

Sensitivity-driven adaptive sparse stochastic approximations in plasma microinstability analysis

Ionuț-Gabriel Farcaș^a, Tobias Görler^b, Hans-Joachim Bungartz^c, Frank Jenko^{b,c,a}, Tobias Neckel^c

^a*Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, 201 E 24th St, Austin, TX 78712, United States*

^b*Max Planck Institute for Plasma Physics, Boltzmannstr. 2, 85748 Garching, Germany*

^c*Department of Informatics, Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany*

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. Adaptive sparse grid approximations are an established approach to overcome these challenges. However, standard adaptivity is based on global information, thus properties such as lower intrinsic stochastic dimensionality or anisotropic coupling of the input directions, which are common in practical applications, are not fully exploited. We propose a novel structure-exploiting dimension-adaptive sparse grid approximation methodology using Sobol' decompositions in each subspace to introduce a sensitivity scoring system to drive the adaptive process. By employing sensitivity information, we explore and exploit the anisotropic coupling of the stochastic inputs as well as the lower intrinsic stochastic dimensionality. The proposed approach is generic, i.e., it can be formulated in terms of arbitrary approximation operators and point sets. In particular, we consider sparse grid interpolation and pseudo-spectral projection constructed on (L)-Leja sequences. The power and usefulness of the proposed method is demonstrated by applying it to the analysis of gyrokinetic microinstabilities in fusion plasmas, one of the key scientific problems in plasma physics and fusion research. In this context, it is shown that a 12D parameter space can be scanned very efficiently, gaining more than an order of magnitude in computational cost over the standard adaptive approach. Moreover, it allows for the uncertainty propagation and sensitivity analysis in higher-dimensional plasma microturbulence problems, which would be almost impossible to tackle with standard screening approaches.

Keywords: uncertainty propagation, sensitivity analysis, adaptivity, sparse grid approximations, plasma microturbulence simulation

2000 MSC: 62P35, 65C60, 65D05, 65D15, 65N25, 65Y05, 65Z05, 68U20

Email addresses: ionut.farcas@austin.utexas.edu (Ionuț-Gabriel Farcaș), tobias.goerler@ipp.mpg.de (Tobias Görler), bungartz@in.tum.de (Hans-Joachim Bungartz), frank.jenko@ipp.mpg.de (Frank Jenko), neckel@in.tum.de (Tobias Neckel)

Preprint submitted to J. Comp. Phys.

March 10, 2020

1. Introduction

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*.

In this paper, we assume the considered phenomenon is model-driven, i.e., represented using a mathematical model, such as a differential or an integral equation. Quantifying and reducing uncertainties in such problems is done within the framework of *uncertainty quantification* (UQ). Specifically, we focus on *uncertainty propagation*, also known as forward UQ. In forward UQ, the input uncertainty is typically modelled via a probability density function stemming from e.g., estimations from measurement data or expert opinion. By sampling this density, the underlying model is evaluated once for each sample. These samples can be, for example, pseudo-random numbers or deterministic collocation points. In this paper, we employ the latter. After performing all simulations, the ensemble of model evaluations is used to assess statistics such as expectation and standard deviation. Note that uncertainty propagation is an example of an *outer-loop* scenario, i.e., scenarios in which a specific quantity is computed using multiple evaluations of the underlying model (see [26]). To summarize, we focus on forward UQ in which the modelling of the uncertain parameters is based on expert opinion supported by experimental measurements.

The real-world application under consideration here is the simulation of microturbulence in magnetized fusion plasmas. This problem is of high practical relevance for efforts such as the ITER¹ experiment, which will aim at creating – for the very first time – a self-sustained (“burning”) plasma in the laboratory. This amounts to a milestone on the way towards a future fusion power plant. A physics obstacle on this route are 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 plasma temperature and density gradients. Unfortunately, the measurement of these gradients, as well as further decisive physics parameters affecting the underlying microinstabilities are subject to uncertainties, thus requiring a UQ framework. In this paper, we employ the established plasma microturbulence simulation code GENE [13, 19] and focus on linear gyrokinetic eigenvalue problems in 5D phase space. Even without the nonlinear terms, the significant computational requirements and large number of stochastic parameters, typically referred to as the *curse of dimensionality* (see, e.g., [4]), render the quantification of uncertainty in plasma problems challenging.

Overcoming or delaying the curse of dimensionality is one of the grand challenges in UQ and in scientific computing in general. Approximations based on sparse grids [4] have been well established as suitable counter-measures. To this end, starting with works such as [21, 25, 37], sparse grid approximations have been extensively used in UQ. In recent years, significant efforts have been made towards designing enhanced approximation strategies for

¹ITER was initially an acronym for International Thermonuclear Experimental Reactor. Nowadays it is mainly referred to the Latin word “iter”, meaning “the way”.

computationally expensive stochastic problems. A non-exhaustive list includes [5, 6], where the so-called sparse pseudo-spectral projection (PSP) method was proposed and analysed. With PSP, one constructs sparse approximations free of internal aliasing error by choosing the maximum degree of the projection basis such that the continuous and discrete inner products coincide when applied to the basis polynomials. Moreover, in [5], several adaptive strategies were proposed. In [34], two adaptive methods – a nested approach and a PSP approximation with directional refinement – were studied in problems having multiple design and stochastic parameters. In addition, there is a growing interest in employing UQ in the numerical simulation of fusion plasmas. In [14, 31], a simple uniform, deterministic parameter scan was used in both linear and nonlinear gyrokinetic simulations. In both papers, the focus was on assessing the sensitivity to changes in the ions temperature gradient to validate the underlying simulation codes. Furthermore, in [32], nonintrusive collocation methods were used to quantify uncertainty in a drift-wave turbulence study from the CSDX linear plasma experiment. However, most existing UQ studies in numerical fusion plasma problems sample the underlying stochastic space using either dense grids, hence suffering from the curse of dimensionality, or a priori chosen sparse grids, whose number of points still grows prohibitively large with the dimensionality. Moreover, most existing adaptive sparse grid strategies employ deterministic or global information, thus not discriminating between individual directions nor exploiting the (usually) available sensitivity information. In addition, in most applications, the intrinsic stochastic dimensionality is smaller than the given stochastic dimensionality, i.e., only a subset of stochastic inputs is important.

To this end, in this paper, we formulate and test a novel structure-exploiting adaptive sparse approximation methodology, as follows. We build our strategy on the dimension-adaptive algorithms of [12, 17], in which the underlying sparse grid is constructed adaptively via a judiciously chosen linear combination of full subspaces with low-cardinality. However, we drive the adaptive process using sensitivity information to preferentially refine the directions rendered important. Specifically, we introduce a *sensitivity scoring system* to assess the importance of the individual stochastic input parameters as well as of their interaction. To obtain these scores, we perform Sobol’ decompositions in each subspace and compute variance surpluses which reveal the importance of each input and of their interaction. We note that the proposed approach is generic, in the sense that it can be formulated in terms of arbitrary approximation operators and point sets. In particular, we consider Lagrange interpolation (see, e.g., [2]), which requires the underlying model solution to be continuous, and PSP (see, e.g., [34, 36]), which requires square-integrability. Moreover, to define these two approximations, we employ two (L)-Leja point constructions (see, e.g., [24]). We remark that besides the formulation of a novel adaptive sparse approximation strategy for stochastic computations, another novelty of this paper is the undertaking, to the best of our knowledge, of one of the first UQ studies in plasma microinstability analysis. In addition, the proposed methodology is model-agnostic, provided that certain smoothness assumptions such as continuity or square-integrability are fulfilled.

The remainder of this paper is organized as follows. In Section 2, we introduce our notation and provide a brief background on plasma microturbulence simulation and dimension-adaptive sparse grids, on which we construct our proposed approach based on sensitivity

scores. We present in detail our proposed algorithm in Section 3. We present our numerical results in two plasma microturbulence test cases in Section 4. In Section 4.1, we consider a modified version of the benchmark [7] with eight stochastic inputs. In Section 4.2, we investigate a more realistic test case based on the study from [9], first considering three and afterwards 12 uncertain inputs. Finally, we summarize and give an outlook in Section 5.

2. Background

2.1. Problem formulation

In this work, we are interested in quantifying uncertainty in complex, real-world phenomena such as plasma microturbulence analysis. These problems are typically governed by a model \mathcal{F} specified by complex and nonlinear ODE/PDE systems. However, the solution to this model is very rarely, if ever, available analytically, hence we resort to numerical approximations. Let us denote the discretization of the underlying continuous model by \mathcal{F}_h such that $\mathcal{E}(\mathcal{F} - \mathcal{F}_h)$ can be made arbitrarily small, where \mathcal{E} is a suitable error function. We 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 a vector of d stochastic inputs $\boldsymbol{\theta} := (\theta_1, \theta_2, \dots, \theta_d)$. We model $\boldsymbol{\theta}$ as a multivariate random vector with independent components θ_i defined in a probability space (Ω, \mathcal{A}, P) , where Ω is the event space, which is equipped with σ -algebra \mathcal{A} and probability measure P . We assume that $\boldsymbol{\theta}$ is a continuous random vector characterized by a probability density function $\boldsymbol{\pi}$ with a product structure, i.e.,

$$\boldsymbol{\pi}(\boldsymbol{\theta}) := \prod_{i=1}^d \pi_i(\theta_i)$$

with image $\mathbf{X} := \boldsymbol{\pi}(\Omega)$ having a product structure as well

$$\mathbf{X} := \bigotimes_{i=1}^d X_i,$$

where X_i is the image of π_i . The independence assumption on the components of $\boldsymbol{\theta}$ can be relaxed if a suitable transformation, e.g., a transport map [22] is employed. In uncertainty propagation, we are interested in computing quantities of interest such as the expectation $\mathbb{E}[\mathcal{F}_h]$ and standard deviation $\sigma[\mathcal{F}_h] := \sqrt{\text{Var}[\mathcal{F}_h]}$ of the output of \mathcal{F}_h . When a single evaluation of \mathcal{F}_h is computationally expensive, as it is assumed in this work, estimating $\mathbb{E}[\mathcal{F}_h]$ and $\sigma[\mathcal{F}_h] := \sqrt{\text{Var}[\mathcal{F}_h]}$ becomes prohibitive. Addressing the challenges of quantifying uncertainty in computationally expensive, real-world problems is the main goal of this work.

2.2. Plasma microturbulence simulations

Fusion devices are subject to plasma microturbulence which is driven by the intrinsically steep density and temperature gradients. The associated turbulent transport determines the energy confinement time which in turn is a key parameter for creating a burning plasma.

Any insight into the nature of plasma microturbulence and ways for avoiding turbulence related confinement degradations are therefore crucial for the design of fusion power plants.

Unfortunately, even turbulence in fluid systems is difficult to assess and considered one of the most important problems in classical physics. In magnetically confined plasmas – basically very hot but dilute ionized gases – the problem is further complicated by the low collisionality which renders fluid descriptions insufficient in many situations. An appropriate description is therefore given by a kinetic model, i.e., 6D Vlasov equations per species, which are coupled via Maxwell’s equations. Computing corresponding solutions in complex geometries is challenging even on present-day computational resources. Even more, performing tasks such as uncertainty propagation which require (large) ensembles of such simulations can quickly become computationally prohibitive. This therefore calls for using surrogate models. The following subsections summarize the most popular approach for plasma microturbulence analysis, emphasize the need for UQ in plasma microturbulence analysis, and introduce the employed plasma turbulence code.

2.2.1. The gyrokinetic approach and its applications

The most popular theory for assessing plasma microturbulence is the so-called gyrokinetic theory [3, 20]. Acknowledging a time-scale separation between the fast gyromotion of the particles around the magnetic field lines and typical turbulence time scales, the knowledge of the exact position of the particles along their orbit is irrelevant. Gyrokinetics therefore effectively removes the gyrophase information and yields a 5D system of equations which better fits nowadays computational resources and is considered a valid approach for a wide range of plasma parameters.

Several numerical implementations have been developed over the last two decades and encouraging progress has been made [11, 20]. Early gyrokinetic studies were often limited to restricted physics, e.g., adiabatic electrons and simplified geometries, which could only yield qualitative results and predictions. However, the complexity has meanwhile dramatically increased and flagship codes aim for quantitative validation with various observables obtained from experimental measurements, see e.g., [33] and references therein. While some codes aim for a full flux-driven setup, where profiles and turbulence are self-consistently developing in response to prescribed heat and particle source, these simulations are usually too costly for regular applications and therefore typically 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 as well as the experimental fluctuation observables can only be measured within some uncertainty. This can be quite troublesome since plasma profiles are often found to be quite stiff, i.e., a small increase in inputs such as the gradients may cause large differences in the resulting turbulent heat fluxes. To this end, the standard practice is to identify the key parameters affecting the scenario at hand and *scan* these within their underlying uncertainty bounds. In the plasma physics community, a *parameter scan* means using a full, Cartesian grid of equidistant points from the input domain to ascertain the importance of the input parameters. Given the enormous computational efforts associated to

high-dimensional parameter scans, the former (identification) step is often being performed without considering the nonlinearities in the Vlasov equation which can take up to 50% of the run time. Such linear simulations are still 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. However, many of such studies are only considering a subset of stochastic input parameters due to the tremendously large required computational cost. This clearly provides motivation to develop and apply modern UQ methods in plasma microinstability analysis, which we address in this paper.

2.2.2. The plasma microturbulence code GENE

The gyrokinetic solver employed in this publication is the GENE code – one of the first grid-based codes in this field which has now been under continuous development for almost two decades [19]. GENE computes the time evolution of gyrocenter distribution functions on a fixed grid in a 5D (3D-2V) position-velocity space. The underlying nonlinear partial differential equations of gyrokinetic theory are solved via a mix of numerical methods also widely used in Computational Fluid Dynamics, including finite difference, spectral, finite element, and finite volume schemes. Details are described, e.g., in [13].

Originally, GENE had been restricted to flux-tube simulation domains [1], i.e., thin magnetic-field-line following boxes which allow for highly efficient simulations if the radial correlation lengths of the turbulent fluctuations are small compared to the profile scale lengths. In this limit, the radial variations of the profiles (as well as their gradients) can be assumed to be constant across the simulation domain. Consequently, one may safely assume periodic boundary conditions in the directions perpendicular to the magnetic field line and apply spectral methods which greatly simplify operators such as gyro-averages. For applications in small devices or with steep profiles, however, locality is not a safe assumption anymore and – by considering full radial profiles – periodicity is at least lost in the radial direction. Correspondingly modified numerical methods, e.g., finite differences instead of spectral methods, have been added as another option in GENE which since then allows for radially global simulations [13] or full flux-surface simulations [35] if toroidal instead of radial background variations are considered, respectively. 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 centimeters in the perpendicular plane. GENE simulations are parallelized by domain decomposition in all five dimensions, typically using MPI [23]. Fully nonlinear simulations may require at least 10k CPU-hours which is why linear local (flux-tube) simulations will be employed for testing the UQ framework. Linear local simulations have a runtime of up to at most a few hours. These simulations enable the quantification of uncertainty of the microinstability character. Our long-term goal, nevertheless, is to quantify the uncertainty in fully nonlinear, turbulent simulations, which we will address in subsequent publications.

2.3. Approximation with dimension-adaptive sparse grids

In this section, we summarize dimension-adaptive sparse grid approximations [12, 17], on which we build our proposed methodology in Section 3. Let $f : X_i \rightarrow \mathbb{R}$ denote a univariate function and let $\mathcal{U}^i[f]$ be a sequence of univariate linear continuous operators depending on the marginals π_i of the input density, $\boldsymbol{\pi}(\boldsymbol{\theta}) := \prod_{i=1}^d \pi_i(\theta_i)$. Moreover, let $\mathcal{U}_\ell^i[f]$ be approximations such that

$$\|\mathcal{U}^i[f] - \mathcal{U}_\ell^i[f]\| \xrightarrow{\ell \rightarrow \infty} 0, \quad i = 1, 2, \dots, d,$$

in a suitable norm $\|\cdot\|$, where $1 \leq \ell \in \mathbb{N}$ is referred to as *level*. $\mathcal{U}_\ell^i[f]$ is obtained from discrete evaluations of $\mathcal{U}^i[f]$ at $m(\ell)$ points in X_i , where $m(\ell) : \mathbb{N} \rightarrow \mathbb{N}$ is called the *level-to-nodes* function. In this work, $\mathcal{U}^i[f]$ is interpolation or PSP. In general, we assume that $\mathcal{U}^i[f]$ are *global* operators, i.e., their support is the entire domain X_i .

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

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

We denote by $\Delta_\ell^i[f] := \mathcal{U}_\ell^i[f] - \mathcal{U}_{\ell-1}^i[f]$ the *hierarchical surpluses*, with the convention $\mathcal{U}_0^i[f] := 0$. Therefore, the above telescoping series becomes

$$\mathcal{U}^i[f] = \sum_{\ell=1}^{\infty} \Delta_\ell^i[f]. \quad (2.1)$$

We are interested in approximations that depend on $d \geq 2$ stochastic parameters. A natural way to lift the 1D operators to d dimensions is via *tensorization*, i.e.,

$$\boldsymbol{\mathcal{U}}^d[\mathcal{F}_h] = (\mathcal{U}^1 \otimes \mathcal{U}^2 \otimes \dots \otimes \mathcal{U}^d)[\mathcal{F}_h] = \left(\bigotimes_{i=1}^d \mathcal{U}^i \right) [\mathcal{F}_h]. \quad (2.2)$$

Plugging the telescoping series (2.1) into the tensorized representation (2.2), we obtain

$$\begin{aligned} \boldsymbol{\mathcal{U}}^d[\mathcal{F}_h] &= \left(\sum_{\ell_1=1}^{\infty} \Delta_{\ell_1}^1 \otimes \sum_{\ell_2=1}^{\infty} \Delta_{\ell_2}^2 \otimes \dots \otimes \sum_{\ell_d=1}^{\infty} \Delta_{\ell_d}^d \right) [\mathcal{F}_h] \\ &= \left(\sum_{\boldsymbol{\ell} \in \mathbb{N}^d} \Delta_{\ell_1}^1 \otimes \Delta_{\ell_2}^2 \otimes \dots \otimes \Delta_{\ell_d}^d \right) [\mathcal{F}_h] = \sum_{\boldsymbol{\ell} \in \mathbb{N}^d} \left(\bigotimes_{i=1}^d \Delta_{\ell_i}^i \right) [\mathcal{F}_h] = \sum_{\boldsymbol{\ell} \in \mathbb{N}^d} \boldsymbol{\Delta}_{\boldsymbol{\ell}}^d[\mathcal{F}_h], \end{aligned} \quad (2.3)$$

where $\boldsymbol{\ell} = (\ell_1, \ell_2, \dots, \ell_d) \in \mathbb{N}^d$ denotes a *multiindex* and

$$\boldsymbol{\Delta}_{\boldsymbol{\ell}}^d[\mathcal{F}_h] = \sum_{\mathbf{z} \in \{0,1\}^d} (-1)^{|\mathbf{z}|_1} \boldsymbol{\mathcal{U}}_{\boldsymbol{\ell}}^d[\mathcal{F}_h],$$

where $\mathbf{u}_\ell^d[\mathcal{F}_h] = \left(\bigotimes_{i=1}^d \mathcal{U}_{\ell_i}^i\right)[\mathcal{F}_h]$ and $|\mathbf{z}|_1 := \sum_{i=1}^d z_i$. However, (2.3) is a representation of $\mathbf{u}^d[\mathcal{F}_h]$ with an infinite number of terms, thus unsuitable for computations. Therefore, we restrict the multiindices ℓ to a finite set $\mathcal{L} \subset \mathbb{N}^d$ and define

$$\mathbf{u}_{\mathcal{L}}^d[\mathcal{F}_h] = \sum_{\ell \in \mathcal{L}} \Delta_\ell^d[\mathcal{F}_h]. \quad (2.4)$$

\mathcal{L} must be constructed such that the summation in (2.4) telescopes correctly. Such suitable sets are called *admissible* or *downward closed* (see [12]). In particular, for an admissible set \mathcal{L} it holds that $\ell \in \mathcal{L} \Rightarrow \ell - \mathbf{e}_i \in \mathcal{L}$ for $i = 1, 2, \dots, d$, where \mathbf{e}_i denotes the i th unit vector in \mathbb{R}^d . Note that (2.4) can be rewritten as a *combination scheme* (see, e.g., [15, 27]),

$$\mathbf{u}_{\mathcal{L}}^d[\mathcal{F}_h] = \sum_{\ell \in \mathcal{L}} a_\ell \mathbf{u}_\ell^d[\mathcal{F}_h], \quad (2.5)$$

where $a_\ell = \sum_{\mathbf{z} \in \{0,1\}^d} (-1)^{|\mathbf{z}|_1} \chi_{\mathcal{L}}(\ell + \mathbf{z})$ and $\chi_{\mathcal{L}}$ is the characteristic function on \mathcal{L} , i.e., $\chi_{\mathcal{L}}(\ell) = 1$ if $\ell \in \mathcal{L}$ and $\chi_{\mathcal{L}}(\ell) = 0$ otherwise. Since (2.5) is computationally more convenient, we use it throughout our numerical experiments.

2.3.1. Approximation operators

Without loss of generality, we employ the same approximation operators in all d directions in (2.5). Therefore, for simplicity we drop the superscript i in \mathcal{U}^i and \mathcal{U}_ℓ^i and use the notation \mathcal{U} and \mathcal{U}_ℓ instead. In this work, we consider two approximation operators, interpolation and PSP. Let \mathbb{P}_{P_ℓ} be the space of univariate polynomials of degree $P_\ell \in \mathbb{N}$. Further, let $C^0(X_i)$ denote the space of continuous functions $f : X_i \rightarrow \mathbb{R}$ and $L^2(X_i)$ be the separable Hilbert space of square-integrable functions $f : X_i \rightarrow \mathbb{R}$. To distinguish between interpolation and PSP, we use the superscripts *in* and *psp*. In addition, to refer to either one of the two approximations, we use the superscript *op*.

A popular interpolation approach used in UQ is Lagrange interpolation (see, e.g., [2, 6, 8, 24, 25]), which we also employ in this work. The univariate Lagrange interpolation operator is defined as:

$$\mathcal{U}_\ell^{\text{in}} : C^0(X_i) \rightarrow \mathbb{P}_{P_\ell}, \quad \mathcal{U}_\ell^{\text{in}}[f] := \sum_{n=1}^{m(\ell)} f(\theta_n) L_n(\theta), \quad (2.6)$$

where $\{\theta_n\}_{n=1}^{m(\ell)}$ are interpolation nodes computed w.r.t. the density π_i and $\{L_n(\theta)\}_{n=1}^{m(\ell)}$ are Lagrange polynomials of degree $n - 1$ satisfying the interpolation condition $L_n(\theta_m) = \delta_{nm}$, where δ_{nm} is Kronecker's delta function. For improved numerical stability, we implement (2.6) in terms of the barycentric formula (see, e.g., [2]).

Another commonly used approximation operation in UQ is PSP (see, e.g., [5, 6, 36]). The PSP operator is defined as a series expansion of the form:

$$\mathcal{U}_\ell^{\text{psp}} : L^2(X_i) \rightarrow \mathbb{P}_{P_\ell}, \quad \mathcal{U}_\ell^{\text{psp}}[f] := \sum_{p=0}^{P_\ell} c_p \phi_p(\theta), \quad (2.7)$$

where $\{\phi_p\}_{p=0}^{P_\ell}$ are orthonormal polynomials satisfying

$$\langle \phi_p, \phi_q \rangle := \int_{X_i} \phi_p(\theta) \phi_q(\theta) \pi(\theta) d\theta = \delta_{pq}.$$

Moreover, c_p are the *pseudo-spectral coefficients* defined via projection and computed as:

$$c_p := \int_{X_i} f(\theta) \phi_p(\theta) \pi_i(\theta) d\theta \approx \sum_{n=1}^{m(\ell)} f(\theta_n) \phi_p(\theta_n) w_n, \quad (2.8)$$

where $\{\theta_n\}_{n=1}^{m(\ell)}$ are quadrature nodes and w_n are normalized weights, i.e. $\sum_{n=1}^{m(\ell)} w_n = 1$ computed w.r.t. π_i . In this work, the quadrature rule used to compute the coefficients c_p in (2.8) is exact² for polynomials of degree at most $m(\ell) - 1$. With this in mind, we choose P_ℓ in (2.7) such that $\int_0^1 \phi_p(\theta) \phi_q(\theta) du = \sum_{n=1}^{m(\ell)} \phi_p(\theta_n) \phi_q(\theta_n) w_n$ for all $p, q \leq P_\ell$, that is, there is no internal aliasing error in the underlying projection space (see [5, 6, 34]). The maximal degree that guarantees this is $P_\ell = \lfloor (m(\ell) - 1)/2 \rfloor$. No internal aliasing error in the projection spaces eliminates potential quadrature errors that drastically affect the approximation quality of PSP, as has been numerically shown in [6] and later proved in [5].

2.3.2. (L)-Leja sequences

In this section, we summarize weighted (L)-Leja sequences, which we use to construct the dimension-adaptive sparse grid interpolation and PSP. 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 $i = 1, 2, \dots, d$.

The choice of point sets to construct sparse grid interpolation or PSP for uncertainty propagation has a critical impact on the overall computational cost, since for each grid point an evaluation of the forward model is needed. When these evaluations are computationally expensive, we seek a point set that:

- is nested – so that we can reuse computations from previous levels;
- grows slowly with the level – so that $m(\ell)$ is not very large;
- has good approximation properties – so that we obtain accurate approximations.

A point set having all aforementioned properties is the weighted (L)-Leja sequence (see, e.g., [18]), which we consider in this paper.

For interpolation, we employ standard weighted (L)-Leja points constructed as:

$$\begin{aligned} \theta_1^{\text{in}} &= \operatorname{argmax}_{\theta \in X_i} \pi_i(\theta) \\ \theta_n^{\text{in}} &= \operatorname{argmax}_{\theta \in X_i} \pi_i(\theta) \prod_{m=1}^{n-1} |(\theta - \theta_m^{\text{in}})|, \quad n = 2, 3, \dots \end{aligned} \quad (2.9)$$

²Up to the employed arithmetic precision

Note that the above point sequence is in general not uniquely defined, because (2.9) might have multiple solutions. In that case, we simply pick one of the maximizers. In addition, (2.9) indicates that to increase the interpolation degree from $j - 1$ to j , we need to add only θ_j^{in} to $\{\theta_1^{\text{in}}, \theta_2^{\text{in}}, \dots, \theta_j^{\text{in}}, \}$. Therefore, we add one new (L)-Leja point per level, i.e., $m(\ell) = \ell$.

For PSP, we need a quadrature node sequence to evaluate the PSP coefficients (see (2.8)). When the input density is symmetric and compactly supported, which is the case, for example, for uniform densities which we consider in this work, (2.9) places the first (L)-Leja point in the center of the domain, while the next two points are placed on the boundary of the weight function's support. In this way, adding only one (L)-Leja point at a time will lead to a zero quadrature weight at level $\ell = 2$, causing the employed adaptive algorithms to stop prematurely. Therefore, for PSP, we employ symmetrized (L)-Leja points [28]:

$$\begin{aligned} \theta_1^{\text{PSP}} &:= \operatorname{argmax}_{\theta \in X_i} \pi_i(\theta) \\ \theta_{2n}^{\text{PSP}} &:= \operatorname{argmax}_{\theta \in X_i} \pi_i(\theta) \left| \prod_{\ell=1}^{2n-1} (\theta - \theta_\ell^{\text{PSP}}) \right|, \quad \theta_{2n+1}^{\text{qu}} = 2\theta_1^{\text{PSP}} - \theta_{2n}^{\text{PSP}}, \quad n = 1, 2, \dots \end{aligned} \quad (2.10)$$

In this construction, we add two additional points per level. The first point is obtained via the standard (L)-Leja construction (2.9) and the second is its symmetric point w.r.t. θ_1^{PSP} . Therefore, the level-to-nodes mapping is $m(\ell) = 2\ell - 1$.

Remark 1. The symmetrized Leja construction (2.10) is designed for densities π_i which are symmetric w.r.t. θ_1^{PSP} . If π_i is not symmetric w.r.t. θ_1^{PSP} , we can instead employ, for example, standard (L)-Leja points (2.9) with level-to-nodes mapping $m(\ell) = 2\ell - 1$.

For more details on weighed (L)-Leja points and their properties, we refer to [18, 24].

2.3.3. Standard dimension-adaptivity

To define the sparse grid approximation (2.5), we need an approximation operator (in this work, this is interpolation or PSP), a (discrete) point set to compute these approximations (for this purpose, we employ weighted (L)-Leja points) and a finite multiindex set, \mathcal{L} . To determine \mathcal{L} , we employ dimension-adaptivity [12, 17]. In the following, we summarize the basics of this algorithm and refer the reader to [5, 12, 17, 24, 34] for more details.

In the dimension-adaptive approach, the multiindex set \mathcal{L} is split into two subsets, \mathcal{O} (the old-index set) and \mathcal{A} (the active set) such that $\mathcal{L} := \mathcal{O} \cup \mathcal{A}$ is admissible. \mathcal{O} comprises the already visited multiindices, while the multiindices in \mathcal{A} are used to drive the adaptive process. In the first step, $\mathcal{O} = \{\mathbf{1}_d\}$, because the corresponding number of points is $1^d = 1$, and $\mathcal{A} = \{\mathbf{1}_d + \mathbf{e}_i, i = 1, 2, \dots, d\}$. In the remaining steps, the algorithm employs the following principle: if a multiindex $\ell \in \mathcal{A}$ contributes significantly to the current solution, its adjacent neighbours are likely to contribute as well. To this end, based on a refinement indicator $\epsilon(\cdot)$, $\epsilon(\ell)$ is computed for each $\ell \in \mathcal{A}$. Afterwards, the multiindex with the largest refinement indicator is moved to \mathcal{O} and all its forward neighbors that preserve the admissibility of \mathcal{L} are added to the active set \mathcal{A} . Additionally, a surrogate $\rho := \sum_{\ell \in \mathcal{A}} \epsilon(\ell)$ of the global error is computed at each step. Note that although ρ is a heuristic, it was proven

to be an acceptable surrogate of the global error in a large number of numerical studies (see, e.g., [5, 12, 34]). The adaptivity stops if $\rho < tol^{op}$, for a user-defined tolerance tol^{op} , if $\mathcal{A} = \emptyset$, or if a user-defined maximum level, L_{\max}^{op} , is reached.

The essential ingredient in dimension adaptivity is the refinement indicator $\epsilon^{op}(\cdot)$. A standard indicator (see [5]) which we also employ in this work reads

$$\epsilon^{op}(\ell) := \|\Delta_{\ell}^{op}[\mathcal{F}_h]\|_{L^2}/C_{\ell}^{op}, \quad (2.11)$$

where C_{ℓ}^{op} is the cost, usually the number of grid points, needed to compute $\Delta_{\ell}^{op}[\mathcal{F}_h]$. In Section 3, we explain why we use the L^2 norm of the surplus $\Delta_{\ell}^{op}[\mathcal{F}_h]$ in (2.11) and how does this lead to our proposed refinement indicator based on sensitivity scores.

2.4. Computing quantities of interest

Our goal in this work is to quantify uncertainty in complex, computationally expensive real-world problems such as plasma microturbulence analysis. To this end, we employ dimension-adaptive sparse grid approximations, summarized in Section 2.3. Although these approximations yield a surrogate for the forward model, \mathcal{F}_h , in uncertainty propagation we are interested in computed quantities such as the expectation, standard deviation or Sobol' indices for sensitivity analysis corresponding to outputs of \mathcal{F}_h . In this section, we summarize how these quantities can be estimated at no additional cost from PSP coefficients. In Section 3, we show how PSP coefficients are computed for interpolation as well.

Let

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

denote a multivariate PSP approximation, where \mathcal{P}^{PSP} is a set containing the multivariate PSP polynomial degrees. For example, if dimension-adaptive sparse grid PSP based on symmetrized (L)-Leja points is employed (see Section 2.3), then

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

where $\mathbf{p}_{\max} = (\ell_{1,\max} - 1, \ell_{2,\max} - 1, \dots, \ell_{d,\max} - 1)$, since for PSP the univariate degrees are $P_{\ell_i} = \lfloor (m(\ell_i) - 1)/2 \rfloor = \lfloor (2\ell_i - 2)/2 \rfloor = \ell_i - 1$, $i = 1, 2, \dots, d$, and $(\ell_{1,\max}, \ell_{2,\max}, \dots, \ell_{d,\max})$ is the maximum reached multiindex by the dimension-adaptive algorithm.

In [30, 36] it was shown that the expectation of the forward model, $\mathbb{E}[\mathcal{F}_h]$, its standard deviation, $\sigma[\mathcal{F}_h]$, and total Sobol' indices for variance-based global sensitivity analysis, S_i^T , can be estimated from the pseudo-spectral coefficients. We have

$$\begin{aligned} \mathbb{E}[\mathcal{F}_h] &\approx \hat{\mathbb{E}}[\mathcal{F}_h] = c_{\mathbf{0}} \\ \sigma[\mathcal{F}_h] &\approx \hat{\sigma}[\mathcal{F}_h] = \sqrt{\sum_{\mathbf{p} \in \mathcal{P}^{PSP} \setminus \{\mathbf{0}\}} c_{\mathbf{p}}^2} \\ S_i^T &\approx \hat{S}_i^T = \frac{\sum_{\mathbf{p} \in \mathcal{J}_i^{PSP}} c_{\mathbf{p}}^2}{\sigma^2[\mathcal{F}_h]}, \end{aligned} \quad (2.12)$$

where $\mathcal{J}_i^{\text{PSP}} := \{\mathbf{p} \in \mathcal{P}^{\text{PSP}} : \mathbf{p}_i \neq 0\}$. Total Sobol' indices comprise the total contribution of a stochastic input to the resulting variance, i.e., its individual contribution as well as contributions due to interactions with other inputs. For more details on Sobol' indices, we refer the reader to [30, 36].

3. Sensitivity-driven adaptive refinement

In this section, we present in detail the major algorithmic contribution of this work. We build our method on the adaptive algorithm of [12, 17]. Our novelty is a context-aware refinement policy based on sensitivity information such that the important directions from a stochastic perspective are preferentially refined. We show in Section 3.1 how we obtain directional variances from Sobol' decompositions. In Section 3.2, we describe in detail our proposed adaptive strategy based on sensitivity scores. Finally, in Section 3.3, we showcase the proposed adaptive strategy in a simple example.

3.1. Directional variances from Sobol' decompositions

The essential ingredient in the dimension adaptive algorithm [12, 17] is the refinement indicator $\epsilon^{\text{op}}(\cdot)$. In this work, we propose an indicator based on sensitivity information: starting from the L^2 norm of the surpluses $\Delta_{\ell}^{\text{op}}[\mathcal{F}_h]$, we assess the importance of each stochastic parameter as well as of their interaction. We first summarize how we obtain sensitivity information for PSP, and then we show how to connect interpolation and PSP.

For a multiindex ℓ in the active set \mathcal{A} , the surplus PSP expansion reads:

$$\Delta_{\ell}^{\text{PSP}}[\mathcal{F}_h] := \sum_{\mathbf{p}=0}^{\mathbf{P}_{\ell}} \Delta_{c_{\mathbf{p}}} \Phi_{\mathbf{p}}(\boldsymbol{\theta}), \quad (3.1)$$

where $\Phi_{\mathbf{p}}(\boldsymbol{\theta}) := \prod_{i=1}^d \Phi_{p_i}(\theta_i)$ are the multivariate orthonormal PSP basis polynomials and $\mathbf{P}_{\ell} := (\lfloor (2\ell_1 - 2)/2 \rfloor, \lfloor (2\ell_2 - 2)/2 \rfloor, \dots, \lfloor (2\ell_d - 2)/2 \rfloor) = (\ell_1 - 1, \ell_2 - 1, \dots, \ell_d - 1)$ is the total multivariate degree, since $m(\ell) = 2\ell - 1$ for PSP (cf. Section 2.3.1). Moreover,

$$\Delta_{c_{\mathbf{p}}} := \sum_{\mathbf{z} \in \{0,1\}^d} (-1)^{|\mathbf{z}|_1} c_{\mathbf{p}-\mathbf{z}}, \quad (3.2)$$

where $\Delta_{c_{\mathbf{0}}} := c_{\mathbf{0}}$. From Parseval's identity, we have that

$$\left\| \Delta_{\ell}^{\text{PSP}}[\mathcal{F}_h] \right\|_{L^2}^2 = \sum_{\mathbf{p}=0}^{\mathbf{P}_{\ell}} \Delta_{c_{\mathbf{p}}}^2. \quad (3.3)$$

In [30], it was shown that a PSP expansion is equivalent to a Sobol' decomposition [29] of the same degree, i.e., the terms in the two decompositions are equivalent. A Sobol' expansion is used to represent a d -variate function as a summation of the function's expectation and variances due to each individual input, which reveal the relative importance of each input direction, and variances stemming from all possible interactions between inputs, which

quantify the relative importance of input interactions. Using the equivalence between PSP and Sobol' decompositions, we can rewrite (3.3) as

$$\left\| \Delta_{\ell}^{\text{PSP}}[\mathcal{F}_h] \right\|_{L^2}^2 = \sum_{\mathbf{p}=\mathbf{0}}^{\mathbf{P}_{\ell}} \Delta c_{\mathbf{p}}^2 = \Delta \text{Var}_{\ell}^{\text{PSP},0}[\mathcal{F}_h] + \sum_{i=1}^d \Delta \text{Var}_{\ell}^{\text{PSP},i}[\mathcal{F}_h] + \Delta \text{Var}_{\ell}^{\text{PSP},\text{inter}}[\mathcal{F}_h], \quad (3.4)$$

where $\Delta \text{Var}_{\ell}^{\text{PSP},0}[\mathcal{F}_h] := \Delta c_{\mathbf{0}}^2$ is the expectation surplus, usually small (see, e.g., [34]),

$$\Delta \text{Var}_{\ell}^{\text{PSP},i}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{I}_i^{\text{PSP}}} \Delta c_{\mathbf{p}}^2,$$

are the surplus contributions to the d individual directional variances for $i = 1, 2, \dots, d$, where $\mathcal{I}_i^{\text{PSP}} = \{\mathbf{p} \leq \mathbf{P}_{\ell} : \mathbf{p}_i \neq 0 \wedge \mathbf{p}_j = 0, \forall j \neq i\}$, and $\Delta \text{Var}_{\ell}^{\text{PSP},\text{inter}}[\mathcal{F}_h] := \sum_{\mathbf{p} \in \mathcal{P}_{\text{inter}}^{\text{PSP}}} \Delta c_{\mathbf{p}}^2$, where $\mathcal{P}_{\text{inter}}^{\text{PSP}} = \bigcup_{i=1}^d \{\mathbf{p} \leq \mathbf{P}_{\ell} : \mathbf{p} \notin \mathcal{I}_i^{\text{PSP}}\}$, refers to the surplus variance due to all possible interactions. Therefore, using the L^2 norm of the surplus in the standard indicator (2.11) means using stochastic information to drive the adaptive process, which is desirable in UQ contexts. We can also write $\Delta_{\ell}^{\text{PSP}}[\mathcal{F}_h]$ in (3.1) in terms of a scalar index $j = 1, 2, \dots, N_{\ell} := \prod_{i=1}^d \ell_i$ representing the position of \mathbf{p} in $\{\mathbf{0}, \dots, \mathbf{P}_{\ell}\}$.

For Lagrange interpolation, there is no direct connection to Sobol' expansions. To this end, we 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., [8, 10]). Let $\mathcal{U}_{\ell}^{\text{in}}$ denote the full-grid interpolation operator for multiindex ℓ . The change of basis from Lagrange interpolation and PSP is done as:

$$\mathcal{U}_{\ell}^{\text{in}}[\mathcal{F}_h] = \sum_{\mathbf{p}=\mathbf{0}}^{\mathbf{I}_{\ell}} c_{\mathbf{p}} \Phi_{\mathbf{p}}(\boldsymbol{\theta}) = \sum_{j=1}^{M_{\ell}} c_j \Phi_j(\boldsymbol{\theta}), \quad (3.5)$$

where $\{\Phi_{\mathbf{p}}(\boldsymbol{\theta})\}_{\mathbf{p}=\mathbf{0}}^{\mathbf{I}_{\ell}}$ is the equivalent PSP basis, $\{c_{\mathbf{p}}(\boldsymbol{\theta})\}_{\mathbf{p}=\mathbf{0}}^{\mathbf{I}_{\ell}}$ are the PSP coefficients and $\mathbf{I}_{\ell} = (\ell_1 - 1, \ell_2 - 1, \dots, \ell_d - 1)$, since $m(\ell) = \ell$ for interpolation. We also use the scalar index $j = 1, 2, \dots, M_{\ell} := \prod_{i=1}^d \ell_i$. To obtain the PSP coefficients $\{c_j\}_{j=1}^{M_{\ell}}$, we solve $\sum_{j=1}^{M_{\ell}} c_j \Phi_j(\boldsymbol{\theta}_k) = \mathcal{U}_{\ell}^{\text{in}}[\mathcal{F}_h(\boldsymbol{\theta}_k)]$ for all line (L)-Leja points $\boldsymbol{\theta}_k$ associated to the multiindex ℓ (see, e.g., [8]). Then, we compute the surpluses (3.2) and the squared L^2 norm (3.4).

Remark 2. For the employed level-to-nodes mappings, the bases for interpolation and PSP have the same size $M_{\ell} = N_{\ell}$, therefore, the number of pseudo-spectral coefficients is the same in both cases. However, for PSP, $m(\ell) = 2\ell - 1$, whereas $m(\ell) = \ell$ for interpolation. Therefore, to achieve the same accuracy, we expect interpolation to require fewer grid points.

3.2. Sensitivity-driven dimension adaptivity

Next, we present in detail our proposed adaptive strategy. In Section 3.2.1, we describe how to assess what we call *sensitivity scores*. We depart from standard adaptive techniques which use global information to guide the refinement process and use instead sensitivity

information in each subspace to preferentially refine the directions rendered important. Since several sensitivity scores can be equal, we introduce a classification algorithm in Section 3.2.2 to ensure that one maximum score is found per refinement step. We present the proposed dimension-adaptive algorithm based on sensitivity scores in Section 3.2.3.

3.2.1. Sensitivity scores

The standard error indicator (2.11) employed in this work depends on the norm of the surpluses, $\|\Delta_{\ell}^{\text{op}}[\mathcal{F}_h]\|_{L^2}$, which represents global information. In this way the adaptive algorithm does not discriminate between the individual directions nor their interactions. Having a finer control over the individual input directions as well as their interaction is especially desired in problems in which these directions are anisotropically coupled. In addition, in the majority of practical applications, such as plasma microturbulence simulation, the *intrinsic stochastic dimensionality* is usually smaller than the total number of stochastic parameters, d , i.e., only a subset of inputs are *important* in the uncertainty propagation problem.

To this end, we propose an adaptive approach to better exploit the structure of the underlying problem. We introduce a *sensitivity scoring system* to quantify the contribution of each subspace to the individual variances as well as to the variance due to the parameters' interaction. Specifically, for each multiindex ℓ in the active set \mathcal{A} , we first compute $\|\Delta_{\ell}^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2$ using (3.4)³. Recall that the summation terms in $\|\Delta_{\ell}^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2$ comprise the directional variance surpluses for each input parameter as well as the contribution due to the inputs' interaction. Hence, using $d+1$ user-defined tolerances $\boldsymbol{\tau}^{\text{op}} = (\tau_1^{\text{op}}, \tau_2^{\text{op}}, \dots, \tau_d^{\text{op}}, \tau_{d+1}^{\text{op}})$, we compute the *sensitivity score* of the multiindex ℓ , $s_{\ell}^{\text{op}} \in \mathbb{N}$. Initially, $s_{\ell}^{\text{op}} = 0$. For $i = 1, 2, \dots, d$, we increase s_{ℓ}^{op} by one if the individual variance surpluses satisfy $\Delta\text{Var}_{\ell}^{\text{op},i}[\mathcal{F}_h] \geq \tau_i^{\text{op}}$. Hence, after this step, s_{ℓ}^{op} can be at most d . Finally, if the variance surplus due to the interaction of the stochastic parameters, $\Delta\text{Var}_{\ell}^{\text{op,inter}}[\mathcal{F}_h]$, satisfies $\Delta\text{Var}_{\ell}^{\text{op,inter}}[\mathcal{F}_h] \geq \tau_{d+1}^{\text{op}}$, we increase s_{ℓ}^{op} by one as well. Therefore, s_{ℓ}^{op} can take integer values between 0 and $d+1$.

Since the introduced sensitivity scoring system is based on the d individual stochastic input directions as well as their interaction, it filters the important directions and ensures that they are preferentially refined. Therefore, when the underlying inputs are anisotropically coupled or when the intrinsic stochastic dimensionality is smaller than d , the value of the sensitivity score will reflect these properties. Note, in addition, that our scoring system distinguishes between individual and interaction variance surpluses. In this way, we ensure that if the mixed directions are stochastically unimportant, we prevent the algorithm to refine too much these directions, which are computationally expensive due to their larger number of grid points. In summary, when the underlying stochastic problem has a rich structure, the proposed sensitivity scoring system will *explore* and *exploit* that structure.

We summarize the computation of the sensitivity scores in Algorithm 1. The inputs are the probability density, $\boldsymbol{\pi}$, w.r.t. which all computations are performed, the tolerances, $\boldsymbol{\tau}^{\text{op}}$ and the hierarchical surplus $\Delta_{\ell}^{\text{op}}[\mathcal{F}_h]$. Note that the choice of $\boldsymbol{\tau}^{\text{op}}$ 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 $\boldsymbol{\tau}^{\text{op}}$

³For interpolation, we do the basis transformation (3.5) and compute the squared L^2 norm using (3.4).

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

Algorithm 1 Sensitivity Scores Computation

- 1: **procedure** COMPUTESENSITIVITYSCORE($\boldsymbol{\pi}, \boldsymbol{\tau}^{\text{op}}, \boldsymbol{\Delta}_\ell^{\text{op}}[\mathcal{F}_h]$)
- 2: $s_\ell^{\text{op}} := 0$
- 3: Compute $\|\boldsymbol{\Delta}_\ell^{\text{op}}[\mathcal{F}_h]\|_{L^2}^2$ to obtain the variance $\Delta\text{Var}_\ell^{\text{op}}[\mathcal{F}_h]$ w.r.t. $\boldsymbol{\pi}$
- 4: Decompose $\Delta\text{Var}_\ell^{\text{op}}[\mathcal{F}_h]$ via (3.4) to obtain all unnormalized Sobol' indices

$$\Delta\text{Var}_\ell^{\text{op}}[\mathcal{F}_h] = \sum_{i=1}^d \Delta\text{Var}_\ell^{\text{op},i}[\mathcal{F}_h] + \Delta\text{Var}_\ell^{\text{op,inter}}[\mathcal{F}_h]$$

- 5: **for** $i \leftarrow 1, 2, \dots, d$ **do**
 - 6: **if** $\Delta\text{Var}_\ell^{\text{op},i}[\mathcal{F}_h] \geq \tau_i^{\text{op}}$ **then**
 - 7: $s_\ell^{\text{op}} = s_\ell^{\text{op}} + 1$
 - 8: **if** $\Delta\text{Var}_\ell^{\text{op,inter}}[\mathcal{F}_h] \geq \tau_{d+1}^{\text{op}}$ **then**
 - 9: $s_\ell^{\text{op}} = s_\ell^{\text{op}} + 1$
 - 10: **return** s_ℓ^{op}
-

3.2.2. Maximum sensitivity score

The adaptive process in our proposed approach is driven by the multiindex from the active set \mathcal{A} with the largest sensitivity score. To this end, if several subspaces have equal scores we need an additional step to distinguish between them. Note that two or more subspaces have the same sensitivity score if the same number of directional variances are larger than the prescribed tolerances. However, these directions do not need to be necessarily the same. With this in mind, we distinguish between subspaces having equal sensitivity scores via the following *classification* algorithm.

First, we compute

$$\Delta\text{Var}_\ell^{\text{op,tot}}[\mathcal{F}_h] := \sum_{i=1}^d \Delta\text{Var}_\ell^{\text{op},i}[\mathcal{F}_h] + \Delta\text{Var}_\ell^{\text{op,inter}}[\mathcal{F}_h],$$

and then select the subspace with the largest $\Delta\text{Var}_\ell^{\text{op,tot}}[\mathcal{F}_h]$. In this way, if two or more subspaces contribute significantly in the same number of directions, we select the subspace with the largest global contribution. We summarize the aforementioned strategy for finding the subspace with the maximum score in Algorithm 2. The inputs are two sets: \mathcal{S} , comprising all current scores, and \mathcal{D} , which contains all current surpluses. Between lines 4 – 5, we find the maximum sensitivity scores. If only one maximum score exists, then the algorithm ends (lines 7 – 9). In the case when several maximum scores exist, we select the subspace with the largest $\Delta\text{Var}_\ell^{\text{op,tot}}$ (lines 11 – 16). In the unlikely case that $\Delta\text{Var}_\ell^{\text{op,tot}}$ is the same for several scores, the algorithm returns the multiindex corresponding to the last score.

Algorithm 2 Maximum Sensitivity Score Computation

```
1: procedure FINDINDEXMAXIMUMSCORE( $\mathcal{S}, \mathcal{D}$ )
2:    $I := []$ 
3:    $\Delta\text{Var}_{\max}^{\text{op}}[\mathcal{F}_h] := 0, \quad \ell_{\max}^{\text{op}} := \mathbf{1}$ 
4:   Find the maximum sensitivity scores  $s_{\ell_n}^{\text{op}}$  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:      $\ell_{\max}^{\text{op}} := \ell_p$ 
10:  else
11:    for  $m \leftarrow 1, 2, \dots, n_{\max}$  do
12:       $q := I[m]$ 
13:      Take  $\Delta_{\ell_q}^{\text{op}}[\mathcal{F}_h]$  from  $\mathcal{D}$  and compute  $\Delta\text{Var}_{\ell_q}^{\text{op}}[\mathcal{F}_h]$ 
14:      if  $\Delta\text{Var}_{\ell_q}^{\text{op}}[\mathcal{F}_h] \geq \Delta\text{Var}_{\max}^{\text{op}}[\mathcal{F}_h]$  then
15:         $\Delta\text{Var}_{\max}^{\text{op}}[\mathcal{F}_h] := \Delta\text{Var}_{\ell_q}^{\text{op}}[\mathcal{F}_h]$ 
16:         $\ell_{\max}^{\text{op}} := \ell_q$ 
17:  return  $\ell_{\max}^{\text{op}}$ 
```

3.2.3. Sensitivity-driven dimension-adaptive sparse grid approximations

We hereby summarize our proposed dimension-adaptive strategy based on sensitivity scores in Algorithm 3. The inputs are the forward model, \mathcal{F}_h , the input probability density, $\boldsymbol{\pi}$, the vector $\boldsymbol{\tau}^{\text{op}}$ of $d + 1$ tolerances and the maximum grid level that can be reached, L_{\max}^{op} . At lines 2 – 3, we initialize \mathcal{O} and \mathcal{A} as in the standard algorithm. In addition, we initialize the two new data structures, \mathcal{S} and \mathcal{D} , 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 Algorithm 1 and update \mathcal{S} and \mathcal{D} accordingly. Afterwards, in each refinement step, we determine the multiindex with the maximum score (line 8) based on Algorithm 2 and update \mathcal{A} and \mathcal{O} . Moreover, we also append the largest score and its associated surplus to \mathcal{S} and \mathcal{D} , respectively. The algorithm continues with adding the forward neighbours of the multiindex with the largest sensitivity score provided that the total multiindex set remains admissible. Further, we assess the sensitivity score for each of these neighbours via Algorithm 1 and update the sets \mathcal{O} , \mathcal{A} , \mathcal{S} and \mathcal{D} accordingly. Our algorithm terminates if all scores in \mathcal{S} are zero, i.e., if the contributions of all multiindices in \mathcal{A} are rendered stochastically unimportant. Moreover, the algorithm stops as well if $\mathcal{A} = \emptyset$ or if the maximum level L_{\max}^{op} is reached. Note that we do not employ a surrogate for the global error as in the standard dimension-adaptive algorithm of Section 2.3.3, since, first of all, our algorithm relies on error indicators in each subspace. Second of all, the algorithm stops, by design, when all subspaces from \mathcal{A} have a zero sensitivity score, i.e., when their variance contribution in all directions becomes insignificant. At the end, we determine the PSP coefficients c_{ℓ} for all $\ell \in \mathcal{L}$ and we use these coefficients to assess the expectation,

standard deviation, and total Sobol' indices for global sensitivity analysis via (2.12).

Algorithm 3 Sensitivity-driven Dimension-adaptive Sparse Grid Algorithm

```

1: procedure SENSITIVITYDRIVENADAPTIVESPARSEGRIDAPPROX( $\mathcal{F}_h, \boldsymbol{\pi}, \boldsymbol{\tau}^{\text{op}}, L_{\max}^{\text{op}}$ )
2:    $\mathbf{1}_d := (1, 1, \dots, 1)$ 
3:    $\mathcal{O} := \emptyset, \quad \mathcal{A} := \{\mathbf{1}_d\}$ 
4:    $\mathcal{S} := \emptyset, \quad \mathcal{D} := \emptyset$ 
5:    $s_{\mathbf{1}}^{\text{op}} := \text{COMPUTESENSITIVITYSCORE}(\boldsymbol{\pi}, \boldsymbol{\tau}^{\text{op}}, \boldsymbol{\Delta}_{\mathbf{1}}^{\text{op}}[\mathcal{F}_h])$ 
6:    $\mathcal{S} = \mathcal{S} \cup \{s_{\mathbf{1}}^{\text{op}}\}, \quad \mathcal{D} = \mathcal{D} \cup \{\boldsymbol{\Delta}_{\mathbf{1}}^{\text{op}}[\mathcal{F}_h]\}$ 
7:   while all( $\mathcal{S}$ )  $\neq 0$  or  $\mathcal{A} \neq \emptyset$  or  $\max(\mathcal{L}) < L_{\max}^{\text{op}}$  do
8:      $\mathbf{k} := \text{FINDINDEXMAXIMUMSCORE}(\mathcal{S}, \mathcal{D})$ 
9:      $\mathcal{A} = \mathcal{A} \setminus \{\mathbf{k}\}, \quad \mathcal{O} = \mathcal{O} \cup \{\mathbf{k}\}$ 
10:     $\mathcal{S} = \mathcal{S} \setminus \{s_{\mathbf{k}}^{\text{op}}\}, \quad \mathcal{D} = \mathcal{D} \setminus \{\boldsymbol{\Delta}_{\mathbf{k}}^{\text{op}}[\mathcal{F}_h]\}$ 
11:    for  $i \leftarrow 1, 2, \dots, d$  do
12:       $\mathbf{r} \leftarrow \mathbf{k} + \mathbf{e}_i$ 
13:      if  $\mathbf{r} - \mathbf{e}_q \in \mathcal{O}$  for all  $q = 1, 2, \dots, d$  then
14:         $\mathcal{A} = \mathcal{A} \cup \{\mathbf{r}\}$ 
15:         $s_{\mathbf{r}}^{\text{op}} := \text{COMPUTESENSITIVITYSCORE}(\boldsymbol{\pi}, \boldsymbol{\tau}^{\text{op}}, \boldsymbol{\Delta}_{\mathbf{r}}^{\text{op}}[\mathcal{F}_h])$ 
16:         $\mathcal{S} = \mathcal{S} \cup \{s_{\mathbf{r}}^{\text{op}}\}, \quad \mathcal{D} = \mathcal{D} \cup \{\boldsymbol{\Delta}_{\mathbf{r}}^{\text{op}}[\mathcal{F}_h]\}$ 
17:     $\mathcal{L} = \mathcal{O} \cup \mathcal{A}$ 
18:    Determine the PSP coefficients  $\{c_{\ell}\}_{\ell \in \mathcal{L}}$ 
19:    Compute  $\hat{\mathbb{E}}[\mathcal{F}_h], \hat{\sigma}[\mathcal{F}_h], \hat{S}_1^T, \hat{S}_2^T, \dots, \hat{S}_d^T$  from  $\{c_{\ell}\}_{\ell \in \mathcal{L}}$  using (2.12)
20:  return  $\hat{\mathbb{E}}[\mathcal{F}_h], \hat{\sigma}[\mathcal{F}_h], \hat{S}_1^T, \hat{S}_2^T, \dots, \hat{S}_d^T$ 

```

3.3. 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 (2.11).

For a fair comparison, we employ small tolerances ($tol^{\text{in}} = 10^{-8}$ in the standard version and $\boldsymbol{\tau}^{\text{in}} = 10^{-16} \cdot \mathbf{1}_{d+1}$ in our approach) and compute the L^2 approximation error

$$\mathcal{E}^2(\mathcal{F} - \mathbf{u}_{\mathcal{L}}[\mathcal{F}]) := \sqrt{\int_{\mathbf{X}} (\mathcal{F}(\boldsymbol{\theta}) - \mathbf{u}_{\mathcal{L}}[\mathcal{F}(\boldsymbol{\theta})])^2 d\boldsymbol{\theta}} \quad (3.6)$$

and the *relative error* of the expectation approximation

$$\mathcal{E}_{\text{rel}}(\mathbb{E}[\mathcal{F}] - \hat{\mathbb{E}}[\mathcal{F}]) := |1 - \mathbb{E}[\mathcal{F}]/\hat{\mathbb{E}}[\mathcal{F}]|, \quad (3.7)$$

where $\hat{\mathbb{E}}[\mathcal{F}]$ is estimated using the first PSP coefficient as shown in (2.12).

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

$$\mathcal{F}(\boldsymbol{\theta}) = 1 + \cos(\pi + 1.5\theta_1 + 0.5\theta_2 + 0.05\theta_3 + 0.1\theta_4 + 0.002\theta_5), \quad (3.8)$$

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 as in (2.12) using 10^5 Gauss-Legendre nodes and obtain

$$\hat{S}_1^T = 0.9034, \quad \hat{S}_2^T = 0.0098, \quad \hat{S}_3^T = 9.8216 \cdot 10^{-4}, \quad \hat{S}_4^T = 3.9287 \cdot 10^{-3}, \quad \hat{S}_5^T = 1.5714 \cdot 10^{-6}.$$

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, $\mathcal{E}^2(\mathbb{E}[\mathcal{F}] - \hat{\mathbb{E}}[\mathcal{F}])$, and relative error, $\mathcal{E}_{\text{rel}}(\mathbb{E}[\mathcal{F}] - \hat{\mathbb{E}}[\mathcal{F}])$, cannot be computed analytically, we estimate them numerically using 1000 Monte Carlo samples for the L^2 error and the expectation estimate on the Gauss-Legendre grid with 10^5 points for the relative error.

We depict the error estimations in Figure 1. 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 $3 \cdot 10^{-8}$, our approach requires 376 Leja points, whereas the standard approach needs 578 points. 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 see again that our approach is, for a lower computational cost, as accurate as the standard adaptive method. Therefore, for the con-




Figure 1: Monte Carlo estimate of the L^2 approximation error (3.6) (left) and the estimate of the relative error of the expectation approximation (3.7) (right) for the example in (3.8).

sidered 5D test case (3.8), we showed that our approach is computationally cheaper than the dimension-adaptive approach based on the standard refinement indicator (2.11).

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 $\mathcal{F} : [0, 1]^2 \rightarrow \mathbb{R}$,

$$\mathcal{F}(\boldsymbol{\theta}) = \cos(\theta_1 + 0.1\theta_2), \quad (3.9)$$

with $\boldsymbol{\pi}(\boldsymbol{\theta}) = U(0, 1)^2$. The associated total Sobol' indices estimated on a Gauss-Legendre grid comprising $16^2 = 256$ nodes are

$$\hat{S}_1^T = 0.9908, \quad \hat{S}_2^T = 0.0112,$$

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

The two errors are estimated using 1000 Monte Carlo samples as previously. We depict the results in Figure 2. 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,

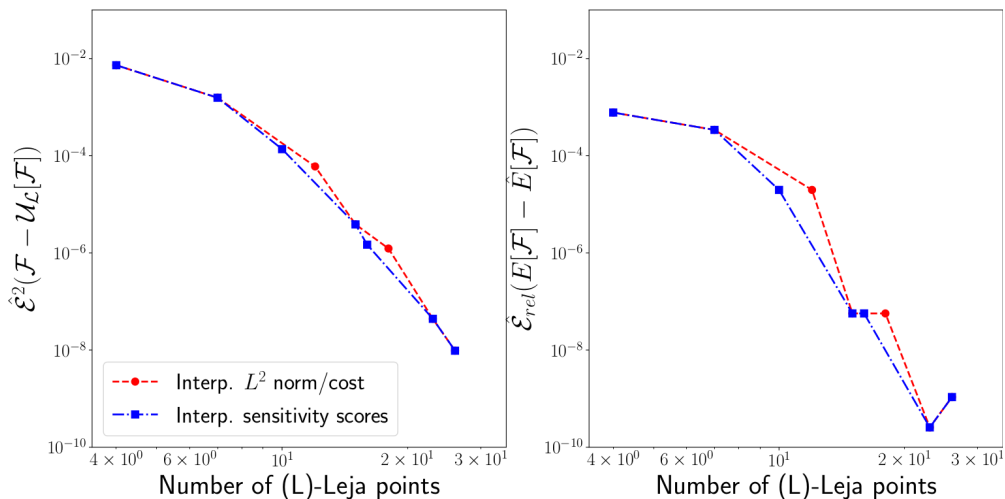


Figure 2: Monte Carlo estimate of the L^2 approximation error (3.6) (left) and the estimate of the relative error of the expectation approximation (3.7) (right) for the example in (3.9).

underlines that our proposed sensitivity 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 approximations in uncertainty propagation: we obtained accurate approximation with at most 576 Leja points in the 5D test case and at most 32 points in the 2D example.

4. Application to plasma microinstability analysis

In this section, we present results of the proposed UQ method applied to two plasma microinstability test cases. In addition, we also employ the standard adaptive strategy summarized in Section 2.3.3, which we compare with our method. Both test cases represent linear local (flux-tube) simulations in which microinstabilities are characterised using the linear eigenvalue solver from the gyrokinetic code GENE (recall Section 2.2.2). Therefore,

\mathcal{F}_h is the linear eigenvalue solver from GENE. The uncertain inputs in both test cases are independent uniform random variables, i.e., $\boldsymbol{\pi}$ is a multivariate uniform density, with left and right bounds stemming from expert opinion supported by experimental measurements. All linear eigenvalue simulations are performed using 32 cores on two Intel Xeon E5-2697 nodes from the CoolMUC-2 cluster ⁴ at the Leibniz Supercomputer Center. The runtime of one simulation varies roughly between five and 90 minutes, depending on the used setup.

In Section 4.1, we consider a modified gyrokinetic benchmark 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 4.2 a particular discharge of the ASDEX Upgrade experiment. For this second test case, the analysis is performed step-wisely – first, only three uncertain input parameters are considered in Section 4.2.1, whereas 12 stochastic parameters are taken into account in Section 4.2.2.

4.1. Modified Cyclone Base Case

The first test case is based on the so-called Cyclone Base Case (CBC) of [7]. CBC has been selected since it is a popular benchmark in the gyrokinetic community and since its parameters are known to display a significant sensitivity to changes of the temperature and density gradients, which calls for quantify uncertainty *ab initio*. The original CBC benchmarks have been restricted to one gyrokinetic ion (i) species, assuming an adiabatic electron response and thus only electrostatic fluctuations. However, we modify the setting to allow for more choices of stochastic parameters to better resemble realistic applications.

The extensions and deviations from the parameters from [7] and the additionally educated assumptions for the uncertainties are summarized in Table 1. The electrons (e) are treated fully gyrokinetically such that their logarithmic temperature gradients, $-L_s \partial_x \ln T_e$, and density gradients, $-L_s \partial_x \ln n_e$, need to be considered as well. While the former is taken in the same range as the ion temperature gradient, $-L_s \partial_x \ln T_i$, but can be varied independently, the logarithmic density gradient is forced to the exact value of the ion counterpart due to the quasi-neutrality constraint in plasma physics. The logarithmic density gradient also fixes the density ratio to 1 while the temperature ratio, T_i/T_e , can be varied. Adding an electron species furthermore allows to consider electromagnetic effects in the gyrokinetic simulations. Their strength 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 1. Uncertainties in the circular magnetic ($\hat{s} - \alpha$) equilibrium can be considered by attributing lower and upper limits to the safety factor, q , the ratio of toroidal turns of a magnetic field line per poloidal turn, and its normalized radial derivative, the magnetic shear, $\hat{s} = r/q \partial q / \partial r$, where r is the radial coordinate labeling flux surfaces. To summarize, we consider eight uncertain input parameters: the first six characterize the underlying particle species, ions and electrons, whereas the last two parameters are associated to the magnetic geometry.

⁴<https://www.lrz.de/services/compute/linux-cluster/>

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 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.

The bounds of the eight uniformly distributed stochastic parameters, found in the last two columns in Table 1, are symmetric around a nominal value employed in deterministic simulations. The output of interest is the growth rate, $\gamma[c_s/L_s]$, of the dominant eigenmode, where $c_s = \sqrt{T_e/m_i}$ is the *ion sound speed* and L_s is the *characteristic length*, which we compute with six digits of precision. For a more in-depth overview of the influence of the eight stochastic parameters, we perform uncertainty propagation for multiple *normalized perpendicular wavenumbers*, $k_y \rho_s$:

$$k_y \rho_s = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9. \quad (4.1)$$

We treat $k_y \rho_s$ as a deterministic, free parameter and for each of its values in (4.1), we propagate the uncertainty in the eight parameters listed in Table 1 using Algorithm 3. Note that since simulations corresponding to different $k_y \rho_s$ are independent of each other, they can be performed embarrassingly parallel, thus adding a second layer of parallelism.

Prior to uncertainty propagation, a deterministic simulation using the nominal values of the eight parameters from Table 1 is performed to obtain more insights into the underlying microinstabilities. We compute the linear growth rates and real frequencies, $\omega[c_s/L_s]$, of the dominant and first subdominant eigenmodes for all for all $k_y \rho_s$ in (4.1). We depict the results in Figure 3. 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 to 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., [16] and the references therein. Given the steep increase of the latter at large wavenumbers, the transition is marked by a very sharp growth rate gradient at $k_y \rho_s = 0.6$. Since in this work we employ sparse grid approximations formulated in terms of global polynomial basis functions, these approximations were ill-conditioned at $k_y \rho_s = 0.6$. Therefore, we discard $k_y \rho_s = 0.6$ and present in what follows results at the remaining values of the normalized perpendicular wavenumber in (4.1).




Figure 3: 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 uncertain parameters listed in Table 1.

We consider both interpolation and PSP to construct adaptive sparse grid surrogates for the 8D stochastic model. In addition, we compare our approach with the standard adaptive technique summarized in Section 2.3.3, which we use to obtain reference results. In our approach, we prescribe $\boldsymbol{\tau}_{\text{op}} = 10^{-6} \cdot \mathbf{1}_9$ and $L_{\text{max}}^{\text{op}} = 20$, whereas in the standard adaptive approach we employ $\text{tol}_{\text{op}} = 10^{-3}$ and $L_{\text{max}}^{\text{op}} = 20$. After computing the pseudo-spectral coefficients, the expectation, standard deviation and total Sobol' indices (recall Section 2.4) are evaluated in the postprocessing for each considered $k_y \rho_s$.

Figure 4 illustrates the expected value and one standard deviation of $\gamma[c_s/L_s]$ for the different adaptive surrogates as well as the deterministic growth rate results. On the one hand, a good agreement between all adaptive sparse approximation approaches can be stated. This demonstrates that for identical setups, interpolation and PSP perform comparatively well. In addition, our sensitivity scores approach has a similar accuracy as the standard adaptive approach. On the other hand, the deterministic results overlap almost perfectly with the expectation of the stochastic results, which is no surprise, given the relative simplicity of this test case. Therefore, the most likely value of the growth rate yielded by uncertainty propagation is similar to the deterministic result. The novelty of using uncertainty propagation is that we also obtain uncertainty bars which represent a quantitative measure of the uncertainty associated to the expected value of each stochastic input. Based on these results, the ITG mode appears to be robust w.r.t. the uncertainty in the eight stochastic parameters while the onset of the TEM/ETG-branch could be quite different given the much larger uncertainty bars. Already these results can be very interesting for physicists trying to compare and predict microturbulence in plasmas. However, the analysis can be taken even further. To quantitatively understand the total contribution of each stochastic input to this uncertainty, we analyse in a next step the total Sobol' indices for global sensitivity analysis.

We show the total Sobol' indices in Figure 5. The top two subfigures contain the results using adaptive interpolation, whereas in the bottom subfigures total Sobol' indices obtained by adaptive PSP are presented. In addition, the left figures correspond to the standard




Figure 4: Expected value and one standard deviation as well as the deterministic growth rates for the CBC test case with eight uncertain parameters.

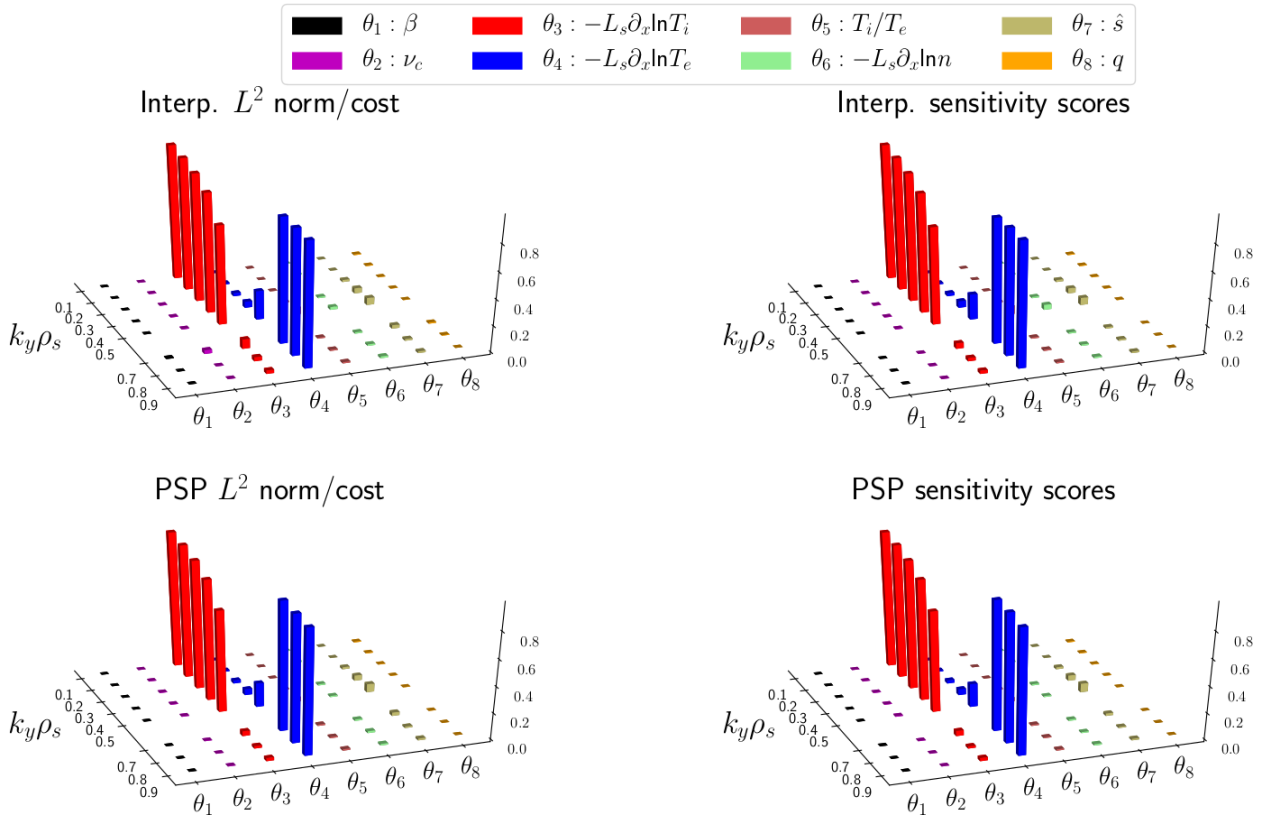


Figure 5: Total Sobol' indices for global sensitivity analysis via adaptive sparse grid interpolation (top two subfigures) and PSP (bottom two subfigures) based on the standard adaptive approach (left plots) and our proposed approach (right plots) for the CBC test case using $k_y \rho_s \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.8, 0.9\}$.

adaptive approach, whereas the right plots depict the results corresponding to the proposed sensitivity scores approach. The similarity of all results indicates again that interpolation and PSP perform similarly, and that our proposed adaptive approach 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 inputs have negligible contributions. This indicates that although we considered a total of eight stochastic inputs, only two of them are important for uncertainty propagation in the considered $k_y\rho_s$ domain. Such findings are very important for the physics community. The other stochastic parameters are very well known to have an impact on the modes themselves but here apparently not within the assumed uncertainty bars. This could, for instance, motivate to restrict the much more costly nonlinear studies to the two temperature gradients if a compromise between computational resources and sensitivity studies needs to be found. In addition, the two temperature gradients have complementary total Sobol' indices: the ion temperature gradient dominates whereas 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 $k_y\rho_s \geq 0.7$. In other words, 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 in the deterministic case: 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, i.e., the stochastic model can be well approximated by a linear model. Such information can be extremely helpful in constructing reduced, e.g., quasi-linear models in the plasma physics community.

Finally, the cost in terms of total number of simulations for all four employed adaptive approaches is illustrated in Figure 6. First, we observe that the most expensive problem

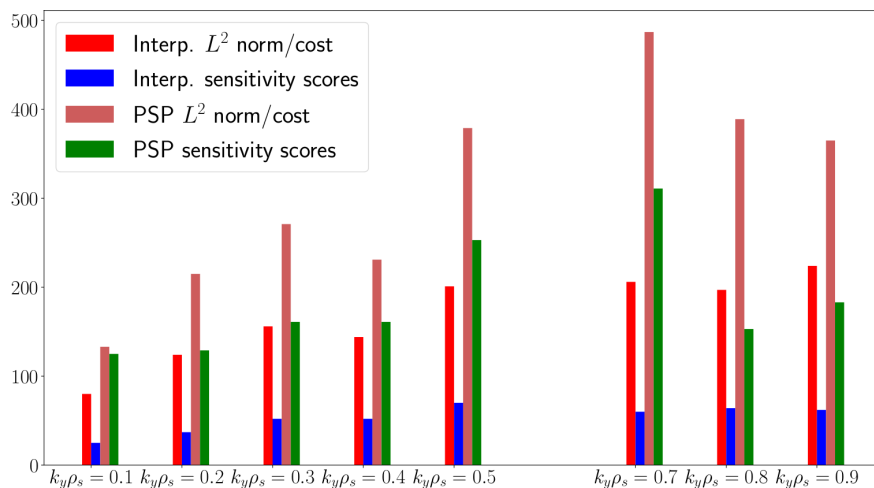


Figure 6: Total number of evaluations needed to construct the sparse interpolation and PSP surrogates using our proposed approach based on sensitivity scores vs. the standard adaptive approach.

required only 473 simulations in total; this was mainly possible because the stochastic in-

puts are anisotropically coupled and the intrinsic stochastic dimensionality is significantly lower than eight. To put this in perspective, a standard parameter scan using 10 points in each direction would require a total of 10^8 simulations. In addition, as expected, for either standard or sensitivity scores adaptivity, PSP is more expensive than interpolation (see Remark 2); even the adaptive PSP based on sensitivity scores is generally more expensive than standard adaptive interpolation. By far the cheapest approach is the adaptive interpolation using our proposed sensitivity scores approach. For example, at $k_y \rho_s = 0.7$, it is 3.4 times cheaper than interpolation using standard adaptivity, 8.1 times cheaper than adaptive PSP based on the standard method, and 5.2 times cheaper than PSP with adaptivity based on sensitivity scores. As a conclusion, although PSP and interpolation have very similar accuracies, the significantly lower cost of interpolation makes it our method of choice in the following realistic and computationally more expensive test case.

4.2. Realistic test case

While the CBC benchmark is a popular parameter set, corresponding nonlinear simulations are found to dramatically overestimate the transport levels obtained in the experiment that inspired this exercise. This can to some degree be explained by idealizations in the choice of parameters and by missing physics such as external shear flows. The level of realism shall therefore be increased by considering a particular GENE validation study from [9], performed for the ASDEX Upgrade experiment. Therein, nonlinear simulation results have been confronted 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 which became accessible via a new diagnostic. Linear simulations furthermore revealed that the parameters associated to the discharge of interest are very close to the dominant mode transitions and therefore represent a challenging set which could (i) be influenced by further stochastic parameters and (ii) turn out to be subject to interactions among these parameters. We took these concerns as motivation to perform a two-step UQ analysis.

4.2.1. Realistic test case with three stochastic parameters

First, we perform uncertainty propagation using three stochastic parameters as listed in Table 2. The three stochastic parameters are the logarithmic density and temperature gradients for ions and electrons. As for the CBC test case, the left and right bounds are symmetric around a nominal value used in deterministic simulations (see the last two columns in Table 2). However, for this test case the nominal values stem from experimental measurements.

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

We perform uncertainty propagation for multiple values of $k_y\rho_s$. Because this test case is more representative for real-world problems, wave-numbers up to the hyper-fine electron gyroradius scales are considered:

$$k_y\rho_s = 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.9, 1.0, \\ 2.0, 3.0, 4.0, 5.0, 7.5, 10.0, 12.5, 15.0, 17.5, 20.0, 22.5, 25.0, 27.5, 30.0. \quad (4.2)$$

As in the CBC test case, we treat $k_y\rho_s$ as a deterministic, free parameter and perform uncertainty propagation using Algorithm 3 for each value in (4.2).

To obtain some first insights into the underlying microinstabilities, deterministic runs using the nominal values of all input parameters, which stem from experimental data, are performed to compute 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 7 (note the negative sign of the frequency). Both curves are smooth and monotonous, suggesting that the dominant mode does neither change its (electron diamagnetic) drift direction nor its character in a dramatic way for the parameters at hand. The slightly different slopes in frequency, i.e., dispersion relations, and the first hump in growth rate imply that nevertheless two or more different microinstabilities such as pure TEM and TEM/ETG-hybrids are excited.




Figure 7: Growth rate (left) and frequency spectra (right) for the ASDEX Upgrade test case obtained using the nominal values of the input parameters. Note that for this test case, these values stem from experimental measurements.

In Section 4.1, we concluded that adaptive sparse grid interpolation is our method of choice for this realistic test case. Hence, a surrogate for the 3D stochastic gyrokinetic model is constructed via adaptive interpolation using both the proposed methodology based on sensitivity scores as well as the standard adaptive technique summarized in Section 2.3.3. In our proposed approach we prescribe $\tau_{\text{in}} = 10^{-6} \cdot \mathbf{1}_4$ and $L_{\text{max}}^{\text{in}} = 20$, whereas in the standard adaptive approach, $tol_{\text{in}} = 10^{-3}$ and $L_{\text{max}}^{\text{in}} = 20$ are employed. As in the CBC test case, the output of interest are the growth rate, $\gamma[c_s/L_s]$, of the dominant eigenmode with six digits of precision. Additionally, we compute its expectation, standard deviation and total Sobol' indices for each $k_y\rho_s$ in postprocessing (recall Section 2.4). Due to the larger growth rates and more simple mode structure, the runtime significantly decreased with $k_y\rho_s$.

The resulting expectation and one standard deviation as well as the deterministic growth rates are displayed in Figure 8. An overlap between the results obtained with the two adap-





Figure 8: Expected value 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 (4.2).

tive interpolation approaches is observed, showing that our proposed approach is similarly accurate as the standard adaptive method for this test case as well. In addition, the deterministic result is similar to the expectation of the stochastic simulation; the absolute difference varies uniformly roughly between 0.01 and 0.1. However, these differences are more significant for $k_y\rho_s \leq 1.0$ as in this region they represent a larger fraction of the growth rate's magnitude. Since this wave number range is usually considered to be most important for a correct assessment of the *ion heat flux*, results from corresponding non-linear simulations with nominal values should be carefully analysed. These results imply that the more likely results considering uncertainties may be somewhat different which is an important piece of information for the plasma turbulence modeling community.


For a more in-depth understanding of the results, the total Sobol' indices are furthermore depicted in Figure 9 (top two subfigures: standard adaptive approach, bottom subplots: our proposed approach). For a clear illustration, we split the results in two figures corresponding to $k_y\rho_s \leq 1$ and $k_y\rho_s > 1$, respectively. Again a good agreement between 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, whereas for $k_y\rho_s = 0.9, 1$, the two logarithmic temperature gradients are the most important, i.e., the intrinsic stochastic dimensionality is two. For $k_y\rho_s \leq 1$, the stochastic model exhibits a nonlinear behaviour due to the non-negligible interaction between the inputs. When $k_y\rho_s > 1$, on the other hand, the uncertainty in the logarithmic electron temperature gradient is the most important stochastic parameter. Moreover, the sum of the total Sobol' indices is close to 1, which shows that the stochastic model can be well approximated by a linear model. In other words, when $k_y\rho_s > 1$ the intrinsic stochastic dimension is one. Figure 9 also provides information on the underlying microinstabilities. For $k_y\rho_s > 1$, we clearly have an ETG




Interp. L^2 norm/cost: $k_y \rho_s \leq 1.0$



Interp. L^2 norm/cost: $k_y \rho_s \geq 2.0$



Interp. sensitivity scores: $k_y \rho_s \leq 1.0$



Interp. sensitivity scores: $k_y \rho_s \geq 2.0$




Figure 9: Total Sobol' indices for global sensitivity analysis via adaptive interpolation with the standard approach (top two subfigures) and our proposed approach (bottom two subfigures) for the ASDEX Upgrade test case with three uncertain parameters and $k_y \rho_s$ as in (4.2).

mode: the electron temperature gradient is most important parameter. For $k_y\rho_s \leq 1$, on the other hand, 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 mode for $k_y\rho_s \leq 1$ which is, however, also affected by the ion temperature gradient and may be in competition with subdominant ITG modes.

In Figure 10, we visualize the cost of the two adaptive strategies in terms of total number of grid points, i.e., GENE evaluations (left: $k_y\rho_s \leq 1$, right: $k_y\rho_s \geq 2$). Note that the




Figure 10: 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 for the ASDEX Upgrade test case with three uncertain parameters.

behaviour observed in the total Sobol' indices is reflected in the cost as well. It is computationally more expensive to perform uncertainty propagation for $k_y\rho_s \leq 1$ than for $k_y\rho_s > 1$ due to the higher intrinsic stochastic dimension. In addition, our approach requires fewer GENE runs than the standard adaptive approach for all normalized perpendicular wavenumbers $k_y\rho_s$. For example, at $k_y\rho_s = 3$, we need about 3.3 times fewer GENE evaluations than the standard approach. 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 difference between the given stochastic dimensionality, three, is not significantly larger than the intrinsic dimensionality, one.

To further illustrate the computational savings due to our approach, we depict, in Figure 11, the multiindices and the associated sparse grids at $k_y\rho_s = 1.0$ (left: standard adaptive approach, right: sensitivity scores approach). The three estimated total Sobol' indices are $\hat{S}_1^T = 0.014$, $\hat{S}_2^T = 0.120$, and $\hat{S}_3^T = 0.900$. Therefore, the logarithmic electron density gradient is the most important parameter, the logarithmic ion density gradient is the second, while the third important parameter is the logarithmic density gradient. The total Sobol' indices sum up to 1.033, which indicates that the interactions between the three inputs are negligible. The multiindex sets in the top subfigures in Figure 11 show that both algorithms detect the third stochastic direction as the most important. However, the standard adaptive algorithm adds many (unnecessary) interaction multiindices, while our approach better exploits the fact that one direction is significantly more important than the other two combined and that the coupling between the three inputs is negligible. At the end of

the refinement process, the standard approach yields 83 grid points, whereas our approach needs a total of only 32 points to produce similar results.

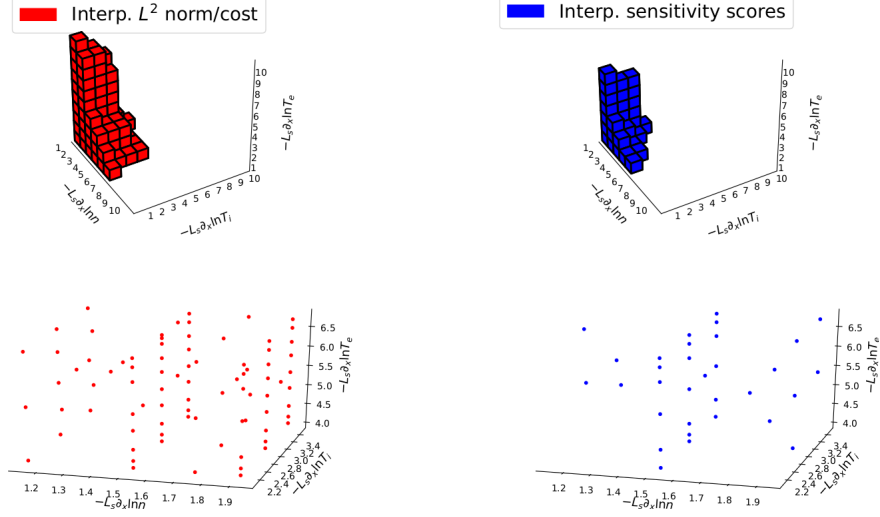


Figure 11: Multiindices and their associated sparse grids at $k_y \rho_s = 1.0$ for adaptive interpolation for the 3D ASDEX Upgrade scenario (left: standard adaptive approach, right: our proposed approach).

4.2.2. 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 3, resulting in a total of 12 stochastic inputs for the ASDEX Upgrade test case. The first seven characterize the two particle species, ions and

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 3: Summary of the 12 stochastic parameters considered for the ASDEX Upgrade test case.

electrons, while the latter five parameters are associated to the magnetic geometry. As for the previous test case, the uniform bounds of the stochastic inputs are symmetric around

nominal values stemming from experimental measurements. Note that the deterministic growth rates and frequencies are the same as for the 3D case in Section 4.2.1 as we have the same nominal input parameters.

We extend the stochastic setup from the 3D scenario for 12 uncertain inputs: $\boldsymbol{\tau}_{\text{in}} = 10^{-6} \cdot \mathbf{1}_{13}$, while the other tolerance and the maximum levels remain the same. We visualize the expectation, one standard deviation and the deterministic growth rates for both the 3D and 12D scenarios, for comparison, in Figure 12. On the one hand, we see that for this





Figure 12: Expected value and 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 (4.2).



scenario too, our proposed approach yields results very similar to the standard adaptive strategy. Moreover, the difference between the deterministic results and the error plots 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, we assume that the extra 9 uncertain parameters contribute insignificantly to the overall results which would be an important piece of information for on-going validation studies which may now mainly focus on the logarithmic gradient sensitivities in nonlinear turbulence simulations.

To ascertain the above assumption, we analyse the total Sobol' indices for sensitivity analysis in Figure 13 (top: standard adaptive interpolation, bottom: adaptive interpolation with our proposed sensitivity scores approach). Note that as for the 3D scenario, our proposed approach produces results very similar to the standard 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 9 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. In addition, throughout the considered $k_y\rho_s$ domain, the total Sobol' indices of the aforementioned important parameters are qualitatively and quantitatively very similar to the indices from the 3D scenario. Hence, we conclude that



Interp. L^2 norm/cost: $k_y \rho_s \leq 1.0$

Interp. L^2 norm/cost: $k_y \rho_s \geq 2.0$



Interp. sensitivity scores: $k_y \rho_s \leq 1.0$

Interp. sensitivity scores: $k_y \rho_s \geq 2.0$





Figure 13: Total Sobol' indices for global sensitivity analysis via adaptive interpolation with the standard approach (top two subfigures) and our proposed approach (bottom two subfigures) for the ASDEX Upgrade test case with 12 uncertain parameters and $k_y \rho_s$ as in (4.2).

the 9 extra stochastic input parameters contribute insignificantly to the overall results. In addition, the underlying microinstabilities are the same as for the 3D scenario.

In Figure 14, we visualize the cost comparison between our proposed approach and the standard adaptive approach (left: $k_y\rho_s \leq 1$, right: $k_y\rho_s \geq 2$). We observe that for all $k_y\rho_s$ the proposed method requires significantly fewer GENE evaluations than the standard adaptive approach. For example, at $k_y\rho_s = 15$ our approach requires 13.3 times fewer evaluations than the standard adaptive approach. Note, on the one hand, that for $k_y\rho_s > 1$ the cost of the sensitivity scores approach is similar for all $k_y\rho_s$, which is due to having the intrinsic stochastic dimensionality one in this region: our methodology detects that only one stochastic direction is important, hence it refines predominantly that direction. To put the cost savings for this scenario into perspective, a full-tensor grid typically used in parameter scans constructed using q_0 points in each direction would require 10^{12} GENE evaluations in total, which is computationally prohibitive. On the other hand, the standard adaptive strategy required at most 2283 GENE evaluations, whereas the largest number of simulations in the proposed approach was 546.




Figure 14: 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 the two stochastic scenarios considered for the ASDEX Upgrade test case, the presented adaptive sparse approximation strategies enable uncertainty propagation and sensitivity analysis in problems in which standard approaches would be almost impossible to use. In addition, our proposed approach proved to be comparatively accurate to the standard adaptive approach, but at a significantly reduced computational cost. Finally, when the intrinsic stochastic dimensionality is smaller than the total number of stochastic inputs, which is usually the case in practice, our approach explores and exploits this structure to preferentially refine the important directions.

5. Conclusions and outlook

In this work, we introduced a novel adaptive sparse grid approximation methodology for uncertainty propagation and sensitivity analysis in computationally expensive real-world applications. We leveraged Sobol' decompositions and introduced a sensitivity scoring system to drive the adaptive process. By employing sensitivity information, we explored and

exploited the anisotropic coupling of the stochastic inputs as well as the lower intrinsic stochastic dimensionality which are typical in practical applications.

The proposed approach can be formulated in terms of arbitrary approximation operators and point sets. Here, we considered Lagrange interpolation and pseudo-spectral projection operators which we defined using (L)-Leja points. Although the proposed methodology is model-agnostic, a real-world problem arising in plasma turbulence research has been considered. Therefore, another novelty was that this work represents, to the best of our knowledge, one of the first uncertainty propagation studies in plasma microinstability analysis.

We tested and compared the proposed approach with a standard adaptive strategy in two plasma microturbulence simulation test cases. In both test cases, the gyrokinetic code GENE was employed to characterize the microinstabilities. The first test case was based on the popular Cyclone Base Case benchmark of the gyrokinetic community. Therein, we considered eight uncertain inputs characterizing both the underlying particle species, and the magnetic geometry. The second test case was a real-world example stemming from a particular validation study for the ASDEX Upgrade experiment. For this test case, we carried out a two step analysis, initially considering three uncertain inputs characterizing the ions and electrons, and then 12 stochastic parameters associated to the particle species and the magnetic geometry. Our results showed that the proposed approach is comparatively accurate to the standard adaptive approach while requiring a significantly reduced computational cost; for example, for the 12D scenario, we obtained factors of up to 13.3 fewer GENE evaluations. The presented adaptive sparse grid approximations enable the uncertainty propagation and sensitivity analysis in higher-dimensional plasma microturbulence problems, which would be almost impossible to tackle with standard screening approaches. It furthermore demonstrated its capability to provide deep and new insights on the parametric dependencies which will be very helpful for further comparison with experiments and construction of reduced plasma turbulence models.

For future research, we want to extend the current methodology to a multifidelity framework for uncertainty propagation. Our long-term goal with respect to plasma turbulence research is to quantify uncertainty in fully nonlinear gyrokinetic simulations.

Acknowledgements

The first author thankfully acknowledges the support of the German Academic Exchange Service (<http://daad.de/>) and the support of the German Research Foundation and Technical University of Munich through the TUM International Graduate School of Science and Engineering (<http://igsse.tum.de/>) within the project 10.02 BAYES. Moreover, the authors gratefully acknowledge the compute and data resources provided by the Leibniz Supercomputing Centre (www.lrz.de). Part of this research was performed while the first, third and fourth authors were visiting the Institute for Pure and Applied Mathematics (IPAM), which is supported by the National Science Foundation.

References

- [1] Beer, M. A., Cowley, S. C., Hammett, G. W., 1995. Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Phys. Plasmas* 2, 2687–2700.
- [2] Berrut, J.-P., Trefethen, L. N., 2004. Barycentric lagrange interpolation. *SIAM Review* 46 (3), 501–517.
- [3] Brizard, A. J., Hahm, T. S., 2007. Foundations of nonlinear gyrokinetic theory. *Rev. Mod. Phys.* 79, 421–468.
- [4] Bungartz, H.-J., Griebel, M., 2004. Sparse grids. *Acta Numer.* 13, 147–269.
- [5] Conrad, P. R., Marzouk, Y. M., 2013. Adaptive smolyak pseudospectral approximations. *SIAM J. Sci. Comput.* 35 (6), A2643–A2670.
- [6] Constantine, P. G., Eldred, M. S., Phipps, E. T., 2012. Sparse pseudospectral approximation method. *Comput. Methods in Appl. Mech. Eng.* 229-232, 1–12.
- [7] Dimits, A. M., Bateman, G., Beer, M. A., Cohen, B. I., Dorland, W., Hammett, G. W., Kim, C., Kinsey, J. E., Kotschenreuther, M., Kritz, A. H., Lao, L. L., Mandrekas, J., Nevins, W. M., Parker, S. E., Redd, A. J., Shumaker, D. E., Sydora, R., Weiland, J., 2000. Comparisons and physics basis of tokamak transport models and turbulence simulations. *Phys. Plasmas* 7 (3), 969–983.
- [8] Formaggia, L., Guadagnini, A., Imperiali, I., Lever, V., Porta, G., Riva, M., Scotti, A., Tamellini, L., 2013. Global sensitivity analysis through polynomial chaos expansion of a basin-scale geochemical compaction model. *Comput. Geosci.* 17 (1), 25–42.
- [9] Freethy, S. J., Grler, T., Creely, A. J., Conway, G. D., Denk, S. S., Happel, T., Koenen, C., Hennequin, P., White, A. E., 2018. Validation of gyrokinetic simulations with measurements of electron temperature fluctuations and density-temperature phase angles on asdex upgrade. *Phys. Plasmas* 25 (5), 055903.
- [10] Gander, W., 2005. Change of basis in polynomial interpolation. *Numer. Linear Algebra Appl.* 12 (8), 769–778.
- [11] Garbet, X., Idomura, Y., Villard, L., Watanabe, T. H., 2010. Topical Review: Gyrokinetic simulations of turbulent transport. *Nucl. Fusion* 50, 043002.
- [12] Gerstner, T., Griebel, M., 2003. Dimension-adaptive tensor-product quadrature. *Computing* 71 (1), 65–87.
- [13] Görler, T., Lapillonne, X., Brunner, S., Dannert, T., Jenko, F., Merz, F., Told, D., 2011. The global version of the gyrokinetic turbulence code gene. *J. Comput. Phys.* 230 (18), 7053–7071.
- [14] Görler, T., White, A. E., Told, D., Jenko, F., Holland, C., Rhodes, T. L., 2014. A flux-matched gyrokinetic analysis of diii-d l-mode turbulence. *Phys. Plasmas* 21 (12), 122307.
- [15] Griebel, M., Schneider, M., Zenger, C., 1992. A combination technique for the solution of sparse grid problems. In: de Groen, P., Beauwens, R. (Eds.), *Iterative Methods in Linear Algebra*. IMACS, Elsevier, North Holland, pp. 263–281.
- [16] Güler, T., Tronko, N., Hornsby, W. A., Bottino, A., Kleiber, R., Norscini, C., Grandgirard, V., Jenko, F., Sonnendrcker, E., 2016. Intercode comparison of gyrokinetic global electromagnetic modes. *Phys. Plasmas* 23 (7), 072503.
- [17] Hegland, M., 2003. Adaptive sparse grids. In: Burrage, K., Sidje, R. B. (Eds.), *Proc. of 10th Computational Techniques and Applications Conference CTAC-2001*. Vol. 44. pp. C335–C353.
- [18] Jantsch, P., Webster, C. G., Zhang, G., 06 2018. On the Lebesgue constant of weighted Leja points for Lagrange interpolation on unbounded domains. *IMA J. Numer. Anal.* 39 (2), 1039–1057.
- [19] Jenko, F., Dorland, W., Kotschenreuther, M., Rogers, B. N., 2000. Electron temperature gradient driven turbulence. *Phys. Plasmas* 7 (5), 1904–1910.
- [20] Krommes, J. A., 2012. The Gyrokinetic Description of Microturbulence in Magnetized Plasmas. *Annu. Rev. Fluid Mech.* 44, 175–201.
- [21] Ma, X., Zabaras, N., 2009. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *J. Comput. Phys.* 228 (8), 3084–3113.
- [22] Marzouk, Y., Moselhy, T., Parno, M., Spantini, A., 2016. *Sampling via Measure Transport: An Introduction*. Springer International Publishing, Cham, pp. 1–41.
- [23] MPI Forum, 2017. MPI: A Message-Passing Interface Standard. Version 4.0. Available at: <http://www.mpi-forum.org>.

- [24] Narayan, A., Jakeman, J. D., 2014. Adaptive Leja sparse grid constructions for stochastic collocation and high-dimensional approximation. *SIAM J. Sci. Comput.* 36 (6), A2952–A2983.
- [25] Nobile, F., Tempone, R., Webster, C. G., 2008. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* 46 (5), 2309–2345.
- [26] Peherstorfer, B., Willcox, K., Gunzburger, M., 2018. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review* 60 (3), 550–591.
- [27] Rüttgers, A., Griebel, M., 2018. Multiscale simulation of polymeric fluids using the sparse grid combination technique. *Appl. Math. Comput.* 319, 425–443.
- [28] Schillings, C., Schwab, C., 2013. Sparse, adaptive smolyak quadratures for bayesian inverse problems. *Inverse Probl.* 29 (6), 065011.
- [29] Sobol, I. M., 2001. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Math. Comput. Simul.* 55 (1), 271–280.
- [30] Sudret, B., 2008. Global sensitivity analysis using polynomial chaos expansions. *Reliab. Eng. Syst. Safe.* 93 (7), 964–979.
- [31] Told, D., Jenko, F., Görler, T., Casson, F. J., Fable, E., 2013. Characterizing turbulent transport in asdex upgrade l-mode plasmas via nonlinear gyrokinetic simulations. *Phys. Plasmas* 20 (12), 122312.
- [32] Vaezi, P., Holland, C., 2018. An improved approach to uncertainty quantification for plasma turbulence validation studies. *Fusion Sci. Technol.* 74 (1-2), 77–88.
- [33] White, A. E., Görler, T., 2017. Special issue on comparing gyrokinetic simulations to experiments. *Plasma Phys. Controlled Fusion* 59 (5), 050101.
- [34] Winokur, J., Kim, D., Bisetti, F., Le Maître, O. P., Knio, O. M., 2016. Sparse pseudo spectral projection methods with directional adaptation for uncertainty quantification. *J. Sci. Comput.* 68 (2), 596–623.
- [35] Xanthopoulos, P., Mynick, H. E., Helander, P., Turkin, Y., Plunk, G. G., Jenko, F., Görler, T., Told, D., Bird, T., Proll, J. H. E., 2014. Controlling Turbulence in Present and Future Stellarators. *Phys. Rev. Lett.* 113 (15), 155001.
- [36] Xiu, D., 2010. *Numerical methods for stochastic computations : a spectral method approach*. Princeton, N.J. Princeton University Press.
- [37] Xiu, D., Hesthaven, J. S., 2005. High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.* 27 (3), 1118–1139.