymbol{\theta}; \mathbf{y})$ is used to perform Bayesian inference, i.e., in computations w.r.t. the posterior, which typically involve integration. In this work, integration w.r.t. the posterior are performed via adaptive sparse grid quadrature. Both stages are performed as *outer-loops*. Note however, that both stages depend on adaptive algorithms, thus the sizes of the outer-loops is not known a priori.

we have a surrogate of the likelihood function and hence a surrogate of the posterior density as well. Finally, the proposed approach is *sampling-free* since we do not explore the posterior approximation via traditional sampling methods such as MCMC or SMC, but we employ *dimension-adaptive* sparse grid quadrature in which the quadrature nodes are weighted (L)-Leja sequences constructed w.r.t. the updated density. In this way, we place the quadrature points in the region of high posterior probability and hence ignore the region of low-posterior probability, which traditional quadrature techniques are unable to do in Bayesian inversion. Thus, the proposed strategy addresses the difficult tasks of efficiently finding surrogates and performing deterministic quadrature in inference problems. We provide, in Figure 7.3, a visual summary of the proposed approach. We split the Bayesian inference process into two outer-loops. In the first outer-loop, we employ the proposed Multilevel adaptive sparse Leja algorithm to find context-aware surrogates of the potential function. This step is computationally most expensive since it involves evaluations of the forward operator. After we have the sparse grid surrogate, we enter the second outer-loop, in which we (i) compute the Bayesian posterior depending on the surrogate of the potential function and (ii) perform all computations w.r.t. the posterior, which typically involve integration. All integration operations are done via dimension-adaptive sparse grid quadrature. We note that the proposed approach is not designed to address Bayesian inverse problems with arbitrary posterior densities, but rather posterior densities with a clearly defined mode. The numerical example in Section 7.4.3 emphasize this limitation. At the end of this chapter, we discuss provide an outlook for future research directions to address this issue.

Remark: The specific implementations for adaptive sparse grid interpolation and quadrature do not influence the formulation of our proposed approach. Thus, we assume we have available two adaptive strategies, $\text{AdaptSGInterp}(tol^{\text{in-g}}, K_{\text{max}}^{\text{in-g}}, \mathcal{M}, \boldsymbol{\pi})$ and $\text{AdaptSGQuad}(tol^{\text{qu}}, K_{\text{max}}^{\text{qu}}, \mathcal{M}, \boldsymbol{\pi})$. Adaptive sparse grid algorithms typically employ a *tolerance*, tol^{op} and a *maximum*

attainable grid level, K_{\max} . Moreover, we also need the function \mathcal{M} for which we perform the sparse grid approximation, and the weight (density) function w.r.t. which the approximation is performed, π , which needs to have a product structure, since sparse grid approximations are constructed via tensorizations (see Section 2.5). Note that specific implementations might have additional input arguments, however the four inputs considered here are sufficient to illustrate these algorithms. The employed adaptive strategies are presented in Section 7.3.

Next, we present in detail the algorithmic steps of the proposed Multilevel adaptive sparse Leja methodology. We begin with a generic overview on the multilevel setup and introduce some notation in Section 7.2.1. In Section 7.2.2, we summarize the computations performed at the first level in our approach. Note that these computations are performed in a standard way, i.e., w.r.t. the given prior density. The novelty of our approach comes in the remaining levels, which we detail in Section 7.2.3. We end by discussing the computational cost of our algorithm in Section 7.2.4 and its termination criteria in Section 7.2.5.

7.2.1 Multilevel setup

Let $J \in \mathbb{N}$ denote the number of levels in our multilevel formulation such that we have at least two levels, i.e., $J \geq 2$, and let $j = 1, 2, \dots, J$. In addition, we denote by $\mathcal{M} \in \{\mathcal{F}_h, \Phi, L, Z, \pi^y\}$ a generic continuous quantity (forward operator, potential function, likelihood, evidence, posterior density) depending on both physical and stochastic parameters. Note that Φ, L, Z and π^y depend implicitly on the forward operator, \mathcal{F}_h . To enable the numerical simulation of \mathcal{M} , we hence need two discretization parameters, one for the physical domain and the other parameter for the stochastic domain discretization. The discretization of the physical domain of \mathcal{M} is characterized by a parameter h_j such that h_1 is the coarsest and h_J the finest discretization level. For example, h_j could denote the mesh width in a FE approximation. Therefore, by \mathcal{M}_j we denote the semi-discrete, parametric stochastic approximation of \mathcal{M} depending on h_j , whereas $\mathcal{M}_{\delta j}$ refers to either $\mathcal{M}_j - \mathcal{M}_{j-1}$ or $\mathcal{M}_j/\mathcal{M}_{j-1}$. We note that the employed spatial discretization is classical, and is not our major concern here. Building a useful and accurate surrogate for \mathcal{M}_j is the central task ahead. To this end, we employ dimension-adaptive sparse grid approximations with user-defined tolerance tol_j^{op} such that

$$tol_1^{\text{op}} \leq tol_2^{\text{op}} \leq \dots \leq tol_J^{\text{op}}$$

to construct approximations of \mathcal{M}_j . Therefore, by $\mathcal{M}_{j,s}$ we denote the sparse grid approximation of \mathcal{M}_j depending on tol_s^{op} . To be able to employ sparse grid approximations in this context, we need Assumptions 2.1 and 2.2 either to hold true or we need suitable transformations such that they are satisfied. From Assumption 2.2, the given prior π_0 has a product structure, which is usually the common setup employed in Bayesian inverse problems.

Recall from Section 2.7, that in multilevel decomposition, the goal is to determine $\mathcal{M}_{j,J-j+1}$ for all $j = 1, 2, \dots, J$. To simplify the notation we use the subscript $\ell(j)$ to refer to $(h_j, tol_{j-j+1}^{\text{op}})$ and the subscript $\ell(\delta j)$ to denote approximations $\mathcal{M}_{\ell(\delta j)} \approx \mathcal{M}_{\delta j}$, where the subscript δj was used to denote either differences or ratios in terms of \mathcal{M}_j and \mathcal{M}_{j-1} . Hence, we use the notation $\ell(j)$ or $\ell(\delta j)$ to refer to the *level* j in our multilevel approach (see also Section 5.3.2). Note that levels are used to characterize both sparse grid and multilevel formulations. This is because both formulations are intrinsically multilevel (recall Sections 2.5 and 2.10). Nevertheless, to avoid confusion we will explicitly specify what is meant by level in each context.

7.2.2 Level $\ell(1)$

We begin the multilevel algorithm with level $\ell(1)$. At this level, the adaptive sparse grid interpolation surrogate of the potential function and all subsequent quadrature operations involving this approximation are performed w.r.t. the given prior density, π_0 , which by Assumption 2.2, has a product structure.

To reduce the overall computational effort, at $\ell(1)$ we employ the coarsest spatial discretization characterized by h_1 , which is combined with the most accurate sparse grid interpolation tolerance, $tol_J^{\text{in-g}}$. Thus, at $\ell(1)$ we compute the approximation $\Phi_{\ell(1)}(\boldsymbol{\theta}; \mathbf{y}) \approx \Phi_1(\boldsymbol{\theta}; \mathbf{y})$ using adaptive sparse grid interpolation w.r.t. the prior density π_0 . We use the sparse grid approximation $\Phi_{\ell(1)}(\boldsymbol{\theta}; \mathbf{y})$ to obtain the first level likelihood surrogate, $L_1(\boldsymbol{\theta}|\mathbf{y}) \approx L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y}) := \exp(-\Phi_{\ell(1)}(\boldsymbol{\theta}; \mathbf{y}))$. We then employ $L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y})$ in numerical integration operations which we perform via adaptive sparse grid quadrature w.r.t. the prior density π_0 . We assess the evidence $Z_{\ell(1)}(\mathbf{y}) \approx Z_1(\mathbf{y})$ using (2.4). Having obtained the approximations of the likelihood, $L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y})$, and of the evidence, $Z_{\ell(1)}(\mathbf{y})$, we plug these approximations into Bayes' formula (2.3) and obtain the posterior approximation at level $\ell(1)$, $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta})$.

In addition, as we will see next, at level $\ell(2)$, we also need the *Gaussian biasing density* $\hat{\pi}_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta}) := \text{N}(\mathbf{m}_{\ell(1)}, \mathbf{C}_{\ell(1)})$ of $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta})$, defined using the *expectation* and *covariance* of the posterior $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta})$. Thus, we make use of the likelihood approximation $L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y})$ and compute the expectation and the covariance matrix of the posterior density via dimension-adaptive sparse quadrature w.r.t. the prior (see (2.6)). We estimate the posterior expectation $\mathbf{m}_{\ell(1)} \in \mathbb{R}^{d_{\text{sto}}}$ as

$$\mathbf{m}_{\ell(1)}^i := Z_{\ell(1)}^{-1}(\mathbf{y}) \int_{\mathbf{X}} \theta_i L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y}) \pi_0(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad i = 1, \dots, d_{\text{sto}}.$$

The approximation of the covariance matrix, $\mathbf{C}_{\ell(1)} \in \mathbb{R}^{d_{\text{sto}} \times d_{\text{sto}}}$, reads

$$\mathbf{C}_{\ell(1)}^{ij} := Z_{\ell(1)}^{-1}(\mathbf{y}) \int_{\mathbf{X}} \theta_i \theta_j L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y}) \pi_0(\boldsymbol{\theta}) d\boldsymbol{\theta} - \mathbf{m}_{\ell(1)}^i \mathbf{m}_{\ell(1)}^j, \quad i, j = 1, \dots, d_{\text{sto}}.$$

Remark: When the underlying posterior density is concentrated in the support of the prior, adaptive sparse grid quadrature w.r.t. the prior typically fails to converge since most sparse grid points are placed outside of the region of high (and concentrated) posterior probability. When this happens, the adaptive quadrature at level $\ell(1)$ in our proposed approach would fail to provide a Gaussian biasing density. Possible remedies include using a static sparse grid of a large enough level to ensure that grid points are placed within the region of high posterior probability (we used this remedy of one of the test cases; see Section 7.4.4). When using nested point sets, such as (L)-Leja points, we can proceed as following: we can begin with a static grid of a small level and then increase the level, reusing the computations from the previous level, until, e.g., a desired accuracy is reached or the difference between two steps is small.

In the same direction, another remedy is to first compute approximations using a static sparse grid such that grid points are placed inside or close to the region of high-posterior probability, and then to start the adaptive refining from this approximation. However, it is rather unclear how to choose the level of the initial grid and when to start the refinement process. Other solutions include employing a Laplace approximation of the posterior around the MAP point [30], which was proven to be a suitable choice in high-dimensional settings. Moreover, using tempering (see, e.g., [50]) or sparse grids with local (spatial) adaptivity [120] could also be alternative ways to alleviate the challenges of dealing with concentrated posteriors.

The steps corresponding to level $\ell(1)$ are summarized in Algorithm 7.9. The first three inputs are the coarsest spatial discretization parameter, h_1 , and the smallest tolerances for sparse grid

7.2. MULTILEVEL ADAPTIVE SPARSE LEJA ALGORITHM

interpolation and quadrature, $tol_J^{\text{in-g}}$ and tol_J^{qu} . In addition, for both dimension-adaptive sparse grid operations, a maximum reachable level $\mathbf{K}_{\max} := (K_{\max}^{\text{in-g}}, K_{\max}^{\text{qu}})$ is needed as well. The last two inputs are the potential function, $\Phi(\boldsymbol{\theta}; \mathbf{y})$, and the prior density, $\pi_0(\boldsymbol{\theta})$. At the end, the algorithm returns the posterior approximation, $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta})$, and its expectation, $\mathbf{m}_{\ell(1)}$, and covariance matrix, $\mathbf{C}_{\ell(1)}$, needed for the Gaussian biasing density at the second level.

Algorithm 7.9 Level One Adaptive Sparse Leja Algorithm for Bayesian Inversion

- 1: **procedure** LEVELONESPARSELEJA($h_1, tol_J^{\text{in-g}}, tol_J^{\text{qu}}, \mathbf{K}_{\max}, \Phi, \pi_0$)
- 2: Compute the potential's surrogate via adaptive sparse grid interpolation

$$\Phi_{\ell(1)}(\boldsymbol{\theta}; \mathbf{y}) = \text{AdaptSGInterp}(tol_J^{\text{in-g}}, K_{\max}^{\text{in-g}}, \Phi_1, \pi_0)$$

- 3: Construct the likelihood surrogate $L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y}) := \exp(-\Phi_{\ell(1)}(\boldsymbol{\theta}; \mathbf{y}))$
- 4: Compute the evidence via adaptive sparse grid quadrature

$$Z_{\ell(1)}(\mathbf{y}) = \text{AdaptSGQuad}(tol_J^{\text{qu}}, K_{\max}^{\text{qu}}, L_{\ell(1)}, \pi_0)$$

- 5: Compute the posterior

$$\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta}) := \frac{\pi_0(\boldsymbol{\theta})L_{\ell(1)}(\boldsymbol{\theta}|\mathbf{y})}{Z_{\ell(1)}(\mathbf{y})}$$

- 6: Compute the expectation $\mathbf{m}_{\ell(1)}$ of $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta})$ w.r.t. $\pi_0(\boldsymbol{\theta})$
 - 7: Compute the covariance $\mathbf{C}_{\ell(1)}$ of $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta})$ w.r.t. $\pi_0(\boldsymbol{\theta})$
 - 8: **return** $\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta}), \mathbf{m}_{\ell(1)}, \mathbf{C}_{\ell(1)}$
 - 9: **end procedure**
-

7.2.3 Levels $\ell(j)$ with $j \geq 2$

The novelty of the proposed Multilevel adaptive sparse Leja algorithm begins at levels $\ell(j)$ with $j \geq 2$ in the multilevel decomposition. At these levels, we sequentially update the reference information such that the posterior approximation from the previous level, $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$, is reused as prior density at the current level; we detail the sequential update next. In this way, we construct the adaptive sparse grid approximations w.r.t. $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$. Recall that $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ needs to have a product structure, i.e., the underlying (stochastic) space needs to have a tensor structure to allow the construction of sparse grid approximations (see also Assumption 2.1). However, $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ has the desired product structure only in very limited settings, e.g., when the underlying forward model is linear. Since our proposed approach for Bayesian parameter estimation is targeted to problems characterized by nonlinear, complex models, $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ is not generally in a product form. Therefore, we introduce a biasing density $\hat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$, i.e., an approximation of the posterior $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$, which allows to obtain the desired product structure.

Gaussian biasing density

We use the Gaussian density $\hat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ defined as

$$\hat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) := \text{N}(\mathbf{m}_{\ell(j-1)}, \mathbf{C}_{\ell(j-1)}) = \frac{\exp\left(-\frac{1}{2}(\boldsymbol{\theta} - \mathbf{m}_{\ell(j-1)})^T \mathbf{C}_{\ell(j-1)}^{-1}(\boldsymbol{\theta} - \mathbf{m}_{\ell(j-1)})\right)}{\sqrt{(2\pi)^{d_{\text{sto}}} \det(\mathbf{C}_{\ell(j-1)})}} \quad (7.1)$$

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

as biasing density of the posterior $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$, where $\mathbf{m}_{\ell(j-1)}$ and $\mathbf{C}_{\ell(j-1)}$ are the expectation and covariance matrix of $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$.

To arrive at the desired product structure, we write the generic multivariate Gaussian random variable distributed according to (7.1) as an affine mapping of a standard multivariate Gaussian random variable. Specifically, from the spectral decomposition $\mathbf{C}_{\ell(j-1)} = VDV^{-1}$, we have $\mathbf{C}_{\ell(j-1)}^{1/2} = VD^{1/2}V^{-1}$. Using this decomposition, we arrive at

$$\boldsymbol{\theta} = T_{\ell(j-1)}(\boldsymbol{\zeta}) := \mathbf{m}_{\ell(j-1)} + \mathbf{C}_{\ell(j-1)}^{1/2}\boldsymbol{\zeta} \Rightarrow \boldsymbol{\theta} \sim \mathcal{N}(\mathbf{m}_{\ell(j-1)}, \mathbf{C}_{\ell(j-1)}), \quad (7.2)$$

where $\boldsymbol{\zeta}$ is a standard Gaussian random variable, i.e., $\boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, I)$.

To summarize, (7.2) allows to write a general multivariate Gaussian random variable with correlated components as an affine mapping of a standard multivariate Gaussian random variable, which has the desired product structure since the components of $\boldsymbol{\zeta}$ are uncorrelated and thus independent. Note that in a broader context, the mapping $T_{\ell(j-1)}$ in (7.2) can be seen as an *affine transport map* (see [121]).

Remark: The covariance matrix of the posterior density is generally dense. Estimating it in high-dimensional settings can become prohibitive, since in d_{sto} dimensions, $d_{\text{sto}}(d_{\text{sto}} + 1)/2$ integrals need to be computed. Moreover, the accuracy of these estimates usually degrades as the dimensionality increases. To this end, instead of using the Gaussian approximation of the posterior as biasing density, we could use, e.g., the Laplace approximation of the posterior, which was proven to be a feasible and scalable choice in high-dimensional settings (see, e.g., [94]).

Next, we detail how the sparse grid operations are performed at levels $\ell(j)$ with $j \geq 2$.

Level update on tensor domain

Starting with level $\ell(2)$, we sequentially update the auxiliary reference density in Bayes' formula (2.3) such that we reuse the sparse grid approximation of the posterior density from the previous level, $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$. Since $\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ does not have the product structure needed in sparse grid approximations, we employ the Gaussian biasing density $\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ introduced above. Mathematically, the sequential update of the reference density in Bayes' formula reads

$$\begin{aligned} \pi_j^{\mathbf{y}}(\boldsymbol{\theta}) &:= \frac{L_j(\boldsymbol{\theta}|\mathbf{y})\pi_0(\boldsymbol{\theta})}{Z_j(\mathbf{y})} = \frac{L_j(\boldsymbol{\theta}|\mathbf{y})\pi_0(\boldsymbol{\theta})}{Z_j(\mathbf{y})} \frac{L_{j-1}(\boldsymbol{\theta}|\mathbf{y})}{L_{j-1}(\boldsymbol{\theta}|\mathbf{y})} \frac{Z_{j-1}(\mathbf{y})}{Z_{j-1}(\mathbf{y})} \\ &= \frac{L_{j-1}(\boldsymbol{\theta}|\mathbf{y})\pi_0(\boldsymbol{\theta})}{Z_{j-1}(\mathbf{y})} \frac{L_j(\boldsymbol{\theta}|\mathbf{y})}{L_{j-1}(\boldsymbol{\theta}|\mathbf{y})} \frac{Z_{j-1}(\mathbf{y})}{Z_j(\mathbf{y})} = \frac{\pi_{j-1}^{\mathbf{y}}(\boldsymbol{\theta}) \frac{L_j(\boldsymbol{\theta}|\mathbf{y})}{L_{j-1}(\boldsymbol{\theta}|\mathbf{y})}}{\frac{Z_j(\mathbf{y})}{Z_{j-1}(\mathbf{y})}} \\ &\approx \frac{\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) L_{\delta_j}(\boldsymbol{\theta}|\mathbf{y})}{Z_{\delta_j}(\mathbf{y})} = \frac{\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) L_{\delta_j}(\boldsymbol{\theta}|\mathbf{y}) \frac{\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})}{\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})}}{Z_{\delta_j}(\mathbf{y})}, \end{aligned} \quad (7.3)$$

where $L_{\delta_j}(\mathbf{y}) := L_j(\mathbf{y})/L_{j-1}(\mathbf{y})$ and $Z_{\delta_j}(\mathbf{y}) := Z_j(\mathbf{y})/Z_{j-1}(\mathbf{y})$. Thus, (7.3) allows to write $\pi_j^{\mathbf{y}}$ depending on the sparse grid approximation of the posterior from the previous level, $\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$, the likelihood ratio, L_{δ_j} and the bias correcting ratio, $\pi_{\ell(j-1)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$. Moreover, to obtain a probability density function, we normalize using the evidences ratio $Z_{\delta_j}(\mathbf{y})$. In computations, we first find a dimension-adaptive interpolation surrogate of the potential w.r.t. the Gaussian biasing density $\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$. We then employ this surrogate in adaptive quadrature operations, in which we correct the bias introduced by the Gaussian density $\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$ using the ratio $\pi_{\ell(j-1)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$.

Remark: The sequential update of the reference density defined in (7.3) means that the posterior density from level $\ell(1)$ is used at level $\ell(2)$, the posterior from level $\ell(2)$ is employed at level $\ell(3)$, and so forth until level $\ell(J)$, which employs the posterior from level $\ell(J-1)$. Therefore, the posterior density obtained at the last level, $\pi_{\ell(J)}^{\mathbf{y}}(\boldsymbol{\theta})$, represents the posterior solution of the multilevel algorithm, and the quantities of interest are computed w.r.t. this posterior density. Moreover, in the proposed approach, we do not explicitly evaluate *differences* between quantities from adjacent levels, but instead we work with likelihood *ratios* at each step in the multilevel hierarchy. In this way, the multilevel decomposition in the proposed approach is performed *implicitly*, not *explicitly* via a telescoping sum as in standard multilevel methods (see Section 2.7). We illustrate the difference between a standard (explicit) multilevel decomposition based on telescoping sums and our implicit multilevel approach in Figure 7.4.




Figure 7.4: Left: example of a three-level standard multilevel decomposition. In the context of Bayesian inference, in standard multilevel methods, the computations at each level are performed w.r.t. the prior density. At the end, a quantity of interest, \mathcal{Q} , is assessed via telescoping sums (2.35). Here, \mathcal{Q} is general an integral, thus we approximate it, at each level, via dimension-adaptive sparse grid quadrature. Note that five quantities are assessed in total. Right: illustration of the proposed multilevel approach for $J = 3$. We depict the interpolation part in our method, with which we approximate the potential function. At level $\ell(1)$ we proceed as in the standard approach and perform all operations w.r.t. the prior. However, starting with level $\ell(2)$, we update the reference information and find a surrogate for the potentials' difference, i.e., likelihoods' ratio depending on h_2 and h_1 . The level $\ell(2)$ posterior is used to update the reference information at level $\ell(3)$, where again we find a surrogate of the potentials' difference depending on h_3 and h_2 . Observe that only three quantities, that is, three surrogates are computed here. Moreover, \mathcal{Q} is assessed w.r.t. the posterior approximation at level $\ell(3)$, without using telescoping sums.

Moreover, standard multilevel techniques are targeted to assess linear quantities of interest, e.g., expectations, whereas with our approach we can also assess nonlinear quantities of interest, such as the covariance. Besides assessing quantities of interest, our approach additionally determines an approximation $\pi_{\ell(J)}^{\mathbf{y}}(\boldsymbol{\theta})$ of the posterior density, which could be used, for example, in an uncertainty propagation setting.

We begin with the adaptive sparse grid interpolation surrogate of the potential function. Specifically, we construct the surrogate $\Phi_{\ell(\delta_j)}(\boldsymbol{\theta}; \mathbf{y}) := \Phi_{\ell(j)}(\boldsymbol{\theta}; \mathbf{y}) - \Phi_{\ell(j-1)}(\boldsymbol{\theta}; \mathbf{y})$ of $\Phi_{\delta_j}(\boldsymbol{\theta}; \mathbf{y})$ w.r.t. the density $\hat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$ because the sequential update (7.3) depends on the likelihood ratio

$$L_{\delta_j}(\boldsymbol{\theta}|\mathbf{y}) := \frac{\exp(-\Phi_j(\boldsymbol{\theta}; \mathbf{y}))}{\exp(-\Phi_{j-1}(\boldsymbol{\theta}; \mathbf{y}))} = \exp(-\Phi_{\delta_j}(\boldsymbol{\theta}; \mathbf{y})).$$

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

Recall that the mapping $T_{\ell(j-1)}$ defined in (7.2) allows us to use adaptive sparse grid interpolation w.r.t. the Gaussian biasing density $\hat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta})$. Hence, to construct the surrogate for the potential function $\Phi_{\delta_j}(\boldsymbol{\theta}; \mathbf{y})$ we employ the mapping $T_{\ell(j-1)}$ and obtain

$$\Phi_{\delta_j}(\boldsymbol{\theta}; \mathbf{y}) \approx \Phi_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta}); \mathbf{y}).$$

$\Phi_{\ell(\delta_j)}$ gives the approximation $L_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta})|\mathbf{y}) := \exp(-\Phi_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta}); \mathbf{y}))$.

Remark: From a geometric perspective, sparse grid approximations are targeted to functions aligned with the standard Cartesian coordinate system (see [179] for a more detailed discussion on the geometric perspective of the regularity assumptions in sparse grid, low-rank, quasi-Monte Carlo and non-tensorized approximations based on, e.g., Padua points). In our proposed approach, with the exception of the first level, the sparse grid approximations are performed w.r.t. a Gaussian biasing density whose coordinate system is typically non-aligned with the standard Cartesian axes. However, our algorithm overcomes this issue implicitly by employing the affine mapping $T_{\ell(j-1)}$: $T_{\ell(j-1)}$ shifts and rotates the standard coordinate system to the coordinate system of the Gaussian biasing density. We depict this behaviour in Figure 7.5. In the left plot, we have a function and a sparse grid aligned with the standard coordinate system; this is the setup assumed by standard sparse grid approximations. The center plot depicts the situation in which an axis-aligned sparse grid is used to approximate a non axis-aligned target function. Specifically, the target function is shifted and rotated, which is what the Gaussian biasing density does. However, the sparse grid has not been transformed as well, and we can see that only a few sparse grid points reside in the support of the target function. Even worse, if an adaptive algorithm would be employed, it could terminate very early because the corresponding error indicators would be zero. In the right figure, we see again the shifted and rotated target function but this time with an aligned sparse grid: herein, we employed the affine mapping (7.2) to transform the coordinates of the sparse grid points as well. Note that the right figure illustrates what happens in our starting with the second level.




Figure 7.5: In the left plot, the standard setup for sparse grid approximations is depicted: the function and the sparse grid are aligned with the standard Cartesian coordinate system. In the center plot, we visualize the case in which the underlying function is not aligned with the standard coordinate system. This underlines what happens with all target functions in our proposed multilevel algorithm when employing the Gaussian biasing density. If standard, axis-aligned sparse grid are used in this case, then the resulting approximation will have a very poor quality. In the right plot, we depict the desired layout for sparse grid approximations: when the target function is non-aligned with the Cartesian system, we transform the sparse grid accordingly to obtain accurate approximations.

Addressing the issue of approximation of functions not aligned with the standard Cartesian coordinate system is an active area of research. For example, Trefethen proposed to use approximations depending on the so-called *Euclidean degree* (see, e.g., [179]). In addition, the recent work of [15] approached this problem in the context of least-squares regression in large data sets via sparse grid approximations. The authors perform a preprocessing step in which an optimized, problem-dependent coordinate system is determined which reduces the effective dimensionality of a given data set in the ANOVA sense.

Before explaining how quadrature is performed, we prove the following lemma, which is related to Lemma 5.3.1. Let

$$\pi_G(\zeta) := N(\mathbf{0}, I) = \frac{\exp\left(-\frac{1}{2}\zeta^T \zeta\right)}{\sqrt{(2\pi)^{d_{\text{sto}}}}}$$

denote the standard multivariate normal density. Recall the the Gaussian biasing density $\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$ defined in (7.1) and the mapping $T_{\ell(j-1)}(\zeta) = \mathbf{m}_{\ell(j-1)} + \mathbf{C}_{\ell(j-1)}^{1/2} \zeta$ from (7.2).

Lemma: 7.2.1 *For an integrable function $g : \mathbf{X} \rightarrow \mathbb{R}$ we have that*

$$\int_{\mathbf{X}} g(\boldsymbol{\theta}) \widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathbb{R}^{d_{\text{sto}}}} g(T_{\ell(j-1)}(\zeta)) \pi_G(\zeta) d\zeta.$$

Proof: By definition, $T_{\ell(j-1)}$ is injective and differentiable with Jacobian $DT_{\ell(j-1)}(\zeta) = \mathbf{C}_{\ell(j-1)}^{1/2}$ which is nonzero for all ζ (recall (7.2)). Therefore, from $\boldsymbol{\theta} = T_{\ell(j-1)}(\zeta) = \mathbf{m}_{\ell(j-1)} + \mathbf{C}_{\ell(j-1)}^{1/2} \zeta$,

$$d\boldsymbol{\theta} = |\det(DT_{\ell(j-1)})(\zeta)| d\zeta = \det(\mathbf{C}_{\ell(j-1)}^{1/2}) d\zeta.$$

Thus, if we make the substitution $\boldsymbol{\theta} = T_{\ell(j-1)}(\zeta)$ in $\int_{\mathbf{X}} g(\boldsymbol{\theta}) \widehat{\pi}_{j-1}^{\mathbf{y}}(\boldsymbol{\theta}) d\boldsymbol{\theta}$, we obtain

$$\int_{\mathbf{X}} g(\boldsymbol{\theta}) \widehat{\pi}_{j-1}^{\mathbf{y}}(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathbb{R}^{d_{\text{sto}}}} g(T_{\ell(j-1)}(\zeta)) \widehat{\pi}_{j-1}^{\mathbf{y}}(T_{\ell(j-1)}(\zeta)) \det(\mathbf{C}_{\ell(j-1)}^{1/2}) d\zeta. \quad (7.4)$$

Next, we analyse the product $\widehat{\pi}_{j-1}^{\mathbf{y}}(T_{\ell(j-1)}(\zeta)) \det(\mathbf{C}_{\ell(j-1)}^{1/2})$ using (7.1):

$$\begin{aligned} \widehat{\pi}_{j-1}^{\mathbf{y}}(T_{\ell(j-1)}(\zeta)) \det(\mathbf{C}_{\ell(j-1)}^{1/2}) &= \widehat{\pi}_{j-1}^{\mathbf{y}}\left(\mathbf{m}_{\ell(j-1)} + \mathbf{C}_{\ell(j-1)}^{1/2} \zeta\right) \det(\mathbf{C}_{\ell(j-1)}^{1/2}) \\ &= \frac{\exp\left(-\frac{1}{2}\zeta^T \zeta\right)}{\sqrt{(2\pi)^{d_{\text{sto}}} \det(\mathbf{C}_{\ell(j-1)})}} \det(\mathbf{C}_{\ell(j-1)}^{1/2}) = \frac{\exp\left(-\frac{1}{2}\zeta^T \zeta\right)}{\sqrt{(2\pi)^{d_{\text{sto}}}}} = \pi_G(\zeta), \end{aligned} \quad (7.5)$$

where we used $\sqrt{\det(\mathbf{C}_{\ell(j-1)})} = \det(\mathbf{C}_{\ell(j-1)}^{1/2})$ because $\mathbf{C}_{\ell(j-1)}$ is a covariance matrix, hence symmetric and positive definite. From (7.4) and (7.5) we thus get

$$\int_{\mathbf{X}} g(\boldsymbol{\theta}) \widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathbb{R}^{d_{\text{sto}}}} g(T_{\ell(j-1)}(\zeta)) \pi_G(\zeta) d\zeta. \quad \blacksquare$$

Thus, Lemma 7.2.1 shows that employing the mapping $T_{\ell(j-1)}(\zeta)$ leads to integration w.r.t. the standard multivariate Gaussian density π_G .

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

To evaluate the ratio of evidences $Z_{\delta_j}(\mathbf{y})$ we make use of the sequential update formula (7.3) and Lemma 7.2.1, in which we also employ the bias correcting ratio $\pi_{\ell(j-1)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}$ to obtain

$$\begin{aligned} Z_{\delta_j}(\mathbf{y}) &= \int_{\mathbf{X}} \widehat{\pi}_{j-1}^{\mathbf{y}}(\boldsymbol{\theta}) L_{\delta_j}(\boldsymbol{\theta}|\mathbf{y}) \frac{\pi_{j-1}^{\mathbf{y}}(\boldsymbol{\theta})}{\widehat{\pi}_{j-1}^{\mathbf{y}}(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ &\approx \int_{\mathbb{R}^{d_{\text{sto}}}} L_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta})|\mathbf{y}) \frac{\pi_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta}))}{\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta}))} \pi_G(\boldsymbol{\zeta}) d\boldsymbol{\zeta}. \end{aligned}$$

Therefore, to obtain the evidence approximation $Z_{\ell(\delta_j)}(\mathbf{y}) \approx Z_{\delta_j}(\mathbf{y})$, we numerically integrate $\left(L_{\ell(\delta_j)}\pi_{\ell(j-1)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}\right) \circ T_{\ell(j-1)}$ w.r.t. the density π_G via adaptive sparse grid quadrature. For all other quadrature computations, such as assessing the mean and covariance of the posterior densities, or computing quantities of interest w.r.t. the posterior, we employ Lemma 7.2.1.

Having the approximations for the likelihood ratio, $L_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta})|\mathbf{y})$, and for the evidences ratio, $Z_{\ell(\delta_j)}(\mathbf{y})$, at each level $\ell(j)$ with $j \geq 2$, we obtain the posterior density approximation

$$\pi_j^{\mathbf{y}}(\boldsymbol{\theta}) \approx \pi_{\ell(j)}^{\mathbf{y}}(\boldsymbol{\theta}) := Z_{\ell(\delta_j)}^{-1}(\mathbf{y}) \pi_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta})) L_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta})|\mathbf{y}).$$

We summarize the proposed Multilevel adaptive sparse Leja methodology in Algorithm 7.10. The first three inputs are the number of levels, $J \geq 2$, the sequence of mesh sizes, \mathbf{h} , and \mathbf{tol} , which comprises the tolerances for adaptive sparse grid interpolation and quadrature at all levels. The fourth input denotes the maximum reachable levels for the adaptive sparse grid algorithms, $\mathbf{K}_{\max} := (K_{\max}^{\text{qu}}, K_{\max}^{\text{in-g}})$. Lastly, π_0 is the prior density, $\Phi(\boldsymbol{\theta}; \mathbf{y})$ is the potential function and \mathcal{Q} is the quantity of interest. In the first step (line 2), we perform the level $\ell(1)$ operations using Algorithm 7.9. The steps corresponding to levels greater than two begin at line 4. At the end, we compute and return the quantity of interest w.r.t. the posterior from the last level. Note that besides computing quantities of interest, our algorithm also yields an approximation of the posterior density, i.e., the posterior at the last level.

7.2.4 Computational cost

At each level $\ell(1), \ell(2), \dots, \ell(J)$ in the proposed multilevel approach, two types of operations are performed, i.e., dimension-adaptive sparse grid interpolation and quadrature. Interpolation is employed to find a surrogate of the potential function, which involves evaluations of the forward model. Thus, the largest computational effort in the proposed multilevel methodology is spent in adaptive interpolation. Assume that one evaluation of the forward operator is performed on P processes, where $P = 1$ if the simulation is performed serially. To this end, for $j = 1, 2, \dots, J$, let \mathcal{C}_{h_j} denote the cost of one evaluation of $\mathcal{F}_{h_j}(\boldsymbol{\theta})$ on P processes, where $\mathcal{F}_{h_j}(\boldsymbol{\theta})$ denotes the spatial discretization of the forward operator, $\mathcal{F}(\boldsymbol{\theta})$, depending on the mesh width h_j . Additionally, let $N_{J-j+1}^{\text{in-g}}$ denote the total number of $\mathcal{F}_{h_j}(\boldsymbol{\theta})$ evaluations needed to find the dimension-adaptive interpolation surrogate of the potential function depending on the tolerance $\text{tol}_{J-j+1}^{\text{in-g}}$ and the maximum reachable level $K_{\max, J-j+1}^{\text{in-g}}$. Recall that the sequential update of the reference density in the proposed approach necessitates a surrogate for potentials differences (see (7.3)). Thus, at each level except the last, the same potential function and with that the same forward model enters in two different potentials differences. Putting everything together, the total *interpolation cost*, which we denote by $\mathcal{C}_J^{\text{in-g}}$, of the proposed approach when using J levels amounts to

$$\mathcal{C}_J^{\text{in-g}} = \sum_{j=1}^{J-1} (N_{J-j+1}^{\text{in-g}} + N_{J-j}^{\text{in-g}}) \mathcal{C}_{h_j} + N_1^{\text{in-g}} \mathcal{C}_{h_J}.$$

Algorithm 7.10 Multilevel Adaptive Sparse Leja Algorithm

- 1: **procedure** MULTILEVELADAPTSPARSELEJA($J, \mathbf{h}, \mathbf{tol}, \mathbf{K}_{\max}, \Phi, \boldsymbol{\pi}_0, \mathcal{Q}$)
 - 2: Use Algorithm 7.9 to obtain

$$\pi_{\ell(1)}^{\mathbf{y}}(\boldsymbol{\theta}), \mathbf{m}_{\ell(1)}, \mathbf{C}_{\ell(1)} = \text{LevelOneSparseLeja}(h_1, \mathbf{tol}_J^{\text{in-g}}, \mathbf{tol}_J^{\text{qu}}, \mathbf{K}_{\max}, \Phi, \boldsymbol{\pi}_0)$$
 - 3: **for** $j \leftarrow 2, 3, \dots, J$ **do**
 - 4: Construct the Gaussian biasing density (7.1) and the affine mapping (7.2)

$$\widehat{\boldsymbol{\pi}}_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) := \mathbf{N}(\mathbf{m}_{\ell(j-1)}, \mathbf{C}_{\ell(j-1)}), \quad T_{\ell(j-1)}(\boldsymbol{\zeta}) := \mathbf{m}_{\ell(j-1)} + \mathbf{C}_{\ell(j-1)}^{1/2} \boldsymbol{\zeta}$$
 - 5: Compute the potentials ratio surrogate via adaptive interpolation

$$\Phi_{\ell(\delta_j)}(\boldsymbol{\theta}; \mathbf{y}) = \text{AdaptSGInterp}(\mathbf{tol}_{J-j+1}^{\text{in-g}}, K_{\max}^{\text{in-g}}, \Phi_{\delta_j} \circ T_{\ell(j-1)}(\boldsymbol{\zeta}), \widehat{\boldsymbol{\pi}}_{\ell(j-1)}^{\mathbf{y}})$$
 - 6: Construct $L_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta})|\mathbf{y}) := \exp(-\Phi_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta}); \mathbf{y}))$
 - 7: Compute the evidence ratio via adaptive quadrature

$$Z_{\ell(\delta_j)}(\mathbf{y}) = \text{AdaptSGQuad}(\mathbf{tol}_{J-j+1}^{\text{qu}}, K_{\max}^{\text{qu}}, L_{\ell(\delta_j)}, \widehat{\boldsymbol{\pi}}_{\ell(j-1)}^{\mathbf{y}})$$
 - 8: Compute the updated posterior

$$\pi_{\ell(j)}^{\mathbf{y}}(\boldsymbol{\theta}) := \frac{\pi_{\ell(j-1)}^{\mathbf{y}}(\boldsymbol{\theta}) L_{\ell(\delta_j)}(\boldsymbol{\theta}|\mathbf{y})}{Z_{\ell(\delta_j)}(\mathbf{y})}$$
 - 9: Compute the expectation $\mathbf{m}_{\ell(j)}$ of $\pi_{\ell(j)}^{\mathbf{y}}(\boldsymbol{\theta})$ w.r.t. $\widehat{\boldsymbol{\pi}}_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta}))$
 - 10: Compute the covariance $\mathbf{C}_{\ell(j)}$ of $\pi_{\ell(j)}^{\mathbf{y}}(\boldsymbol{\theta})$ w.r.t. $\widehat{\boldsymbol{\pi}}_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta}))$
 - 11: **end for**
 - 12: Compute $\mathcal{Q} \approx \mathcal{Q}_{\ell(J)}$
 - 13: **return** $\mathcal{Q}_{\ell(J)}$
 - 14: **end procedure**
-

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

As the level increases the interpolation is performed w.r.t. a more informed density function, typically with a lower-dimensional structure. Assuming that (i) Φ_{δ_j} is getting close to zero as j increases and (ii) the Gaussian approximation of the posterior from level $j - 1$ is accurate, we expect the number of forward model evaluations needed for interpolation to decrease significantly with the level.

The remaining costs, which are typically due to adaptive sparse grid quadrature, involve evaluations of the interpolation surrogates. Therefore these costs are significantly smaller compared to the interpolation costs, $\mathcal{C}_J^{\text{in-g}}$. Let $\mathcal{C}_{\ell(j)}^{\text{in-g,approx}}$ denote the cost of evaluating once the surrogate of the potential function at level $\ell(j)$ on P processes (since evaluating this surrogate is computationally cheap, usually we have $P = 1$). Additionally, let $N_{J-j+1}^{\text{qu},\mathcal{Q}}$ be the number of surrogate evaluations needed by adaptive sparse grid quadrature to assess a quantity \mathcal{Q} when employing the tolerance $\text{tol}_{J-j+1}^{\text{qu}}$ and the maximum reachable level $K_{\max, J-j+1}^{\text{qu}}$. \mathcal{Q} is, for example, the mean or covariance matrix of the posterior density at level $\ell(j)$, or the quantity of interest computed at the end of the multilevel algorithm. Since all quadrature operations at level $\ell(j)$ are performed w.r.t. the same weight function, the same nodes are likely to be employed in the assessment of different quantities \mathcal{Q} . Thus, we keep a lookup table with all existing evaluations and only evaluate the surrogate at new quadrature nodes. The total *quadrature cost*, denoted by $\mathcal{C}_J^{\text{qu}}$, of the proposed approach on J levels reads

$$\mathcal{C}_J^{\text{qu}} = \sum_{j=1}^J \bigcap_{\mathcal{Q}} N_{J-j+1}^{\text{qu},\mathcal{Q}} \mathcal{C}_{\ell(j)}^{\text{in-g,approx}}.$$

7.2.5 Termination of the multilevel algorithm

One possible stopping criterion of the proposed multilevel approach is to impose a maximum number of levels, J , which has an associated computational cost estimated in Section 7.2.4. Let $p \in \mathbb{R}$ denote a user defined computational budget. Assuming that the sparse grid interpolation costs are significantly larger than the costs of sparse grid quadrature, which depend on evaluating the interpolation surrogates, one criterion to stop the proposed algorithm is

$$\lfloor \mathcal{C}_J^{\text{in-g}} - p \rfloor = 0.$$

In contrast, if the quadrature costs are nonnegligible, then we stop the proposed algorithm when

$$\lfloor (\mathcal{C}_J^{\text{in-g}} + \mathcal{C}_J^{\text{qu}}) - p \rfloor = 0.$$

At each level we find an adaptive sparse grid approximation of the potentials difference, $\Phi_{\delta_j}(\boldsymbol{\theta}; \mathbf{y})$. As the level increases, this difference will get closer and closer to zero. Thus, $L_{\delta_j}(\boldsymbol{\theta}|\mathbf{y}) := \exp(-\Phi_{\delta_j}(\boldsymbol{\theta}; \mathbf{y})) \rightarrow 1$. With this in mind, another possible stopping criterion is to first assess, using Lemma 7.2.1,

$$I_{L_{\ell(\delta_j),1}} := \int_{\mathbb{R}^{d_{\text{sto}}}} \exp(-\Phi_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta}); \mathbf{y})) \pi_G(\boldsymbol{\zeta}) d\boldsymbol{\zeta} \quad (7.6)$$

and check if $|I_{L_{\ell(\delta_j),1}} - 1| \leq \delta_1$ for a user-defined tolerance δ_1 . In other words, if the expectation of the adaptive sparse grid approximation of the likelihood ratio w.r.t. the standard multivariate Gaussian density gets close to 1, the proposed algorithm can be stopped. A stronger stopping criterion would be to compute

$$I_{L_{\ell(\delta_j),2}} := \int_{\mathbb{R}^{d_{\text{sto}}}} \exp(-\Phi_{\ell(\delta_j)}(T_{\ell(j-1)}(\boldsymbol{\zeta}); \mathbf{y})) \frac{\pi_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta}))}{\widehat{\pi}_{\ell(j-1)}^{\mathbf{y}}(T_{\ell(j-1)}(\boldsymbol{\zeta}))} \pi_G(\boldsymbol{\zeta}) d\boldsymbol{\zeta}, \quad (7.7)$$

that is, the evidence ratio from the sequential update (7.3) of the reference information, and then to verify if $|I_{L_{\ell(\delta_j),2}} - 1| \leq \delta_2$ for a prescribed tolerance δ_2 . Put differently, (7.7) evaluates close to 1 if the expectation of the product between the surrogate of the likelihood ratio, $\exp(-\Phi_{\ell(\delta_j)} \circ T_{\ell(j-1)})$, and the bias correcting ratio, $(\pi_{\ell(j-1)}^y / \widehat{\pi}_{\ell(j-1)}^y) \circ T_{\ell(j-1)}$, is close to 1.

7.3 Adaptive strategies for surrogate constructions

Our proposed Multilevel adaptive sparse Leja algorithm is independent of the specific implementations of adaptive sparse grid interpolation and quadrature, i.e., `AdaptSGInterp` and `AdaptSGQuad`. In the following, the considered refinement strategies are presented. To exploit the sparsity of the input parameters w.r.t. the sequentially updated reference information, we employ dimension-adaptivity [71, 90] summarized in Section 2.5.3. We consider a standard dimension-adaptive strategy for both interpolation and quadrature. Since the cost-critical operation in our approach is the sparse grid interpolation (it involves evaluations of the forward map, assumed to be computationally expensive), we formulate an additional context-aware strategy for sparse grid interpolation based on directional variance information. Recall that by $\mathcal{M} \in \{\mathcal{F}, \Phi, L, Z, \pi^y\}$ we denote a generic continuous quantity (forward operator, potential function, likelihood, evidence, posterior density),

7.3.1 Standard dimension-adaptive interpolation and quadrature

We employ the standard dimension-adaptive Algorithm 2.1 for both sparse grid interpolation and quadrature. The essential ingredient of this strategy is the refinement indicator. For sparse grid quadrature, we consider the standard refinement indicator (2.25):

$$\mathcal{I}(\Delta_{\mathbf{k}}^{\text{op}}[\mathcal{M}], \mathcal{C}_{\mathbf{k}}^{\text{op}}) := \|\Delta_{\mathbf{k}}^{\text{qu}}[\mathcal{M}]\|_{L^1} / \Delta L_{\mathbf{k}}^{\text{qu}} = |\Delta_{\mathbf{k}}^{\text{qu}}[\mathcal{M}]| / \Delta L_{\mathbf{k}}^{\text{qu}},$$

which is a surrogate for quadrature error in the subspace associated with \mathbf{k} . For sparse interpolation, we use

$$\mathcal{I}(\Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}], \mathcal{C}_{\mathbf{k}}^{\text{in-g}}) = \|\Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}]\|_{L^2} / \Delta L_{\mathbf{k}}^{\text{in-g}} = \|\Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}]\|_{L^2}.$$

To assess the L^2 norm of surpluses, $\|\Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}]\|_{L^2}$, we proceed as in Section 4.2.2. First, we transform the Lagrange basis into the equivalent PSP basis and find the associated PSP coefficients via (4.6). We then find the L^2 norm of the surpluses using (4.3).

7.3.2 Directional variance dimension-adaptive sparse interpolation

The cost-critical operation in our approach is the sparse grid interpolation since it involves evaluations of the forward map, assumed to be computationally expensive. We hereby formulate an additional context-aware adaptive strategy based on directional variance information.

We discussed in Chapter 4 that although standard refinement indicators for dimension-adaptivity proved to an acceptable heuristic numerous numerical experiments (see, e.g., [35, 71, 90, 193]), they do not inherently distinguish between the individual input parameters. This is particularly relevant in our proposed multilevel approach since we perform the sparse grid operations w.r.t. more informed biasing densities and we thus have information about the sparsity of the stochastic input parameters w.r.t. this biasing density. To better exploit this information in dimension-adaptive interpolation, we enhance the standard dimension-adaptive Algorithm 2.1 with an additional context-aware *filtering* step, as follows.

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

In each refinement step, we additionally compute a *global measure of importance of each input parameter* via *total directional variances*, and we stop refining the input directions for which the associated global measure of importance becomes insignificant; the goal is to preferentially refine the directions rendered stochastically important and thus to exploit the sparsity of the underlying problem in a context-aware manner. The linearity of the global sparse grid approximation formula (2.10) implies that

$$\mathbf{u}_{\mathcal{K}}^{\text{in-g}}[\mathcal{M}] := \sum_{\mathbf{k} \in \mathcal{K}} \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}] = \sum_{\mathbf{k} \in \mathcal{O}} \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}] + \sum_{\mathbf{k} \in \mathcal{A}} \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}],$$

that is, the approximation can be split into the contribution due to the already visited subspaces from the old index set and the contribution of the active set, \mathcal{A} , which contains the subspaces available for refinement. To this end, to enhance the standard adaptive approach for interpolation, we proceed as follows. First, we find the equivalent PSP coefficients using (4.6). Afterwards, we proceed analogously to [193] and (i) perform the Sobol' decomposition [166] of the active set and (ii) make use of the equivalence between PSP coefficients and Sobol' indices [173] (see Section 2.9) to obtain *directional variance surpluses*:

$$\left\| \sum_{\mathbf{k} \in \mathcal{A}} \Delta_{\mathbf{k}}^{\text{in-g}}[\mathcal{M}] \right\|_{L^2}^2 = \sum_{\mathbf{p} \in \mathcal{P}_{\mathcal{A}}} \Delta c_{\mathbf{p}}^2 = (\mathbb{E}_{\mathcal{A}}^{\text{in-g}}[\mathcal{M}])^2 + \sum_{i=1}^{d_{\text{sto}}} \Delta \text{Var}_{\mathcal{A}}^{\text{in-g},i}[\mathcal{M}] + \Delta \text{Var}_{\mathcal{A}}^{\text{in-g,inter}}[\mathcal{M}],$$

where $\mathcal{P}_{\mathcal{A}} := \bigcup_{\mathbf{k} \in \mathcal{A}} \mathcal{P}_{\mathbf{k}}$ is the union of maximum PSP degrees corresponding to all multiindices in \mathcal{A} , i.e., $\mathcal{P}_{\mathbf{k}} := \{\mathbf{p} \in \mathbb{N}^{d_{\text{sto}}} : \mathbf{0} \leq \mathbf{p} \leq \mathbf{P}_{\mathbf{k}}\}$ and $\mathbf{P}_{\mathbf{k}} = (k_1 - 1, k_2 - 1, \dots, k_{d_{\text{sto}}} - 1)$ (recall Section 4.2.2). Moreover, $\mathbb{E}_{\mathcal{A}}^{\text{in-g}}[\mathcal{M}]$ refers to the expectation contribution, corresponding to $\mathbf{p} = \mathbf{0}$,

$$\Delta \text{Var}_{\mathcal{A}}^{\text{in-g},(i)}[\mathcal{M}] := \sum_{\mathbf{p} \in \mathcal{J}_i} \Delta c_{\mathbf{p}}^2, \quad i = 1, 2, \dots, d_{\text{sto}},$$

where $\mathcal{J}_i = \{\mathbf{p} \in \mathcal{P}_{\mathcal{A}} \setminus \{\mathbf{0}\} : \mathbf{p}_i \neq 0 \wedge \mathbf{p}_j = 0, \forall j \neq i\}$, are the *individual directional variance surpluses* corresponding to the d_{sto} directions and $\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,inter}}[\mathcal{M}] := \sum_{\mathbf{p} \in \mathcal{J}_{\text{inter}}} \Delta c_{\mathbf{p}}^2$ refers to the directional variance surplus due to all possible interactions between the d_{sto} stochastic inputs. $\mathcal{J}_{\text{inter}} = \bigcup_{i=1}^{d_{\text{sto}}} \mathcal{J}_{i,\text{inter}}$ such that $\mathcal{J}_{i,\text{inter}}$ comprises all interactions involving the i th input, $\mathcal{J}_{i,\text{inter}} = \{\mathbf{p} \in \mathcal{P}_{\mathcal{A}} \setminus \{\mathbf{0}\} : \mathbf{p}_i \neq 0 \wedge \exists j \neq i, \mathbf{p}_j \neq 0\}$, i.e., all multivariate PSP degrees for which the i th component is nonzero and at least one other component (and at most $d_{\text{sto}} - 1$) is also nonzero. We assess the *total directional variance surpluses* associated to the i th input as

$$\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot}_i}[\mathcal{M}] := \Delta \text{Var}_{\mathcal{A}}^{\text{in-g},i}[\mathcal{M}] + \Delta \text{Var}_{\mathcal{A}}^{\text{in-g,inter}_i}[\mathcal{M}], \quad (7.8)$$

where $\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,inter}_i}[\mathcal{M}] := \sum_{\mathbf{p} \in \mathcal{J}_{i,\text{inter}}} \Delta c_{\mathbf{p}}^2$ denotes the contribution due to all interactions involving direction i . Therefore, $\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot}_i}[\mathcal{M}]$ can be seen as a surrogate used to ascertain the *global* importance for each stochastic input: a large $\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot}_i}[\mathcal{M}]$ implies that the i th stochastic parameter is stochastically significant.

To this end, we prescribe d_{sto} user-defined *directional tolerances* $\boldsymbol{\delta}^{\text{in-g}} := (\delta_1^2, \delta_2^2, \dots, \delta_{d_{\text{sto}}}^2)$ and compare $\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot}_i}$ with δ_i^2 for $i = 1, 2, \dots, d_{\text{sto}}$. When the stochastic direction i is rendered unimportant, i.e., when $\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot}_i}[\mathcal{M}] < \delta_i^2$, we simply stop adding multiindices with the i th component exceeding the maximum i th index in the current multiindex set \mathcal{K} , in all remaining refinement steps. In this way, the algorithm preferentially refines only the important directions and stops adding multiindices in the directions rendered unimportant, which decreases the overall interpolation cost. Note that in the worst case scenario, when neither of the directional tolerances are met, the enhanced algorithm reduces to the standard approach summarized in Algorithm 2.1. We illustrate the presented context-aware filtering in the following example.

Example: Let $d_{\text{sto}} = 3$ and assume that in the current refinement step, the active set is $\mathcal{A} = \{(4, 1, 1), (1, 6, 1), (3, 4, 1), (1, 1, 3), (3, 1, 2), (2, 2, 2), (1, 4, 2)\}$. The multiindex with the largest error indicator is $(4, 1, 1)$. The dimension-adaptive procedure moves $(4, 1, 1)$ to the old-index set, \mathcal{O} , and adds to \mathcal{A} the forward neighbours of $(4, 1, 1)$ which are \mathcal{O} -admissible. These neighbours are $(5, 1, 1)$ and $(4, 2, 1)$.

In the directional variance-enhanced approach, we additionally perform the Sobol' decomposition of the active set. Assume we obtain

$$\Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot1}}[\mathcal{M}] = 10^{-4}, \quad \Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot2}}[\mathcal{M}] = 0.1, \quad \Delta \text{Var}_{\mathcal{A}}^{\text{in-g,tot3}}[\mathcal{M}] = 10^{-2}.$$

Prescribing, for example, the directional tolerances $\boldsymbol{\delta}^{\text{in-g}} = (10^{-3}, 10^{-3}, 10^{-3})$ in the enhanced approach, the first direction would be rendered unimportant. Hence, we stop adding multi-indices with the first index exceeding the maximum reached indices so far. Looking at the active set, \mathcal{A} , this index is $k_1 = 4$. Therefore, from the two \mathcal{O} -admissible forward neighbours of $(4, 1, 1)$, that is, $(5, 1, 1)$ and $(4, 2, 1)$, we would only add $(4, 2, 1)$ to \mathcal{A} .

Remark: (i) The total directional variances defined in (7.8) are nothing else but the numerators of total Sobol' indices (see Section 2.9). Since the denominators of total Sobol' indices are the same, i.e., the output variance, their ordering is the same as the ordering of total directional variance. Total Sobol' indices have also been employed to enhance standard multilevel decompositions with online dimensionality reduction in Chapter 5.

(ii) In the Sensitivity-driven dimension-adaptive sparse grid algorithm proposed in Chapter 4, we performed a Sobol' decomposition in *each* subspace in the active set, \mathcal{A} , i.e., at a *subspace level*, from which we assessed the sensitivity scores. Here, in contrast, we perform a Sobol' decomposition for the *entire* active set to obtain a surrogate for the importance of each parameter at the level of the active set, \mathcal{A} , that is, at a *global level*.

We remark that the presented directional variance-based dimension adaptivity is similar to the algorithm proposed in [193]. However, there are a few differences. Notably, [193] focused on settings in which the model has an input comprising a set of *design* variables and a set of *stochastic* variables, the goal being to distinguish between the two directions. Our adaptive approach here, in contrast, focuses only on the stochastic domain of the problem. Our goal is to exploit the anisotropic coupling of the stochastic inputs at a *global level* to enhance the standard dimension-adaptive algorithm. Furthermore, we formulate our strategy for sparse grid interpolation constructed using (L)-Leja sequences, whereas [193] focuses on PSP constructed using other point sequences, such as Gaussian. Moreover, in this work, we compute total directional variances for the entire active set, whereas [193] computes only upper bounds for the total directional variances.

We summarize the directional variance-enhanced dimension-adaptive algorithm for sparse grid interpolation in Algorithm 7.11. The inputs are the user-defined global interpolation tolerance $tol^{\text{in-g}}$, the maximal attainable level $K_{\text{max}}^{\text{in-g}}$, the multivariate function \mathcal{M} for which we perform the interpolation, and the density $\boldsymbol{\pi}(\boldsymbol{\theta})$ w.r.t. which the interpolation is performed and the associated weighted (L)-Leja points are constructed. The last input is $\boldsymbol{\delta}^{\text{in-g}} := (\delta_1^2, \delta_2^2, \dots, \delta_{d_{\text{sto}}}^2)$, comprising the d_{sto} directional tolerances. In this algorithm, we introduce a new data structure, $\mathcal{V}_{\mathcal{A}}^{\text{tot}}$, to hold the d_{sto} total directional variances computed using (7.8). Note that the total directional variances need to be recomputed at each refinement step as the active set and thus the corresponding PSP basis and coefficients change. The total directional variances are computed at lines 8 – 9 for the first adaptive step and in steps 24 – 25 for the remaining steps. A termination criterion specific to using total directional variances is for all directional variances

Algorithm 7.11 Directional variance dimension-adaptive sparse grid interpolation

```

1: procedure DIRVARADAPTIVITY( $tol^{\text{in-g}}, K_{\text{max}}^{\text{in-g}}, \mathcal{M}, \boldsymbol{\pi}, \boldsymbol{\delta}^{\text{in-g}}$ )
2:    $\mathbf{1} = (1, 1, \dots, 1)$ ,  $\mathcal{O} = \emptyset$ ,  $\mathcal{A} = \{\mathbf{1}\}$ 
3:    $\mathcal{V}_{\mathcal{A}}^{\text{tot}} = \emptyset$ 
4:    $a = \Delta_{\mathbf{1}}^{\text{in-g}}[\mathcal{M}]$ 
5:   Compute  $\epsilon(\mathbf{1}) = \|\Delta_{\mathbf{1}}^{\text{in-g}}[\mathcal{M}]\|_{L^2}$  based on  $\boldsymbol{\pi}(\boldsymbol{\theta})$ 
6:    $\rho = \epsilon(\mathbf{1})$ 
7:   for  $i \leftarrow 1, 2, \dots, d_{\text{sto}}$  do
8:     Compute  $\Delta\text{Var}_{\mathcal{A}}^{\text{in-g}, \text{tot}_i}$  using (7.8)
9:      $\mathcal{V}_{\mathcal{A}}^{\text{tot}} = \mathcal{V}_{\mathcal{A}}^{\text{tot}} \cup \Delta\text{Var}_{\mathcal{A}}^{i, \text{tot}}$ 
10:  end for
11:  while  $\rho \geq tol^{\text{in-g}}$  or  $\mathcal{V}_{\mathcal{A}}^{\text{tot}} \geq \boldsymbol{\delta}^{\text{in-g}}$  or  $\mathcal{A} \neq \emptyset$  or  $\max(\mathcal{K}) \leq K_{\text{max}}^{\text{in-g}}$  do
12:    Select  $\mathbf{k}$  from  $\mathcal{A}$  with the largest  $\epsilon(\mathbf{k})$ 
13:     $\mathcal{A} = \mathcal{A} \setminus \{\mathbf{k}\}$ ,  $\mathcal{O} = \mathcal{O} \cup \{\mathbf{k}\}$ 
14:     $\rho = \rho - \epsilon(\mathbf{k})$ 
15:     $\mathcal{V}_{\mathcal{A}}^{\text{tot}} = \emptyset$ 
16:    for  $i \leftarrow 1, 2, \dots, d_{\text{sto}}$  do
17:       $\mathbf{r} = \mathbf{k} + \mathbf{e}_i$ 
18:      if  $\mathbf{r} - \mathbf{e}_q \in \mathcal{O}$  for all  $q = 1, 2, \dots, d_{\text{sto}}$  then
19:         $\mathcal{A} = \mathcal{A} \cup \{\mathbf{r}\}$ 
20:         $a = a + \Delta_{\mathbf{r}}^{\text{in-g}}[\mathcal{M}]$ 
21:        Compute  $\epsilon(\mathbf{r}) = \|\Delta_{\mathbf{r}}^{\text{in-g}}[\mathcal{M}]\|_{L^2}$  based on  $\boldsymbol{\pi}(\boldsymbol{\theta})$ 
22:         $\rho = \rho + \epsilon(\mathbf{r})$ 
23:      end if
24:      Compute  $\Delta\text{Var}_{\mathcal{A}}^{\text{in-g}, \text{tot}_i}$  using (7.8)
25:       $\mathcal{V}_{\mathcal{A}}^{\text{tot}} = \mathcal{V}_{\mathcal{A}}^{\text{tot}} \cup \Delta\text{Var}_{\mathcal{A}}^{i, \text{tot}}$ 
26:    end for
27:  end while
28:   $\mathcal{K} = \mathcal{O} \cup \mathcal{A}$ 
29:  return  $a$ 
30: end procedure

```

to fall below the tolerances $\delta^{\text{in-g}}$, i.e., $\mathcal{V}_{\mathcal{A}}^{\text{tot}} \geq \delta^{\text{in-g}}$. All other steps, including other termination criteria, are as in standard dimension-adaptive Algorithm 2.1.

The degrees of freedom in dimension-adaptive sparse grid algorithms are the prescribed tolerances. On the one hand, in the standard version, only one (global) tolerance is prescribed. When the heuristic for the global error, ρ , falls below this tolerance, it is assumed that the algorithm converged. Thus, this tolerance serves as a heuristic for the targeted accuracy. On the other hand, in the directional variance-enhanced version presented in Section 7.3.2, we additionally prescribe d_{sto} tolerances, $\delta^{\text{in-g}} = (\delta_1^2, \delta_2^2, \dots, \delta_{d_{\text{sto}}}^2)$. These directional tolerances are used to ascertain the global stochastic importance of an input direction. When one of these tolerances is met, we assume that the associated direction becomes unimportant. The chosen numerical values of the directional variances should reflect any a priori available information about the stochastic importance of each input. For example, in a two-dimensional setting, if the first direction is a priori known to be more important, δ_1^2 should be smaller than δ_2^2 . However, such information is rarely available and we therefore generally recommend prescribing equal directional tolerances, i.e., $\delta^{\text{in-g}} = \delta^2 \cdot \mathbf{1}$.

7.3.3 Illustrative examples

Before presenting our results in Bayesian inverse problems, we first employ the directional variance-enhanced dimension-adaptive interpolation and compare it to interpolation with standard dimension-adaptivity in three forward problems with analytically available models.

Two-dimensional examples

First, we provide the reader with a visual illustration about how the enhanced adaptivity exploits the anisotropy of the stochastic inputs. To this end, we consider two simple models with two stochastic parameters which are uniformly distributed in $[0, 1]^2$. For the standard dimension-adaptive strategy, we consider $\text{tol}^{\text{in-g}} = 10^{-5}$ and $K_{\text{max}}^{\text{in-g}} = 20$. Additionally, we prescribe $\delta^{\text{in-g}} = (10^{-8}, 10^{-8})$ in the directional variance-enhanced algorithm.

Let

$$\mathcal{F}_1(\boldsymbol{\theta}) = \sin(3\theta_1 + 2\theta_2)$$

and

$$\mathcal{F}_2(\boldsymbol{\theta}) = \sin(3\theta_1 + 0.2\theta_2).$$

By construction, we see that the first parameter is the most important in both models. Moreover, in \mathcal{F}_2 , we expect θ_2 to be rather insignificant. To ascertain the stochastic importance of θ_1 and θ_2 , we employ a Gauss-Legendre grid comprising $20^2 = 400$ points and estimate the total Sobol' indices from spectral projection coefficients (see Section 2.9). We obtain

$$\mathcal{F}_1 : \hat{S}_1^T = 0.7909, \quad \hat{S}_2^T = 0.3836, \quad \mathcal{F}_2 : \hat{S}_1^T = 0.9999, \quad \hat{S}_2^T = 0.0190.$$

Thus, we see that indeed θ_1 is more important than θ_2 in both models and in the second model, the importance of θ_2 is negligible. To illustrate how this information has been exploited by the dimension-adaptive algorithms, we visualize, in Figures 7.6 and 7.7, the multiindex sets obtained at the end of the adaptive process. In the left figures, we depict the multiindex sets corresponding to the standard algorithm, whereas in the right figures, we see the sets at the end of the directional variance-enhanced version. First, we observe that the shape of the multiindex sets reflects the coupling of the input parameters: we have more multiindices in the first direction, which is the most important in both models. We also see that the second direction is important as well in the first model. Moreover, since the sum of the two Sobol' indices is about




Figure 7.6: Multiindex sets obtained with dimension-adaptive sparse interpolation of function $f(\theta) = \sin(3\theta_1 + 2\theta_2)$ using the standard (left) and the directional variances-based (right) strategies. While both algorithms detect that θ_1 is more important than θ_2 , the enhanced algorithm stops adding multiindices beyond 8 – for the first component – and larger than 7 – for the second component.




Figure 7.7: Multiindex sets obtained with dimension-adaptive sparse interpolation of function $f(\theta) = \sin(3\theta_1 + 0.2\theta_2)$ using the standard (left) and the directional variances-based (right) strategies. While both algorithms detect that θ_1 is significantly more important than θ_2 , the enhanced algorithm stops adding multiindices beyond 8, for the first component.

1.745 in the first model, the inputs interactions are nonnegligible, thus the algorithm invests effort also in the “mixed” directions, i.e., where both multiindex components are greater than one. In contrast, when considering \mathcal{F}_2 , the algorithms preferentially refine the first direction. Moreover, we observe that for both models, the directional variance-enhanced algorithm led to an earlier stopping of the adaptive refinement. For \mathcal{F}_1 , the enhanced algorithms stopped adding multiindices with the first component larger than eight and the second component larger than six. For \mathcal{F}_2 , the enhanced dimension-adaptivity stopped adding multiindices with the first component exceeding eight.

Five-dimensional example

Let us now look more in depth at the approximation accuracy of the two dimension-adaptive sparse grid strategies. We consider now a forward model with five stochastic inputs uniformly distributed in $[0, 1]^5$, which is again analytically available,

$$\mathcal{F}(\boldsymbol{\theta}) = 1 + \cos(\pi + 1.5\theta_1 + 3.5\theta_2 + 0.05\theta_3 + 0.1\theta_4 + 0.002\theta_5). \quad (7.9)$$

This model is very similar to the five-dimensional example considered in Section 4.2.7, but here we consider a larger multiplication constant for θ_2 (this was 0.5 in Section 4.2.7).

With this setup, the five estimated total Sobol’ indices for global sensitivity analysis estimated using 10^5 Gauss-Legendre nodes read

$$\hat{S}_1^T = 0.3011, \quad \hat{S}_2^T = 0.9439, \quad \hat{S}_3^T = 0.0003, \quad \hat{S}_4^T = 0.0014, \quad \hat{S}_5^T = 5.8162 \cdot 10^{-9}.$$

We see that as the multiplication weights of the five inputs suggest, the first two parameters are significantly more important than the remaining three, with θ_2 being the most important.

We compare the directional variance-enhanced dimension-adaptive sparse grid interpolation with the standard approach w.r.t. their L^2 error

$$\mathcal{E}^2(\mathcal{F} - \mathcal{U}_{\mathcal{L}}[\mathcal{F}]) := \sqrt{\int_{\mathbf{X}} (\mathcal{F}(\boldsymbol{\theta}) - \mathcal{U}_{\mathcal{L}}[\mathcal{F}(\boldsymbol{\theta})])^2 d\boldsymbol{\theta}}$$

which we estimate numerically using 10000 MC samples. For both approaches, $K_{\max}^{\text{in-g}} = 20$ and $\text{tol}^{\text{in-g}} \in [10^{-2}, 10^{-5}]$. Moreover, for the directional variance-based algorithm, we consider the additional tolerances $\boldsymbol{\delta}^{\text{in-g}} = (\text{tol}^{\text{in-g}})^2 \cdot \mathbf{1}_5$. We depict our results in Figure 7.8. We see that for this test case, the directional variance-based adaptivity leads to a more efficient approximation than the standard approach. The enhanced algorithm exploits the fact that three out of five directions are stochastically unimportant and it thus stops refining them early on.

To summarize, we illustrated that directional variance-enhanced dimension adaptivity can lead to further cost reduction as compared to the standard algorithm which is due to exploiting the anisotropic coupling of the stochastic inputs. It is important, however, to remark that the behaviour of these dimension-adaptive methods depends on the structure of the underlying problem: the more structure the model has, the more savings we expect. Moreover, they are based on heuristics, it can happen that in some problems they fail to provide accurate approximations because, e.g., they heuristic fails and the algorithm stops too early.

7.4 Numerical results

We apply the proposed Multilevel adaptive sparse Leja algorithm in four test cases. First, we focus on (L)-Leja quadrature in a simple inversion problem in which the forward model is




Figure 7.8: MC estimate of the L^2 approximation error of dimension-adaptive sparse grid interpolation of the five-dimensional example (7.9) using (i) the standard refinement indicator based on the ratio between the L^2 norm of the surpluses and their cost and (ii) the enhanced version based on directional variances. The estimation was done using 10000 samples.

known analytically (Section 7.4.1). Specifically, we consider an integration problem in which (i) the prior is used as integration weight function and (ii) when the biasing density, i.e., the Gaussian approximation of the posterior, is employed. The following three test cases are characterized by elliptic forward operators in 2D and 3D spatial domains. Therein, we compare the proposed multilevel approach with standard multilevel approximations based on telescoping sums (see Section 2.7), in which all sparse grid operations are performed w.r.t. the prior density. Moreover, since sampling methods are the most popular in Bayesian inversion, we also employ MCMC sampling, but only to obtain a rather qualitative reference. In Section 7.4.2, we apply the proposed approach in a source inversion problem defined in a 2D spatial domain. To test the behaviour of our method in a problem with multimodal posterior, we investigate a source inversion problem with two sources in Section 7.4.3. Finally, in Section 7.4.4, the proposed approach is employed in a more complex and computationally more expensive numerical test case defined in a three-dimensional spatial domain, in which we consider eight stochastic parameters. In all considered test cases, we have non-degenerate Gaussian additive noise and finite dimensional data space. The recent work [110] established that Bayesian inverse problems with this setup are *always* well-posed. All numerical experiments have been performed on a desktop computer with a four core Intel i7-4790 CPU and 24 GB of RAM using standard double precision arithmetics. The FE functionality used for spatial discretization in all four considered test case was provided by the open-source computing platform FEniCS [2].

7.4.1 Simple quadrature showcase

In the first test case, we study the behaviour of weighted (L)-Leja points in integration problems w.r.t. the posterior solution of the Bayesian parameter estimation problem. Consider the

following forward model, which is analytically available:

$$\mathcal{F}(\boldsymbol{\theta}) := \frac{A(\theta_1)}{(w(\theta_2)\pi)^2} (\sin(w(\theta_2)\pi x) - \sin(w(\theta_2)\pi)),$$

where $x \in \Omega := [0, 1]$, $A(\theta_1) = 20\theta_1 + 1$ and $w(\theta_2) = \theta_2 + 1.2$. $\mathcal{F}(\boldsymbol{\theta})$ is the analytical solution to the boundary value problem

$$\begin{aligned} -u_{xx}(x) &= A(\theta_1) \sin(w(\theta_2)\pi x), & x \in \Omega \\ u(x) &= 0, & x \in \partial\Omega, \end{aligned}$$

where u_{xx} denotes the second derivative of u .

The task of the Bayesian inverse problem is to infer (θ_1, θ_2) from noisy observations. Assume that $(\theta_1, \theta_2)_{\text{true}} = (0.45, 0.65)$. The observation data \mathbf{y} are generated synthetically. We take $N_{\text{obs}} = 9$ measurements at locations $0.1, 0.2, \dots, 0.9$, corrupted by additive Gaussian noise $\boldsymbol{\eta} \sim \mathbf{N}(\mathbf{0}, 0.1^2 I)$. Further, we assume a non-informative prior, i.e., the uniform density in $[0, 1]^2$. That is, $\boldsymbol{\pi}_0 = U(0, 1)^2$. To obtain the posterior density, $\pi^{\mathbf{y}}$, we simply apply Bayes' formula (2.3); the evidence $Z(\mathbf{y})$ was computed via dimension adaptive sparse grid quadrature w.r.t. the prior, using a tolerance of 10^{-11} . We depict the posterior density in the left plot in Figure 7.9. Observe that the posterior is unimodal and non-symmetric, but it can be well approximated with a Gaussian density.

Consider the following integration problem

$$\mathcal{I}(g) := \int g(\boldsymbol{\theta}) \pi^{\mathbf{y}}(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (7.10)$$

where $g(\boldsymbol{\theta})$ is an integrable function and $\pi^{\mathbf{y}}(\boldsymbol{\theta})$ is the posterior density. To assess (7.10) we employ adaptive sparse grid quadrature using two different weight functions. In the first case, we employ a tolerance $tol_{\boldsymbol{\pi}_0}^{\text{qu}}$ and perform the quadrature w.r.t. the prior density, $\boldsymbol{\pi}_0(\boldsymbol{\theta})$:

$$(i) \quad \mathcal{I}(g) = \int g(\boldsymbol{\theta}) \frac{L(\boldsymbol{\theta}|\mathbf{y})\boldsymbol{\pi}_0(\boldsymbol{\theta})}{Z(\mathbf{y})} d\boldsymbol{\theta} \approx Z^{-1}(\mathbf{y}) \left(\sum_{n=1}^{N_{\text{pr}}} g(\boldsymbol{\theta}_{n,\text{pr}}) L(\boldsymbol{\theta}_{n,\text{pr}}|\mathbf{y}) w_{n,\text{pr}} \right), \quad (7.11)$$

where $\{\boldsymbol{\theta}_{n,\text{pr}}\}_{n=1}^{N_{\text{pr}}}$ denote the (L)-Leja nodes computed w.r.t. $\boldsymbol{\pi}_0(\boldsymbol{\theta})$ and tolerance $tol_{\boldsymbol{\pi}_0}^{\text{qu}}$.

With the second strategy we assess (7.10) using our proposed approach. Specifically, we employ Lemma 7.2.1 and perform the adaptive sparse grid quadrature based on a prescribed tolerance $tol_{\hat{\boldsymbol{\pi}}^{\mathbf{y}}}^{\text{qu}}$ w.r.t. the biasing density $\hat{\boldsymbol{\pi}}^{\mathbf{y}}(\boldsymbol{\theta})$, which is the Gaussian approximation of the posterior density. The second strategy for integration thus reads

$$\begin{aligned} (ii) \quad \mathcal{I}(g) &= \int g(\boldsymbol{\theta}) \frac{\pi^{\mathbf{y}}(\boldsymbol{\theta})}{\hat{\boldsymbol{\pi}}^{\mathbf{y}}(\boldsymbol{\theta})} \hat{\boldsymbol{\pi}}^{\mathbf{y}}(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int g(T(\boldsymbol{\zeta})) \frac{\pi^{\mathbf{y}}(T(\boldsymbol{\zeta}))}{\hat{\boldsymbol{\pi}}^{\mathbf{y}}(T(\boldsymbol{\zeta}))} \boldsymbol{\pi}_G(\boldsymbol{\zeta}) d\boldsymbol{\zeta} \\ &\approx \sum_{n=1}^{N_{\text{post}}} g(T(\boldsymbol{\zeta}_{n,\text{post}})) \frac{\pi^{\mathbf{y}}(T(\boldsymbol{\zeta}_{n,\text{post}}))}{\hat{\boldsymbol{\pi}}^{\mathbf{y}}(T(\boldsymbol{\zeta}_{n,\text{post}}))} w_{n,\text{post}}, \end{aligned} \quad (7.12)$$

where $\{\boldsymbol{\zeta}_{n,\text{post}}\}_{n=1}^{N_{\text{post}}}$ are (L)-Leja nodes computed based on the standard multivariate normal density, $\boldsymbol{\pi}_G$, and $T(\boldsymbol{\zeta}) := \mathbf{m} + \mathbf{C}^{1/2}\boldsymbol{\zeta}$, where \mathbf{m} and \mathbf{C} are the expectation and covariance matrix associated with the biasing density $\hat{\boldsymbol{\pi}}^{\mathbf{y}}(\boldsymbol{\theta})$.

Remark: \mathbf{m} and \mathbf{C} in (7.12) are assessed via sparse grid quadrature as well. However, since the focus of this test case is on integration w.r.t. the posterior density, we are not concerned with the cost of estimating these two quantities.

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

In the numerical experiments, we consider $g(\boldsymbol{\theta}) := \exp(-\theta_1 - \theta_2)$. First, we compute a sampling solution using $3 \cdot 10^5$ Metropolis-Hastings (MH) MCMC samples stemming from a random walk with initial sample $\boldsymbol{\theta}_1 = (1, 1)$ and Gaussian proposal density with covariance matrix $C_{\text{MH}} = 7 \cdot 10^{-3}I$. The employed setup yields an acceptance rate of 44%. We then employ adaptive sparse grid quadrature w.r.t. the prior density, i.e., (7.11), prescribing the tolerance $\text{tol}_{\pi_0}^{\text{qu}} = 10^{-11}$. Additionally, we employ a tolerance $\text{tol}_{\hat{\pi}^y}^{\text{qu}} = 10^{-5}$ in our approach, (7.12). We compare the two integration strategies in terms of accuracy w.r.t. the sampling result and number of sparse grid quadrature points, i.e., N_{pr} and N_{post} . The results are summarized in Table 7.1. With the employed tolerances, the two integration strategies match four digits of

Table 7.1: Results for the quadrature problem (7.10) using a MH sampling solution with $3 \cdot 10^5$ samples, integration w.r.t. the prior density as in (7.11) and our proposed approach in which we integrate w.r.t. the Gaussian approximation of the posterior, as showed in (7.12).

Method	No. quadrature points	Result
MH	$3 \cdot 10^5$	0.33813
Integration w.r.t. π_0	1603	0.33813
Integration w.r.t. $\hat{\pi}^y(\boldsymbol{\theta})$	49	0.33811

the sampling results; the two employed tolerances were the largest that could ensure four digits of accuracy. However, if we look at the number of sparse grid quadrature nodes, the standard approach requires a total of $N_{\text{pr}} = 1603$ (L)-Leja nodes. In contrast, our approach which uses the Gaussian approximation of the posterior as biasing density requires only $N_{\text{post}} = 49$ quadrature nodes, i.e., about 33 times fewer nodes. This is because the support of the posterior density is significantly smaller than the support of the prior. Thus, when the integration weight function is the prior, the adaptive algorithm places a large number of quadrature points outside of the support of the region of high posterior probability. The quadrature nodes corresponding to (7.11) and (7.12) are visualized in the center and right figures in Figure 7.9, respectively.




Figure 7.9: Left: Posterior density $\pi^y(\boldsymbol{\theta})$. Center: (L)-Leja points computed w.r.t. the uniform prior density used in the integration problem (7.11). Right: Gaussian approximation $\hat{\pi}^y(\boldsymbol{\theta})$ of the posterior and the associated weighted (L)-Leja points used in (7.12).

In this example, we saw that employing the Gaussian approximation of the posterior as biasing density in quadrature computations reduces the overall quadrature effort while maintaining the accuracy of the standard approach. This is because the adaptive algorithm places quadrature nodes in the region of high posterior probability and ignores the regions of low pos-

terior probability, which standard, prior-based approaches are unable to do. In the following test cases, we apply the proposed approach in more complex test cases governed by forward models given in terms of more complex PDEs.

7.4.2 Source inversion with one source in a 2D spatial domain

In the second test case, we apply the proposed multilevel approach in a source inversion problem with one source defined on a two-dimensional spatial domain. Let the forward model $\mathcal{F}_h(\boldsymbol{\theta})$ be the numerical solution of the following elliptic PDE defined on $\Omega := [0, 1]^2$,

$$\begin{aligned} -(u_{xx} + u_{yy})(x, y) &= A(\alpha) \exp(-[(x - \theta_1)^2 + (y - \theta_2)^2]/2\alpha^2), & (x, y) \in \Omega \\ u(x, y) &= 0, & (x, y) \in \partial\Omega, \end{aligned} \quad (7.13)$$

with $A(\alpha) = 5/(2\pi\alpha^2)$ and $\alpha = 0.2$. $\mathcal{F}_h(\boldsymbol{\theta})$ is found using FE.

The goal of the Bayesian parameter estimation is to infer the coordinates (θ_1, θ_2) of the source term on the right-hand side, thus seeking the solution to a source inversion problem. We are therefore solving a 2D stochastic inverse problem defined on a two-dimensional spatial domain. To this end, we perform the multilevel Bayesian inversion as described in Algorithm 7.10 using three levels, i.e., $J = 3$ and $j = 1, 2, 3$. We summarize the setup at the three levels in Table 7.2. For spatial discretization, we employ standard two-dimensional, triangular FE with mesh widths h_j such that $h_2 = h_1/2$ and $h_3 = h_2/2 = h_1/4$. To find surrogates for the potential function at each level we employ dimension-adaptive sparse interpolation. Moreover, for a comprehensive overview of the proposed approach, we employ both the standard and the directional variance-based strategies for adaptive refinement, detailed in Section 7.3. Recall that for the standard adaptive strategy we have tolerances $tol_1^{\text{in-g}}, tol_2^{\text{in-g}}, tol_3^{\text{in-g}}$ whereas for directional variances-based adaptivity from Section 7.3.2 we additionally need the directional tolerances $\{\boldsymbol{\delta}_j^{\text{in-g}}\}_{j=1}^3$. In the multilevel decomposition, we combine h_1 with $tol_3^{\text{in-g}}$ and $\boldsymbol{\delta}_3^{\text{in-g}}$ at level $\ell(1)$, at level $\ell(2)$, h_2 is combined with $tol_2^{\text{in-g}}$ and $\boldsymbol{\delta}_2^{\text{in-g}}$, and at level $\ell(3)$ we employ h_3 together with $tol_1^{\text{in-g}}$ and $\boldsymbol{\delta}_1^{\text{in-g}}$. We choose the FE mesh width $h_1 = \sqrt{2}/2^4$ and the tolerances for adaptive sparse grid interpolation such that the FE and sparse grid approximation errors are quantitatively similar. To perform the integration operations at each level, e.g., the computation of the evidence or of the mean and covariance of the posterior density, we employ dimension-adaptive sparse grid quadrature based on the standard refinement scheme summarized in Algorithm 2.1. To prevent the adaptive quadrature from stopping prematurely, especially when integrating w.r.t. the prior density, we employ small tolerances $tol_1^{\text{qu}}, tol_2^{\text{qu}}, tol_3^{\text{qu}}$. For comparison purposes, we compute

Table 7.2: Multilevel setup for the 2D inversion problem with forward model (7.13).

Level	h	$tol^{\text{in-g}}$	$\boldsymbol{\delta}^{\text{in-g}}$	tol^{qu}
$\ell(1)$	$h_1 = \sqrt{2}/2^4$	$tol_3^{\text{in-g}} = 10^{-5}$	$\boldsymbol{\delta}_3^{\text{in-g}} = (10^{-7}, 10^{-7})$	$tol_3^{\text{qu}} = 10^{-12}$
$\ell(2)$	$h_2 = \sqrt{2}/2^5$	$tol_2^{\text{in-g}} = 10^{-4}$	$\boldsymbol{\delta}_2^{\text{in-g}} = (10^{-6}, 10^{-6})$	$tol_2^{\text{qu}} = 10^{-11}$
$\ell(3)$	$h_3 = \sqrt{2}/2^6$	$tol_1^{\text{in-g}} = 10^{-3}$	$\boldsymbol{\delta}_1^{\text{in-g}} = (10^{-5}, 10^{-5})$	$tol_1^{\text{qu}} = 10^{-10}$

a reference solution using a standard multilevel scheme based on telescoping sums in which the dimension-adaptive sparse grid interpolation and quadrature at all levels are performed w.r.t. the prior density. Finally, we also compute a sampling solution using MH sampling.

We assume $\boldsymbol{\theta}_{\text{true}} = (0.35, 0.65)$, which gives rise to a non-symmetric posterior density about the standard Cartesian coordinate system. To obtain the observation data, 16 sensors are placed

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

at equidistant locations $(0.2n, 0.2m)$ for $n, m = 1, 2, 3, 4$. The 16 measurements are obtained synthetically by evaluating the forward model on a finer mesh of width $h = \sqrt{2}/2^7$ to avoid committing a “Bayesian crime”, and assuming additive Gaussian noise $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, 0.2^2 I)$. The parameters to be inferred, (θ_1, θ_2) , need to reside inside the domain of the forward, $\Omega = [0, 1]^2$. Thus, we choose the prior to be the uniform density in $[0, 1]^2$, i.e., $\boldsymbol{\pi}_0 = U(0, 1)^2$.

To allow for an easier comparison between the proposed approach and the standard multilevel approach, we assess a linear quantity of interest, namely the posterior expectation, i.e., $\mathcal{Q} := \mathbb{E}_{\pi^y}[\boldsymbol{\theta}]$. We first compute a qualitative reference sampling solution using $2 \cdot 10^5$ samples on the finest FE mesh obtained from a random walk MH algorithm with Gaussian proposal having covariance matrix $C_{\text{MH}} = 4 \cdot 10^{-3} I$. The chain starts from $\boldsymbol{\theta}_1 = (1, 1)$. The employed setup yields an acceptance rate of 64%. Moreover, we assess the quantity of interest via the standard three-level approach to obtain a reference multilevel solution. Finally, we apply the proposed multilevel approach using the two adaptive strategies for sparse grid interpolation. To distinguish between the three employed multilevel variants, we use the abbreviation StdML to refer to the standard multilevel approach. Moreover, MLLejaStd stands for our approach in which standard dimension-adaptive interpolation is used at each level, and MLLejaDV is used to refer to our approach combined with directional variance-based adaptive interpolation. The posterior expectation estimates are presented in Table 7.3. Observe that the two variants of our proposed approach, MLLejaStd and MLLejaDV, yield results which are identical to up to the four displayed digits with the results produced by the standard multilevel approach. Moreover, our results are similar to the sampling estimate as well.

Table 7.3: Estimates of $\mathbb{E}[\pi^y(\boldsymbol{\theta})]$ for the source inversion problem with forward model (7.13). We first compute a sampling solution using $2 \cdot 10^5$ MH samples. Afterwards, we employ StdML and the two variants of our proposed multilevel approach, MLLejaStd and MLLejaDV.

Method	$\mathbb{E}_{\pi^y}[\boldsymbol{\theta}]$
MH	(0.3628, 0.6370)
StdML	(0.3631, 0.6368)
MLLejaStd	(0.3630, 0.6369)
MLLejaDV	(0.3630, 0.6369)

Before discussing the computational costs, we take a closer look at each level of the three employed multilevel variants. In Figure 7.10, we illustrate the behaviour of the adaptive sparse grid interpolation. In the left plots, we show the results for StdML, whereas the center and right figures depict the results for MLLejaStd and MLLejaDV. Furthermore, the prior as well as the corresponding (L)-Leja points used in sparse grid interpolation at each level are visualized on the top part. In the bottom figures, we depict the resulting posterior densities. First, note that at the first level, $\ell(1)$, all three multilevel strategies have the same behaviour since the same weight function, i.e., the prior, $\boldsymbol{\pi}_0$, is employed. Observe that a large number of interpolation nodes are placed outside of the region of high posterior probability. In contrast, starting with level $\ell(2)$ the sequential update of the reference density in the proposed approach leads to significantly fewer interpolation points compared to StdML: since the reference density is information, the interpolation nodes are placed in the region of high posterior probability. Comparing the two variants of our proposed approach, MLLejaDV requires fewer (L)-Leja points than MLLejaStd. This is because at both levels $\ell(2)$ and $\ell(3)$, the two total directional variances in MLLejaDV fall below the imposed tolerances. In other words, the values of the surrogates for the global measure of importance of the two stochastic parameters become insignificant. Thus MLLejaDV better exploits the update of the reference information in the underlying test case.




Figure 7.10: Results obtained using StdML (left), MLLejaStd (center) and MLLejaDV (right) to find the adaptive sparse grid interpolation surrogate for the potential function in the source inversion problem with forward model (7.13). At each of the three levels, in the top plots we depict the prior density and the corresponding weighted (L)-Leja points used to find the surrogate. At levels $\ell(2)$ and $\ell(3)$, the prior is the Gaussian approximation of the posterior from the posterior level. In the bottom plots we depict the corresponding posterior density solution.

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

We visualize the multiindex sets for the three multilevel variants in Figure 7.11. At level $\ell(1)$, all three multiindex sets have a symmetric structure since the two stochastic parameters have equal importance relative to the symmetric (and uninformative) uniform prior density. Nevertheless, MLLejaDV leads to a smaller multiindex set since the directional variances $\delta_3^{\text{in-g}}$ fall below $(10^{-7}, 10^{-7})$, but we also see no clear distinction between the two inputs. As expected, at levels $\ell(2)$ and $\ell(3)$ we observe a different behaviour in the two variants of our proposed approach, MLLejaStd and MLLejaDV. Recall that at level $\ell(2)$, interpolation is performed w.r.t. the biasing density which is the Gaussian approximation of the posterior from level $\ell(1)$. The eigenvalues (λ_1, λ_2) of its covariance matrix are $\lambda_1 = 0.0097$ and $\lambda_2 = 0.0055$, indicating that the first direction is more important than the second. This information is reflected in the multiindex sets corresponding to MLLejaStd and MLLejaDV. On the other hand, the multiindex set for the standard multilevel approach remains symmetric because the prior density is unchanged. Finally, at level $\ell(3)$, we see low-cardinality multiindex sets for both MLLejaStd and MLLejaDV. This is because at this stage we have a very informative prior, thus the likelihood ratio is close 1, which requires little approximation effort. This suggests that going beyond level $\ell(3)$ might not be necessary. To ascertain this claim, we compute the two suggested termination criteria from Section 7.2.3, $|1 - I_{L_{\ell(\delta_j),1}}|$ and $|1 - I_{L_{\ell(\delta_j),2}}|$, with $I_{L_{\ell(\delta_j),1}}$ defined in (7.6) and $I_{L_{\ell(\delta_j),2}}$ defined in (7.7). The results are showed in Table 7.4. We observe that indeed, at $\ell(3)$ both $|1 - I_{L_{\ell(\delta_j),1}}|$ and $|1 - I_{L_{\ell(\delta_j),2}}|$ are in $O(10^{-3})$ for MLLejaStd and MLLejaDV.




Figure 7.11: Multiindex sets corresponding to adaptive sparse grid interpolation in the source inversion problem (7.13) for StdML (left), MLLejaStd (center) and MLLejaDV (right).

The total costs of all multilevel methods are visualized in Figure 7.12. For a clearer overview, the costs for both interpolation and quadrature are depicted. The total number of forward model evaluations needed to find the adaptive sparse grid interpolant of the potential function are illustrated in the left figure. In the right figure, we show the total number of surrogate evaluations in all quadrature computations. Note that in all multilevel variants we need quadrature

Table 7.4: Estimates of the two stopping criteria, $|1 - I_{L_{\ell(\delta_j),1}}|$ and $|1 - I_{L_{\ell(\delta_j),2}}|$, suggested in Section 7.2.3, for MLLejaStd and MLLejaDV at levels $\ell(2)$ and $\ell(3)$ in the second test case.

Level	$ 1 - I_{L_{\ell(\delta_j),1}} _{\text{MLLejaStd}}$	$ 1 - I_{L_{\ell(\delta_j),2}} _{\text{MLLejaStd}}$	$ 1 - I_{L_{\ell(\delta_j),1}} _{\text{MLLejaDV}}$	$ 1 - I_{L_{\ell(\delta_j),2}} _{\text{MLLejaDV}}$
$\ell(2)$	0.0106	0.0091	0.0108	0.0093
$\ell(3)$	0.0020	0.0019	0.0019	0.0022

to assess the evidences and expectations at all of the three levels. Additionally, in MLLejaStd and MLLejaDV, we need to compute the covariance matrices at levels $\ell(1)$ and $\ell(2)$, which are needed in the Gaussian biasing density of the associated posteriors and in the affine mapping (recall (7.1) and (7.2)). Observe, first of all, that at level $\ell(1)$, our proposed approach is slightly more expensive for interpolation, which is expected since the FE solver on level $\ell(1)$ is evaluated at level $\ell(2)$ as well: recall that at level $\ell(2)$ we construct a sparse grid surrogate for the ratio of potential functions. However, the increased cost is not significant since it involves evaluations of the coarsest FE solver, which is the cheapest. Starting with level $\ell(2)$, the update of the reference density in the proposed approach leads to significant cost savings for both interpolation and quadrature. For example, at level $\ell(2)$, MLLejaStd leads to about 2 times fewer forward model evaluations for sparse grid interpolation, and around 12.5 fewer surrogate evaluations for sparse grid quadrature. Moreover, at level $\ell(3)$, about 15 times fewer interpolation points and about 7.5 times fewer quadrature nodes are needed. In addition, MLLejaDV leads to about 5 times fewer interpolation nodes and about 12.5 times fewer quadrature evaluations, and the interpolation cost at level $\ell(3)$ is reduced by a factor of about 20, while the quadrature cost by a factor of around 9.5. These results clearly show that updating the reference information in our multilevel approach for Bayesian inversion makes the proposed approach significantly more efficient than the standard multilevel approach based on telescoping sums, which employs the prior density in all computations.




Figure 7.12: Left: total number of forward model evaluations needed in the adaptive sparse interpolation of the potential using the three multilevel variants in the source inversion problem (7.13). Right: total number of quadrature nodes.

7.4.3 Source inversion with two sources in a 2D spatial domain

In the first two test cases, we dealt with unimodal posterior densities for which the corresponding Gaussian bias density (7.1) was an accurate approximation. For a more comprehensive overview of the proposed multilevel approach, we consider now a test case with a multimodal posterior density. In particular, we consider another source inversion test case in which we use two sources to generate the observations, to have bimodal data and thus a bimodal posterior, and one source

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

to perform the Bayesian inference.

The forward operator is again elliptic, defined on $\Omega := [0, 1]^2$:

$$\begin{aligned} -(u_{xx} + u_{yy})(x, y) &= A(\alpha) \left(\exp(-[(x - \theta_1)^2 + (y - \theta_2)^2]/2\alpha^2) \right. \\ &\quad \left. + b \exp(-[(x - \theta_3)^2 + (y - \theta_4)^2]/2\alpha^2) \right), \quad (x, y) \in \Omega \quad (7.14) \\ u(x, y) &= 0, \quad (x, y) \in \partial\Omega, \end{aligned}$$

where $A(\alpha) = 5/(2\pi\alpha^2)$, $\alpha = 0.15$ and the binary parameter $b = 1$ when generating the data and $b = 0$ when performing the inference. Therefore, we are solving a source inversion similar to the one in Section 7.4.2, but starting from bimodal data.

For Bayesian inference we employ StdML, MLLejaStd and MLLejaDV using three levels as well as the standard MH MCMC algorithm, to obtain a sampling solution. The multilevel setup is outlined in Table 7.5. We choose the true locations of the two sources far apart, i.e.,

Table 7.5: Multilevel setup for the 2D inversion problem with forward model (7.14).

Level	h	$tol^{\text{in-g}}$	$\delta^{\text{in-g}}$	tol^{qu}
$\ell(1)$	$h_1 = \sqrt{2}/2^5$	$tol_3^{\text{in-g}} = 10^{-6}$	$\delta_3^{\text{in-g}} = (10^{-8}, 10^{-8})$	$tol_3^{\text{qu}} = 10^{-13}$
$\ell(2)$	$h_2 = \sqrt{2}/2^6$	$tol_2^{\text{in-g}} = 10^{-5}$	$\delta_2^{\text{in-g}} = (10^{-7}, 10^{-7})$	$tol_2^{\text{qu}} = 10^{-12}$
$\ell(3)$	$h_3 = \sqrt{2}/2^7$	$tol_1^{\text{in-g}} = 10^{-4}$	$\delta_1^{\text{in-g}} = (10^{-6}, 10^{-6})$	$tol_1^{\text{qu}} = 10^{-11}$

$(\theta_1, \theta_2)_{\text{true}} = (0.15, 0.15)$ and $(\theta_3, \theta_4)_{\text{true}} = (0.85, 0.85)$. The observation data are generated synthetically. As in the previous test case, 16 sensors are considered at equidistant locations $(0.2n, 0.2m)$ for $n, m = 1, 2, 3, 4$, and to obtain the 16 measurements we evaluate the forward model discretized on a finer mesh with width $h = \sqrt{2}/2^8$. Moreover, additive Gaussian noise $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, 0.2^2 I)$ is assumed. Finally, the prior is uniform, i.e., $\boldsymbol{\pi}_0 = U(0, 1)^2$.

Figure 7.13 depicts the bimodal posterior density obtained via standard Bayes' formula (2.3) for which $50^2 = 2500$ Gauss-Legendre nodes are used to assess the evidence. Note that the two




Figure 7.13: Posterior density for the source inversion problem with forward model (7.14). We used $50^2 = 2500$ Gauss-Legendre points to assess the evidence.

peaks are symmetric around $(0.5, 0.5)$. Therefore, Bayesian inference using only once source in the forward model can yield, in the best case, a posterior density centered at $(0.5, 0.5)$.

The quantity of interest is the posterior mean, i.e., $\mathcal{Q} := \mathbb{E}[\boldsymbol{\pi}^y(\boldsymbol{\theta})]$. We show the obtained estimates in Table 7.6. The MH sampling estimate was computed using $2 \cdot 10^5$ samples on the finest FE mesh, using a random walk Gaussian proposal. Due to having a bimodal posterior with

far apart peaks, we choose a “wider” Gaussian proposal with covariance matrix $C_{\text{MH}} = 10^{-1}I$. The chain starts from $\boldsymbol{\theta}_1 = (1.0, 1.0)$. The employed setup yields an acceptance rate of 45%. We observe that the MCMC and the standard multilevel solutions are close to $(0.5, 0.5)$. However, the estimates given by MLLejaStd and MLLejaDV are far away from this value.

Table 7.6: Comparison of estimates of $\mathbb{E}[\pi^{\mathbf{y}}(\boldsymbol{\theta})]$ for the source inversion problem with forward model (7.14). We first, we compute a sampling solution using $2 \cdot 10^5$ MH samples. Afterwards, we employ StdML and the two variants of our multilevel approach, MLLejaStd and MLLejaDV.

Method	$\mathbb{E}_{\pi^{\mathbf{y}}}[\boldsymbol{\theta}]$
MH	(0.5032, 0.5068)
StdML	(0.5002, 0.5002)
MLLejaStd	(0.6688, 0.6548)
MLLejaDV	(0.6648, 0.6548)

To better understand why our proposed approach yielded inaccurate posterior expectation estimates, we depict, in Figure 7.14, the prior and posterior densities as well as the weighted (L)-Leja points used to construct the adaptive sparse grid interpolation surrogate of the potential function for all employed multilevel methods. We observe that at levels $\ell(2)$ and $\ell(3)$, the region of high probability in the Gaussian approximation is much larger than the region of high probability in the corresponding posterior. This leads to a bias-correcting ratio in quadrature operations, $\pi_{\ell(j)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j)}^{\mathbf{y}}$ (recall (7.3)), that is significantly different from 1 in most regions of the domain. The large variations of this ratio cause large variations of the error indicators in dimension-adaptive sparse grid quadrature which prevent the algorithm to attain the imposed tolerance and to thus yield accurate estimates. In this way, the estimates of the expectation and covariance matrix of the posteriors from levels $\ell(1)$ and $\ell(2)$, and with that, the Gaussian biases employed at levels $\ell(2)$ and $\ell(3)$, are inaccurate. This is also visible by looking at level $\ell(2)$ and especially at level $\ell(3)$, where the (L)-Leja points are unaligned with the coordinate system of the associated Gaussian approximations and very few of them lie in the region of high posterior probability. Note that the large spread of the biasing densities leads to weighted (L)-Leja points outside of the domain of the uniform prior, thus outside the domain of the forward operator, Ω (see (7.14)). However, since these points represent the coordinates of the first source, they need to reside within Ω . Thus, whenever a weighted (L)-Leja point falls outside of Ω , we impose the corresponding likelihood evaluation to be zero.

These results call for a further investigation of the effect of the bias correcting ratio $\pi_{\ell(j)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j)}^{\mathbf{y}}$ in quadrature operations. To this end, we recompute the quantity of interest using the two variants of our proposed approach, MLLejaStd and MLLejaDV without using the bias correcting ratio in any quadrature computation. The quantity of interest estimates are showed in Table 7.7. We observe that both estimates are significantly closer to the MH sampling result and the StdML estimation from Table 7.6. Note also that the two estimates are identical. In fact, in this numerical experiment, the two directional tolerances in MLLejaDV are not attained, thus MLLejaStd and MLLejaDV behave the same.

In Figure 7.15, we depict the prior density and the associated (L)-Leja points used for sparse grid interpolation, as well as the posterior density at levels $\ell(2)$ and $\ell(3)$ for MLLejaStd (left) and MLLejaDV (right). We observe that at level $\ell(2)$, the layout of (L)-Leja points is similar to Figure 7.14, but because we do not use the bias correcting ratio in quadrature computations, the adaptive sparse grid algorithm provides reasonable approximations of the mean and covariance matrix of the posterior. This impacts the computations at the third level, where we observe a different layout of the (L)-Leja points and of the Gaussian biasing density, as compared

Table 7.7: Comparison of estimates of $\mathbb{E}[\pi^{\mathbf{y}}(\boldsymbol{\theta})]$ for the source inversion problem with forward model (7.14) MLEjaStd and MLEjaDV without employing the bias correcting ratio $\pi_{\ell(j)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j)}^{\mathbf{y}}$ in any computation.

Method	$\mathbb{E}_{\pi^{\mathbf{y}}}[\boldsymbol{\theta}]$
MLEjaStd	(0.4694, 0.5097)
MLEjaDV	(0.4694, 0.5097)




Figure 7.15: Results obtained using MLEjaStd (left) and MLEjaDV (right) to find the adaptive sparse grid interpolation surrogate for the potential function in the source inversion problem with forward model (7.14). Different from Figure 7.14, we did not use the bias correcting ratio, $\pi_{\ell(j)}^{\mathbf{y}}/\widehat{\pi}_{\ell(j)}^{\mathbf{y}}$, in any quadrature computation. At each of the three levels, in the top plots we depict the prior density and the corresponding weighted (L)-Leja points used to find the surrogate. In the bottom plots we depict the corresponding posterior density solution.

to what is seen in Figure 7.14. Thus, without using the bias correcting ratio, $\pi_{\ell(j)}^y / \widehat{\pi}_{\ell(j)}^y$, in quadrature computations, our proposed multilevel approach behaves similarly to the sampling method and the reference solution given by StdML. However, this is not a viable solution for multimodal inverse problems in general. In fact, this test case underlines that the proposed multilevel approach might perform poorly in problems with multimodal observation data and thus multimodal posterior. We stress out, nevertheless, that our approach was not designed for multimodal problems, since the current implementation relies on unimodal Gaussian biasing densities. We discuss these aspects and provide an outlook at the end of this chapter.

In summary, in Section 7.4.2, we showed that in a source inversion problem with unimodal posterior, the proposed multilevel approach is significantly more efficient than the standard multilevel approach based on telescoping sums, for both sparse grid interpolation and quadrature operators. Moreover, from the two variants of adaptive interpolation considered in our approach, MLLejaDV proved to be the most accurate. In the source inversion problem in this section, in which two sources were used to generate bimodal observation data, we underlined the main limitation of the proposed multilevel method.

7.4.4 Higher-dimensional problem in a 3D spatial domain

In the fourth test case, we apply the proposed approach in a more challenging and computationally expensive inverse problem in which we infer eight parameters. The forward model is elliptic and defined on $\Omega := [0, 1]^3$,

$$\begin{aligned} -\nabla \cdot (k(x, y, z, \boldsymbol{\theta}) \nabla u(x, y, z)) &= f, & (x, y, z) \in \Omega \\ u(x, y, z) &= 0, & (x, y, z) \in \partial\Omega, \end{aligned} \quad (7.15)$$

where $f \equiv 5$ with a permeability field projected onto a Fourier basis:

$$k(x, y, z, \boldsymbol{\theta}) := \exp \left(\sum_{i=1}^8 s_i \theta_i \sin(p_{i,1}\pi x) \sin(p_{i,2}\pi y) \sin(p_{i,3}\pi z) \right), \quad (7.16)$$

where $s_1 = 0.1785$, $s_2 = s_3 = s_4 = 0.1428$, $s_5 = s_6 = s_7 = 0.1071$ and $s_8 = 0.0714$ are normalized scaling factors, i.e., $\sum_{i=1}^8 s_i = 1$, and $(p_{i,1}, p_{i,2}, p_{i,3}) \in \{1, 2\}^3$. With the above setup, θ_1 is the most important parameter, $\theta_2, \theta_3, \theta_4$ are the second most important parameters etc. under the prior density. Note that (7.16) resembles a Karhunen-Loève (KL) approximation in which $\{\sqrt{s_i}\}_{i=1}^8$ are the eigenvalues and $\{\sin(p_{i,1}\pi x) \sin(p_{i,2}\pi y) \sin(p_{i,3}\pi z)\}_{i=1}^8$ the eigenfunctions. We note that these models are relevant, e.g., in geotechnical engineering (steady-state groundwater flow problems): the elliptic forward model couples the permeability field, k , with the hydrostatic pressure, which is denoted by u in (7.15). This elliptic problem is used extensively in the Bayesian inversion literature, see, e.g., [27, 30, 111, 160, 161, 171, 172] and the references therein.

Bayesian inference is carried out for the weights $(\theta_1, \theta_2, \dots, \theta_8)$ of the permeability field $k(x, y, z, \boldsymbol{\theta})$. As in the previous test cases, we first perform MH MCMC sampling to obtain a sampling solution and we then apply the standard multilevel approach based on telescoping sum as well as the two variants of the proposed approach, MLLejaStd and MLLejaDV. In the multilevel setting, we consider again $J = 3$. The multilevel setup is outlined in Table 7.8. The forward model (7.15) is discretized via standard tetrahedral FE meshes characterized by widths h_j . We choose $h_1 = \sqrt{3}/2^4$ and $h_2 = h_1/2$ and $h_3 = h_2/2 = h_1/4$. Accordingly, the sparse grid interpolation tolerances $tol_j^{\text{in-g}}$ and $\delta_j^{\text{in-g}}$ are chosen to yield quantitatively similar errors to the FE approximation for $j = 1, 2, 3$. Finally, we choose small tolerances for quadrature to prevent the adaptive algorithm from stopping prematurely.

Table 7.8: Multilevel setup for the 8D inversion problem with forward model (7.15).

Level	h	$tol^{\text{in-g}}$	$\delta^{\text{in-g}}$	tol^{qu}
$\ell(1)$	$h_1 = \sqrt{3}/2^4$	$tol_3^{\text{in-g}} = 10^{-5}$	$\delta_3^{\text{in-g}} = 10^{-7} \cdot \mathbf{1}_8$	$tol_3^{\text{qu}} = 10^{-9}$
$\ell(2)$	$h_2 = \sqrt{3}/2^5$	$tol_2^{\text{in-g}} = 10^{-4}$	$\delta_2^{\text{in-g}} = 10^{-6} \cdot \mathbf{1}_8$	$tol_2^{\text{qu}} = 10^{-8}$
$\ell(3)$	$h_3 = \sqrt{3}/2^6$	$tol_1^{\text{in-g}} = 10^{-3}$	$\delta_1^{\text{in-g}} = 10^{-5} \cdot \mathbf{1}_8$	$tol_1^{\text{qu}} = 10^{-7}$

In a KL expansion, $(\theta_1, \theta_2, \dots, \theta_8)$ would be standard (independent) Gaussian random variables. With this in mind, we choose a standard normal prior distribution, i.e., $\pi_0 = \mathbf{N}(\mathbf{0}, I)$. The observation data consists of $N_{\text{obs}} = 729$ measurements at locations $\{0.1, 0.2, \dots, 0.9\}^3 \in \Omega$. These measurements stem from the FE solution to the forward model discretized using a finer mesh width $h = \sqrt{3}/2^7$, and assuming measurement noise $\boldsymbol{\eta} \sim \mathbf{N}(\mathbf{0}, 0.1^2 I)$. In addition, to avoid further biasing, $\boldsymbol{\theta}_{\text{true}}$ is a realization of a standard 8D Gaussian random variable, i.e.,

$$\boldsymbol{\theta}_{\text{true}} = (0.3015, 0.6578, -0.5002, 0.4608, 1.1345, 0.5447, -1.5353, -0.1689).$$

The quantity of interest \mathcal{Q} is again the expectation of the posterior density, $\mathbb{E}_{\pi^y}[\boldsymbol{\theta}]$. The MH sampling solution is computed using 10^5 samples using the finest FE mesh. The Gaussian proposal has covariance matrix $C_{\text{MH}} = 0.5I$. To reduce the burn-in and to obtain an accurate expectation estimate, the chain starts from $\boldsymbol{\theta}_{\text{true}}$. This setup yields a 24% acceptance rate, which is satisfactory given the higher-dimensional setting. Afterwards, we apply the standard multilevel approach based on telescoping sums. We begin at level $\ell(1)$, where all operations are performed w.r.t. the prior density. The sparse grid interpolation surrogate of the potential function is computed adaptively. The surrogate is then used to assess the evidence via dimension-adaptive sparse grid quadrature. Moreover, both the expectation and the covariance matrix of the corresponding posterior are computed as well, since we need them at $\ell(2)$ in our proposed multilevel approach. We obtain, however, an indefinite covariance matrix with a negative variance for θ_1 . This is mainly due to the limitations of performing dimension-adaptivity w.r.t. the prior density, especially when the complexity and dimensionality of the Bayesian inverse problem increase: a large number of sparse grid points are placed outside of the region of high posterior probability, making the approximation inaccurate even when small tolerances are imposed. To overcome this limitation, we employ instead standard sparse grids of a priori fixed levels having sufficiently many points to guarantee a positive definite covariance matrix. In particular, we consider a static sparse grid of level ten for interpolation (24310 grid points) and a quadrature grid of the same level comprising 598417 points.

Remark: Observe the large difference in the number of grid points between interpolation and quadrature. For a grid of same level, ten, in an 8D setting, we have about 25 times more quadrature points. This is because for interpolation, the number of surplus (L)-Leja points per level is always one, whereas for quadrature it varies between one and $2^{d_{\text{sto}}}$ (recall Section 2.5.2).

Since our goal is to compare multilevel methods based on adaptive sparse grid algorithms, we do not perform StdML on $\ell(2)$ and $\ell(3)$ using a priori chosen sparse grids, but rather focus only on the two variants of our proposed approach starting from the Gaussian approximation of the posterior at level $\ell(1)$ obtained with the aforementioned static sparse grids. Moreover, evaluating the FE discretizations depending on h_2 and h_3 on the standard sparse grid of level ten, i.e., 24310 evaluations, requires a rather significant computational cost.

We employ MLLejaStd and MLLejaDV as described above. The results are showed in Table 7.9. The two variants of our approach produce results comparable to the MH solution, thus

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

making the proposed approach comparably accurate to sampling methods in this test case. Observe, however, that the accuracy of the expectation estimates deteriorates compared to θ_{true} as the index increases from one to eight. This is because the likelihood is not informative in all eight directions: the likelihood updates the reference information very well for the first direction, relatively well for the next three, and almost not at all for the last four directions. Nevertheless, since the last four directions are the least important by construction (7.16), we expect that having inaccurate corresponding mean estimates is not too significant.

Table 7.9: Estimation of the posterior’s mean value for the fourth test case using a reference MH solution with 10^5 samples and the two variants of our proposed multilevel approach for Bayesian inversion.

Method	$\mathbb{E}_{\pi^y}[\theta_1]$	$\mathbb{E}_{\pi^y}[\theta_2]$	$\mathbb{E}_{\pi^y}[\theta_3]$	$\mathbb{E}_{\pi^y}[\theta_4]$	$\mathbb{E}_{\pi^y}[\theta_5]$	$\mathbb{E}_{\pi^y}[\theta_6]$	$\mathbb{E}_{\pi^y}[\theta_7]$	$\mathbb{E}_{\pi^y}[\theta_8]$
MH	0.2532	0.2123	-0.1363	0.1326	0.1486	0.0753	-0.1584	0.0066
MLStd	0.2642	0.2111	-0.1630	0.1539	0.1429	0.0670	-0.1816	-0.0053
MLDV	0.2620	0.2114	-0.1600	0.1542	0.1448	0.0689	-0.1808	-0.0050

To assess the quality of the expectation estimates, we construct instances of the permeability field, $k(x, y, z, \theta)$, as follows. First, we use θ_{true} to obtain the reference representation of the permeability field, $k(x, y, z, \theta_{\text{true}})$. We then compare $k(x, y, z, \theta_{\text{true}})$ with representations in which we use the MH MCMC expectation estimate as well as the estimates obtained with the two variants of our multilevel approach. We depict, in Figure 7.16, 2D slices of the field in which the spatial coordinates x (top), y (middle) and z (bottom) are fixed respectively to 0.5. We observe that having good estimates for the first components of θ_{true} is sufficient to obtain good representations of the permeability field: the inaccurate estimates for the latter four (uninformed) components of θ do not significantly affect the estimation, as expected.




Figure 7.16: 2D slices of the 3D permeability field from (7.16) parametrized using θ_{true} and the mean estimates from Table 7.9. Top: yz slice taking $x = 0.5$. Middle: xz slice when $y = 0.5$. Bottom: xy slice fixing $z = 0.5$.

Table 7.10 shows the two stopping criteria estimates, $|1 - I_{L_{\ell(\delta_j),1}}|$ and $|1 - I_{L_{\ell(\delta_j),2}}|$, for both MLLejaStd and MLLejaDV at levels $\ell(2)$ and $\ell(3)$. We observe, first of all, that at level $\ell(2)$, these estimates are quite large, in $O(10^{-1})$. Hence, the expectation of the likelihood ratio is significantly different from 1, indicating that we need to go beyond level $\ell(2)$. At level $\ell(3)$ the estimates are in $O(10^{-3})$, which suggests that we can stop the multilevel algorithm here. Furthermore, $|1 - I_{L_{\ell(\delta_j),2}}|$, which incorporates the bias correcting ratio, is small as well, which tells us that the Gaussian biasing density approximates well the region of high posterior probability, hence its expectation and covariance matrix have been accurately estimated.

Table 7.10: Estimates of the two stopping criteria, $|1 - I_{L_{\ell(\delta_j),1}}|$ and $|1 - I_{L_{\ell(\delta_j),2}}|$, suggested in Section 7.2.3, for MLLejaStd and MLLejaDV at levels $\ell(2)$ and $\ell(3)$ in the fourth test case.

Level	$ 1 - I_{L_{\ell(\delta_j),1}} _{\text{MLLejaStd}}$	$ 1 - I_{L_{\ell(\delta_j),2}} _{\text{MLLejaStd}}$	$ 1 - I_{L_{\ell(\delta_j),1}} _{\text{MLLejaDV}}$	$ 1 - I_{L_{\ell(\delta_j),2}} _{\text{MLLejaDV}}$
$\ell(2)$	0.1182	0.1203	0.1203	0.1154
$\ell(3)$	0.0139	0.0021	0.0081	0.0082

In Figure 7.17, we depict the costs of MLLejaStd and MLLejaDV in terms of number of (L)-Leja points for interpolation (left) and quadrature (right). Note that for interpolation we also show costs for level $\ell(1)$ because evaluations of the forward operator discretization depending on h_1 are employed at level $\ell(2)$ as well. For interpolation, MLLejaDV is cheaper than MLLejaStd at all of the three levels, requiring about 4.3 times fewer evaluations at level $\ell(1)$, 4.8 times fewer evaluations at level $\ell(2)$, and 7 times fewer evaluations at $\ell(3)$. Interestingly, the overall interpolation costs are very small given that we have an 8D inversion problem at hand. For example, at level $\ell(3)$, only 22 PDE evaluations are sufficient for MLLejaDV: the maximally reached level in the corresponding multiindex set is 4, and all multiindices have components larger than 1 only in the first four directions. Thus, the directional variances-based algorithm detects early on that the latter four directions are uninformed and unimportant and therefore preferentially refines only the first four directions. Note that the reduced costs at level $\ell(3)$ are in accordance to the small stopping criteria estimates from Table 7.10 as well. We see once again that using the enhanced adaptive algorithm for adaptive sparse grid interpolation leads to a nonnegligible cost reductions. For quadrature, the total number of surrogate evaluations are similar. Nevertheless, given the large difference in runtime between evaluating the 3D FE discretization of the forward model and the interpolation surrogates, the quadrature costs are smaller, as compared to the interpolation costs.




Figure 7.17: Left: total number of forward model evaluations needed by the two variants of our proposed approach, MLLejaStd and MLLejaDV, at all three levels in the 8D inversion problem with forward model (7.15). Right: total number of quadrature nodes used in all three in MLLejaStd and MLLejaDV at levels $\ell(2)$ and $\ell(3)$.

In summary, we employed the proposed multilevel adaptive approach for Bayesian inference

CHAPTER 7. MULTILEVEL ADAPTIVE SPARSE GRID LEJA APPROXIMATIONS FOR BAYESIAN INFERENCE

in a problem defined on a 3D spatial domain, with eight uncertain inputs. We underlined the main limitation of using dimension-adaptivity w.r.t. the prior density in higher-dimensional Bayesian inference problems: as the dimensionality increases and when the region of high posterior probability is (significantly) smaller than the region of high prior probability, standard dimension-adaptive algorithms can yield inaccurate results. This limitation is not inherent in our approach, since alternative methods, e.g., a Laplace approximation of the posterior, can be used at the first level. We saw that the sequential update of the reference information in the proposed approach led to efficient approximations at levels $\ell(2)$ and $\ell(3)$ with an accuracy comparable to the considerably more expensive MH solution. Moreover, the MLLejaDV variant of our approach proved again to be the most efficient method.

Conclusion and remarks

We presented a Multilevel adaptive sparse Leja algorithm for efficient Bayesian inference in higher-dimensional, computationally expensive scenarios. Our method is based on two layers of *model hierarchies*, one based on global sparse grid approximations and another hierarchy based on sequentially-updated multilevel decompositions, which were performed implicitly, not explicitly via telescoping sums. At each level in the multilevel decompositions, sparse grid interpolation is employed to adaptively find a surrogate of the potential or data misfit function. Hereafter, this surrogate is used in all other operations, typically involving quadrature, which we perform via dimension-adaptive sparse grid quadrature. The main novelty of the proposed approach is the sequential update of the reference information such that, starting with the second level, the previous level posterior is reused as the prior at the current level. In this way, the adaptive algorithms can place the sparse grid points in the region of high posterior probability. Moreover, the properties of the potential function w.r.t. the updated reference density, such as sparsity, can be exploited by the adaptive algorithms to further reduce the cost of interpolation and quadrature. For a broader overview, we considered two adaptive strategies for sparse grid interpolation, one based on standard refinement indicators and the other context-aware strategy employing total directional variances to preferentially refine only the important stochastic directions. Since the posterior density does not generally have a product structure, which is required by sparse grid approximations, at each level we employed a biasing density, i.e., the Gaussian approximation of the posterior. The four presented numerical test cases showed the usefulness but also the main limitation of the proposed method. In the first test case, we studied quadrature w.r.t. the prior density and the Gaussian biasing density, showing that our proposed approach for quadrature can lead to significant cost savings. In the second test case, a source inversion in a two-dimensional spatial domain with one source, we applied the proposed multilevel method and we compared it with standard MCMC sampling and a standard multilevel strategy employing telescoping sums. The results showed that our multilevel methodology leads to significant cost savings for both interpolation and quadrature, while maintaining the accuracy of the standard multilevel scheme. The third test case underlines that the proposed approach is likely to yield inaccurate results when the observation data is multimodal. This is because by construction, our multilevel algorithm is not designed to handle multimodal densities: the Gaussian biasing of the posterior employed at levels greater than or equal to two in our approach performs well in unimodal settings. There are several possibilities to address this limitation. A first idea would be to employ a (Gaussian) mixture model instead of a biasing in terms of a single Gaussian density. However, it is unclear how to then use this mixture model with adaptive sparse grid interpolation, for example. A more promising alternative is to instead employ transport maps [121], which provide a more versatile framework for complex, nonlinear Bayesian inverse problems. Note, however, that standard transport maps

are typically parametrized in terms of global basis functions, such as polynomials, which would fail to capture the multimodal nature of the density. To generalize the proposed approach to multimodal Bayesian inverse problems, part of our future research is to design transport maps parametrized in terms of basis functions with localized support, such as hat functions or radial basis functions, or, alternatively to develop transport map-based strategies alongside the recently developed Stein variational Newton method (see [31, 41, 117]). Finally, in the fourth test case, we applied the proposed approach in an eight-dimensional inference problem defined on a 3D spatial domain and we compared it with MCMC sampling and the standard multilevel approach as well. Therein, the likelihood was informative for only the first few directions. The standard multilevel decomposition yielded an indefinite covariance matrix for the biasing density at the first level, underlying the main limitations of standard sparse grid approximations in Bayesian inversion: performing sparse grid approximations w.r.t. the given prior quickly becomes prohibitive as the dimensionality and complexity of the inference problem increase. We thus decided to only focus on the two variants of our proposed approach. Of the two variants, the strategy employing directional variances for adaptive interpolation proved to better exploit the low-dimensional structure given by the non-informative likelihood to find the surrogates at a very low computational cost.

This concludes the fourth and last contribution chapter of this work. In the last part, we give a summary of all our contributions and results, and an outlook for future research directions.

PART V

CONCLUSION

"The end of a melody is not its goal: but nonetheless, had the melody not reached its end it would not have reached its goal either. A parable. "

Friedrich Nietzsche

8

Summary and conclusion

We come to an end and summarize again the main contributions of this work, in Section 8.1, and give an outlook for the next research steps, in Section 8.2.

8.1 What was achieved

In summary, our main motivation in this work was to overcome, at least to a certain extent, the computational challenges of quantifying uncertainty in *higher-dimensional*, cost intensive, practically relevant applications. To this end, we formulated and tested four novel algorithms for efficient black box uncertainty propagation (first three) and Bayesian inference (the fourth approach) based on approximation with sparse grids, multilevel decompositions and multifidelity sampling methods. We showed that intuitively, multilevel and multifidelity methods can be viewed as (generalized, weighted) combination schemes and, moreover, that all three strategies are based on *model hierarchies*. In particular, in global sparse grid approximation, the given, high-fidelity model is used in all computations, the hierarchy appearing in the stochastic space, where different levels of approximation are used. Going one step further, in multilevel decompositions, the forward model is discretized on different grids, and these model discretizations are combined with a (complementary) sampling hierarchy of the stochastic space. Here, the high-fidelity model is the forward model discretization with the highest resolution. Finally, in multifidelity sampling methods, heterogeneous models with different fidelities are used to estimate statistics of the high-fidelity model output. The high-fidelity model is given, whereas the other models in the hierarchy, the so-called low-fidelity models can be heterogeneous, such as data-fit, projection-based or reduced physics models.

The four proposed approaches are *context-aware*, in the sense that to overcome the computational challenges of *higher-dimensional uncertainty quantification*, they exploit the structure of the underlying problem. In uncertainty propagation settings, by structure we meant properties such as anisotropic coupling of the stochastic inputs, the lower stochastic dimensionality or the accuracy and cost rates of reduced models. In the context of Bayesian inference, we employed a sequential update of the reference information to construct posterior-focused approximations of the posterior. To find these approximations efficiently, we exploited the anisotropic coupling of the input parameters under the updated reference information.

We divided our contributions into three parts (Part II – IV). In Part II, we formulated two context-aware algorithms for uncertainty propagation based on sparse grid approximations and multilevel decompositions:

- In Chapter 4, we formulated the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm. In this algorithm, we employed model hierarchies appearing in sparse grid interpolation and pseudo-spectral projections. The main novelty was a context-aware refinement indicator for dimension-adaptivity based on sensitivity information, whose goal

was to explore and exploit the anisotropic coupling of the stochastic inputs and their lower intrinsic dimensionality. We applied this algorithm in a real-world problem, the analysis of plasma microturbulence, and considered two stochastic scenarios with three, eight or 12 uncertain inputs. Moreover, for a broader overview, we considered two output of interest. To the best of our knowledge, this represented one of the first uncertainty quantification studies in plasma microinstability analysis. Our results showed that the proposed approach was more efficient than dimension-adaptivity based on standard refinement indicators. Moreover, the performed uncertainty and sensitivity analyses aided the physical interpretation of the results and revealed insights about the underlying microturbulence.

- The Multilevel adaptive spectral projection with stochastic dimensionality reduction was proposed in Chapter 5. This algorithm was based on hierarchies appearing in multilevel decompositions. The main novelty was the enhancement of standard multilevel decompositions based on telescoping sums with the capability to perform context-aware, online stochastic dimensionality reduction at no additional computational cost. Moreover, at each level in the multilevel decomposition, we formulated a novel approach to compute pseudo-spectral coefficients via spatially adaptive sparse grids and products of one-dimensional integrals. As an example of complex and computationally expensive problem, we applied the proposed approach in a fluid-structure interaction simulation with five uncertain inputs. We showed that online stochastic dimensionality reduction leads to significant cost savings by reducing the number of necessary forward runs. Moreover, in a model problem with five and then with six uncertain parameters, we showed that our algorithm maintains the accuracy of equivalent approaches with no dimensionality reduction, while reducing the computational cost as well.

In Part III, we formulated our third and last contribution for uncertainty propagation. Therein, we formulated the

- Context-aware multifidelity Monte Carlo sampling algorithm, in Chapter 6. This algorithm was based on multifidelity model hierarchies in which the low-fidelity models were explicitly constructed for being used together with the high-fidelity model. The main contribution of this method was the learning of context-aware low-fidelity models for the estimation of statistics of high-fidelity model outputs. The surrogates were constructed by quasi-optimally trading off their refinement, to improve their deterministic approximation quality, with sampling the models, to reduce the statistical error. We first considered the bifidelity setting. Our analysis showed that the quasi-optimal computational effort to spend on improving the surrogates is bounded, which meant that low-fidelity models can become too accurate for multifidelity methods, which is in stark contrast to traditional model reduction. We next considered the setting with an arbitrary number of surrogates, in which we extended, numerically, the results from the bifidelity setup. We applied the proposed context-aware approach in a thermal block problem with nice uncertain inputs in which we considered two surrogates with heterogeneous accuracy and cost rates. We showed that our algorithm is more efficient than both standard Monte Carlo sampling and static multifidelity Monte Carlo sampling in which the surrogates were constructed with either fewer or more high-fidelity model evaluations than were used to find the context-aware surrogate in our approach. We then applied the new approach to a plasma microturbulence simulation scenario with 12 uncertain inputs in which we showed again that our algorithm is more efficient than standard multifidelity Monte Carlo sampling. To the best of our knowledge, this was one of the first times when multifidelity methods are employed to quantify uncertainty in plasma microturbulence analysis.

Finally, in Part IV, we switched focus from uncertainty propagation to Bayesian inference and formulated our fourth and last contribution in this work:

- The Multilevel adaptive sparse Leja algorithm was proposed in Chapter 7. It is a fully deterministic, sampling-free, adaptive approach based on model hierarchies appearing in multilevel decompositions. Moreover, at each level in the multilevel hierarchy, we employed dimension-adaptive sparse grid interpolation and quadrature, hence, a second layer of model hierarchies. The main novelty of this algorithm was the sequential update of the reference information. This allowed us to place interpolation and quadrature points in the region of high-posterior probability and hence construct posterior-focused surrogates at no additional cost. For interpolation, besides standard dimension-adaptivity, we also considered an enhanced version in which a filtering step was used to prevent the algorithm to further refine the directions rendered unimportant. Our results in three test cases with elliptic forward models in 2D and 3D spatial domains and two or eight uncertain inputs, showed that the proposed approach is more efficient than an equivalent multilevel method based on telescoping sums. In addition, the enhanced dimension-adaptive approach proved to be the most cost-effective for adaptive sparse grid interpolation. We also considered a test case to underline the main limitation of the current approach.

8.2 Outlook

In our future research, we aim to enhance the proposed algorithms and to apply them to other realistic and practically relevant applications, or in other scenarios involving the applications considered in this work. Specifically, we plan to

- enhance the Sensitivity-driven dimension-adaptive sparse grid approximation algorithm to settings in which the uncertain inputs are dependent. A promising approach to this end is employing transport maps [121], which already gained significant attention in the Bayesian inference community.
- extend the Multilevel adaptive spectral projection with online stochastic dimensionality reduction also to settings with dependent stochastic inputs. Recent years witnessed significant progress in the development of methods for sensitivity analysis of systems with dependent inputs, such as active subspaces [36]. These methods can be easily be incorporated in our approach.
- extend the mathematical analysis from the bifidelity setting to the setup with an arbitrary number of (heterogeneous) low-fidelity models in the Context aware multifidelity Monte Carlo sampling algorithm. We want to find the conditions for which the number of high-fidelity model evaluations to construct and improve the surrogates exists, is unique and is bounded for any budget.
- enhance the Multilevel adaptive sparse Leja algorithm to settings with multimodal posteriors. One solution would be use locally-parametrized transport maps or the recently developed Stein Variational Newton methods (see [31, 41, 117]). Moreover, in situations in which the number of posterior modes is a priori known, we could use, e.g., clustering techniques to determine the location of the cluster and then apply a domain decomposition-like approach in which we construct a separate transport map for each mode.

Moreover, some of the proposed algorithms can be also combined in future test cases. For example, in plasma microturbulence analysis, a complete simulation pipeline roughly includes:

1. gathering the experimental data;
2. analysing the experimental data, typically via Bayesian approaches,;
3. creating temperature and density gradients profiles from these data;
4. input these profiles into the simulation code, e.g., GENE and perform subsequent forward simulations to understand the nature of the microinstability.

In current simulations, the third step is usually deterministic, in the sense that the temperature and density gradient profiles are usually the mean or maximum a posteriori estimate of the Bayesian posterior density stemming from the second step. To exploit that the second step provides, in fact, a quantification of uncertainty in the input parameters, we could, for example

1. employ the presented or an enhanced version of the Multilevel adaptive sparse Leja algorithm to aid the Bayesian data analysis step;
2. make use of either of the three other (enhanced) algorithms to quantify the effects of the input uncertainty in the underlying microturbulence modes.

Finally, one milestone for plasma physics simulations is to quantify uncertainty in fully nonlinear gyrokinetic simulations. As the high-fidelity model in this settings typically requires days, maybe weeks for a single run, the most promising approach to quantify uncertainty in such problems is to employ multifidelity methods. We already showed that our proposed Context-aware multifidelity sampling algorithm is superior to the standard version. Thus, in a next step, we will apply it to the setting in which the high-fidelity model is quasi-linear, and then to fully nonlinear simulations. That being said, this thesis comes to an end.

Bibliography

- [1] A. ALEXANDERIAN, *On spectral methods for variance based sensitivity analysis*, Probability Surveys, 10 (2013), pp. 51–68.
- [2] M. S. ALNÆS, J. BLECHTA, J. HAKE, A. JOHANSSON, B. KEHLET, A. LOGG, C. RICHARDSON, J. RING, M. E. ROGNES, AND G. N. WELLS, *The fenics project version 1.5*, Archive of Numerical Software, 3 (2015).
- [3] S. AMARAL, D. ALLAIRE, AND K. WILLCOX, *Optimal l_2 -norm empirical importance weights for the change of probability measure*, Statistics and Computing, 27 (2017), pp. 625–643.
- [4] F. AXISA AND J. ANTUNES, *Modelling of Mechanical Systems: Fluid-Structure Interaction*, vol. 3, Elsevier, 2007.
- [5] I. BABUSKA, F. NOBILE, AND R. TEMPONE, *A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data*, SIAM Journal on Numerical Analysis, 45 (2007), pp. 1005–1034.
- [6] A. BARTH, C. SCHWAB, AND N. ZOLLINGER, *Multilevel monte carlo finite element method for elliptic pdes with stochastic coefficients*, Numerische Mathematik, 119 (2011), pp. 123–161.
- [7] G. BATCHELOR, *An introduction to fluid mechanics*, Cambridge University Press, 1967.
- [8] Y. BAZILEVS, K. TAKIZAWA, AND T. TEZDUYAR, *Computational Fluid-Structure Interaction: Methods and Applications*, Wiley, 2013.
- [9] M. A. BEER, S. C. COWLEY, AND G. W. HAMMETT, *Field-aligned coordinates for nonlinear simulations of tokamak turbulence*, Physics of Plasmas, 2 (1995), pp. 2687–2700.
- [10] P. BENNER, M. OHLBERGER, A. COHEN, AND K. WILLCOX, *Model Reduction and Approximation*, Society for Industrial and Applied Mathematics (SIAM), 2017.
- [11] J.-P. BERRUT AND L. N. TREFETHEN, *Barycentric Lagrange interpolation*, SIAM Review, 46 (2004), pp. 501–517.
- [12] A. BESKOS, M. GIROLAMI, S. LAN, P. E. FARRELL, AND A. M. STUART, *Geometric mcmc for infinite-dimensional inverse problems*, Journal of Computational Physics, 335 (2017), pp. 327 – 351.
- [13] C. BIERIG AND A. CHERNOV, *Estimation of arbitrary order central statistical moments by the multilevel monte carlo method*, Stochastics and Partial Differential Equations Analysis and Computations, 4 (2016), pp. 3–40.
- [14] D. BIGONI, O. ZAHM, A. SPANTINI, AND Y. MARZOUK, *Greedy inference with layers of lazy maps*, arXiv e-prints, (2019).
- [15] B. BOHN, M. GRIEBEL, AND J. OETTERSCHAGEN, *Optimally rotated coordinate systems for adaptive least-squares regression on sparse grids*, in Proceedings of the 2019 SIAM International Conference on Data Mining, 2019, pp. 163–171.
- [16] L. M. BREKHOVSKIKH AND V. GONCHAROV, *Mechanics of Continua and Wave Dynamics*, Springer, 1993.

BIBLIOGRAPHY

- [17] A. J. BRIZARD AND T. S. HAHM, *Foundations of nonlinear gyrokinetic theory*, *Reviews of Modern Physics*, 79 (2007), pp. 421–468.
- [18] S. BROOKS, A. GELMAN, G. L. JONES, AND X.-L. MENG, *Handbook of Markov chain Monte Carlo*, Chapman & Hall/CRC Handbooks of Modern Statistical Methods, CRC Press, Boca Raton, FL, 2011.
- [19] H.-J. BUNGARTZ, *Finite Elements of Higher Order on Sparse Grids*, habilitationsschrift, Fakultät für Informatik, Technische Universität München, Nov. 1998.
- [20] H.-J. BUNGARTZ AND M. GRIEBEL, *Sparse grids*, *Acta Numerica*, 13 (2004), pp. 147–269.
- [21] H.-J. BUNGARTZ, F. LINDNER, B. GATZHAMMER, M. MEHL, K. SCHEUFELE, A. SHUKAEV, AND B. UEKERMANN, *precice – a fully parallel library for multi-physics surface coupling*, *Computers & Fluids*, (2016).
- [22] T. BUTLER, J. JAKEMAN, AND T. WILDEY, *Combining push-forward measures and bayes’ rule to construct consistent solutions to stochastic inverse problems*, *SIAM Journal on Scientific Computing*, 40 (2018), pp. A984–A1011.
- [23] T. BUTLER, J. D. JAKEMAN, AND T. WILDEY, *A Consistent Bayesian Formulation for Stochastic Inverse Problems Based on Push-forward Measures*, arXiv e-prints, (2017).
- [24] R. CAMERON AND W. MARTIN, *The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals*, *Annals of Mathematics*, 48 (1947), pp. 385–392.
- [25] P. CHEN, A. QUARTERONI, AND G. ROZZA, *Reduced basis methods for uncertainty quantification*, *SIAM/ASA Journal on Uncertainty Quantification*, 5 (2017), pp. 813–869.
- [26] P. CHEN AND C. SCHWAB, *Sparse-grid, reduced-basis Bayesian inversion*, *Computer Methods in Applied Mechanics and Engineering*, 297 (2015), pp. 84–115.
- [27] ———, *Adaptive sparse grid model order reduction for fast Bayesian estimation and inversion*, in *Sparse grids and applications—Stuttgart 2014*, vol. 109 of *Lecture Notes in Computational Science and Engineering*, Springer, Cham, 2016, pp. 1–27.
- [28] ———, *Sparse-grid, reduced-basis Bayesian inversion: nonaffine-parametric nonlinear equations*, *Journal of Computational Physics*, 316 (2016), pp. 470–503.
- [29] P. CHEN AND C. SCHWAB, *Model Order Reduction Methods in Computational Uncertainty Quantification*, Springer International Publishing, Cham, 2017, pp. 937–990.
- [30] P. CHEN, U. VILLA, AND O. GHATTAS, *Hessian-based adaptive sparse quadrature for infinite-dimensional Bayesian inverse problems*, *Computer Methods in Applied Mechanics and Engineering*, 327 (2017), pp. 147–172.
- [31] P. CHEN, K. WU, J. CHEN, T. O’LEARY-ROSEBERRY, AND O. GHATTAS, *Projected Stein Variational Newton: A Fast and Scalable Bayesian Inference Method in High Dimensions*, arXiv e-prints, (2019).
- [32] C. CLENSHAW AND A. CURTIS, *A method for numerical integration on an automatic computer*, *Numerische Mathematik*, 2 (1960), pp. 197–205.

- [33] K. A. CLIFFE, M. B. GILES, R. SCHEICHL, AND A. L. TECKENTRUP, *Multilevel monte carlo methods and applications to elliptic pdes with random coefficients*, Computing and Visualization in Science, 14 (2011).
- [34] N. COLLIER, A.-L. HAJI-ALI, F. NOBILE, E. VON SCHWERIN, AND R. TEMPONE, *A continuation multilevel monte carlo algorithm*, BIT Numerical Mathematics, 55 (2015), pp. 399–432.
- [35] P. CONRAD AND Y. MARZOUK, *Adaptive Smolyak pseudospectral approximations*, SIAM Journal on Scientific Computing, 35 (2013), pp. A2643 – A2670.
- [36] P. G. CONSTANTINE, *Active Subspaces*, Society for Industrial and Applied Mathematics (SIAM), 2015.
- [37] P. G. CONSTANTINE, M. S. ELDRED, AND E. T. PHIPPS, *Sparse pseudospectral approximation method*, Computer Methods in Applied Mechanics and Engineering, 229-232 (2012), pp. 1–12.
- [38] R. CRISOVAN, D. TORLO, R. ABGRALL, AND S. TOKAREVA, *Model order reduction for parametrized nonlinear hyperbolic problems as an application to uncertainty quantification*, Journal of Computational and Applied Mathematics, 348 (2019), pp. 466 – 489.
- [39] T. CUI, K. J. LAW, AND Y. M. MARZOUK, *Dimension-independent likelihood-informed mcmc*, Journal of Computational Physics, 304 (2016), pp. 109 – 137.
- [40] P. DEL MORAL, A. DOUCET, AND A. JASRA, *Sequential Monte Carlo samplers*, Journal of the Royal Statistical Society. Series B. Statistical Methodology, 68 (2006).
- [41] G. DETOMMASO, T. CUI, A. SPANTINI, Y. MARZOUK, AND R. SCHEICHL, *A stein variational newton method*, in Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS’18, USA, 2018, Curran Associates Inc., pp. 9187–9197.
- [42] G. DETOMMASO, T. DODWELL, AND R. SCHEICHL, *Continuous level monte carlo and sample-adaptive model hierarchies*, SIAM/ASA Journal on Uncertainty Quantification, 7 (2019), pp. 93–116.
- [43] G. DETOMMASO, J. KRUSE, L. ARDIZZONE, C. ROTHER, U. KÖTHE, AND R. SCHEICHL, *HINT: Hierarchical Invertible Neural Transport for General and Sequential Bayesian inference*, arXiv e-prints, (2019).
- [44] J. DICK, R. N. GANTNER, Q. T. LE GIA, AND C. SCHWAB, *Higher order Quasi-Monte Carlo integration for Bayesian Estimation*, (2016).
- [45] J. DICK, R. N. GANTNER, Q. T. LE GIA, AND C. SCHWAB, *Multilevel higher-order quasi-Monte Carlo Bayesian estimation*, Mathematical Models and Methods in Applied Sciences, 27 (2017), pp. 953–995.
- [46] J. DICK, F. KUO, Q. T. LE GIA, AND C. SCHWAB, *Multilevel higher order QMC Petrov–Galerkin discretization for affine parametric operator equations*, SIAM Journal on Numerical Analysis, 54 (2016), pp. 2541–2568.

BIBLIOGRAPHY

- [47] A. M. DIMITS, G. BATEMAN, M. A. BEER, B. I. COHEN, W. DORLAND, G. W. HAMMETT, C. KIM, J. E. KINSEY, M. KOTSCHENREUTHER, A. H. KRITZ, L. L. LAO, J. MANDREKAS, W. M. NEVINS, S. E. PARKER, A. J. REDD, D. E. SHUMAKER, R. SYDORA, AND J. WEILAND, *Comparisons and physics basis of tokamak transport models and turbulence simulations*, *Physics of Plasmas*, 7 (2000), pp. 969–983.
- [48] J. DONEA, A. HUERTA, J.-P. PONTHOT, AND A. RODRIGUEZ-FERRAN, *Arbitrary Lagrangian-Eulerian Methods*, John Wiley & Sons, Ltd, 2004, ch. 14, pp. 413–437.
- [49] D. L. DONOHO, *High-dimensional data analysis: The curses and blessings of dimensionality*, in *AMS CONFERENCE ON MATH CHALLENGES OF THE 21ST CENTURY*, 2000.
- [50] S. E. DOSSO, C. W. HOLLAND, AND M. SAMBRIDGE, *Parallel tempering for strongly nonlinear geoaoustic inversion*, *The Journal of the Acoustical Society of America*, 132 (2012), pp. 3030–3040.
- [51] M. EIGEL, M. MARSCHALL, AND R. SCHNEIDER, *Sampling-free Bayesian inversion with adaptive hierarchical tensor representations*, *Inverse Problems*, 34 (2018), pp. 035010, 29.
- [52] M. EIGEL, C. MERDON, AND J. NEUMANN, *An adaptive multilevel monte carlo method with stochastic bounds for quantities of interest with uncertain data*, *SIAM/ASA Journal on Uncertainty Quantification*, 4 (2016), pp. 1219–1245.
- [53] O. G. ERNST, B. SPRUNGK, AND L. TAMELLINI, *On Expansions and Nodes for Sparse Grid Collocation of Lognormal Elliptic PDEs*, (2019).
- [54] I.-G. FARÇAŞ, T. GÖRLER, H.-J. BUNGARTZ, F. JENKO, AND T. NECKEL, *Sensitivity-driven adaptive sparse stochastic approximations in plasma microinstability analysis*, arXiv e-prints, (2018).
- [55] I.-G. FARÇAŞ, J. LATZ, E. ULLMANN, T. NECKEL, AND H.-J. BUNGARTZ, *Multilevel adaptive sparse Leja approximations for Bayesian inverse problems*, arXiv e-prints, (2019). Accepted for publication in *SIAM Journal on Scientific Computing*.
- [56] I.-G. FARÇAŞ, P. C. SÂRBU, H.-J. BUNGARTZ, T. NECKEL, AND B. UEKERMANN, *Multilevel adaptive stochastic collocation with dimensionality reduction*, in *Sparse Grids and Applications - Miami 2016*, J. Garcke, D. Pflüger, C. G. Webster, and G. Zhang, eds., Cham, 2018, Springer International Publishing, pp. 43–68.
- [57] I.-G. FARÇAŞ, B. UEKERMANN, T. NECKEL, AND H.-J. BUNGARTZ, *Nonintrusive uncertainty analysis of fluid-structure interaction with spatially adaptive sparse grids and polynomial chaos expansion*, *SIAM Journal on Scientific Computing*, 40 (2018), pp. B457–B482.
- [58] C. FARHAT, P. GEUZAINÉ, AND G. BROWN, *Application of a three-field nonlinear fluid-structure formulation to the prediction of the aeroelastic parameters of an F-16 fighter*, *Computers & Fluids*, 32 (2003), pp. 3–29.
- [59] C. A. J. FLETCHER, *Computational Techniques for Fluid Dynamics*, vol. 2, Springer-Verlag, 1991.

-
- [60] L. FORMAGGIA, A. GUADAGNINI, I. IMPERIALI, V. LEVER, G. PORTA, M. RIVA, A. SCOTTI, AND L. TAMELLINI, *Global sensitivity analysis through polynomial chaos expansion of a basin-scale geochemical compaction model*, Computational Geosciences, 17 (2013), pp. 25–42.
- [61] M. FRANGOS, Y. MARZOUK, K. WILLCOX, AND B. VAN BLOEMEN WAANDERS, *Surrogate and reduced-order modeling: a comparison of approaches for large-scale statistical inverse problems*, in Large-scale inverse problems and quantification of uncertainty, Wiley Series in Computational Statistics, Wiley, Chichester, 2011, pp. 123–149.
- [62] F. FRANZELIN, P. DIEHL, AND D. PFLÜGER, *Non-intrusive Uncertainty Quantification with Sparse Grids for Multivariate Peridynamic Simulations*, in Meshfree Methods for Partial Differential Equations VII, M. Griebel and M. A. Schweitzer, eds., vol. 100 of Lecture Notes in Computational Science and Engineering, Springer International Publishing, 2015, pp. 115–143.
- [63] S. J. FREETHY, T. GÖRLER, A. J. CREELY, G. D. CONWAY, S. S. DENK, T. HAPPEL, C. KOENEN, P. HENNEQUIN, AND A. E. WHITE, *Validation of gyrokinetic simulations with measurements of electron temperature fluctuations and density-temperature phase angles on asdex upgrade*, Physics of Plasmas, 25 (2018), p. 055903.
- [64] D. GALBALLY, K. FIDKOWSKI, K. WILLCOX, AND O. GHATTAS, *Non-linear model reduction for uncertainty quantification in large-scale inverse problems*, International Journal for Numerical Methods in Engineering, 81 (2010), pp. 1581–1608.
- [65] D. GALINDO, P. JANTSCH, C. WEBSTER, AND G. ZHANG, *Accelerating stochastic collocation methods for partial differential equations with random input data*, SIAM/ASA Journal on Uncertainty Quantification, 4 (2016), pp. 1111–1137.
- [66] W. GANDER, *Change of basis in polynomial interpolation*, Numerical Linear Algebra with Applications, 12 (2005), pp. 769–778.
- [67] R. GANTNER, L. HERRMANN, AND C. SCHWAB, *Quasi-Monte Carlo integration for affine-parametric, elliptic PDEs: Local supports and product weights*, SIAM Journal on Numerical Analysis, 56 (2018), pp. 111–135.
- [68] X. GARBET, Y. IDOMURA, L. VILLARD, AND T. H. WATANABE, *Topical Review: Gyrokinetic simulations of turbulent transport*, Nuclear Fusion, 50 (2010), p. 043002.
- [69] B. GATZHAMMER, *Efficient and Flexible Partitioned Simulation of Fluid-Structure Interactions*, phd thesis, Technische Universität München, 2014.
- [70] T. GERSTNER AND M. GRIEBEL, *Numerical integration using sparse grids*, Numerical Algorithms, 18 (1998), pp. 209–232.
- [71] T. GERSTNER AND M. GRIEBEL, *Dimension-adaptive tensor-product quadrature*, Computing, 71 (2003), pp. 65–87.
- [72] T. GERSTNER AND S. HEINZ, *Dimension- and time-adaptive multilevel monte carlo methods*, in Sparse Grids and Applications, J. Garcke and M. Griebel, eds., Berlin, Heidelberg, 2013, Springer Berlin Heidelberg, pp. 107–120.
- [73] R. G. GHANEM AND P. D. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, Berlin, Heidelberg, 1991.
-

BIBLIOGRAPHY

- [74] M. B. GILES, *Multilevel monte carlo path simulation*, Operations Research, 56 (2008), pp. 607–617.
- [75] ———, *Multilevel monte carlo methods*, Acta Numerica, 24 (2015), pp. 259–328.
- [76] M. GIROLAMI AND B. CALDERHEAD, *Riemann manifold langevin and hamiltonian monte carlo methods*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 73 (2011), pp. 123–214.
- [77] T. GÖRLER, X. LAPILLONNE, S. BRUNNER, T. DANNERT, F. JENKO, F. MERZ, AND D. TOLD, *The global version of the gyrokinetic turbulence code gene*, Journal of Computational Physics, 230 (2011), pp. 7053–7071.
- [78] T. GÖRLER, A. D. SIENA, H. DOERK, T. HAPPEL, S. FREETHY, I.-G. FARÇAŞ, A. B. NAVARRO, R. BILATO, A. BOCK, J. CITRIN, G. CONWAY, A. CREELY, P. HENNEQUIN, F. JENKO, T. JOHNSON, C. LECHTE, T. NECKEL, E. POLI, M. SCHNEIDER, E. SONNENDRUECKER, J. STOBER, A. WHITE, ASDEX-UPGRADE-TEAM, AND JET-CONTRIBUTORS, *En route to high-performance discharges: Insights & guidance from high-realism gyrokinetics*, in 27th IAEA Fusion Energy Conference (FEC 2018), International Atomic Energy Agency (IAEA), Oct 2018.
- [79] T. GÖRLER, N. TRONKO, W. A. HORNSBY, A. BOTTINO, R. KLEIBER, C. NORSCINI, V. GRANDGIRARD, F. JENKO, AND E. SONNENDRUECKER, *Intercode comparison of gyrokinetic global electromagnetic modes*, Physics of Plasmas, 23 (2016), p. 072503.
- [80] T. GÖRLER, A. E. WHITE, D. TOLD, F. JENKO, C. HOLLAND, AND T. L. RHODES, *A flux-matched gyrokinetic analysis of diii-d l-mode turbulence*, Physics of Plasmas, 21 (2014), p. 122307.
- [81] M. GRIEBEL, H. HARBRECHT, AND M. D. PETERS, *The sparse grid combination technique for parametric partial differential equations*, 2017.
- [82] M. GRIEBEL AND J. OETTERSCHAGEN, *On tensor product approximation of analytic functions*, Journal of Complexity, 207 (2016), pp. 348 – 379.
- [83] M. GRIEBEL, M. SCHNEIDER, AND C. ZENGER, *A combination technique for the solution of sparse grid problems*, in Iterative Methods in Linear Algebra, P. de Groen and R. Beauwens, eds., Elsevier, North Holland, 1992, IMACS, pp. 263–281.
- [84] J. HADAMARD, *Sur les problèmes aux dérivées partielles et leur signification physique*, Princeton University Bulletin, 13 (1902), pp. 49–52.
- [85] A.-L. HAJI-ALI, F. NOBILE, L. TAMELLINI, AND R. TEMPONE, *Multi-index stochastic collocation for random PDEs*, Computer Methods in Applied Mechanics and Engineering, 306 (2016), pp. 95–122.
- [86] A.-L. HAJI-ALI, F. NOBILE, AND R. TEMPONE, *Multi-index monte carlo: When sparsity meets sampling*, Numer. Math., 132 (2016), pp. 767–806.
- [87] J. HAMPTON, H. FAIRBANKS, A. NARAYAN, AND A. DOOSTAN, *Parametric/stochastic model reduction: Low-rank representation, non-intrusive bi-fidelity approximation, and convergence analysis*, arXiv e-prints, (2017).

-
- [88] B. HARDING AND M. HEGLAND, *A parallel fault tolerant combination technique*, Advances in Parallel Computing, 25 (2014), pp. 584–592.
- [89] B. HARDING, M. HEGLAND, J. LARSON, AND J. SOUTHERN, *Scalable and fault tolerant computation with the sparse grid combination technique*, arXiv e-prints, (2014).
- [90] M. HEGLAND, *Adaptive sparse grids*, in Proceedings of 10th Computational Techniques and Applications Conference CTAC-2001, K. Burrage and R. B. Sidje, eds., vol. 44, 2003, pp. C335–C353.
- [91] S. HEINRICH, *Multilevel monte carlo methods*, in Large-Scale Scientific Computing, S. Margenov, J. Waśniewski, and P. Yalamov, eds., Berlin, Heidelberg, 2001, Springer Berlin Heidelberg, pp. 58–67.
- [92] H. HOEL, E. VON SCHWERIN, A. SZEPESSY, AND R. TEMPONE, *Adaptive multilevel monte carlo simulation*, in Numerical Analysis of Multiscale Computations, B. Engquist, O. Runborg, and Y.-H. R. Tsai, eds., Berlin, Heidelberg, 2012, Springer Berlin Heidelberg, pp. 217–234.
- [93] J. M. HOKANSON AND P. G. CONSTANTINE, *The Lipschitz Matrix: A Tool for Parameter Space Dimension Reduction*, arXiv e-prints, (2019).
- [94] T. ISAAC, N. PETRA, G. STADLER, AND O. GHATTAS, *Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the antarctic ice sheet*, Journal of Computational Physics, 296 (2015), pp. 348 – 368.
- [95] K. ITO AND B. JIN, *Inverse problems : Tikhonov theory and algorithms*, World Scientific, 2014.
- [96] P. JANTSCH, C. G. WEBSTER, AND G. ZHANG, *On the Lebesgue constant of weighted Leja points for Lagrange interpolation on unbounded domains*, IMA Journal of Numerical Analysis, 39 (2018), pp. 1039–1057.
- [97] F. JENKO, W. DORLAND, M. KOTSCHENREUTHER, AND B. N. ROGERS, *Electron temperature gradient driven turbulence*, Physics of Plasmas, 7 (2000), pp. 1904–1910.
- [98] L. JIANG AND N. OU, *Bayesian inference using intermediate distribution based on coarse multiscale model for time fractional diffusion equations*, Multiscale Modeling & Simulation, 16 (2018), pp. 327–355.
- [99] J. KAIPIO AND E. SOMERSALO, *Statistical and computational inverse problems*, vol. 160 of Applied Mathematical Sciences, Springer-Verlag, New York, 2005.
- [100] M. C. KENNEDY AND A. O’HAGAN, *Bayesian calibration of computer models*, Journal of the Royal Statistical Society. Series B. Statistical Methodology, 63 (2001), pp. 425–464.
- [101] D. KEYES, L. MCINNES, C. WOODWARD, W. GROPP, E. MYRA, M. PERNICE, J. BELL, J. BROWN, A. CLO, J. CONNORS, E. CONSTANTINESCU, D. ESTEP, K. EVANS, C. FARHAT, A. HAKIM, G. HAMMOND, G. HANSEN, J. HILL, T. ISAAC, X. JIAO, K. JORDAN, D. KAUSHIK, E. KAXIRAS, A. KONIGES, K. LEE, A. LOTT, Q. LU, J. MAGERLEIN, R. MAXWELL, M. MCCOURT, M. MEHL, R. PAWLOWSKI, A. RANDELES, D. REYNOLDS, B. RIVIERE, U. RUDE, T. SCHEIBE, J. SHADID, B. S. AMD

BIBLIOGRAPHY

- M. SHEPHARD, A. SIEGEL, B. SMITH, X. TANG, C. WILSON, AND B. WOHLMUTH, *Special issue of multiphysics simulations: Challenges and opportunities*, The International Journal of High Performance Computing Applications, 27 (2013), pp. 4–83.
- [102] A. KHODADADIAN, M. PARVIZI, AND C. HEITZINGER, *An adaptive multilevel monte carlo algorithm for the stochastic drift-diffusion-poisson system*, arXiv e-prints, (2019).
- [103] D. C. KNILL AND A. POUGET, *The bayesian brain: the role of uncertainty in neural coding and computation*, Trends in Neurosciences, 27 (2004), pp. 712 – 719.
- [104] D. KOHLER, J. MUELLER, AND U. WEVER, *Cellular probabilistic automata a novel method for uncertainty propagation*, SIAM/ASA Uncertainty Quantification, 2 (2014), pp. 29–54.
- [105] J. KONRAD, *Multifidelity monte carlo sampling in plasma microturbulence analysis*, bachelorarbeit, Technical University of Munich, 2019.
- [106] B. KRAMER, A. N. MARQUES, B. PEHERSTORFER, U. VILLA, AND K. WILLCOX, *Multifidelity probability estimation via fusion of estimators*, Journal of Computational Physics, 392 (2019), pp. 385 – 402.
- [107] J. A. KROMMES, *The Gyrokinetic Description of Microturbulence in Magnetized Plasmas*, Annual Review of Fluid Mechanics, 44 (2012), pp. 175–201.
- [108] F. KÜNZNER, T. NECKEL, AND H.-J. BUNGARTZ, *Prediction and reduction of runtime in non-intrusive forward uq simulations*, Springer Nature Applied Sciences, (2019).
- [109] R. LAM, O. ZAHM, Y. MARZOUK, AND K. WILLCOX, *Multifidelity dimension reduction via active subspaces*, arXiv e-prints, (2018).
- [110] J. LATZ, *On the well-posedness of Bayesian inverse problems*, arXiv e-prints, (2019).
- [111] J. LATZ, M. EISENBERGER, AND E. ULLMANN, *Fast sampling of parameterised Gaussian random fields*, Computer Methods in Applied Mechanics and Engineering, 348 (2019), pp. 978–1012.
- [112] J. LATZ, I. PAPAIOANNOU, AND E. ULLMANN, *Multilevel sequential² Monte Carlo for Bayesian inverse problems*, Journal of Computational Physics, 368 (2018), pp. 154 – 178.
- [113] P. LE TALLEC AND J. MOURO, *Fluid structure interaction with large structural displacements*, Computer Methods in Applied Mechanics and Engineering, 190 (2001), pp. 3039–3067. Advances in Computational Methods for Fluid-Structure Interaction.
- [114] C. LESSING, *Some remarks on monte carlo integration and the curse of dimensionality*, 2017.
- [115] J. LI AND Y. M. MARZOUK, *Adaptive construction of surrogates for the Bayesian solution of inverse problems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1163–A1186.
- [116] C. LIEBERMAN, K. WILLCOX, AND O. GHATTAS, *Parameter and state model reduction for large-scale statistical inverse problems*, SIAM Journal on Scientific Computing, 32 (2010), pp. 2523–2542.

- [117] Q. LIU AND D. WANG, *Stein variational gradient descent: A general purpose bayesian inference algorithm*, in Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16, USA, 2016, Curran Associates Inc., pp. 2378–2386.
- [118] A. LOEVEN, J. WITTEVEEN, AND H. BIJL, *Efficient uncertainty quantification using a two-step approach with chaos collocation*, in ECCOMAS CFD 2006: Proceedings of the European Conference on Computational Fluid Dynamics, Egmond aan Zee, The Netherlands, September 5-8, 2006, Delft University of Technology; European Community on Computational Methods in Applied Sciences (ECCOMAS), 2006.
- [119] X. MA AND N. ZABARAS, *An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations*, Journal of Computational Physics, 228 (2009), pp. 3084–3113.
- [120] X. MA AND N. ZABARAS, *An efficient Bayesian inference approach to inverse problems based on an adaptive sparse grid collocation method*, Inverse Problems, 25 (2009), p. 035013.
- [121] Y. MARZOUK, T. MOSELHY, M. PARNO, AND A. SPANTINI, *Sampling via Measure Transport: An Introduction*, Springer International Publishing, Cham, 2016, pp. 1–41.
- [122] Y. MARZOUK AND D. XIU, *A stochastic collocation approach to Bayesian inference in inverse problems*, Communications in Computational Physics, 6 (2009), pp. 826–847.
- [123] Y. M. MARZOUK AND H. N. NAJM, *Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems*, Journal of Computational Physics, 228 (2009), pp. 1862–1902.
- [124] Y. M. MARZOUK, H. N. NAJM, AND L. A. RAHN, *Stochastic spectral methods for efficient Bayesian solution of inverse problems*, Journal of Computational Physics, 224 (2007), pp. 560–586.
- [125] H. G. MATTHIES, A. LITVINENKO, B. V. ROSIC, AND E. ZANDER, *Bayesian Parameter Estimation via Filtering and Functional Approximations*, arXiv e-prints, (2016).
- [126] H. G. MATTHIES, E. ZANDER, B. V. ROSIC, AND A. LITVINENKO, *Parameter estimation via conditional expectation: a bayesian inversion*, Advanced Modeling and Simulation in Engineering Sciences, 3 (2016), p. 24.
- [127] H. G. MATTHIES, E. ZANDER, B. V. ROSIĆ, A. LITVINENKO, AND O. PAJONK, *Inverse Problems in a Bayesian Setting*, Springer International Publishing, Cham, 2016, pp. 245–286.
- [128] S. A. MATTIS AND B. WOHLMUTH, *Goal-oriented adaptive surrogate construction for stochastic inversion*, Computer Methods in Applied Mechanics and Engineering, 339 (2018), pp. 36–60.
- [129] M. MEHL, B. UEKERMANN, H. BIJL, D. BLOM, B. GATZHAMMER, AND A. VAN ZUIJLEN, *Parallel coupling numerics for partitioned fluid-structure interaction simulations*, Computers and Mathematics with Applications, 71 (2016), pp. 869–891.
- [130] N. METROPOLIS AND S. ULAM, *The monte carlo method*, Journal of American Statistical Association, 44 (1949), pp. 335–341.

BIBLIOGRAPHY

- [131] M. D. MORRIS, *Factorial sampling plans for preliminary computational experiments*, Technometrics, (1991), pp. 161–174.
- [132] Y. NAKATSUKASA, *Approximate and integrate: Variance reduction in Monte Carlo integration via function approximation*, arXiv e-prints, (2018).
- [133] A. NARAYAN, C. GITTELSON, AND D. XIU, *A stochastic collocation algorithm with multifidelity models*, SIAM Journal on Scientific Computing, 36 (2014), pp. A495–A521.
- [134] A. NARAYAN AND J. D. JAKEMAN, *Adaptive Leja sparse grid constructions for stochastic collocation and high-dimensional approximation*, SIAM Journal on Scientific Computing, 36 (2014), pp. A2952 – A2983.
- [135] B. L. NELSON, *On control variate estimators*, Computers & Operations Research, 14 (1987), pp. 219 – 225.
- [136] N. M. NEWMARK, *A method of computation for structural dynamics*, Journal of the Engineering Mechanics Division, 85 (1959), pp. 67–94.
- [137] L. W.-T. NG AND M. ELDRED, *Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation*, 2012.
- [138] N. C. NGUYEN, G. ROZZA, D. B. P. HUYNH, AND A. T. PATERA, *Reduced Basis Approximation and a Posteriori Error Estimation for Parametrized Parabolic PDEs: Application to Real-Time Bayesian Parameter Estimation*, Wiley-Blackwell, 2010, ch. 8, pp. 151–177.
- [139] F. NOBILE, R. TEMPONE, AND C. G. WEBSTER, *A sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM Journal on Numerical Analysis, 46 (2008), pp. 2309–2345.
- [140] M. OBERPARLEITER, H. NORDMAN, G. VERDOOLAEGE, AND F. JENKO, *Uncertainty estimation and a stopping rule in nonlinear gyrokinetic simulations*, Journal of Physics: Conference Series, 775 (2016), p. 012009.
- [141] B. PEHERSTORFER, *Multifidelity monte carlo estimation with adaptive low-fidelity models*, SIAM/ASA Journal on Uncertainty Quantification, 7 (2019), pp. 579–603.
- [142] B. PEHERSTORFER, M. GUNZBURGER, AND K. WILLCOX, *Convergence analysis of multifidelity monte carlo estimation*, Numerische Mathematik, 139 (2018), pp. 683–707.
- [143] B. PEHERSTORFER, B. KRAMER, AND K. WILLCOX, *Multifidelity preconditioning of the cross-entropy method for rare event simulation and failure probability estimation*, SIAM/ASA Journal on Uncertainty Quantification, 6 (2018), pp. 737–761.
- [144] B. PEHERSTORFER AND Y. MARZOUK, *A transport-based multifidelity preconditioner for Markov chain Monte Carlo*, arXiv e-prints, (2018).
- [145] B. PEHERSTORFER, K. WILLCOX, AND M. GUNZBURGER, *Optimal model management for multifidelity monte carlo estimation*, SIAM Journal on Scientific Computing, 38 (2016), pp. A3163–A3194.
- [146] B. PEHERSTORFER, K. WILLCOX, AND M. GUNZBURGER, *Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization*, SIAM Review, 60 (2018), pp. 550–591.

-
- [147] D. PFÜGER, *Spatially Adaptive Sparse Grids for High-Dimensional Problems*, doctoral dissertation, Institut für Informatik, Technische Universität München, Sept. 2010.
- [148] E. QIAN, B. PEHERSTORFER, D. O'MALLEY, V. VESSELINOV, AND K. WILLCOX, *Multifidelity monte carlo estimation of variance and sensitivity indices*, SIAM/ASA Journal on Uncertainty Quantification, 6 (2018), pp. 683–706.
- [149] C. E. RASMUSSEN, *Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals*, in Bayesian statistics, 7 (Tenerife, 2002), Oxford University Press, New York, 2003, pp. 651–660.
- [150] C. E. RASMUSSEN AND C. K. I. WILLIAMS, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, 2005.
- [151] W. D. REMIGIUS, H. P. LAL, J. K. DUBEY, S. M. GODBOLE, S. GUPTA, AND S. SARKAR, *Stochastic reduced order modelling of a fluid structure interaction system*, Procedia Engineering, 144 (2016), pp. 1213 – 1219. International Conference on Vibration Problems 2015.
- [152] C. P. ROBERT AND G. CASELLA, *Monte Carlo statistical methods*, Springer Texts in Statistics, Springer-Verlag, New York, 1999.
- [153] G. ROZZA, D. B. P. HUYNH, AND A. T. PATERA, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations*, Archives of Computational Methods in Engineering, 15 (2008), pp. 229 – 275.
- [154] ———, *Reduced Basis Approximation and a Posteriori Error Estimation for Affinely Parametrized Elliptic Coercive Partial Differential Equations*, Archives of Computational Methods in Engineering, 15 (2008), pp. 229 – 275.
- [155] A. RÜTTGERS AND M. GRIEBEL, *Multiscale simulation of polymeric fluids using the sparse grid combination technique*, Applied Mathematics and Computation, 319 (2018), pp. 425–443.
- [156] A. SALTELLI, K. CHAN, AND E. M. SCOTT, *Sensitivity Analysis*, Wiley, 1st ed., 2000.
- [157] K. SARGSYAN, C. SAFTA, H. N. NAJM, B. J. DEBUSSCHERE, D. RICCIUTO, AND P. THORNTON, *Dimensionality reduction for complex models via bayesian compressive sensing*, International Journal for Uncertainty Quantification, 4 (2014), pp. 63–93.
- [158] F. SCHÄFER, S. MÜLLER, T. UFFINGER, S. BECKER, J. GRABINGER, AND M. KALTENBACHER, *Fluid-structure-acoustic interaction of the flow past a thin flexible structure*, AIAA Journal, 48 (2010).
- [159] R. SCHEICHL, A. STUART, AND A. TECKENTRUP, *Quasi-Monte Carlo and multilevel Monte Carlo methods for computing posterior expectations in elliptic inverse problems*, SIAM/ASA Journal on Uncertainty Quantification, 5 (2017), pp. 493–518.
- [160] C. SCHILLINGS AND C. SCHWAB, *Sparse, adaptive Smolyak quadratures for Bayesian inverse problems*, Inverse Problems, 29 (2013), pp. 065011, 28.
- [161] ———, *Sparsity in Bayesian inversion of parametric operator equations*, Inverse Problems, 30 (2014), pp. 065007, 30.
-

BIBLIOGRAPHY

- [162] C. SCHWAB AND A. M. STUART, *Sparse deterministic approximation of Bayesian inverse problems*, *Inverse Problems*, 28 (2012), pp. 045003, 32.
- [163] C. SCHWAB AND R. TODOR, *Sparse finite elements for stochastic elliptic problems - higher order moments*, *Computing*, 71(1) (2003), pp. 43–63.
- [164] R. C. SMITH, *Uncertainty Quantification: Theory, Implementation, and Applications*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, USA, 2014.
- [165] S. SMOLYAK, *Quadrature and interpolation formulas for tensor products of certain classes of functions*, *Soviet Mathematics - Doklady*, 4 (1963), pp. 240–243.
- [166] I. M. SOBOL', *Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates*, *Mathematics and Computers in Simulation*, 55 (2001), pp. 271 – 280.
- [167] A. SPANTINI, D. BIGONI, AND Y. MARZOUK, *Inference via low-dimensional couplings*, *Journal of Machine Learning Research*, 19 (2018), pp. 1–71.
- [168] A. SPANTINI, A. SOLONEN, T. CUI, J. MARTIN, L. TENORIO, AND Y. MARZOUK, *Optimal low-rank approximations of bayesian linear inverse problems*, *SIAM Journal on Scientific Computing*, 37 (2015), pp. A2451–A2487.
- [169] I. SRAJ, O. P. L. MATRE, O. M. KNIO, AND I. HOTEIT, *Coordinate transformation and polynomial chaos for the bayesian inference of a gaussian process with parametrized prior covariance function*, *Computer Methods in Applied Mechanics and Engineering*, 298 (2016), pp. 205 – 228.
- [170] D. STIRZAKER, *Elementary Probability*, Cambridge University Press, 2 ed., 2003.
- [171] A. STUART, *Inverse problems: a Bayesian perspective*, *Acta Numerica*, 19 (2010), pp. 451–559.
- [172] A. STUART AND A. TECKENTRUP, *Posterior consistency for Gaussian process approximations of Bayesian posterior distributions*, *Mathematics of Computation*, 87 (2018), pp. 721–753.
- [173] B. SUDRET, *Global sensitivity analysis using polynomial chaos expansions*, *Reliability Engineering & Systems Safety*, 93 (2008), pp. 964 – 979.
- [174] T. J. SULLIVAN, *Introduction to Uncertainty Quantification*, vol. 63, Springer, Cham, 2015.
- [175] A. TARANTOLA, *Inverse problem theory and methods for model parameter estimation*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2005.
- [176] A. L. TECKENTRUP, P. JANTSCH, C. G. WEBSTER, AND M. GUNZBURGER, *A multilevel stochastic collocation method for partial differential equations with random input data*, *SIAM/ASA Journal on Uncertainty Quantification*, 3 (2015), pp. 1046–1074.
- [177] T. E. TEZDUYAR, K. TAKIZAWA, C. MOORMAN, S. WRIGHT, AND J. CHRISTOPHER, *Spacetime finite element computation of complex fluidstructure interactions*, *International Journal for Numerical Methods in Fluids*, 64 (2010), pp. 1201–1218.

-
- [178] D. TOLD, F. JENKO, T. GÖRLER, F. J. CASSON, AND E. FABLE, *Characterizing turbulent transport in asdex upgrade l-mode plasmas via nonlinear gyrokinetic simulations*, Physics of Plasmas, 20 (2013), p. 122312.
- [179] L. TREFETHEN, *Cubature, approximation, and isotropy in the hypercube*, SIAM Review, 59 (2017), pp. 469–491.
- [180] L. N. TREFETHEN, *Is gauss quadrature better than clenshaw-curtis?*, SIAM Review, 50 (2008), pp. 67–87.
- [181] ———, *Approximation Theory and Approximation Practice*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, USA, 2012.
- [182] S. TUREK AND J. HRON, *Propoposal for numerical benchmarking of fluid-structure interaction between an elastic object and a laminar incompressible flow*, in Fluid-Structure Interaction, H.-J. Bungartz and M. Schaefer, eds., vol. 53 of Springer Lecture Notes in Computational Science and Engineering, Springer Verlag, 2006, pp. 371–385.
- [183] B. UEKERMANN, *Partitioned Fluid-Structure Interaction on Massively Parallel Systems*, dissertation, Institut für Informatik, Technische Universität München, 2016.
- [184] B. UEKERMANN, J. C. CAJAS, B. GATZHAMMER, G. HOUZEAUX, M. MEHL, AND M. VAZQUEZ, *Towards Partitioned Fluid-Structure Interaction on Massively Parallel Systems*, in 11th World Congress on Computational Mechanics (WCCM XI), E. Oñate, J. Oliver, and A. Huerta, eds., 2014, pp. 1–12.
- [185] P. VAEZI AND C. HOLLAND, *An improved approach to uncertainty quantification for plasma turbulence validation studies*, Fusion Science and Technology, 74 (2018), pp. 77–88.
- [186] P. VAEZI AND C. HOLLAND, *Quantifying the Temporal Uncertainties of Nonlinear Turbulence Simulations*, arXiv e-prints, (2019).
- [187] H.-W. VAN WYK, *Multilevel sparse grid methods for elliptic partial differential equations with random coefficients*, arXiv e-prints, (2014).
- [188] M. VAZQUEZ, G. HOUZEAUX, S. KORIC, A. ARTIGUES, J. AGUADO-SIERRA, R. ARIS, D. MIRA, C. HADRIEN, F. CUCCHIETTI, H. OWEN, A. TAHA, AND J. M. CELA, *Alya: Towards Exascale for Engineering Simulation Codes*, Journal of Computational Science, 14 (2016), pp. 15–27.
- [189] K. VEROY, C. PRUD’HOMME, D. ROVAS, AND A. PATERA, *A Posteriori Error Bounds for Reduced-Basis Approximation of Parametrized Noncoercive and Nonlinear Elliptic Partial Differential Equations*, 2003.
- [190] C. VILLANI, *Optimal transport: old and new*, vol. 338, Springer Science & Business Media, 2008.
- [191] A. E. WHITE AND T. GÖRLER, *Special issue on comparing gyrokinetic simulations to experiments*, Plasma Physics and Controlled Fusion, 59 (2017), p. 050101.
- [192] N. WIENER, *The homogeneous chaos*, American Journal of Mathematics, 60 (1938), pp. 897–936.
-

BIBLIOGRAPHY

- [193] J. WINOKUR, D. KIM, F. BISETTI, O. P. LE MAÎTRE, AND O. M. KNIO, *Sparse Pseudo Spectral Projection Methods with Directional Adaptation for Uncertainty Quantification*, *Journal of Scientific Computing*, 68 (2016), pp. 596–623.
- [194] J. WITTEVEEN AND H. BIJL, *Modeling Arbitrary Uncertainties Using Gram-Schmidt Polynomial Chaos*, Aerospace Sciences Meetings, American Institute of Aeronautics and Astronautics, Jan 2006.
- [195] ———, *Uncertainty Quantification in Fluid-Structure Interaction Simulations Using a Simplex Elements Stochastic Collocation Approach*, Fluid Dynamics and Co-located Conferences, American Institute of Aeronautics and Astronautics, Jun 2009.
- [196] J. WITTEVEEN, S. SARKAR, AND H. BIJL, *Modeling physical uncertainties in dynamic stall induced fluid–structure interaction of turbine blades using arbitrary polynomial chaos*, *Computers & Structures*, 85 (2007), pp. 866–878. Fourth MIT Conference on Computational Fluid and Solid Mechanics.
- [197] J. A. S. WITTEVEEN, A. DOOSTAN, T. CHANTRASMI, R. PECNIK, AND G. IACCARINO, *Comparison of stochastic collocation methods for uncertainty quantification of the transonic rae 2822 airfoil*, 2009.
- [198] P. XANTHOPOULOS, H. E. MYNICK, P. HELANDER, Y. TURKIN, G. G. PLUNK, F. JENKO, T. GÖRLER, D. TOLD, T. BIRD, AND J. H. E. PROLL, *Controlling Turbulence in Present and Future Stellarators*, *Physical Review Letters*, 113 (2014), p. 155001.
- [199] D. XIU, *Numerical methods for stochastic computations : a spectral method approach*, Princeton, N.J. Princeton University Press, 2010.
- [200] D. XIU AND J. S. HESTHAVEN, *High-order collocation methods for differential equations with random inputs*, *SIAM Journal on Scientific Computing*, 27 (2005), pp. 1118–1139.
- [201] D. XIU AND G. E. KARNIADAKIS, *The Wiener-Askey Polynomial Chaos for Stochastic Differential Equations*, *SIAM Journal on Scientific Computing*, 24 (2002), pp. 619–644.
- [202] D. XIU AND J. SHEN, *Efficient stochastic galerkin methods for random diffusion equations*, *Journal of Computational Physics*, 228 (2009), pp. 266 – 281.
- [203] L. YAN AND Y.-X. ZHANG, *Convergence analysis of surrogate-based methods for Bayesian inverse problems*, *Inverse Problems*, 33 (2017), pp. 125001, 20.
- [204] L. YAN AND T. ZHOU, *Adaptive multi-fidelity polynomial chaos approach to Bayesian inference in inverse problems*, *Journal of Computational Physics*, 381 (2019), pp. 110–128.
- [205] O. ZAHM, T. CUI, K. LAW, A. SPANTINI, AND Y. MARZOUK, *Certified dimension reduction in nonlinear Bayesian inverse problems*, arXiv e-prints, (2018).
- [206] C. ZENGER, *Sparse grids*, in *Parallel Algorithms for Partial Differential Equations*, W. Hackbusch, ed., vol. 31 of *Notes on Numerical Fluid Mechanics*, Vieweg, 1991, pp. 241–251.
- [207] G. ZHANG, D. LU, M. YE, M. GUNZBURGER, AND C. WEBSTER, *An adaptive sparse-grid highorder stochastic collocation method for Bayesian inference in groundwater reactive transport modeling*, *Water Resources Research*, 49 (2013), pp. 6871–6892.

- [208] X. ZHU, E. M. LINEBARGER, AND D. XIU, *Multi-fidelity stochastic collocation method for computation of statistical moments*, Journal of Computational Physics, 341 (2017), pp. 386 – 396.
- [209] X. ZHU, A. NARAYAN, AND D. XIU, *Computational aspects of stochastic collocation with multifidelity models*, SIAM/ASA Journal on Uncertainty Quantification, 2 (2014), pp. 444–463.