

A non-intrusive reduced order model for the characterisation of the spatial power distribution in large thermal reactors

Original

A non-intrusive reduced order model for the characterisation of the spatial power distribution in large thermal reactors / Abrate, N., Dulla, S., Pedroni, N.. - In: ANNALS OF NUCLEAR ENERGY. - ISSN 0306-4549. - ELETTRONICO. - 184:(2023), p. 109674. [10.1016/j.anucene.2022.109674]

Availability:

This version is available at: 11583/2975120 since: 2023-01-25T09:38:16Z

Publisher:

Elsevier

Published

DOI:10.1016/j.anucene.2022.109674

Terms of use:

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

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2023. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.anucene.2022.109674>

(Article begins on next page)

Highlights

A non-intrusive reduced order model for the characterisation of the spatial power distribution in large thermal reactors

Nicolò Abrate, Sandra Dulla, Nicola Pedroni

- Large thermal reactors endowed with heavy reflectors are very sensitive to localised perturbations.
- Legacy code may be inefficient for this kind of parametric analysis, and their improvement and modification would require a new code qualification process.
- A non-intrusive surrogate model capable of delivering approximated yet sufficiently accurate power distributions is proposed.
- The meta-model is composed of two-step, one for the group constant generation and one for the full-core diffusion calculations.
- Due to the spatial dependence of the input perturbations, some strategies are proposed and successfully applied in the metamodel training phase.
- The accuracy of the surrogate model, assessed with the bootstrap technique, is satisfactory.

A non-intrusive reduced order model for the characterisation of the spatial power distribution in large thermal reactors

Nicolò Abrate^a, Sandra Dulla^a, Nicola Pedroni^a

^a*Politecnico di Torino, Dipartimento Energia,
Corso Duca degli Abruzzi, 24 - 10129 Torino (Italy)*

Abstract

Large thermal reactors endowed with heavy reflectors are very sensitive to localised perturbations. To adequately characterise the spatial effects of these localised disturbances on the reactor power and, thus, to study the spatial stability of the core in a computationally efficient manner, without modifying the legacy codes that may be inefficient for this kind of parametric safety analysis, a non-intrusive surrogate model to approximate the power distributions is proposed in this paper. The meta-model is composed of two steps. First, the group constants, originally produced with the Monte Carlo method for the full-core diffusion calculations, are approximated via a polynomial chaos expansion regression model, with an acceptable accuracy. Then, a model combining Proper Orthogonal Decomposition and Radial Basis Functions is trained to replace the expensive full-core diffusion calculations. Due to the spatial dependence of the input perturbations, some *ad hoc* strategies are proposed and successfully applied in the reduced order model training phase. Finally, the bootstrap technique is employed to assess the quality of the meta-model approximations and to provide an estimation of the modelling error distribution.

Keywords: pressurised water reactors; power tilt; flux tilt; spatial perturbations; two-group diffusion; full-core analysis; neutron importance; reduced order modelling; polynomial chaos expansion; proper orthogonal decomposition; bootstrap; radial basis functions;

1. Introduction

Due to the increasingly strict constraints on safety, sustainability and capital cost, Gen-III+ Light Water Reactors (LWRs) are endowed with new design features, which aims at extending the life of the system, reducing the radioactive waste and increasing the safety level of the plant. These improved performance features are mainly obtained increasing the core size, which improves the reactor burn-up and, consequently, the lifetime of the plant, and adopting a heavy reflector, made of stainless steel, for the reduction of neutron leakages from the core.

Compared to traditional reflectors, the heavy reflector tends to absorb more thermal neutrons, reflecting the epithermal ones towards the core. This aspect constitutes an advantage both on the radiological risk, that is reduced, and on the neutron economy, that improves.

In spite of these advantages, the larger core optical size and the heavy reflector increase the overall spatial decoupling of the system, which thus becomes more sensitive to localised perturbations than compact, leakage-dominated cores. The combination of different perturbations, which may be either

Email addresses: nicolo.abrate@polito.it (Nicolò Abrate), sandra.dulla@polito.it (Sandra Dulla), nicola.pedroni@polito.it (Nicola Pedroni)

of operational nature, like the fuel assemblies (FAs) bowing, or of fabrication nature, like the fuel tolerances, can induce local inhomogeneities in the moderation process. As a consequence, when these disturbances do not mutually compensate, some localised perturbations in the flux, known as *flux tilts*, can arise (Sargeni et al., 2016). According to the standard safety practices for PWRs (NRC, 2016), in case the power tilt overcomes a certain threshold, a set of specific actions are to be followed in order to bring the core back to a safe operational state, starting with the thermal power reduction.

Since the effect of the tilt on the overall power distribution is proportional to the spatial decoupling of the system, the heavy reflector and the larger core size amplify the tilt detrimental effect on the core neutronic stability. This unpleasant design feature makes the operation of the EPR more delicate, since the detection of flux and power tilts is practically quite difficult, especially during the cold start-up (Sargeni et al., 2016). In this phase, the tilts can be particularly dangerous, since the thermal feedback is not effective in damping their effects, and the fission chambers used to map the flux distribution in the core are less precise, because of the lower flux with respect to the nominal one. Hence, the ex-core measurements should be complemented by accurate calculations.

Unfortunately, the larger spatial decoupling of these systems have detrimental effects not only on their operation, but also on the efficiency of the numerical techniques used to simulate them. Most of the legacy codes used in LWR industry heavily rely on the so-called *power iteration method*, whose convergence rate depends on the dominance ratio of the system, i.e. the ratio between the fundamental and the first-order eigenvalue of the multiplication eigenvalue problem (Abrate et al., 2019a). This parameter is inversely proportional to the decoupling degree of the reactor, thus the determination of the fundamental eigenpair of the system is extremely inefficient.

One of the best alternative to this old and inefficient approach consists of adopting a method belonging to the class of Krylov-Arnoldi methods (Saad, 1992). In virtue of their non-intrusiveness, these methods can be wrapped around an existing code, minimising the implementation effort and making their validation easier. However, even in the case a method of this class was implemented in the legacy codes, the number of parameters influencing the flux distribution would require a too large number of calculations.

The usual way to reduce the number of expensive calculations in reactor physics is to rely on perturbation methods, like the Generalised Perturbation Theory (GPT) (Gandini, 1978). This approach yields an approximation of the perturbed flux and multiplication eigenvalue, exploiting the knowledge of the perturbation and of some of the higher-order flux harmonics. In this respect, this method would take full advantage of computationally efficient methods like the Krylov-Arnoldi, which are sufficiently robust to yield the leading higher-order harmonics even for system featured by a low dominance ratio.

The GPT is a well-established technique for sensitivity and uncertainty analysis (Aufiero et al., 2015; Abrate et al., 2021), available in most of the legacy codes, but it may not converge if the perturbation magnitude is too large (Abrate et al., 2019b). Since the tilt characterisation may require the analysis of strong perturbations, the use of the GPT for this kind of analysis may not be sufficiently robust and reliable to perform the core qualification.

In light of these limitations concerning the code computational performances and considering that modifying the legacy codes is not viable as a short term solution due to complexity of their qualification, a possible approach to tackle these issues is to resort to non-intrusive reduced order models. These surrogate models allow to mimic either the reality, reproducing the experimental data, or a physico-mathematical-computational model, reproducing its output with a lower computational time, at the price of some controllable approximations (Benner et al., 2015).

When the surrogate is built as a computationally cheaper alternative to the reference model, it is often addressed with the term *meta-model*, since it aims at reproducing the responses of the original model as well as the latter tries to imitate the reality. In this respect, the definition of Science given in (Popper and Barthley III, 1988), "*Science may be described as the art of systematic over-simplification — the art of discerning what we may with advantage omit*", perfectly describes the "art" of meta-modelling: in some specific situations, it is convenient and possible to simplify a complex model that represents the real world, still retaining its most important features in order to map exhaustively its behaviour in a very broad range of conditions. In other words, we omit some physical details in favour

of information concerning the overall system behaviour, which is the fundamental information that cannot be disregarded. Meta-modelling is particularly relevant in the frame of the safety analysis of complex systems, requiring a very large number (e.g., hundreds of thousands) of evaluations of complex, typically long-running models in correspondence of different combinations of many input parameters, as in the characterisation of the spatial distribution of the core power in a reactor subject to localised perturbations.

Reduced Order Models (ROMs) can be roughly divided in physics-driven and data-driven methods. The first ones yield an approximated model by direct manipulation of the Full-Order Model (FOM) (Lorenzi et al., 2016, 2017; Sun et al., 2020). Acting on the model governing equations minimises the biases in the model response and can provide an *a priori* estimate of the ROM error bounds (Rahman et al., 2018). However, on top of the manifest difficulty in accessing the code in case it is not open-source, these interventions would demand a new code verification and validation phase, which is a critical aspect when dealing with codes used for the design verification of nuclear installations.

On the contrary, the second class of methods has the significant advantage of considering the code as a black-box that receives an input parameter and responds with a certain output value. The input-output pair can be used to train a supervised empirical regression model using suitable mathematical devices. The data-driven nature of such methods makes them non-intrusive and application-independent, hence they can be employed effectively even with commercial, validated codes (Casenave et al., 2020).

However, this class of methods also brings some drawbacks. First, the fact that they need data to learn the model behaviour makes their performances sensitive to the training data selection. In order to realise an accurate yet flexible model, a trade-off between the *model variance*, i.e. the sensitivity of the model to the training sample choice, and the *model bias*, i.e. the error between the meta-model and the reference model, is needed. The other main drawback is the fact that the lack of important training data can seriously limit the ROM prediction capability. This aspect is brilliantly summarised in Lassila et al. (2014) with the *caveat*, "If it is not in the snapshots, it is not in the ROM". Finally, it is very difficult to obtain an *a priori* error prediction (Rahman et al., 2018).

Nevertheless, since the advantages of these methods often overcome their shortcomings, they are becoming very popular.

Typically, data-driven methods are composed by an offline and an online stage. In the first one, the computationally expensive FOM is run possibly a small number of times to obtain the model realisation for the ROM training. The number of model executions strongly depends on its complexity and on the number of input parameters which can affect its response. In the online phase, the surrogate is finally adopted to mimic the FOM behaviour for new input parameter values.

A common feature shared by the different ROMs is that they usually rely on dimensionality reduction techniques aiming at the minimisation of the degrees of freedom of the FOM. Several methods are available in the open literature for reducing the dimensionality of FOMs and correspondingly aiding the construction of ROMs. In this paper, a feature extraction technique based on the Proper Orthogonal Decomposition (POD) (Volkwein, 2011; Nagel et al., 2020) has been chosen, due to its suitability and demonstrated ability in efficiently decomposing FOM functional responses (of paramount interest to the present work) into spatio-temporal bases and a set of parameter-dependent coefficients (Marrel et al., 2014; Nanty et al., 2017; Wu et al., 2018b,a).

The major differences in the Non-Intrusive ROMs (NIROMs) are often related to the technique adopted for the evaluation of the POD parameter-dependent coefficients when new parameters are investigated. Among the various methods present in the literature, a SVM with a kernel constituted by radial basis functions (RBF) (Schaback, 1995), which have been proved to be very effective for non-linear, high-dimensionality interpolation in a broad range of applications (Xiao et al., 2017; Dutta et al., 2021; Iuliano and Quagliarella, 2013; Chen et al., 2018) is adopted to interpolate the POD coefficients on new parameters in this work.

The aim of this paper is to present a non-intrusive POD-RBF model for an efficient model reduction in the frame of the characterisation of the spatial effects featuring large thermal reactors, which consists of determining the flux and/or power spatial distributions of the system in presence of localised

perturbations. Due to the peculiarities of this application with respect to other common NIROM applications (e.g., CFD), some modifications to the POD-RBF model will be proposed in the paper. In particular, a two-step ROM will be envisaged, to mimic the usual two-step approach followed in the full-core LWR analyses. The first ROM will be based on a polynomial chaos expansion regression, for the estimation of the perturbed few-group constants, needed for the diffusion calculation, which is the objective of the second model reduction. In this last case, a physics-based importance sampling strategy is defined to explore the most important perturbed configurations of the core, to gather the training data. In this phase, an uncommon metric is adopted for evaluating the distance between the perturbed core configurations in the RBF training process. Finally, the prediction capabilities of the two-step ROM are assessed exploiting the bootstrap method, which allows to get a conservative estimate of the error distribution of the surrogate model.

The paper is organised as follows. In section 4.1, an overview of the POD-RBF model is given, focusing on its structure and features. Then, in section 2, the physico-mathematical statement of the problem is given, addressing the need for two ROMs, one for producing the perturbed group constants and one for evaluating the full-core fission power distribution. In section 3 a polynomial chaos regression model is presented and adopted for the group constant generation, while in section 4 the POD-RBF model is applied to the full-core analysis, using the UAM benchmark as a case study. Finally, in section 5 some conclusions and future perspectives are given.

2. Physico-mathematical statement of the problem

The characterisation of the tilt behaviour according to the different types of localised perturbations occurring during the nominal core operation can be achieved estimating the flux and the associated thermal power distributions, which can be obtained by solving the steady state neutron balance equation.

Since the analysis should involve the full-core, it is not practical nor useful to use a detailed transport model. As it is usually done in this case, a low-transport model is often sufficient to convey an acceptable amount of information regarding the state of the core, provided that a set of homogenised and collapsed cross sections representative of the system under investigation is available. The generation of the few-group data is usually accomplished using a high-fidelity transport calculation on a simplified reactor configuration, e.g., a set of representative fuel assemblies. The usual definition of collapsed cross section yields

$$\Sigma_{y,g}(\vec{r}, \vec{p}) = \frac{\int_{E_{g-1}}^{E_g} dE \Sigma_y(E, \vec{p}) \phi(E, \vec{p})}{\int_{E_{g-1}}^{E_g} dE \phi(E, \vec{p})}, \quad (1)$$

where the terms have their usual meaning and \vec{p} represents a general localised perturbation that can affect the nominal local conditions.

More specifically, for the Gen-III+ reactors as the European Pressurised Reactor (EPR), the localised perturbations indicated by \vec{p} can be assumed to belong to two main families, namely fabrication tolerances and operational disturbances. In the first group, the main fabrication uncertainties can be due to

1. the fuel pellet density, which has an impact on the fission rate
2. the fuel pellet diameter, which influences both the fission rate and the moderation ratio
3. the cladding thickness, which has an impact on the parasitic capture and on the moderation ratio
4. the enrichment, which affects the fission rate.

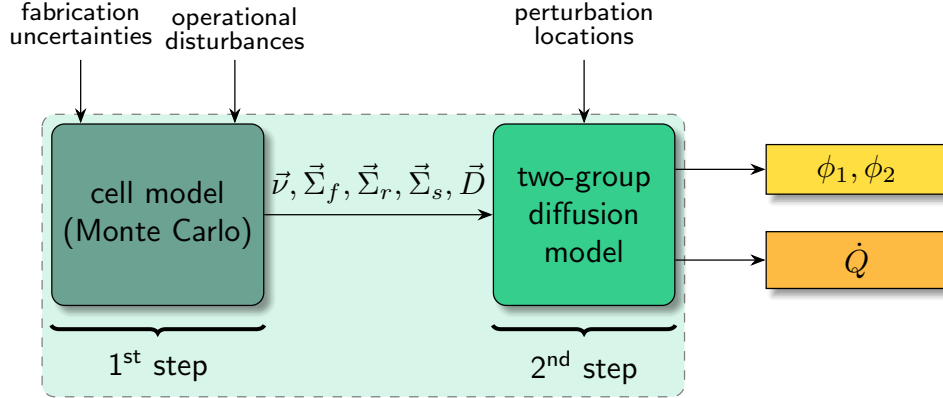


Figure 1: *Two-step calculation path followed for the full-core neutronic simulation.*

Concerning the operational disturbances, it is possible to identify

1. the uncertainty in the behaviour of the circulation pumps (due, e.g., to random variability of their operating conditions), which affects the local water density and, thus, the local moderation
2. the assembly bowing, which alters the flow cross section in the sub-channel and, thus, the local moderation ratio
3. the change in the fuel density, which depends on two competing phenomena, i.e. the fuel densification and the fuel swelling (Cacuci, 2010). The first mechanism, due to the coalescence of the micro-pores in the fuel structure under irradiation, tends to increase the density of the fuel, while the second one, caused by the accumulation of the gaseous fission products, tends to decrease the fuel density because of the pellet ballooning.

These random, localised disturbances introduce some competing phenomena in the neutron balance. Because of the large spatial decoupling of the system, which is a consequence of the short diffusion length of the neutrons, the information of these disturbances is not propagated at the full-core level, but affects the local flux behaviour, inducing the tilt. Since these phenomena occurs at the fuel assembly level, but different assemblies can be affected, the proper characterisation of these disturbances should be carried out considering them since the few-group data generation process.

Once the set of nominal and perturbed few-group constants has been produced, the full-core flux and power distributions can be finally estimated with a few-group diffusion model. Due to the thermal spectrum of the system, a two-group approach can be adopted for this purpose. Consistently with the system features, it is possible to neglect the up-scattering and to consider that all neutrons coming from fissions are emitted in the fast group,

$$\begin{cases} -\nabla \cdot D_1(\vec{r}, \vec{p}) \nabla \phi_1(\vec{r}, \vec{p}) + \Sigma_{r,1}(\vec{r}, \vec{p}) \phi_1(\vec{r}, \vec{p}) = \frac{1}{k_{\text{eff}}} \left[\nu \Sigma_{f,1}(\vec{r}, \vec{p}) \phi_1(\vec{r}, \vec{p}) + \nu \Sigma_{f,2}(\vec{r}, \vec{p}) \phi_2(\vec{r}, \vec{p}) \right] \\ -\nabla \cdot D_2(\vec{r}, \vec{p}) \nabla \phi_1(\vec{r}) + \Sigma_{r,1}(\vec{r}, \vec{p}) \phi_1(\vec{r}) - \Sigma_{s,1 \rightarrow 2}(\vec{r}, \vec{p}) \phi_1(\vec{r}) = 0, \end{cases} \quad (2)$$

where the terms have their usual meaning, as it can be found in classical books (Bell and Glasstone, 1970; Duderstadt and Hamilton, 1976).

The two-step nature of the full-core diffusion calculation, depicted in fig. 1, poses a question concerning the development of a reduced order simulation framework. One possibility is to consider this

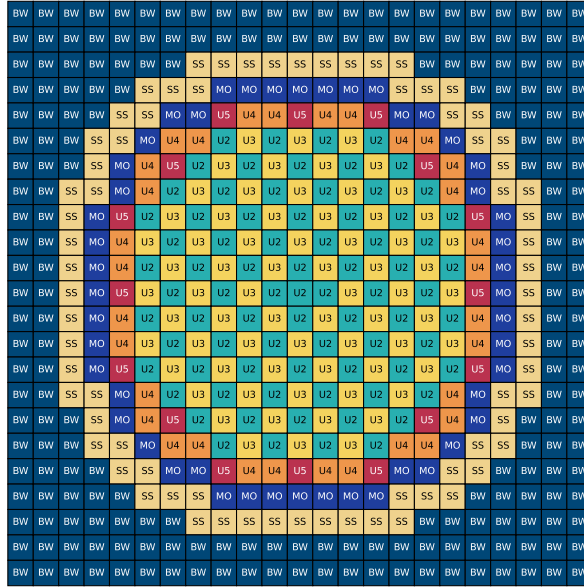


Figure 2: Sketch of the UAM benchmark. BW: borated water; MO: Mixed Oxides of U and Pu; U4: 2.1 % enriched UOX; U3: 3.2 % enriched UOX with burnable poison rods of $UO_2-Gd_2O_3$; SS: stainless steel; U5: 3.2 % enriched UOX with burnable poison rods of $UO_2-Gd_2O_3$; U2: 2.1 % enriched UOX;

calculation chain as a single model \mathcal{M} , sketched as the dashed box in the figure, which takes as input parameters the perturbation magnitudes and their location in the core and outputs the flux and power distributions at the full-core level. Considering the approach as monolithic, the resulting ROM would have the advantage of avoiding the evaluation of the set of few-group data, establishing a direct relationship between input and output. However, due to the large number of input parameters, especially concerning the number and the location of the perturbations, the training phase would be very computationally demanding.

An alternative strategy to this monolithic approach is to consider a reduced order model built in two stages, each one approximating the two steps of the FOM. This two-step ROM would allow to decouple the training phase, taking advantage of the features of the two models, although at the price of the propagation of the modelling errors from the cell ROM to the full-core ROM.

In order to develop an *ad hoc* ROM for the tilt characterisation, a simplified yet still realistic problem is considered, based on the Uncertainty Analysis in Modelling (UAM) benchmark for the Gen-III+ systems (Ivanov et al., 2007), which is sketched in fig. 2. This specific 2D model is featured by 529 squared fuel assemblies with a pitch of 21.42 cm, and includes the different regions of the core, like the heavy reflector and the borated water surrounding the core.

Figure 3 shows the detailed geometry for the fuel assemblies (FAs) constituting the core. For each kind of FAs, a set of energy collapsed and spatial homogenised cross sections and diffusion coefficients are computed with the Serpent 2 Monte Carlo code (Leppänen et al., 2015).

Because of the large number of FAs, the input parameter space of this problem is huge and very complex, since, in principle, different disturbances can affect different FAs at the same time, with different probabilities. Since the aim of this paper is to provide a "proof of concept" for the development of a non-intrusive ROM to characterise the tilt effects on the core operation, the parameter space




Figure 3: Pin cells of the UAM benchmark: UOX with 3.2% enrichment and Gd rods in orange (left), UOX with 2.1% enrichment (centre) and MOX with 9.8% (emerald green), 6.5% (pine green) and 3.7% (mint green) Pu enriched rods (right) .




Figure 4: Distributions considered for the uncertainties in the physical input parameters for the lattice calculation.

variability is properly reduced to limit the computational effort for the model training, still maintaining an adequate level of generality. Therefore, to reduce the number of scenarios, only two disturbances are considered in the following.

The first disturbance is assumed to be the fuel diameter pellet. Since it affects both the fission and the moderation, which depends on the ratio between the fuel and the coolant volume in the elementary cell, this parameter is a good representative for the group of disturbances of fabrication uncertainties. As such, it is assumed that this parameter is normally distributed, following the specifications in [Ivanov et al. \(2007\)](#), i.e. a nominal value of 0.41625 cm and a standard deviation equal to 0.00043. The second disturbance is assumed to be a variation of the water density, which can be assumed originating either from the FAs bowing or from the pump operational uncertainty, as done in [Sargeni et al. \(2016\)](#). The nominal coolant density is $0.700758 \text{ g cm}^{-3}$ and its standard deviation is assumed to be $0.02102274 \text{ g cm}^{-3}$. Due to the difficulty in evaluating the probability of this water density variations, it is assumed that this parameter has a uniform distribution between $\pm 3\%$ of the nominal value, following again [Sargeni et al. \(2016\)](#). The Probability Density Functions (PDFs) for both the parameters is reported in fig. 4.

As briefly discussed above, if the ROM was designed to be "monolithic", the training process would be very expensive, since the input parameter space is very large and the cost of each calculation would include both the cell calculation, which is the most expensive in this specific case, and the full-core simulation. On the contrary, if a two-step approach was employed, it could be possible to split the training phase in two stages, each benefiting of the specific features of each model. The choice of tailoring the ROM according to the peculiarities of the FOM follows a popular principle suggested by Vapnik, who gave fundamental contributions to the development of the Support Vector Machine method ([Cortes and Vapnik, 1995](#)): *when solving a problem of interest, do not solve a more general*

problem as an intermediate step.

Since, in this case, the cell calculation performed with Monte Carlo is the most computationally expensive, this strategy allows to simplify the input parameter space, decoupling the aspect of the perturbation location in the reactor, which is the most challenging aspect of the problem, and the intensity of the perturbations, which are more relevant for the generation of the few-group constants. Of course, the response of the full-core diffusion model would depend both on the kind of perturbation, through the few-group parameters, and on the location of the perturbations, but each set of perturbed cross sections and diffusion coefficients would be generated with a small effort by the cell ROM.

3. Non-intrusive model reduction for the cell calculations

Since both input parameters for the cell calculations follows a statistical rule, the Polynomial Chaos Expansion (PCE) is selected to generate the surrogate model for the generation of the few-group data. The PCE is a very popular technique for Uncertainty Quantification (UQ) and Sensitivity Analysis (SA) that was first suggested by Wiener (1938) and then extended and rigorously formalised by Xiu and Karniadakis (2002). After this seminal work, this technique has been extensively tested for several applications, proving its flexibility in a wide range of applications, including reactor physics (Williams, 2007; Gilli et al., 2013; Santanoceto et al., 2021).

The basic idea of the PCE approach, which belongs to the class of spectral methods (Canuto et al., 2007), is to express a stochastic model \mathcal{M} in terms of an expansion of orthogonal polynomials Ψ ,

$$\mathcal{M}(\vec{s}, t, \vec{p}) = \vec{y} = \sum_{k=0}^{\infty} \vec{a}_k(\vec{s}, t) \Psi_k(\vec{p}) \approx \sum_{k=0}^K \vec{a}_k(\vec{s}, t) \Psi_k(\vec{p}), \quad (3)$$

where \vec{s} is the vector of the state variables (e.g., spatial coordinates and energy), t is the time, \vec{p} is the vector of stochastic parameters (in this case, the pellet diameter and the water density) and a_k is the k -th expansion coefficient. The calculation of these projection coefficients, whose number depends both on the desired polynomial order K and on the number of independent random variable d , can be carried out following different strategies. In this work, they are calculated using the *pseudo-spectral approach*, which consists of approximating the following integral through a quadrature rule,

$$\vec{a}_k(\vec{s}, t) = \int_{-\infty}^{+\infty} d\vec{p} \mathcal{M}(\vec{s}, t, \vec{p}) \Psi_k(\vec{p}) w(\vec{p}) \approx \sum_{i=1}^I \mathcal{M}(\vec{s}, t, \vec{p}_i) \Psi_k(\vec{p}_i) w(\vec{p}_i). \quad (4)$$

When the input vector $\vec{p} \in \mathbb{R}^d$ is constituted by independent random variables, the multivariate polynomials are expressed as a product of univariate polynomials, so the number of their coefficients and, thus, the number of integral evaluations grows exponentially with the number of dimensions, due to the so-called *curse of dimensionality* (Bellman, 1957). To mitigate this issue, the parameter space is sampled using a combination of sparse grids Smolyak (1963), which offer some nice features.

The first one is that, for some specific quadrature rules, they allow to realise nested levels of points, i.e. the higher-order levels contain all the preceding lower-order ones. Despite these rules do not achieve the same accuracy as the Gauss-Legendre quadrature, they allow to perform an adaptive quadrature, which can be improved with more points when needed.

Figure 5 shows the sparse grids constructed to get the model evaluations needed for the quadrature. On the left, it is possible to see the first four levels of quadrature points, sampled with a combination of the Clenshaw-Curtis rule, appropriate for the uniform PDF related to the water density, and of the Genz-Keister rule, appropriate for the normal PDF associated with the pellet density. On the right, the 4th level of the grid is shown distinguishing between positive and negative weights associated with the quadrature rule. In both figures, the size of the points is proportional to the absolute value of their weights. Because of this representation, it is not easy to notice the nodes beyond the dashed lines, which are located at $\pm 3\sigma_{p_D}$, where σ_{p_D} indicates the standard deviation of the normal distribution

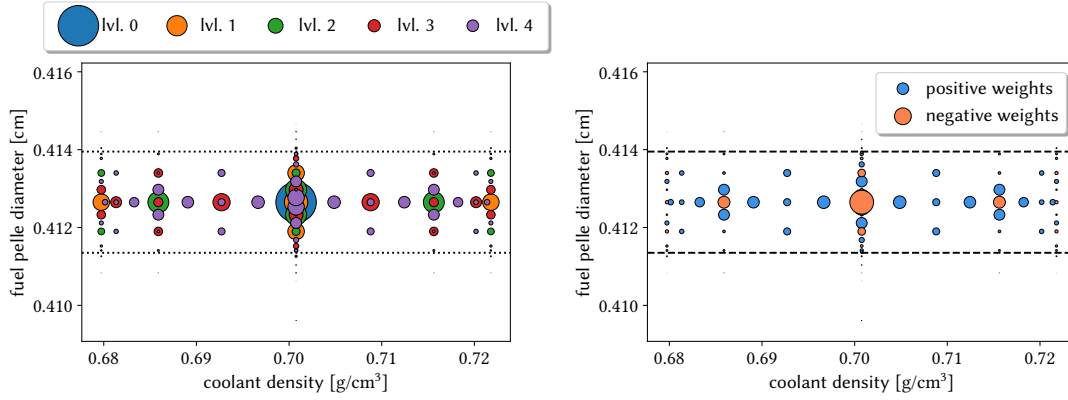


Figure 5: Quadrature weights used in the pseudo-spectral projection. The size of the dots is proportional to their weight, while the dashed, horizontal lines cover the 99.7% of the area, i.e. they are located at $\pm 3\sigma_{p_D}$, where σ_{p_D} is the standard deviation of the pellet diameter.




Figure 6: Quadrature points used to construct the pseudo-spectral PCE model and 160 quasi-random testing points generated using Sobol's rule. The dashed, horizontal lines cover the 99.7% of the area, i.e. they are located at $\pm 3\sigma_{p_D}$, where σ_{p_D} is the standard deviation of the pellet diameter.

for the pellet diameter. The low values of these weights is due to the fact that the region inside the dashed lines covers up to 99.7% of the data, meaning that the values falling outside this range are extremely unlikely.

To better appreciate all the points in the sparse grid, fig. 6 reports both these points and a set of 161 points generated with Sobol's rule for the validation of the PCE model. In this case, the size of the nodes is not proportional to the weights, to clearly show also the points falling outside the range $p_D \leq \pm 3\sigma_{p_D}$. Sobol's rule generates a low-discrepancy, quasi-random sequence of points that obey the statistical rules of each random variable. This sampling rule has been chosen because, although no validation point falls outside the region delimited horizontally by the two dashed lines, the points cover more uniformly the parameter space in the region with largest probability, allowing to check the performances of the PCE model more thoroughly.

3.1. Two-group constants generation

The two-group constants featuring the various perturbed assemblies of the reactor are generated with the Serpent 2 Monte Carlo code by considering a portion of the full-core model in fig. 2 with




Figure 7: Sketch of the 2D system used to generate the two-group constants with Serpent 2. From left to right, it is possible to see the fuel assemblies, the heavy stainless-steel reflector (grey), the water baffle (cyan) and the external boron layer (cerulean) used to simulate the vacuum boundary conditions.




Figure 8: Ratio between the flux energy spectrum computed with the simplified geometry and the spectrum evaluated with the full-core model.

reflective boundary conditions.

This model, sketched in Figure 7, allows to consider all the most relevant FAs and inactive regions featuring the system, reducing as much as possible the core dimension and, thus, its dominance ratio, favouring the convergence of the fission source calculation. Since Serpent does not allow to specify different boundary conditions for different surfaces, an additional layer of assemblies filled with water with a higher boron concentration has been added at the boundary, in such a way that the flux practically vanishes in that region, decoupling the system from its "mirrored" version on the right side. Figure 8 shows the ratio between the energy flux spectrum evaluated with the simplified and the one computed with the full-core models. Except for the borated water region, which is a region that features a low neutron importance, the spectrum of the simplified model is very close to the one evaluated with the full-core calculation, justifying the adoption of such an approach to generate the few-group constants.

A delicate aspect concerning the model reduction in a Monte Carlo simulation framework is that each output is affected by the statistical uncertainty. Since the PCE deals with uncertain input parameters but assumes that the output of the model are deterministic, each Monte Carlo calculation is carried out minimising the statistical uncertainty of the outputs to an acceptable level, so that their impact on the quality of the results is negligible for this specific application. Despite some techniques

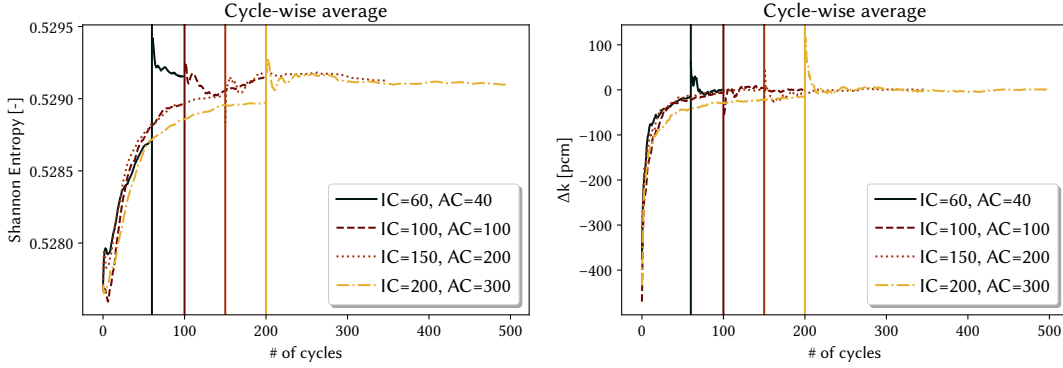


Figure 9: Convergence trend for the Shannon entropy.

may be used to overcome this issue, e.g., bootstrapping (Efron, 1981), their application are not justified considering the scope of the work, i.e. developing a computationally efficient meta-model using well-established codes in the LWR industry, which rely on deterministic cell codes to produce the group constants (Sanchez et al., 2010).

Figure 9 shows the convergence trends for the Shannon entropy and the k_{eff} as a function of the number of neutron cycles for different combinations of active and inactive cycles, which are delimited by the vertical lines in the graph, using 10^6 neutrons per cycle. The convergence behaviour of these parameters and the statistical error of the group constants of interest have been considered acceptable with 200 inactive cycles and 300 active cycles, so each simulation used to train the PCE model relied on these neutron population settings. As a further confirmation, it has been then verified that the output variations due to the physical parameter perturbations are not covered by the statistical noise in the results.

With these settings, each Serpent simulation requires, on average, about 3 hours with 16 CPUs (2x Intel Xeon Scalable Processors Gold 6130 2.10 GHz).

3.2. Model training and validation

The PCE model construction discussed in this section is carried out exploiting the open source Python library *chaospy* (Feinberg and Langtangen, 2015), which offers a complete pipeline for the ROM training. The model evaluations for the training phase are sampled using the nested sparse grid combining the Clenshaw-Curtis and Genz-Kleister rules and appearing in fig. 5, while about 160 model evaluations are performed using the low-discrepancy sequence generated with the Sobol's rule for model validation appearing in fig. 6. As it can be noticed in this figure, all the validation points cover quite uniformly the range $\pm 3\sigma_{pD}$, delimited by the black dashed lines, guaranteeing a uniform model validation as well.

Since the sparse grid for the training phase is nested, the PCE model training is progressive, increasing, when needed, the accuracy level of the grid while minimising the number of new model evaluations. Excluding the 0th order PCE, the first four levels of the sparse grid are employed to generate the PCE models ranging from 1st to 4th order. The four levels of the grid contains 5, 17, 47 and 113 points, respectively. Thank to the fact that the grid is nested, passing from one level to the succeeding one requires only 12, 30 and 66 new model evaluations, respectively.

Figure 10 shows the percentage Root-Mean Square Error (RMSE) computed between the PCE model evaluations and the validation set for different orders of the polynomials. The trend of the RMSE shows that the linear model is the most appropriate in reproducing the Serpent 2 data. (Bargagli Stoffi et al., 2022) so the first-order PCE model is selected as a surrogate for the cell calculations performed with Serpent. Figures 11 and 12 seem to further justify this choice. Despite a slight statistical noise affecting some of the parameters, the linear PCE appears as the most adequate in representing the full-order model response, especially the dependence to the pellet diameter random variable. As




Figure 10: RMSE between the PCE models evaluations and the validation set, sketched in fig. 6.

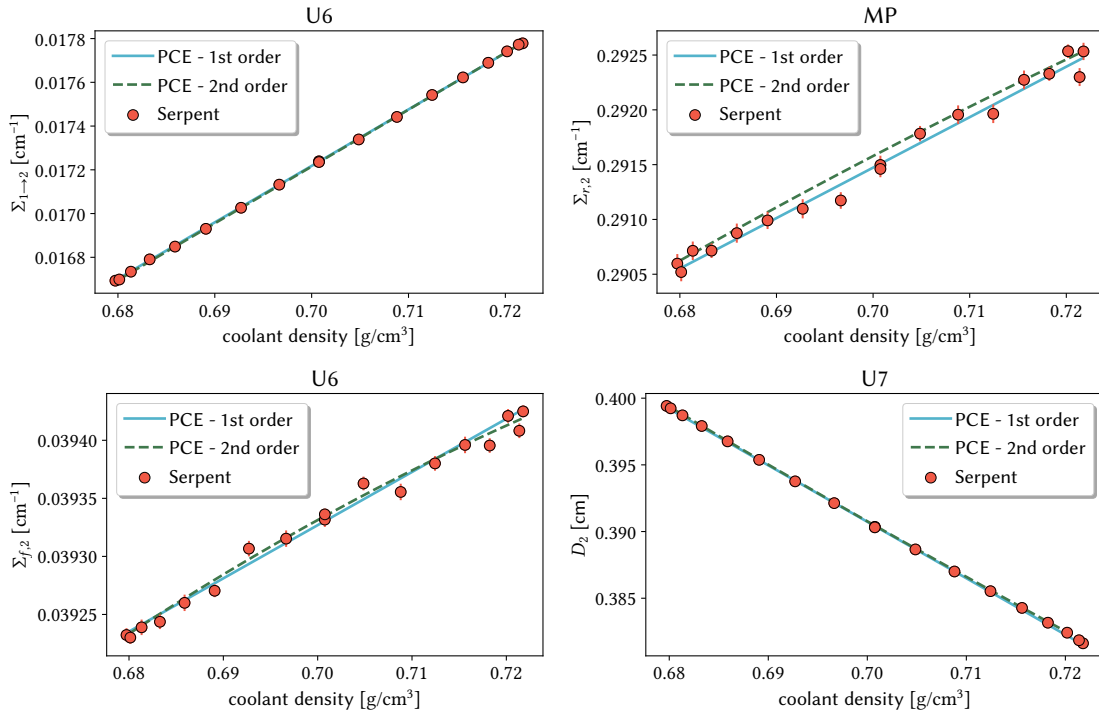


Figure 11: PCE fitting and full-order model evaluations as a function of the coolant density for different FAs. The pellet diameter is considered at its nominal value, 0.41265 cm. U6 indicates the 2.1% enriched UOX FAs, U7 indicates the 3.2% enriched UOX FAs and MP indicates the MOX FAs.




Figure 12: PCE fitting and full-order model evaluations as a function of the pellet diameter for different FAs. The coolant density is assumed to be at its nominal value, $0.700758 \text{ g cm}^{-3}$. U6 indicates the 2.1% enriched UOX FAs, U7 indicates the 3.2% enriched UOX FAs and MP indicates the MOX FAs.

it can be clearly seen in fig. 12, the quadratic PCE fails to represent adequately the validation data, especially those falling outside the range $\pm 2\sigma_{pD}$, represented by the dashed vertical lines. The dependence between the input parameters and the few-group constants is, in general, non-linear, therefore the agreement between the data and the first-order PCE model in this specific case is probably a consequence of the small input variations ($\approx 0.1\%$ for the pellet diameter, $\approx 3\%$ for the coolant density).

In light of these considerations, one could argue that the adoption of a PCE reduced order modelling framework could have been avoided using a much simpler model, e.g., a linear regression. However, it is important to remark that, in a general problem, one does not know *a priori* the general behaviour of a model with respect to a set of input parameters, especially when their number is large and they follow some statistical distributions. Despite this being a simplified problem, the objective of this work is to provide a computationally efficient framework that could be extended to more complex problems. Thus, the fact that the PCE allows to obtain a linear model should be interpreted as a positive feature of this approach. In case the input uncertainty is larger, a first-order PCE may not be adequate anymore. Hence, the PCE model could be naturally extended to a non-linear case, for example if the few-group data dependencies on the burn-up and the temperature were accounted for. Nevertheless, this extension will not be addressed in this paper, leaving it as a future development for this activity.

Concerning the validation phase, the quadrature nodes that are not employed in the model training, i.e. the evaluation of the polynomial coefficients through integration, are added to the validation set. The main advantage of this choice is the fact that, contrarily to the points generated with the Sobol's rule, these nodes cover also the most unlikely parameter region, improving the reliability of the evaluation of the modelling error.

The overall quality of the PCE model is acceptable for the intended application, as it can be noticed in fig. 13, which shows the % relative error between the PCE and the Serpent output for some selected

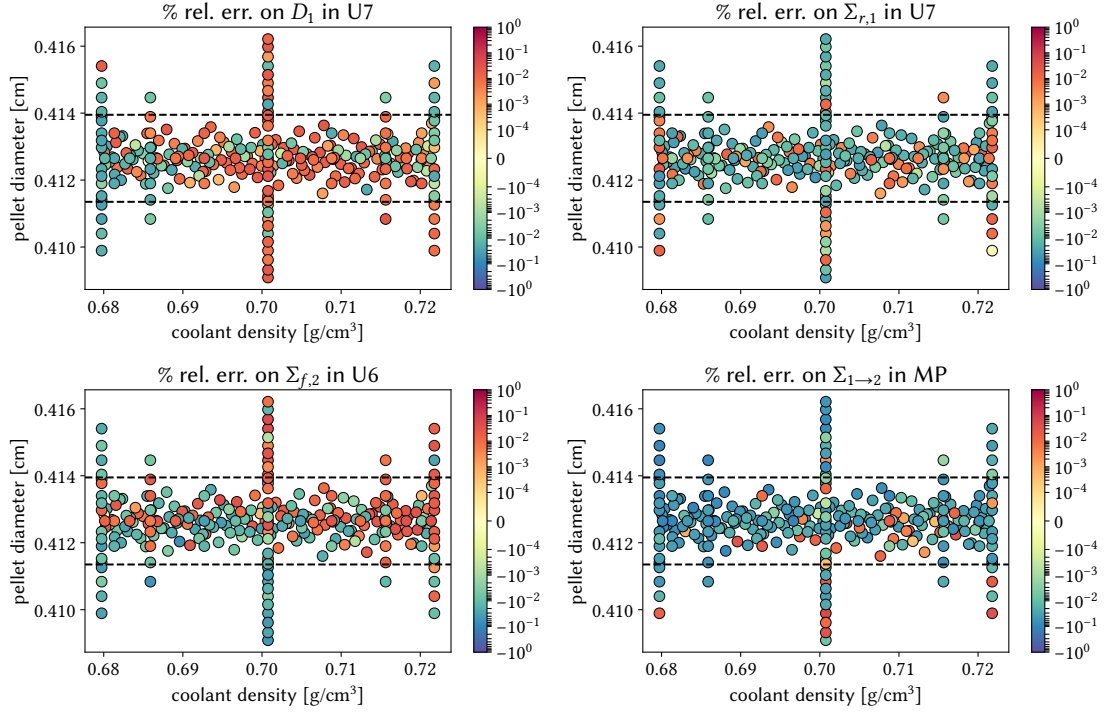


Figure 13: % relative error between PCE and Serpent output for some selected responses. U6 indicates the 2.1% enriched UOX FAs, U7 indicates the 3.2% enriched UOX FAs and MP indicates the MOX FAs.

responses of the validation simulations. Most of the points are featured by an error which is by far lower than 0.1 %, as displayed by the error distributions in fig. 14, where the dashed, vertical lines indicate the values $\pm 0.05\%$. In general, the points featured by the largest error are the ones falling outside the range $\pm 2\sigma_{pD}$, which are thus the less relevant one from a probabilistic point of view.

One of the advantages of the PCE approach is that its coefficients can be used to have a direct evaluation of the first-order Sobol's indices for each output response, defined as

$$S_i = \frac{\sigma_i^2}{\sigma^2[R]}, \quad (5)$$

where σ_i^2 is the variance obtained by varying only the i -th parameter and $\sigma^2[R]$ is the total variance of the output response R . In general, the indices are normalised. In this particular case, they obey the following relationship,

$$S_\rho + S_{pD} + S_{\rho,pD} = 1, \quad (6)$$

where S_ρ is the index associated with the coolant density, S_{pD} is related to the pellet diameter and $S_{\rho,pD}$ is the index accounting for the mutual interaction of the two parameters. These figures of merit, depicted in fig. 15, are fundamental in the framework of the Global Sensitivity Analysis (GSA) (Sudret, 2008), which estimates the relative importance of the various input parameters, and their correlation, in the determination of the output variance.

In this case, except for the diffusion coefficients, the most important uncertain parameter is the coolant density. In addition to its largest relative variation, this parameter is tightly related to the moderation, which strongly alters the flux energy spectrum at the cell level. The overall low sensitivity of the two-group data to the fuel pellet diameter can be explained by its small relative variation with respect to the coolant density. The fact that the diffusion coefficient is more sensitive to the fission cross sections can be explained by a competing phenomena involving the moderation and the spatial






Figure 14: Distribution of the % relative error between PCE and Serpent output for some selected responses. U6 indicates the 2.1% enriched UOX FAs, U7 indicates the 3.2% enriched UOX FAs and MP indicates the MOX FAs.



(a) Sobol's indices for the perturbed MOX FA. (b) Sobol's indices for the perturbed UOX FA enriched at 3.2%.



(c) Sobol's indices for the perturbed UOX FA enriched at 2.1%.

Figure 15: First-order Sobol's indexes for each output model response.

self-shielding featuring the fuel pins (Williams, 2011). For a given coolant density, perturbing the pellet diameter alters the coolant-to-fuel volume ratio inside the cell, which has a non-negligible impact on the mean free path of the particles inside the system and, thus, on their diffusion coefficient. Since the neutron penetration inside the fuel mostly occurs in the outer region, due to the self-shielding, slightly changing the pellet diameter does not significantly impact the fission rate.

Finally, it can be appreciated that the mutual interaction of the two input variables has a negligible impact on the variance of the output responses: the largest index associated with the input parameter correlation is about 0.03 % and, as such, it cannot be appreciated in the figure.

4. Non-intrusive model reduction for the full-core calculations

In this section, the main aspects concerning the training and the validation of the full-core NIROM are discussed, focusing on the peculiarities of the problem under examination.

4.1. A non-intrusive POD-RBF model

For the sake of generality, the two-group diffusion model expressed in eq. (2) will be indicated in the following as a generic model \mathcal{M} taking a P -dimensional vector as input parameter \vec{p} and responding with an m -dimensional output vector \vec{y} , referred to as *model snapshot*,

$$\vec{y} = \mathcal{M}(\vec{p}). \quad (7)$$

In this specific application, the snapshot is the spatial field representing the fission power spatial distribution. If one was interested in more spatial fields, e.g., the group-wise fluxes, this vector could be easily constructed by stacking each spatial field.

Since the model variables (e.g., space, time, energy...) and complexities and the numerical approximations are hidden in \mathcal{M} , this can be thought as a black-box communicating with the external world with only input and output datasets.

The data-driven model reduction, as mentioned, is composed by a set of offline and online stages, summarised in algorithm 1. In the offline phase, a certain sampling strategy is first selected to obtain the high-fidelity responses from some FOM executions. Then, the POD is used to get the POD coefficients, i.e. the the projection of the output datasets onto a reduced order space. These coefficients are then adopted as training data for the construction of the RBF network. Afterwards, the NIROM accuracy is checked against some validation points, i.e. high-fidelity snapshots obtained for parameter values which does not belong to the training set, verifying that the approximation error of the ROM is acceptable for the intended application. The final step of the offline stage involves the application of the bootstrap method for estimating the approximation error of the surrogate induced by the training samples choice. This procedure is performed by constructing a set of different ROMs, each trained sampling with replacement the original training set. The whole procedure is sketched in figs. 16 and 17. After the completion of the offline phase, the ROM can be used to deliver the approximated FOM responses on new parameter values with a fraction of the computational time required by the FOM.

In the following, a more detailed description of the the various steps appearing in algorithm 1 is given.

4.1.1. Model reduction, training and tuning

As briefly outlined above, the performances of non-intrusive surrogates are strongly dependent on the effectiveness of the parameter space sampling, independently on the reduced order modelling approach adopted.

After the selection of a proper sampling strategy, discussed more in detail in section 4.3, the FOM is executed and the snapshots are gathered in the so-called snapshot matrix. The dimensionality reduction is then pursued with the POD algorithm, sketched in fig. 18. Thanks to a singular value decomposition of the FOM snapshot matrix, the set of basis and coefficients is extracted. Since most of the data information is retained by the first modes, the POD basis can be effectively truncated with

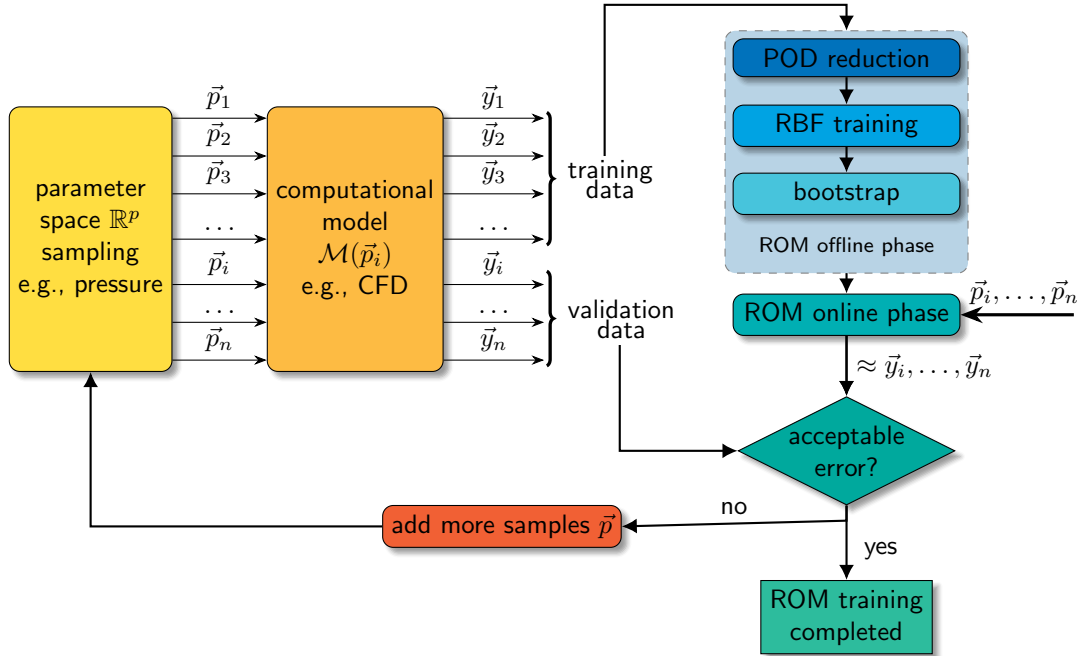
Algorithm 1: POD-RBF with uncertainty estimation

Offline procedures

1. define the p -dimensional parameter space \mathbb{R}^p ;
2. select a parameter space sampling strategy (i.e. sparse grids, random sampling...);
3. generate full-order model snapshots $\vec{y}_i \in \mathbb{R}^m$ for each parameter sample \vec{p}_i ;
4. divide the data into the training and the validation sets;
5. reduce dataset dimensionality, using POD (see Algorithm 2 in appendix Appendix A);
6. train the RBF net with the POD coefficients $\vec{a}_i \in \mathbb{R}^n$ (see Algorithm 3 in appendix Appendix A);
7. apply the bootstrap method (see Algorithm 5) to generate a set of ROMs;
8. compute the error distribution of the ROMs on the validation set;
9. validate the ROM on the validation set. If the average error between validation data and the set of bootstrapped ROMs is not acceptable, go back to step 3, adding more training points;

Online procedure

1. interpolate with RBF the POD coefficients over a new point (i.e. not used during training) \vec{p}_j in \mathbb{R}^p ;
 2. back-project the POD coefficients $\vec{a}_j \in \mathbb{R}^n$ to get the approximated snapshot in the original space, $\vec{y}_j \in \mathbb{R}^m$.
-


 Figure 16: *Offline phase procedure workflow.*

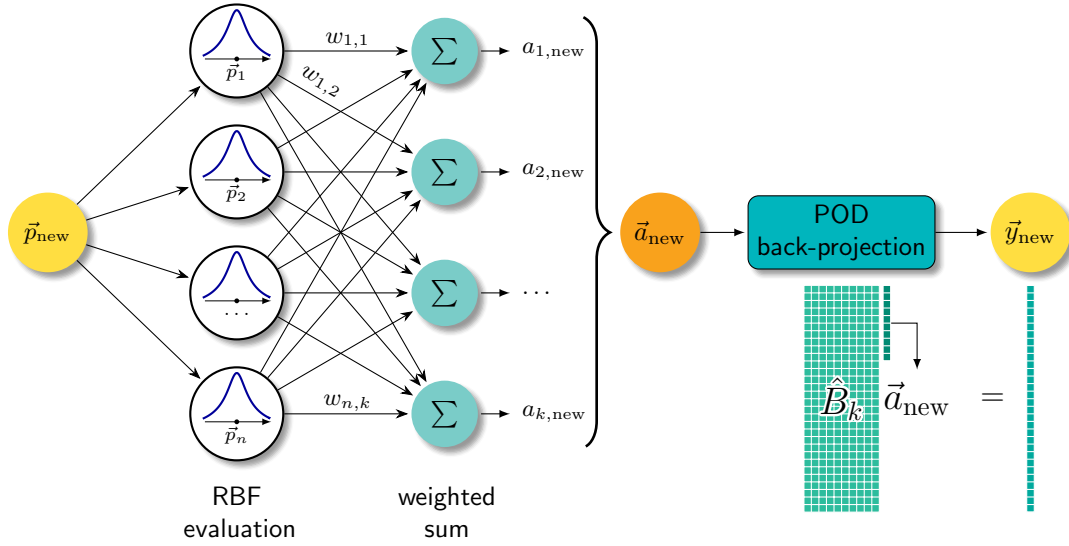


Figure 17: Online phase procedure workflow when a new parameter value \vec{p}_{new} is provided to the ROM. The set of weights $w_{i,j}$ is computed in the training phase, see algorithm 3.




Figure 18: POD-driven dimensionality reduction for the snapshot matrix containing the training model solutions.

a limited loss of information, due to the fact that the POD basis is optimal in a least-square sense (Volkwein, 2011).

The information loss induced by the POD truncation is usually quantified looking at the POD energy \mathcal{E}_k , i.e. the ratio of the sum of the first k squared singular values to the sum of all the squared singular values. The detailed POD algorithm is presented in appendix [Appendix A](#).

The following step consists of training the network of RBFs with the parameter-dependent POD coefficients. The term *radial basis function* indicates each function and distribution whose value depends on the L_2 distance between the RBF network centers, i.e. the training parameter values \vec{p}_i , and the collocation points, i.e. the new parameter \vec{p}_j . A common example of RBF is the Gaussian distribution. In this paper the kind of RBF chosen for the network is the inverse multi-quadrics formulated by Hardy (1971), because of its good numerical stability.

The RBF network requires the selection of a free parameter σ known as *hyperparameter*. A proper tuning of σ is of paramount importance, since it requires a compromise between the interpolation accuracy of the net and its numerical stability. Usually, this parameter is chosen by minimising the interpolation error on some test points, different from the training data. Since this procedure requires a large number of model evaluations, the Leave-One-Out-Cross-Validation (LOOCV) technique is employed. This is a general approach consisting in training the ROM with each of the N sets obtained by taking $N - 1$ training points and using the left one as a test point. The hyperparameter minimising the root-mean square error (RMSE) computed from the N trained models is selected as the best one.

The RBF training phase is summarised in algorithm 3, while the hyperparameter optimisation is presented in algorithm 4.

4.1.2. Model validation and error estimation by a bootstrap-based ensemble of ROMs

The last step of the offline phase has the objective of assessing the surrogate model accuracy with respect to the reference model solutions. To this aim, a set of snapshots on new parameter values, not included in the training set, is usually prepared for validation purposes. Clearly, a thorough validation process would require many FOM evaluations, jeopardising the computational performances of the surrogate. Thus, in the absence of an *a priori* estimation of the ROM error, it is a common practice to select only a few, significant new parameter values not too close to the training points in order to assess the model capability of delivering sufficiently accurate solutions (Rahman et al., 2018; Xiao et al., 2017).

The same strategy is adopted also in this case, but with the addition of a statistical sensitivity study concerning the training set as a complement of the validation phase. Specifically, the analysis is carried out using the bootstrap method, i.e. a non-parametric statistical method consisting in training a large number of ROMs using each time a different set of training samples, obtained by resampling with replacement the original training set. The resulting ensemble of surrogate models is then exploited to construct a distribution of output responses, which can be adopted to estimate the error distribution for the various validation points (Zio, 2006; Secchi et al., 2008; Zio et al., 2010; Pedroni et al., 2010; Marelli and Sudret, 2018). The whole procedure is presented in algorithm 5 in appendix [Appendix A](#) and sketched in fig. 19.

Applying the bootstrap method to the POD-RBF model requires some precautions. First, due to the RBF tuning, training the surrogate hundreds of times can be computationally unaffordable. To address this issue, the ROM bootstrapping may be parallelised exploiting the fact that each model reboot is independent. Also, a special care should be devoted to avoid that the same training data is repeated during the sampling with replacement, since in this case the RBF kernel would become singular. Thus, the resampling is carried out excluding the repetitions, meaning that each training set for the model reboot is always smaller than the original one.

4.2. High-fidelity model of the UAM benchmark

Concerning the high-fidelity full-core model of the UAM benchmark, the finite element code FreeFEM++ (Hecht, 2012) has been selected for the solution of the two-group diffusion equation.

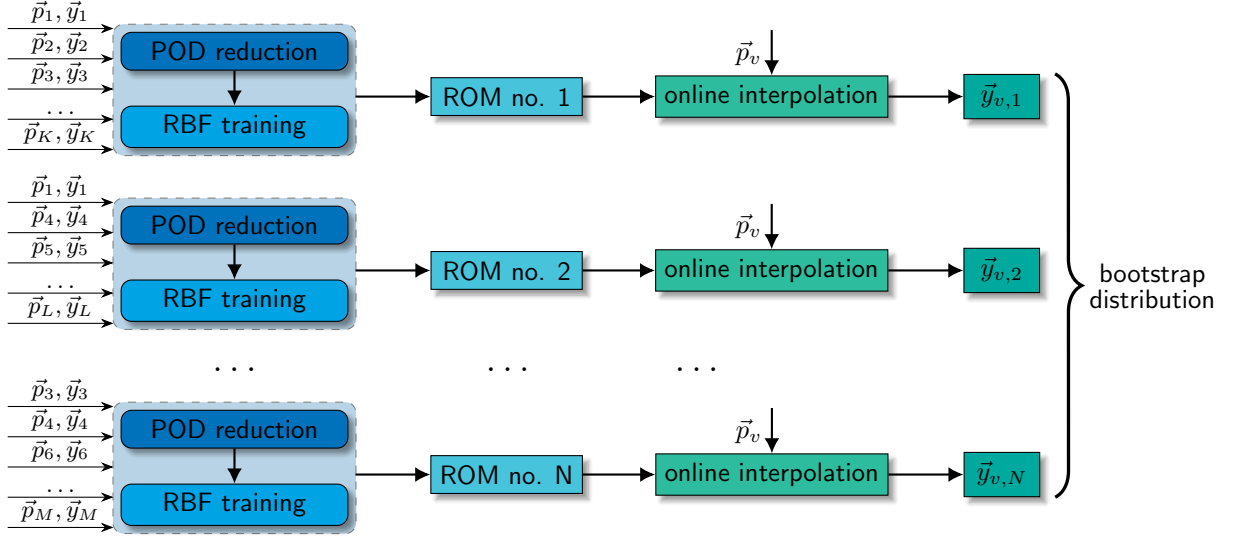


Figure 19: Sketch of the generation of the ROM distribution, and, consequently, of the output spatial fields via bootstrapping.

Despite finite volumes and coarse-mesh methods are usually preferred for the numerical approximation of the neutron diffusion operator, this tool has been selected for several reasons, namely:

- it is open-source;
- it is optimised to handle strongly heterogeneous geometries like the one of the UAM benchmark;
- both the computational mesh generation and the manipulation of the input and output fields are very fast and easy;
- it offers a wide number of wrappers to other computationally efficient libraries, like PETSc (Balay et al., 2020) and SLEPc (Hernandez et al., 2005), which are very efficient for the solution of eigenvalue problems.

On the right of fig. 20, the UAM benchmark sketch with the indication of the numbers used to indicate the various FAs is reported, while on the left the adapted computational mesh employed for the generation of the training and validation solutions is provided. The mesh was generated adding more elements in the central regions of the reactor and using less elements in the external FAs, featured by smoother gradients. All the calculations presented in this work are carried out using piece-wise linear discontinuous finite elements. The number of elements for this mesh was selected according to the results of a grid independence study, reported in fig. 21.

Using this mesh, composed by 632 430 elements, each finite element simulation requires about 3-5 minutes, according to the machine architecture, using 24 CPUs to parallelise the solution of the eigenvalue problem.

4.3. Definition of ad hoc sampling strategies

The selection of an adequate sampling strategy is important for training an efficient and robust meta-model. To the authors' knowledge, a specific sampling strategy adequate to handle the special feature of this application, i.e. the spatial dependence of the perturbations in a reactor core, does not exist in the literature. Therefore, since extending the usual sampling strategies for an input parameter

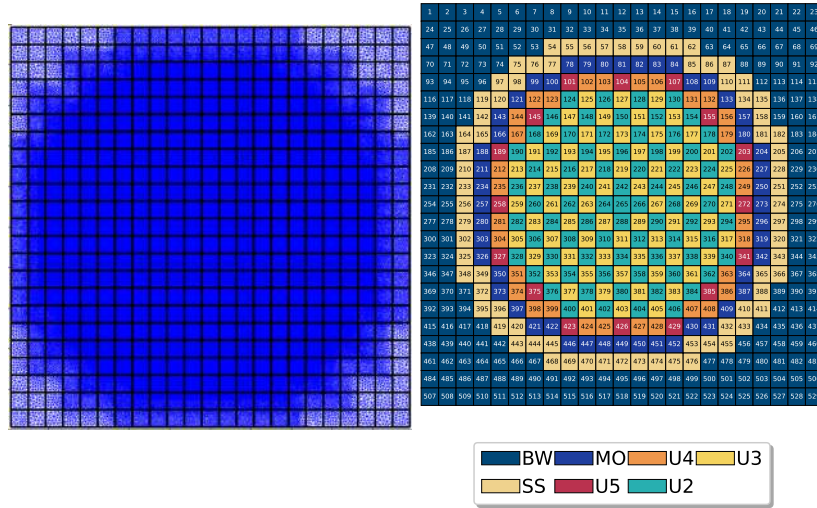


Figure 20: Finite element mesh adopted for the full-core diffusion calculation (left) and sketch of the UAM core with the FA numbers (right).

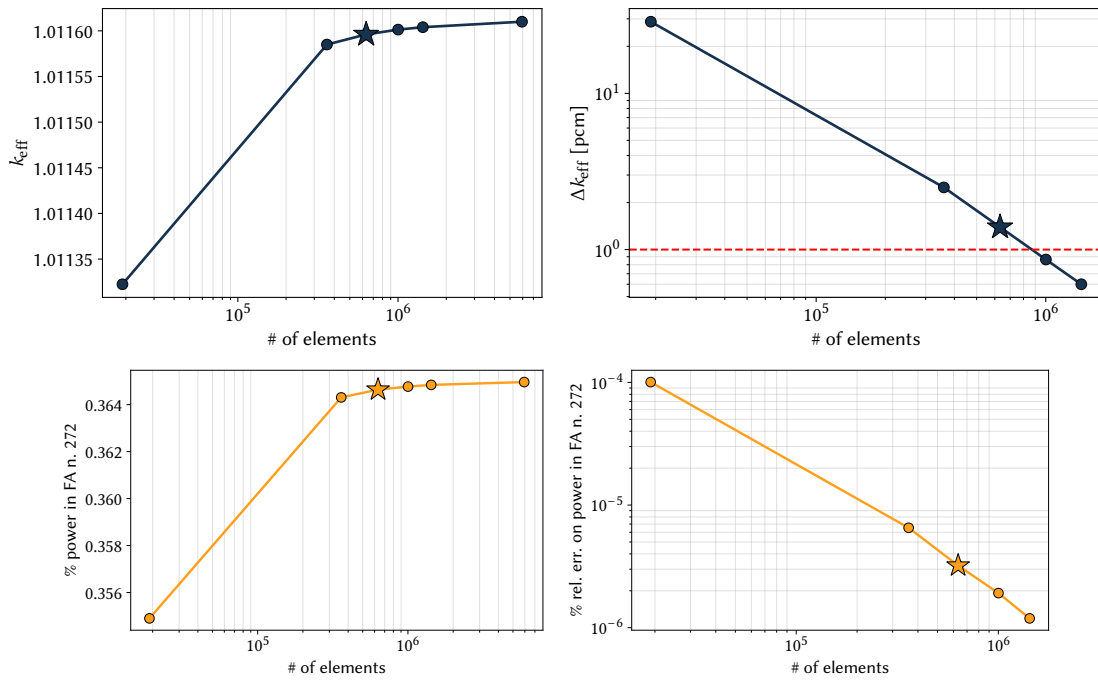


Figure 21: On the left, spatial convergence trend for k_{eff} (top) and the total fission power in the FA number 272 (bottom). On the right, the relative error with respect to a reference calculation with $\approx 5.9 \times 10^6$ elements for k_{eff} (top) and the total fission power in the FA number 272 (bottom).

space made of scalar quantities to this specific case is not straightforward, an *ad hoc* procedure is proposed in the following, trying to exploit the physical features of the problem.

The first aspect that is exploited in the definition of the sampling procedure is the fact that, from the point of view of the power production by fission, not all the neutrons diffusing in the reactor are equal: some of them are more important than others in contributing to the fission chain reaction. Therefore, it would be advisable to introduce the perturbations in the FA to which the fission source is more sensitive.

The concept of neutron importance Ψ , first introduced by [Usachev \(1956\)](#), is tightly related to the adjoint transport equation and its solution ϕ^\dagger , which has been interpreted historically as a measure of the importance of one single neutron in contributing to the response of a certain detector or, for a critical system, to the fission reaction. The interpretation $\Psi = \phi^\dagger$ is the most popular one, and has been successfully used, for instance, as a weighting function to conceive an adjoint-based intrusive ROM ([Lorenzi, 2018](#)). In addition to this interpretation, it is possible to define also other importance functions ([Tal et al., 2019](#)). For example, in the framework of the Generalised Contribution Theory ([Williams, 1991](#)), where the importance function can be defined as the product between the adjoint and the direct fluxes, namely

$$\Psi = \phi\phi^\dagger, \quad (8)$$

or as the product between the adjoint flux and the neutron density, namely

$$\Psi = \frac{\phi}{\mathbf{v}}\phi^\dagger. \quad (9)$$

Figure 22 reports the assembly-wise spatial distributions in each energy group for the direct and adjoint fluxes and for the neutron density, computed ensuring criticality through a proper normalisation of the fission operator by k_{eff} . Since the factor $1/\mathbf{v}_g$ ($g = 1, 2$) is spatially uniform, the spatial distribution of the direct flux and of the neutron density does not change, but, since $\mathbf{v}_1 \gg \mathbf{v}_2$, the ratio between the thermal and the fast densities is reversed: the flux, i.e. the neutron track length per unit volume, is larger in the fast group, but the density, i.e. the number of particles per unit volume, is larger in the thermal group. Since the thermal adjoint in figs. 22e and 22f is larger than the fast one, the resulting importance function would be very different according to its definition.

The three assembly-wise importance function definitions can be appreciated in fig. 23. In all cases, the most important FAs are the peripheral ones featured by 2.1% enriched UOX, without Gd, where the flux achieves its peak. However, the global importance of the other FAs strongly depends on the kind of definition employed. The one coinciding with the adjoint, obtained by summing the group-wise adjoint fluxes, is more spatially uniform than the other ones. Thus, according to this definition, the neutrons generated in the inner FAs have roughly the same importance in sustaining the fission reaction. If the adjoint is weighted with the direct flux, the spatial heterogeneity effects increases, due to the sensitivity of the thermal flux to the FA composition. This effect is exacerbated when the adjoint is weighted with the neutron density, due to the largest weight associated with the thermal density.

Since the last formulation indicates the global importance content in the reactor, this definition is considered to define the probability density function for sampling the location of the perturbed FAs. The assembly-wise distribution of the resulting PDF is reported in 24.

The other physical feature that could be exploited in the training process is the symmetry of the core, which could be used to increase the number of training data at basically no cost by means of proper rotations of the solution. In this way, each FreeFEM++ calculation provides four solutions, one for each multiple of the rotation angle $\theta = \frac{\pi}{2} + k\frac{\pi}{2}$, with $k = 0, 1, 2, 3$. An example of this rotation procedure is visible in fig. 25. In order to emphasise the visualisation of the perturbation clusters, each one is represented by a single colour, although the FAs are featured by different few-group constants according to their type.

To induce the perturbation clusters visible in fig. 25, only the initial FA is sampled according to the importance PDF. Then, an FA among the north, east, south and west ones is selected randomly, and




Figure 22

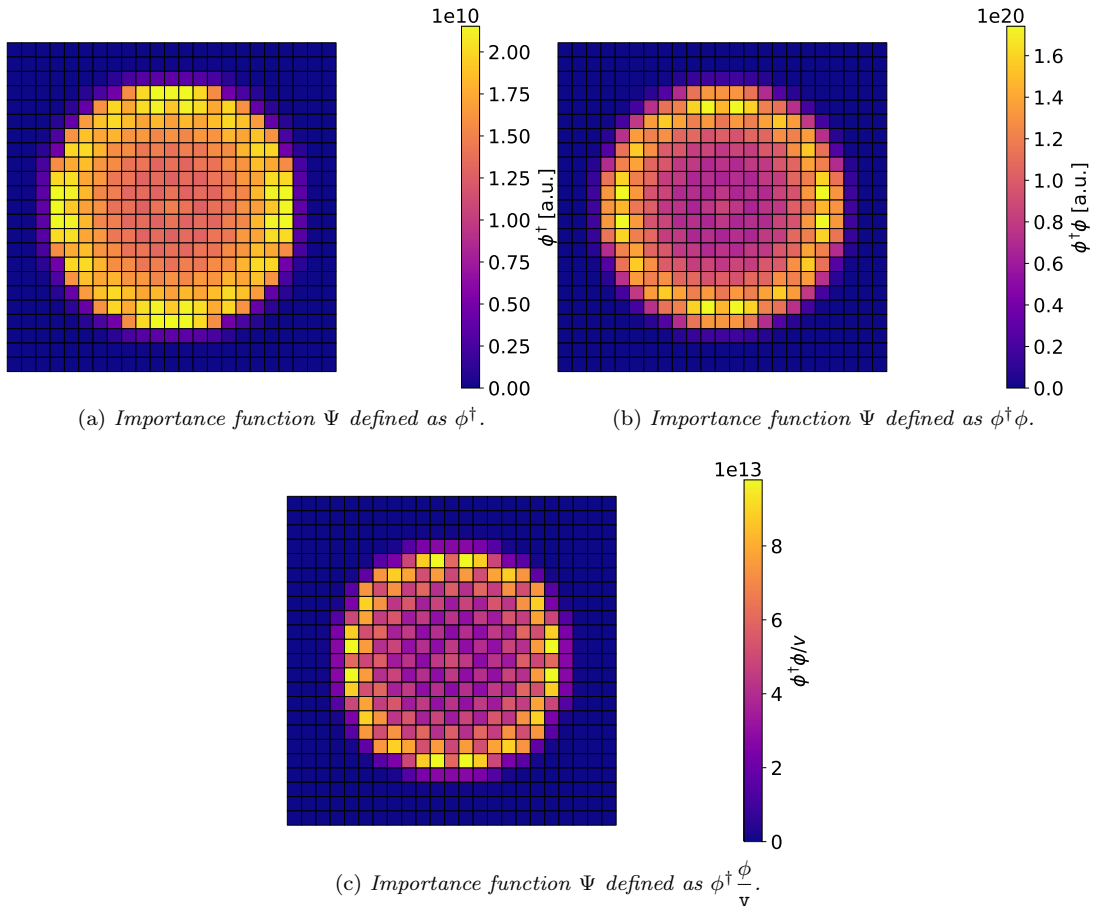


Figure 23: Possible definitions of importance functions adopted during the importance sampling process.




Figure 24: Probability density function to sample the location of the perturbed FA.




Figure 25: Rotated perturbed configurations. The starting configuration (upper left) is progressively rotated with multiple of $\pi/2$ counter-clockwise.

the procedure is iterated until the maximum number of perturbed FAs per each perturbation cluster is achieved. Hence, a kind of perturbation random walk is defined. Figure 26 provides a support to visualise the sampling strategy used to define the block of perturbed FAs. First, the number of perturbed FA is sampled from a uniform distribution defined among 6 and 10, included. Then, the initial FA, 103 in the figure, is sampled according to the PDF in fig. 24. Afterwards, the FA 126 is sampled among the equally probable FAs surrounding the FA 103. The red arrow indicate the direction of the disturbance random walk, while the white arrows indicate the other possible, non-sampled directions. The sampling process goes on until the number of perturbed FAs is achieved, verifying that each new configuration is not contained in the previous samples, to avoid singularities in the RBF kernel.

Each perturbation cluster is featured by the same perturbation magnitudes, which are sampled according to the joint PDF obtained by combining the PDFs showed in fig. 4. Some perturbation arrangements obtained using this algorithm may appear very unlikely in a realistic system, like the ones in fig. 27. However, since one of the objectives of the paper is to study the performances of the NIROM for spatially-dependent perturbations, they are all considered acceptable. The same assumption is made for the fuel pellet diameter uncertainties, in order to reduce the number of possible combinations. For each perturbation, the corresponding set of two-group constants are generated using the PCE model.

In spite of this simplifying assumption, the number of possible scenarios is still dramatic. First, the PDF for the extraction of the first FA is strictly positive for each FA, meaning that potentially every FA could be selected as the origin of the perturbation. Then, even though only some surrounding FAs can be selected to define the perturbation cluster, the total number of perturbed FAs can vary in the range 6-10, increasing the number of potential scenarios. As a further complication, more than one perturbation clusters may appear, each one featured by different values of coolant density and pellet diameter, which are random parameters themselves.




Figure 26: *Example of a perturbation random walk. The white arrows indicate the FAs that could be extracted, while the red arrows indicate the selected FAs. The red X indicates the starting FA.*

Due to the huge size of the parameter space, two alternative strategies are proposed and tested in the following, in combination with the algorithm described above for the generation of the perturbation random walks:

1. a *full random sampling*, which consists of sampling both the spatial features of the perturbation and the perturbation magnitudes;
2. a *hybrid sampling*, which consists of using a random sampling for the spatial features of the perturbations and a deterministic rule, i.e. the Smolyak sparse grid, for the intensity of the perturbations.

The advantages and disadvantages of these approaches will be discussed in more detail in the following; instead, it is important to observe here their fundamental features.

The brute-force random sampling is simpler and more *naive* than the hybrid one, but it can be conveniently used to split the dataset into a training and a validation sets. Moreover, combined with importance sampling, this approach could be still effective in providing a uniform mapping of the parameter space. Thus, the technique is here used to generate a set of around 1125 scenarios (4500 effective scenarios, considering the solution rotations), each featured by a random number of batches, a random number of FAs per batch and a random perturbation intensity, i.e. the values of coolant density and pellet diameter, per batch.

The hybrid sampling is more involved, but it allows a better control of the random physical parameters, enabling also to perform an adaptive training with the nested sparse grids. This feature allows to progressively account for the super-position of coolant density and pellet diameter perturbations. By inspection of fig. 5 it can be noticed that the lower-order levels of the grid are constructed by varying only one parameter and leaving the other unperturbed, meaning that, to have simultaneous, super-imposed perturbations, the minimum level of the grid should be larger than 2. Of course, for a higher number of input parameters this threshold would increase. Notwithstanding its nice features, some precautions are needed to reduce the effects of the curse of dimensionality: since each perturbed batch is featured by an independent perturbation intensity, the number of configurations would scale exponentially with the number of deterministic points used to map the perturbation intensity. For example, if the second level of the Smolyak sparse grid was used to map the pellet diameter and the coolant density, 17 points would be generated, as in fig. 5. Since each perturbation batch would be featured by 17 values of the physical parameters, the number of configurations would scale as 17^b , where b is the number of batches considered.

To avoid this behaviour, it is possible to exploit the polynomial scaling of the Smolyak algorithm, considering a higher-dimensional input space composed by $2b$ parameters, where 2 refers to the density and to the diameter inputs. In this way, only the optimal subset of the possible 17^b combinations would be selected by the Smolyak algorithm. Table 1 provides the number of samples obtained by following the batch-wise and the global sampling. As expected, the batch-wise sampling strategy is

Table 1: *Example of sampling scaling when the second level of the Smolyak sparse grid is defined for each batch independently or considering as input parameters all the density-diameter couples per each batch.*

n. batches	batch-wise sampling	global sampling
1	17	17
2	289	49
3	4913	97
4	83521	161




Figure 27: *Some configurations for a single perturbation in the coolant density distribution in the active FAs.*

not affordable, while the other sampling strategy allows to achieve a massive reduction in the number of samples. Each Smolyak level contains the 0th level, i.e. the reference, unperturbed input, thus excluding these points, the actual number of points is expected to be even lower. Even though, the number of training simulations required to map the core with four batches would be still very large, as 161 simulations are required for a single perturbation configuration. Because of this reason, the hybrid sampling is adopted in the following to analyse a simpler problem, featured by only one batch of perturbed FAs in each configuration, leaving the number of FAs and their location as free, independent parameters. For this problem, 32 independent spatial arrangements are considered, for a total of 1024 cases, considering the rotations and the 16 Smolyak points required per each arrangement.

4.4. Evaluation of the distance between the input parameters

Figure 27 and fig. 28 show some examples of physical perturbations affecting the core, which is supposed to have a uniform density spatial distribution, being at the uniform temperature of 570 K, i.e. at hot zero power.

Since the input parameters (coolant density and pellet diameter) are not scalar quantities but assembly-wise distributions poses the question of taking the spatial aspect into account during the




Figure 28: *Some configurations for multiple perturbations in the coolant density distribution in the active FAs.*

RBF weights assignment. It appears intuitive that the perturbations in figs. 27a and 28a will induce a neutron flux and fission power distributions quite close to the one induced by the perturbations in figs. 27b and 28b, respectively, being the perturbation intensity in each batch of FAs equal. On the contrary, it is reasonable that the perturbation in figs. 27c and 28c would provide substantially different flux and power profiles, due to the location of the L-shape perturbation, which is placed in the northern part of the core instead of the southern one.

However, the RBF weights resulting from the training phase described in Algorithm 3 of appendix Appendix A do not follow this physical feature when the usual L_2 norm is adopted to evaluate the distance $\|\vec{p}_i - \vec{p}_j\|_2$ between the various input parameter configurations. In particular, the Euclidean distance between fig. 27a and fig. 27c turns out to be equal to the one between fig. 27a and fig. 27b, i.e. 0.0442. Despite it does not take into account the physical nature of the problem, this result is fully consistent with the mathematical expression of the L_2 norm,

$$\|\vec{p}_i - \vec{p}_j\|_2 = \sqrt{\sum_{k=0}^K (\vec{p}_{ik} - \vec{p}_{jk})^2}. \quad (10)$$

Since the only non-zero differences would be found for the locations of the perturbed FAs, and both the perturbation intensity and the number of perturbed FAs are the same, also the resulting distance between the two input parameters would be the same. In case the perturbation intensity and the number of perturbed FAs were different, which is probable considering the dimensions of the input parameter space, the distance would certainly change, but it would not still be very sensitive to the spatial distribution of the input. For instance, the Euclidean distance between fig. 28a and fig. 28c would be around 0.0254, while the one between fig. 28a and fig. 28b would be around 0.0442, which are quite similar despite the spatial distributions are noticeably different.

Among the different metrics available in the literature, the so-called IMage Euclidean distance

(IMED) (Wang et al., 2005) seems the most adequate to tackle this issue, due to its robustness to small perturbations and to its simplicity. Contrarily to the traditional L_2 norm, this metric accounts for the relative position of the different pixels composing an image thanks to a proper weighting matrix \hat{G} , which has to be positive definite,

$$\|\vec{p}_i - \vec{p}_j\|_{IMED} = \sqrt{(\vec{p}_i - \vec{p}_j)^\top \hat{G} (\vec{p}_i - \vec{p}_j)}. \quad (11)$$

Given an image defined in $\mathbb{R}^{N \times M}$, the g_{nm} element of the weighting matrix \hat{G} is defined as

$$g_{nm} = h(\|P_n - P_m\|), \quad (12)$$

where P_n and P_m are the n -th and m -th coordinates of the pixel composing the image and h is a continuous function that should be monotonically decreasing when its argument $\|P_n - P_m\|$ increases. As in the reference paper, in this work the h function is chosen to be a Gaussian distribution, namely

$$g_{nm} = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\|P_n - P_m\|}{2\sigma^2}\right). \quad (13)$$

Although the IMED has been developed for image recognition algorithms, it can be easily extended to other multi-dimensional objects for which the relative position in space is important, considering the intensity of the field instead of the pixel colour. In this specific application, each input parameter can be cast as an object belonging to $\mathbb{R}^{X_N \times Y_M}$, where X_N and Y_M are the number of spatial points along the x and y coordinates, respectively. The application of the IMED to one-dimensional fields is trivial, since it would be sufficient to consider $Y_M = 1$, while the application to three dimensional fields, e.g., $\vec{p}(x, y, z)$, which is interesting for 3D full-core analyses, would require an extension of the IMED algorithm to three-dimensional images. This generalisation is not addressed in this work, but it is left as a future development to extend the applicability of this ROM framework to 3D full-core systems.

The IMED distance between the configuration in fig. 27a and the one in fig. 27c yields 0.027, while the one between fig. 27a and fig. 27b amounts to 0.009, proving the higher sensitivity to the perturbation spatial pattern. The better performance of this metric can be verified also comparing the other perturbations depicted in fig. 28. The distance between fig. 28a and fig. 28c yields 0.0267, while the one between fig. 28a and fig. 28b yields, as expected, a lower value, 0.018.

The main disadvantage of the IMED algorithm is its higher computational cost compared to the traditional L_2 norm, since the distance evaluation requires to perform a vector-matrix-vector product, and the RBF approach requires to compute the distance between each pair of training data. However, exploiting the linearity of the algorithm, it is possible to apply a linear transformation to each input that allows to reduce the number of matrix-vector products. Exploiting a factorisation of the \hat{G} matrix (Wang et al., 2005), the IMED turns out to be the classical Euclidean between two transformed variables \vec{r}_i and \vec{r}_j , i.e.

$$\begin{aligned} \|\vec{p}_i - \vec{p}_j\|_{IMED} &= \sqrt{(\vec{p}_i - \vec{p}_j)^\top \hat{G} (\vec{p}_i - \vec{p}_j)} = \sqrt{(\vec{p}_i - \vec{p}_j)^\top \hat{H}^\top \hat{H} (\vec{p}_i - \vec{p}_j)} = \\ &= \sqrt{(\vec{r}_i - \vec{r}_j)^\top (\vec{r}_i - \vec{r}_j)}, \end{aligned} \quad (14)$$

where \hat{H} can be obtained either from SVD or from the eigenvalue decomposition of the weighting matrix \hat{G} .

Thanks to this property, the IMED can be naturally embedded non-intrusively in the ROM framework, just applying the linear transformation to each input before passing them to the POD-RBF algorithm. Figure 29 shows the smoothing effect obtained by applying the \hat{H} matrix to the input coolant density field.

If the input parameters considered to train the ROM were the actual inputs for the diffusion model, i.e. the few-group constants, the smoothing effect would be more relevant for the thermal properties,




Figure 29: *Original (left) and smoothed (right) perturbed coolant density distributions.*

featured by a stronger spatial variation, as visible in fig. 30. Concerning this case, it is possible to notice the appearance of an unphysical "stripe" in the thermal diffusion coefficient distribution. These values are not physical perturbations, but are a consequence of the fact that, to apply the IMED to the few-group data, each group-wise data is stacked in a vector of length $n_{data}N_xN_y$, where n_{data} is the number of data and N_x and N_y are the number of assemblies in the x and y directions. Since each "image" is stacked one below the other to form a unique "image" for the IMED application, the linear transform \hat{H} smooths the passage from one figure to the other. This aspect does not influence at all the distance between the parameters in this application, since each perturbation is applied only in the active region, where this effect does not occur. In view of other applications where this issue may be relevant, it could be possible to avoid it adding some dummy values to each data, which could be interpreted as adding some blank space between each figure.

Another challenging aspect of this application is the large dimension of the input data. In this case, which is featured by a low number of energy groups, $n_{data} = 7$, i.e. the two-group diffusion coefficients, production cross sections, removal cross sections and the slowing down cross section, while N_x and N_y are both equal to 23, for a total of 3703 values. Since the computational cost of the distance matrix evaluation is $\mathcal{O}(n^2)$, where n is the number of training snapshots, and this operation has to be performed several times in the RBF training process (see algorithm 3), the total cost would be prohibitive.

This issue can be effectively mitigated using the POD algorithm to reduce the input dimensionality, exploiting its redundancies. Since the POD is linear, the overall IMED does not change if the reduction is applied to the smoothed input data, making the combination of these two methods straightforward. In fig. 31 it is possible to appreciate the first three POD modes of the thermal fission production cross section $\nu\Sigma_{f,2}$. Concerning the criterion adopted to select the number of POD modes, a very strict truncation error (10^{-12}) is considered, in order to ensure a sufficiently accurate data representation in the reduced order space and, thus, a precise evaluation of the distance between the various input perturbations.

4.5. Model training and validation

In this section, the full-core model results are presented and discussed, considering as output response the assembly-wise power distribution, which is the physical parameter adopted to assess the core conformity in the presence of localised perturbations.

Figure 32 presents the first six modes obtained by applying the POD reduction to the power training snapshots. Contrarily to the POD application to the input data, in this case it is possible to use a much looser tolerance, around $5 \cdot 10^{-6}$, which guarantees an energy of the POD basis larger than 99.999%. Interestingly, the spatial shapes of these modes are very similar to the ones of the higher k -eigenmodes, probably due to the small perturbation magnitudes.

The POD reduction is very important to minimise the computational cost associated with the model training, as one RBF net is trained per each POD coefficient (see algorithm 3), including the RBF




Figure 30: *Original (left) and smoothed (right) fast (top) and thermal (bottom) diffusion coefficient.*

shape parameter optimisation with the LOOCV technique (see algorithm 4). Also in the presence of a significant dimensionality reduction, the computational time of the RBF evaluation strongly depends on the number of training snapshots n , as the RBF weights have to be determined by inverting the distance matrix $\hat{D} \in \mathbb{R}^{n \times n}$. This is a well-known limit of non-parametric techniques like RBF-SVM (Russell, 2010), together with the possible over-fitting behaviour. Since the datasets used to map the input parameter space of the problem are very large, in the following the model sensitivity to the number of training samples will be assessed taking into account both the accuracy of the model and its computational efficiency.

Before moving on with this evaluation, it is important to discuss the input-output association logic. Due to the two-step structure of the high-fidelity model depicted in fig. 1 and to the non-intrusive nature of the ROM, one could be led to assume that the best labelling strategy for the supervised learning model would be to assign the spatial distributions of the coolant density and of the pellet diameter as input to the ROM. This choice could appear legitimate, since, from a physical point of view, the reactor responds with a certain power distribution to a certain perturbation arrangement. However, from the point of view of the model, the same relative density variation affects in a different way each type of FAs in a certain perturbation batch, meaning that the full-core diffusion model would "see" different cross sections and diffusion coefficients in each perturbation batch. This has a non-negligible impact on the effective distance between the training configurations, as visible in fig. 33. The two perturbation batches shown are very similar from a topological point of view, but involve different families of FAs. Therefore, if the ROM input considered were the density and the diameter distributions, the two configurations would turn out to be much closer than they actually are if the two-group data are considered. Figure 34 reports the PDF and the Cumulative Distribution Function (CDF) of the percentage relative error on the assembly-wise power \hat{Q}_i between the FOM and the ROM

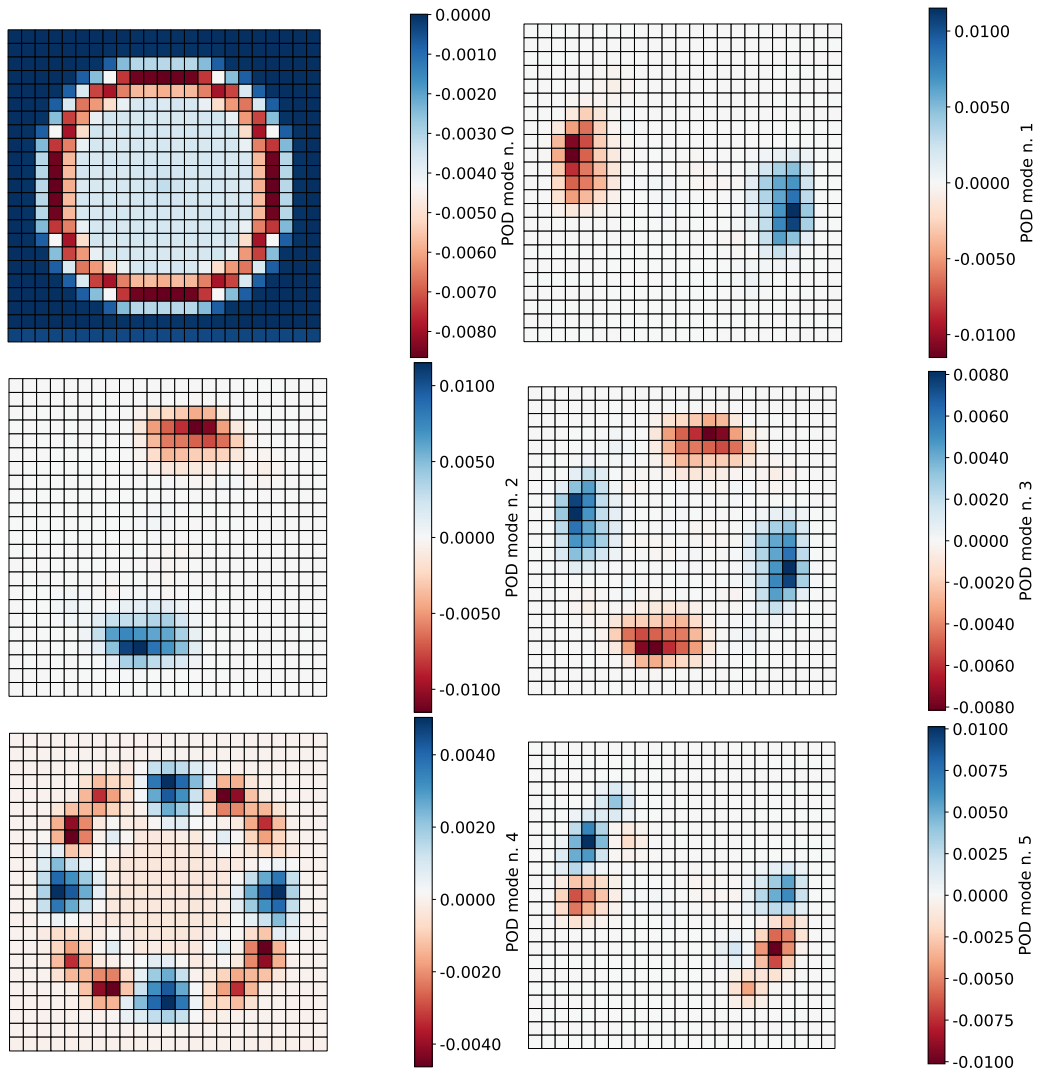


Figure 31: *POD modes for thermal fission production cross section, with increasing spatial frequency moving from top to bottom and from left to right.*




Figure 32: *POD modes for fission power distribution.*

99	100	101	102	103	104	105	106	107	108	109
122	123	124	125	126	127	128	129	130	131	132
145	146	147	148	149	150	151	152	153	154	155
168	169	170	171	172	173	174	175	176	177	178
191	192	193	194	195	196	197	198	199	200	201
214	215	216	217	218	219	220	221	222	223	224
237	238	239	240	241	242	243	244	245	246	247
260	261	262	263	264	265	266	267	268	269	270

Figure 33: *Example of perturbation arrangements which are spatially similar but involve different kinds of FAs.*

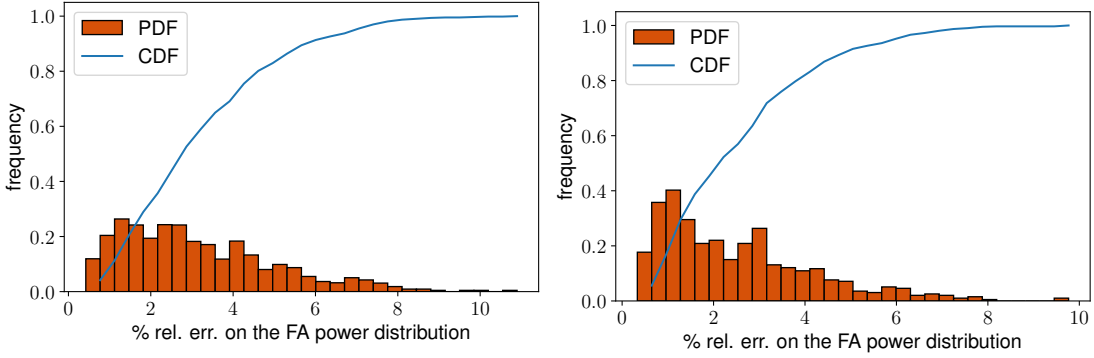


Figure 34: Percentage relative error distribution on the validation samples for the full-random sampling case considering four perturbation batches in the core. On the left, the model is trained considering the coolant density and the pellet diameter as input, on the right the model is trained using the actual few-group data.

Table 2: Computational performances of the ROM trained with different dataset sizes.

n. snapshots	ROM training time [s]	ROM execution time [s]	Normalised memory consumption [-]
500	90	0.035	1
1000	369	0.136	4
1500	860	0.30	9

for 2000 validation cases generated with the full-random sampling approach, namely

$$\epsilon_Q^{\%} = \sum_{i=1}^{N_{FAs}} \left\| \frac{\dot{Q}_{i,FOM} - \dot{Q}_{i,ROM}}{\dot{Q}_{i,FOM}} \right\|_2 100. \quad (15)$$

The assembly-wise power is normalised so that the total power is equal to 1 W, consistently with the Cold Zero Power conditions assumed for the reactor. On the left, the physical perturbations are used as input parameters for the model, while, on the right, the input parameters are the actual two-group constants. At a first glance, it may seem that there is no striking difference between the two cases, but looking at the frequency axis it is easy to notice how the right graph is featured by a peak of around 1% and by a strong reduction of the frequency for the error around 6-10 %. Therefore, these results confirm that, for a proper evaluation of the input parameter distance, which is fundamental for assigning the correct weights to the RBF, it is necessary to assign the actual input data of the FOM to each snapshot.

Figure 35 reports the PDF and CDF of the percentage relative error computed as in eq. (15) for 3000 validation cases generated with the full-random sampling approach using 500, 1000 and 1500 training snapshots. Although the number of samples doubles and then triplicates, the shape of the error distribution for the 3000 validation snapshots does not seem to change significantly, apart from an increase in the frequency of the error by values around 1%. In fact, the model is rather insensitive to the addition of new samples. Since the error seems to decrease slowly, reducing the validation error to an arbitrary value would require a huge number of snapshots.

Finally, it is worth mentioning the behaviour of the other figure of merit, i.e. the computational efficiency of the model. Each snapshot calculation requires about 80 s using 12 CPUs on a workstation. Therefore, the time needed to train the model strongly depends on the number of snapshots. Hence, it is difficult to provide an *a priori* estimate of the time required by the offline calculations. Table 2 proves that both the ROM training, which includes the generation of the input space, the POD reduction and the RBF training, and the ROM execution times increase dramatically with the number of snapshots

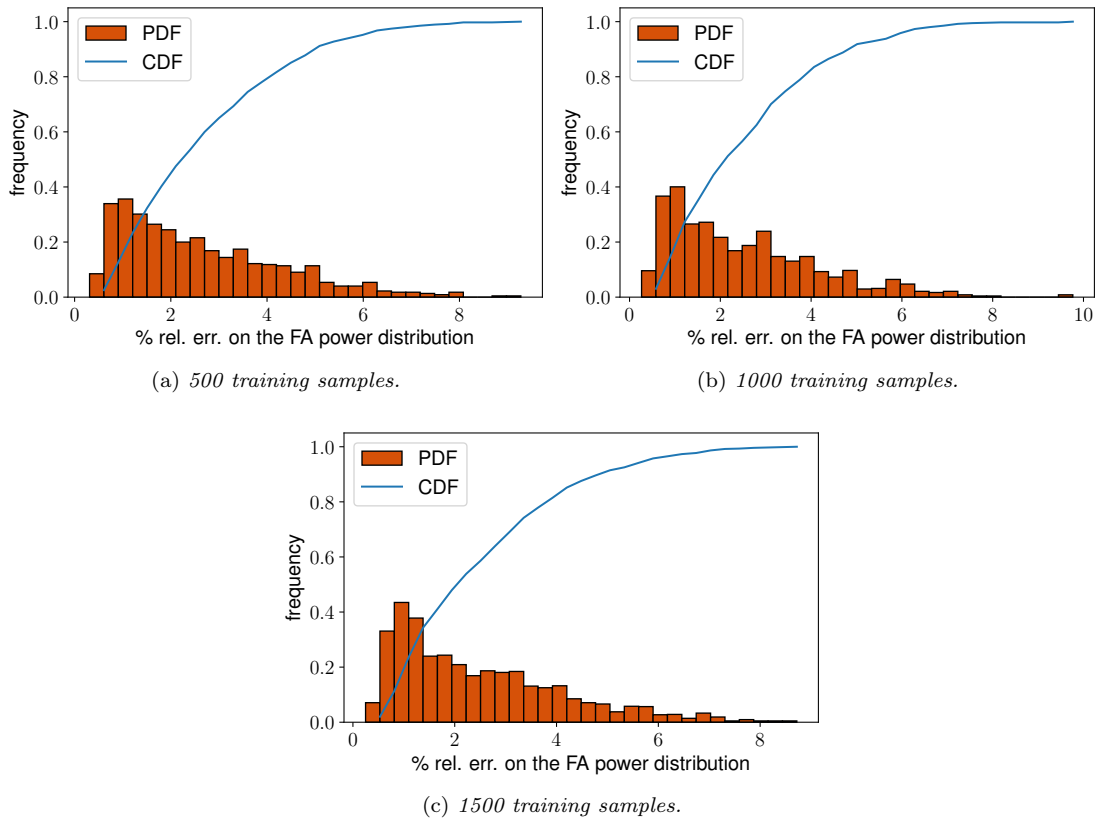


Figure 35: Percentage relative error distribution on 3000 validation samples obtained with the full-random sampling algorithm considering four perturbation batches in the core.




Figure 36: *Input disturbance on the thermal diffusion coefficient (left) and relative relative error on the assembly-wise power (right) for a training case.*

n , as well as the memory resources. Therefore, considering also the behaviour of the error distribution, it is not probably worth to further increase the number of the random training samples.

The final goal of the full-core NIROM described in this section is the accurate and fast estimation of the assembly-wise power to study, for instance, the appearance of the power tilt instability. Hence, in order to complete the evaluation of the ROM performances, it is important to examine more in detail the spatial distribution of the modelling error. Figure 36 reports the perturbation effect on the thermal diffusion coefficient distribution and the relative error between FOM and ROM for a training case. Since the ROM is built in order to reproduce the training data with a very large accuracy, the overall error spatial behaviour is rather uniform, except some local maxima where it can reach values even larger than 1 %. Nevertheless, this accuracy is sufficient to provide an adequate estimate of the quadrant-wise power, which is used in the tilt evaluation (NRC, 2016). The quadrant power is obtained by summing the assembly-wise power in each of the four quadrants, obtained by dividing the core FAs with the lines $x = 0$ and $y = 0$, assuming that the origin of the reference cartesian frame is located in the central FA.

Concerning the validation cases, fig. 37 and fig. 38 show the maps of the largest and lowest power relative errors, respectively. In the first figure, it is possible to appreciate the perturbation effect of the thermal diffusion coefficient and, consequently, on the power distribution, which is clearly skewed towards the left part of the core, due to the increased leakages in the opposite direction. The ROM provides a very good reconstruction of the power, yielding a value of the largest relative error in line with the one characterising the training case (depicted in fig. 36), despite in this case the local spatial distortions are more evident. On the contrary, fig. 38, which indicates the case with the largest relative error $\epsilon\%$, shows that the ROM completely fails in reproducing the power shape. Despite most of the validation points have an error which is less than 5%, the lack of accuracy featuring this case is unacceptable for the intended purpose of the model.

Given that decreasing the modelling error to an acceptable level would require a huge increase in the number of training snapshots, and that a larger dataset could bring the model to the over-fitting condition, a different strategy should be pursued to achieve a higher accuracy.

In this respect, it is convenient to recall the so-called *divide et impera* (divide and conquer) approach (Knuth, 1998), which is a general principle adopted in problem solving. According to this principle, an effective solution strategy consists of dividing recursively the original, complex problem into smaller and smaller sub-problems until a sub-problem that could actually be solved is obtained. Then, one can move backwards solving each sub-problem from the bottom to the top.

In a similar way, smaller, very specialised ROMs could be built for modelling specific situations, e.g., the case in which only one perturbation batch appears. This approach would require to train of a potentially large number of models, but it would offer the advantage of reducing dramatically the parameter space dimension, possibly enabling the adoption of smarter sampling strategies than the brute-force one. As a consequence of the lower input variability, the overall accuracy of the model




Figure 37: *Input disturbance on the thermal diffusion coefficient (top left), power distribution computed by the FOM (top right) and relative relative error on the assembly-wise power (bottom) for the validation case featured by the lowest error.*






Figure 38: *Input disturbance on the thermal diffusion coefficient (left) and relative relative error on the assembly-wise power (right) for the validation case featured by the largest error.*



(a) 512 training samples, composed by 32 different ar-
rangements with 16 different input perturbations each.

(b) 768 training samples, composed by 48 different ar-
rangements with 16 different input perturbations each.



(c) 1024 training samples, composed by 64 different ar-
rangements with 16 different input perturbations each.

Figure 39: Percentage relative error distribution on 1252 validation samples obtained with the full-random sampling algorithm considering one perturbation batch in the core.

would increase. Moreover, for simpler models, increasing the number of training points should be significantly more effective on the model accuracy with respect to the non-specialised ROM.

In the following, the results obtained with a meta-model designed to deal with a single perturbation batch appearing in the full-core are presented. Due to the reduced spatial variability, the hybrid sampling technique described in section 4.3 is employed to map more rigorously the parameter dependence. Concerning the validation samples, the subset of the 3000 validation cases used previously and characterised by one perturbation cluster is selected, for a total of 1252 cases.

As expected, the specialised ROM is much more accurate than the general one, as visible from fig. 39. The training process has been accomplished using 32, 48 and 64 different spatial configurations, each featured by 16 perturbation magnitudes extracted from the Smolyak grid. Despite the lower number of validation samples, the error distribution is clearly peaked towards 1%, with no case going beyond 6%. The superior accuracy of this model can be appreciated looking at the maps for the relative error on the assembly-wise power, reported in fig. 40. The training case is featured by a lower error compared to the one for the general ROM, as well as the validation cases featured by the lowest and the largest errors. Compared to the spatial profile depicted in figs. 37 and 38, the error seems also more spatially uniform, which means that the local deformation of the power distribution due to the modelling approximation is less important.

To assess the sensitivity of the ROM to the training sample selection, the bootstrap method (see algorithm 5) is applied to construct an ensemble of ROMs, each built using 512 snapshots sampled






Figure 40: *Input disturbance on the thermal diffusion coefficient (left) and relative relative error on the assembly-wise power (right) for a training case (top), for the best validation case (centre) and for the worst validation case (bottom).*



(a) Bootstrapped error distribution for a validation case featured by a large variance. The mean error is 6.14 %, featured by a low variance. The mean error is 1.64 %, while the standard deviation is 2.68 %.

(b) Bootstrapped error distribution for a validation case featured by a large variance. The mean error is 6.14 %, featured by a low variance. The mean error is 1.64 %, while the standard deviation is 2.68 %.



(c) Percentage relative error distribution averaged over the ensemble of ROM responses. The mean error is 4.37 %, while the standard deviation is 2.49 %.

Figure 41: Percentage error distributions.

from the set of 1024 training cases available (64 spatial configurations per 16 sets of perturbation magnitude). The computational time required to perform the bootstrap analysis is proportional to the number of ROM reboots and to the computational cost of each model training (see table 2). The choice of limiting the number of training snapshots to 512 is due to a compromise between the computational cost required to train each ROM and the number of ROM reboots needed to provide a sufficient statistics. At a first glance, the bootstrap analysis may appear expensive, especially after the snapshot generation and the training phase. Its application, however, is often justified in practice by the fact that, once properly trained and validated, the ROM is run a very large number of times to perform the parametric safety analysis.

The results of the 1000 ROM reboots are displayed in fig. 41. The upper part of the figure reports the bootstrapped distributions of the percentage relative error for two validation samples. The two cases are, respectively, the ones featured by the largest (left) and lowest (right) errors when the ROM is trained using 1024 points (see the bottom part of fig. 39). In both cases, it can be observed that the error is quite sensitive to the training sample selection, especially for the case characterised by the worst mean error, which also exhibits a very large variance.

The bootstrap analysis also allows to appreciate a very interesting phenomenon: for some training set selection, the validation error of the left case turns out to be smaller than the one obtained if all the 1024 samples were used. This aspect seems to contradict the conclusions drawn observing fig. 39, which shows that adding more and more points has the overall effect of decreasing the percentage

error, despite the maximum error is always around 6%. The solution to this apparent contradiction lies in the fact that, on average, adding more points decreases the overall modelling error, but some cases suffer from over-fitting. In other words, for most of the validation cases, adding more points has a beneficial effect on the modelling error, but, for other points, the opposite situation is verified. In other words, a larger training set may lead, in some cases, to larger modelling errors.

The cause for this "local" *over-fitting* phenomenon is probably due to the fact that some of the validation points fall outside the parameter space covered by the training points. The RBFs are tuned to minimise the error on the training points, hence it may happen that, when this situation occurs, the accuracy of the model becomes extremely poor. Due to this aspect, the points located outside the training points are extremely sensitive to the training dataset selection, thus explaining the unpleasant over-fitting behaviour observed above. These arguments also help to explain the graph on bottom part of fig. 41, which shows the trend of the percentage relative error distribution for the validation set, obtained by averaging the 1000 bootstrapped distributions. In light of these considerations, it can be concluded that, despite it is larger than the average error obtained with 1024 samples, the averaged error computed in this case is a very prudent estimate, which accounts for the sensitivity of the RBF to the selection of the training points.

5. Conclusions and future developments

In this paper, a novel surrogate model has been proposed to reduce the computational cost associated with the full-core diffusion calculations needed to assess the core behaviour in the presence of local disturbances. In the worst case scenario, i.e. the one featured by the most complex ROM, the execution time of the surrogate model after the computationally intensive off-line phase is about 600 times faster than the original diffusion model.

Due to the two-step nature of the full-core calculations, which require a set of few-group homogenised constants, a two-step meta-model has been proposed, combining the polynomial chaos expansion method for the estimation of the multi-group data and the POD-RBF model for the estimation of the assembly-wise power distribution obtained in response to localised perturbations, which are assumed to be variations in the coolant density and in the fuel pellet diameter.

A first-order polynomial chaos expansion has revealed sufficient to reproduce with an acceptable accuracy the cross sections behaviour as a function of the two input perturbations, requiring only a few full-order model evaluations. However, it must be acknowledged that, in the presence of larger variations, probably more training points would be required to match the target accuracy.

Concerning the full-core calculations, some precautions have been taken to handle the spatial arrangement of the perturbations. First of all, an algorithm for the definition of the perturbation and its constraints, e.g., the perturbation shape, has been envisaged. Then, two different sampling techniques are devised to produce the training dataset. The first one draws the first perturbed FA according to an importance probability density function based on the neutron importance, and then randomly selects the surrounding FAs, assigning a random value to the physical perturbation as well. The second sampling technique works in the same way, but the perturbation magnitude is identified according to a deterministic rule.

The last precaution regards the distance evaluation among the input parameter configurations: due to the spatial arrangement of the disturbance, the classical euclidean distance between the various input parameter is not adequate to distinguish similar perturbations, engendering the predicting capabilities of the model. Thus, an alternative algorithm for image recognition has been selected, to guarantee a consistent evaluation of the training samples distance with respect to a new parameter point.

When it is assumed that more perturbations can occur in the core, the parameter space of the problem becomes huge. However, according to the results obtained, it seems that the random sampling technique is adequate to provide good results. Nevertheless, it seems difficult to strongly reduce the error, due to the large variability in the input space.

When a simpler problem is addressed, i.e. only one batch of perturbed FA is considered, the overall accuracy of the model has been shown to improve significantly, also thanks to the adoption

of the hybrid sampling process, which combines a random sampling for the spatial distribution of the perturbation and a deterministic rule for its perturbation. This observation suggests that more specialised meta-models (e.g., only focusing on single batches or single quadrants) could be the most efficient solution to the problem.

Finally, the influence of the training set size and random composition has been assessed quantitatively using the bootstrap method, which has revealed the issue of local over-fitting and allowed to get a conservative estimate of the overall modelling error.

As a future development, hybrid techniques based on data- and physics-driven approaches should be investigated. For example, it would be interesting to use the legacy Generalised Perturbation Theory technique to provide a set of sensitivity coefficients to drive the sampling phase towards the most relevant perturbations. Moreover, the dataset produced in this work could be used to train alternative surrogate models based on other supervised machine learning techniques, in order to take into account also the time evolution of the system subject to localised perturbations.

Data and software availability

The complete datasets and scripts employed to pre- and post-process the calculations presented in this paper are available in the open access Zenodo repository [10.5281/zenodo.7472108](https://doi.org/10.5281/zenodo.7472108).

Acknowledgements

Computational resources were provided by HPC@POLITO, a project of Academic Computing within the Department of Control and Computer Engineering at Politecnico di Torino (<http://www.hpc.polito.it>). The authors would also like to thank Prof. Piero Ravetto for his suggestions concerning the possible definitions of the importance function for neutrons.

References

- N. Abrate, G. Bruna, S. Dulla, and P. Ravetto. Assessment of numerical methods for the evaluation of higher-order harmonics in diffusion theory. *Annals of Nuclear Energy*, 128:455–470, 2019a. URL <https://doi.org/10.1016/j.anucene.2019.01.011>.
- N. Abrate, G. Bruna, S. Dulla, and P. Ravetto. Convergence limits in perturbation theory. In *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, M&C 2019, Portland, OR (U.S.A.)*, pages 2381–2390, 2019b. URL <https://www.ans.org/pubs/proceedings/article-46819/>.
- N. Abrate, S. Dulla, and P. Ravetto. Generalized perturbation techniques for uncertainty quantification in lead-cooled fast reactors. *Annals of Nuclear Energy*, 164:108623, 2021. ISSN 0306-4549. doi: <https://doi.org/10.1016/j.anucene.2021.108623>. URL <https://www.sciencedirect.com/science/article/pii/S0306454921004990>.
- Manuele Aufiero, Adrien Bidaud, Mathieu Hursin, Jaakko Leppänen, Giuseppe Palmiotti, Sandro Peloni, and Pablo Rubiolo. A collision history-based approach to sensitivity/perturbation calculations in the continuous energy monte carlo code serpent. *Annals of Nuclear Energy*, 85:245–258, 2015.
- S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushikand, M. G. Knepley, D. A. Mayand, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, Cass Ave, Lemont, IL (U.S.A.), 2020. URL <https://petsc.org/main/docs/manual/>.
- Falco J Bargagli Stoffi, Gustavo Cevolani, and Giorgio Gnecco. Simple models in complex worlds: Occam’s razor and statistical learning theory. *Minds and Machines*, 32(1):13–42, 2022. URL <https://link.springer.com/article/10.1007/s11023-022-09592-z>.
- G. I. Bell and S. Glasstone. *Nuclear Reactor Theory*. Van Nostrand Reinhold, 1970. URL <https://www.osti.gov/biblio/4074688>.
- R. Bellman. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957. ISBN 9780691079516. URL <https://books.google.it/books?id=wdtoPwAACAAJ>.
- P. Benner, S. Gugercin, and K. Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review*, 57(4):483–531, 2015. doi: 10.1137/130932715. URL <https://doi.org/10.1137/130932715>.
- Dan Gabriel Cacuci. *Handbook of Nuclear Engineering: Vol. 1: Nuclear Engineering Fundamentals; Vol. 2: Reactor Design; Vol. 3: Reactor Analysis; Vol. 4: Reactors of Generations III and IV; Vol. 5: Fuel Cycles, Decommissioning, Waste Disposal and Safeguards*, volume 1. Springer Science & Business Media, 2010. URL <https://link.springer.com/referencework/10.1007/978-0-387-98149-9>.
- Claudio Canuto, M Yousuff Hussaini, Alfio Quarteroni, and Thomas A Zang. *Spectral methods: evolution to complex geometries and applications to fluid dynamics*. Springer Science & Business Media, 2007. URL <https://link.springer.com/book/10.1007/978-3-540-30728-0>.
- F. Casenave, A. Gariah, C. Rey, and F. Feyel. A nonintrusive reduced order model for nonlinear transient thermal problems with nonparametrized variability. *Advanced Modeling and Simulation in Engineering Sciences*, 7(1), 2020. doi: 10.1186/s40323-020-00156-3.
- W. Chen, J. S. Hesthaven, B. Junqiang, Y. Qiu, Z. Yang, and Y. Tihao. Greedy nonintrusive reduced order model for fluid dynamics. *AIAA Journal*, 56(12):4927–4943, 2018. doi: 10.2514/1.J056161. URL <https://doi.org/10.2514/1.J056161>.

- Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995. URL <https://link.springer.com/article/10.1007/BF00994018>.
- J. J. Duderstadt and L. J. Hamilton. *Nuclear Reactor Analysis*. Wiley, New York, 1976. URL <https://www.wiley.com/en-us/Nuclear+Reactor+Analysis-p-9780471223634>.
- S. Dutta, M. W. Farthing, E. Perracchione, G. Savant, and M. Putti. A greedy non-intrusive reduced order model for shallow water equations. *Journal of Computational Physics*, 439:110378, 2021. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2021.110378>. URL <https://www.sciencedirect.com/science/article/pii/S0021999121002734>.
- B. Efron. Nonparametric estimates of standard error: The jackknife, the bootstrap and other methods. *Biometrika*, 68(3):589–599, 1981. ISSN 00063444. URL <http://www.jstor.org/stable/2335441>.
- Jonathan Feinberg and Hans Petter Langtangen. Chaospy: An open source tool for designing methods of uncertainty quantification. *Journal of Computational Science*, 11:46–57, 2015. URL <https://www.sciencedirect.com/science/article/pii/S1877750315300119>.
- A. Gandini. Higher order time-dependent generalized perturbation theory. *Nuclear Science and Engineering*, 67:91–106, 1978. URL <https://doi.org/10.13182/NSE78-A27240>.
- L. Gilli, D. Lathouwers, J.L. Kloosterman, T.H.J.J. van der Hagen, A.J. Koning, and D. Rochman. Uncertainty quantification for criticality problems using non-intrusive and adaptive polynomial chaos techniques. *Annals of Nuclear Energy*, 56:71–80, 2013. ISSN 0306-4549. doi: <https://doi.org/10.1016/j.anucene.2013.01.009>. URL <https://www.sciencedirect.com/science/article/pii/S0306454913000261>.
- R. L. Hardy. Multiquadric equations of topography and other irregular surfaces. *Journal of Geophysical Research (1896-1977)*, 76(8):1905–1915, 1971. doi: <https://doi.org/10.1029/JB076i008p01905>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JB076i008p01905>.
- Frédéric Hecht. New development in freefem++. *Journal of numerical mathematics*, 20(3-4):251–266, 2012. URL <https://www.degruyter.com/document/doi/10.1515/jnum-2012-0013/html>.
- V. Hernandez, J. E. Roman, and V. Vidal. SLEPc: a scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Transactions on Mathematical Software*, 31:351–362, 2005. URL <https://doi.org/10.1145/1089014.1089019>.
- E. Iuliano and D. Quagliarella. Aerodynamic shape optimization via non-intrusive pod-based surrogate modelling. In *2013 IEEE Congress on Evolutionary Computation*, pages 1467–1474, 2013. doi: 10.1109/CEC.2013.6557736.
- Kalcko Ivanov, M Avramova, Ivan Aleksander Kodeli, and E Sartori. *Benchmark for uncertainty analysis in modeling (UAM) for design, operation and safety analysis of LWRs*. Citeseer, 2007. URL https://www.oecd-nea.org/jcms/pl_19290/benchmarks-for-uncertainty-analysis-in-modelling-uam-for-the-design-operation-and-safety-analysis-details=true.
- Donald E. Knuth. *The Art of Computer Programming Volumes 1-3*. Addison-Wesley Longman Publishing Co., Inc., 1998. URL <https://cs.stanford.edu/~knuth/taocp.html>.
- T. Lassila, A. Manzoni, A. Quarteroni, and G. Rozza. Model Order Reduction in Fluid Dynamics: Challenges and Perspectives. In Quarteroni A., Rozza G., editor, *Reduced Order Methods for Modeling and Computational Reduction. MS&A - Modeling, Simulation and Applications*, volume 9, chapter 10. Springer, Cham, 2014. doi: https://doi.org/10.1007/978-3-319-02090-7_9.

- J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho. The Serpent Monte Carlo code: Status, development and applications in 2013. *Annals of Nuclear Energy*, 82:142–150, 2015. doi: 10.1016/j.anucene.2014.08.024.
- S. Lorenzi, A. Cammi, L. Luzzi, and G. Rozza. Pod-galerkin method for finite volume approximation of navier–stokes and rans equations. *Computer Methods in Applied Mechanics and Engineering*, 311:151–179, 2016. ISSN 0045-7825. doi: <https://doi.org/10.1016/j.cma.2016.08.006>. URL <https://www.sciencedirect.com/science/article/pii/S0045782516308829>.
- S. Lorenzi, A. Cammi, L. Luzzi, and G. Rozza. A reduced order model for investigating the dynamics of the gen-iv lfr coolant pool. *Applied Mathematical Modelling*, 46:263–284, 2017. ISSN 0307-904X. doi: <https://doi.org/10.1016/j.apm.2017.01.066>. URL <https://www.sciencedirect.com/science/article/pii/S0307904X17300720>.
- Stefano Lorenzi. An adjoint proper orthogonal decomposition method for a neutronics reduced order model. *Annals of Nuclear Energy*, 114:245–258, 2018.
- S. Marelli and B. Sudret. An active-learning algorithm that combines sparse polynomial chaos expansions and bootstrap for structural reliability analysis. *Structural Safety*, 75:67–74, 2018. ISSN 0167-4730. doi: <https://doi.org/10.1016/j.strusafe.2018.06.003>. URL <https://www.sciencedirect.com/science/article/pii/S0167473017302977>.
- A. Marrel, N. Pérot, and C. Mottet. Development of a surrogate model and sensitivity analysis for spatio-temporal numerical simulators. *Stochastic Environmental Research and Risk Assessment*, 29: 959–974, 2014. doi: 10.1007/s00477-014-0927-y.
- Joseph B Nagel, Jörg Rieckermann, and Bruno Sudret. Principal component analysis and sparse polynomial chaos expansions for global sensitivity analysis and model calibration: Application to urban drainage simulation. *Reliability Engineering & System Safety*, 195:106737, 2020.
- S. Nanty, C. Helbert, A. Marrel, N. Pérot, and C. Prieur. Uncertainty quantification for functional dependent random variables. *Computational Statistics*, 32:559–583, 2017. URL <https://link.springer.com/article/10.1007/s00180-016-0676-0>.
- NRC. Standard technical specifications, westinghouse advanced passive 1000 (ap1000) plants, volume 1: Specifications. Technical Report NUREG-2194, Office of New Reactors, U.S. Nuclear Regulatory Commission, Washington, DC (U.S.A.), 2016. URL <https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr2194/v1/index.html>.
- N. Pedroni, E. Zio, and G. Apostolakis. Comparison of bootstrapped artificial neural networks and quadratic response surfaces for the estimation of the functional failure probability of a thermal-hydraulic passive system. *Reliability Engineering and System Safety*, 95(4):386–395, 2010. URL <https://hal-supelec.archives-ouvertes.fr/hal-00609171>.
- Karl R. Popper and W.W. Barthley III. *The Open universe: an argument for indeterminism from the postscript to the logic of scientific discovery*. Routledge, London, U.K., 1988. URL <https://api.taylorfrancis.com/content/books/mono/download?identifierName=doi&identifierValue=10.4324/9780203713907&type=googlepdf>.
- S. M. Rahman, O. San, and A. Rasheed. A hybrid approach for model order reduction of barotropic quasi-geostrophic turbulence. *Fluids*, 3(4), 2018. ISSN 2311-5521. doi: 10.3390/fluids3040086. URL <https://www.mdpi.com/2311-5521/3/4/86>.
- S. Rippa. An algorithm for selecting a good value for the parameter c in radial basis function interpolation. *Advances in Computational Mathematics*, 34(1):105–126, 2011. ISSN 10197168. doi: 10.1007/s10444-010-9146-3.

- Stuart J Russell. *Artificial intelligence a modern approach*. Pearson Education, Inc., 2010. URL <http://aima.cs.berkeley.edu/global-index.html>.
- Y. Saad. *Numerical methods for large eigenvalue problems*. Manchester University Press, 1992. URL <https://doi.org/10.1137/1.9781611970739>.
- Richard Sanchez, Igor Zmijarevi, Mireille Coste-Delclaux, Emiliano Masiello, Simone Santandrea, Emanuele Martinolli, Laurence Villate, Nadine Schwartz, and Nathalie Guler. Apollo2 year 2010. *Nuclear engineering and technology*, 42(5):474–499, 2010.
- M. Santanoceto, M. Tiberga, Z. Perkó, S. Dulla, and D. Lathouwers. Preliminary uncertainty and sensitivity analysis of the molten salt fast reactor steady-state using a polynomial chaos expansion method. *Annals of Nuclear Energy*, 159:108311, 2021. URL <https://doi.org/10.1016/j.anucene.2021.108311>.
- A Sargeni, KW Burn, and GB Bruna. The impact of heavy reflectors on power distribution perturbations in large pwr reactor cores. *Annals of Nuclear Energy*, 94:566–575, 2016. URL <https://www.sciencedirect.com/science/article/pii/S0306454916301372>.
- R. Schaback. Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics*, 3(3):251–264, 1995. ISSN 10197168. doi: 10.1007/BF02432002.
- P. Secchi, E. Zio, and F. Di Maio. Quantifying uncertainties in the estimation of safety parameters by using bootstrapped artificial neural networks. *Annals of Nuclear Energy*, 35(12):2338–2350, 2008. ISSN 0306-4549. doi: <https://doi.org/10.1016/j.anucene.2008.07.010>. URL <https://www.sciencedirect.com/science/article/pii/S0306454908002119>.
- S. A. Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. *Dokl. Akad. Nauk SSSR*, 148(5):1042–1045, 1963. URL http://www.mathnet.ru/php/archive.phtml?wshow=paper&jrnid=dan&paperid=27586&option_lang=eng.
- Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability engineering & system safety*, 93(7):964–979, 2008. URL <https://www.sciencedirect.com/science/article/pii/S0951832007001329>.
- Yue Sun, Junhe Yang, Yahui Wang, Zhuo Li, and Yu Ma. A pod reduced-order model for resolving the neutron transport problems of nuclear reactor. *Annals of Nuclear Energy*, 149:107799, 2020. URL <https://www.sciencedirect.com/science/article/pii/S0306454920304977>.
- O Tal, E Israeli, P Ravetto, and E Gilad. The adjoint problem as physical heuristic for loading pattern optimization. *Annals of Nuclear Energy*, 134:226–234, 2019. URL <https://www.sciencedirect.com/science/article/pii/S0306454919303329>.
- L. N. Usachev. Equation for the importance of neutrons, reactor kinetics and the theory of perturbations. In *Proc. Int. Conf. on the Peaceful Uses of Atomic Energy, Geneva, Switzerland, Aug. 8-12, 1955*, volume 5, pages 503–510, 1956.
- S. Volkwein. Model reduction using proper orthogonal decomposition. *Lecture notes, Institute of Mathematics and Scientific Computing, University of Graz*. see <http://www.uni-graz.at/imawww/volkwein/POD.pdf>, 1025, 2011. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.616.6416&rep=rep1&type=pdf>.
- Liwei Wang, Yan Zhang, and Jufu Feng. On the euclidean distance of images. *IEEE transactions on pattern analysis and machine intelligence*, 27(8):1334–1339, 2005. URL <https://ieeexplore.ieee.org/abstract/document/1453520/>.

- Norbert Wiener. The homogeneous chaos. *American Journal of Mathematics*, 60(4):897–936, 1938. URL <https://doi.org/10.2307/2371268>.
- M. M. R. Williams. Polynomial chaos functions and neutron diffusion. *Nuclear Science and Engineering*, 155(1):109–118, 2007. doi: 10.13182/NSE05-73TN. URL <https://doi.org/10.13182/NSE05-73TN>.
- Mark L Williams. Resonance self-shielding methodologies in scale 6. *Nuclear Technology*, 174(2):149–168, 2011. URL <https://www.tandfonline.com/doi/abs/10.13182/NT09-104>.
- ML Williams. Generalized contribution response theory. *Nuclear Science and Engineering*, 108(4):355–383, 1991. URL <https://www.tandfonline.com/doi/abs/10.13182/NSE90-33>.
- X. Wu, T. Kozłowski, and H. Meidani. Kriging-based inverse uncertainty quantification of nuclear fuel performance code BISON fission gas release model using time series measurement data. *Reliability Engineering & System Safety*, 169:422–436, 2018a. ISSN 0951-8320. doi: <https://doi.org/10.1016/j.res.2017.09.029>. URL <https://www.sciencedirect.com/science/article/pii/S095183201730532X>.
- X. Wu, T. Kozłowski, H. Meidani, and K. Shirvan. Inverse uncertainty quantification using the modular Bayesian approach based on Gaussian process, Part 1: Theory. *Nuclear Engineering and Design*, 335:339–355, 2018b. ISSN 0029-5493. doi: <https://doi.org/10.1016/j.nucengdes.2018.06.004>. URL <https://www.sciencedirect.com/science/article/pii/S0029549318306423>.
- D. Xiao, F. Fang, C. C. Pain, and I. M. Navon. A parameterized non-intrusive reduced order model and error analysis for general time-dependent nonlinear partial differential equations and its applications. *Computer Methods in Applied Mechanics and Engineering*, 317:868–889, 2017. ISSN 00457825. doi: 10.1016/j.cma.2016.12.033. URL <http://dx.doi.org/10.1016/j.cma.2016.12.033>.
- Dongbin Xiu and George Em Karniadakis. The wiener–askey polynomial chaos for stochastic differential equations. *SIAM journal on scientific computing*, 24(2):619–644, 2002. URL <https://doi.org/10.1137/S1064827501387826>.
- E. Zio. A study of the bootstrap method for estimating the accuracy of artificial neural networks in predicting nuclear transient processes. *IEEE Transactions on Nuclear Science*, 53(3):1460–1478, 2006. doi: 10.1109/TNS.2006.871662.
- E. Zio, G.E. Apostolakis, and N. Pedroni. Quantitative functional failure analysis of a thermal–hydraulic passive system by means of bootstrapped artificial neural networks. *Annals of Nuclear Energy*, 37(5):639–649, 2010. ISSN 0306-4549. doi: <https://doi.org/10.1016/j.anucene.2010.02.012>. URL <https://www.sciencedirect.com/science/article/pii/S0306454910000666>.

Appendix A. POD-RBF algorithms

The appendix reports the detailed algorithms implemented in the ROMpy package.

Algorithm 2 reports the algorithm for the POD reduction, while algorithms 3-4 illustrate the steps to train the RBF net optimising the shape parameter with the Leave-One-Out-Cross-Validation. Finally, algorithm 5 shows the procedure for bootstrapping the ROM.

Algorithm 2: Proper Orthogonal Decomposition algorithm

Input

1. snapshot matrix $\hat{Y} = [\vec{y}_1, \vec{y}_2, \dots, \vec{y}_n]$, with $\vec{y}_i \in \mathbb{R}^m$;
2. snapshot matrix rank r ;
3. truncation error ε ;

Output

1. POD basis $\hat{B}_k = [\vec{b}_1, \vec{b}_2, \dots, \vec{b}_k]$, with $\vec{b}_i \in \mathbb{R}^m$;
2. POD eigenvalues $\lambda_i, \forall i = 1, \dots, k$;
3. POD energy \mathcal{E}_k ;
4. POD coefficients $\hat{A}_k = [\vec{a}_1, \vec{a}_2, \dots, \vec{a}_k]$, with $\vec{a}_i \in \mathbb{R}^n$;

try:

```
[ $\hat{\Psi}, \hat{\Sigma}, \hat{\Phi}$ ] = svd( $\hat{Y}$ ) # Singular Value Decomposition
for  $i = 1, \dots, r$  do
  |  $\vec{b}_i = \vec{\psi}_i$  #  $\vec{\psi}_i$  is the  $\hat{\Psi}$   $i$ -th column
  |  $\lambda_i = \sigma_i^2$  #  $\sigma_i$  is the  $i$ -th diagonal entry of  $\hat{\Sigma}$ 
end
```

except Memory Error:

```
if  $n > m$  then
  | [ $\hat{\Psi}, \hat{\Lambda}$ ] = eig( $\hat{Y}\hat{Y}^T$ ) # Eigenvalue Decomposition
  | for  $i = 1, \dots, r$  do
  | |  $\vec{b}_i = \vec{\psi}_i$  #  $\vec{\psi}_i$  is the  $i$ -th column of  $\hat{\Psi}$ 
  | |  $\lambda_i = \hat{\Lambda}_{i,i}$  #  $\hat{\Lambda}_{i,i}$  is the  $i$ -th diagonal entry of  $\hat{\Lambda}$ 
  | end
```

else

```
[ $\hat{\Phi}, \hat{\Lambda}$ ] = eig( $\hat{Y}^T\hat{Y}$ ) # Eigenvalue Decomposition
for  $i = 1, \dots, r$  do
  |  $\vec{b}_i = \frac{\hat{Y}\vec{\phi}_i}{\sqrt{\lambda_i}}$  #  $\vec{\psi}_i$  is the  $i$ -th column of  $\hat{\Psi}$ 
  |  $\lambda_i = \hat{\Lambda}_{i,i}$  #  $\hat{\Lambda}_{i,i}$  is the  $i$ -th diagonal entry of  $\hat{\Lambda}$ 
end
```

end**end**

choose k such that $1 - \mathcal{E}_k < \varepsilon$

$\hat{A}_k = \hat{Y}^T \hat{B}_k$ # compute reduced order coefficients

Algorithm 3: Radial Basis Function training algorithm.

Input

1. data reduced via POD, $\hat{A}_k \in \mathbb{R}^{n \times k}$;
2. parameter matrix $\hat{P} = [\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n]$, with $\vec{p}_i \in \mathbb{R}^p$;
3. RBF type f (e.g. $f = 1/\sqrt{\|\vec{p}_1 - \vec{p}_j\|_2 + \sigma^2}$);
4. hyperparameter σ via Algorithm 4 or Hardy's formula [Hardy \(1971\)](#);

Output

training matrix $\hat{W} \in \mathbb{R}^{n \times k}$
loop over each column of \hat{A}_k , $\vec{a}_i \in \mathbb{R}^n$
for $i = 1, \dots, k$ **do**
 choose hyperparameter σ_i (optimal selection with algorithm 4)
 # compute distance matrix $\hat{D} = [\vec{d}_1, \dots, \vec{d}_n] \in \mathbb{R}^{n \times n}$
 for $j = 1, \dots, n$ **do**
 $\vec{d}_j = [f(\|\vec{p}_1 - \vec{p}_j\|_2, \sigma_i), \dots, f(\|\vec{p}_n - \vec{p}_j\|_2, \sigma_i)]$
 end
 $\vec{w}_i = \hat{D}^{-1} \vec{a}_i$ # \vec{w}_i is the i -th column of \hat{W}
end

Algorithm 4: Optimal hyperparameter selection.

Input

1. data reduced via POD, $\hat{A}_k \in \mathbb{R}^{n \times k}$;
2. parameter matrix $\hat{P} = [\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n]$, with $\vec{p}_i \in \mathbb{R}^p$;
3. RBF type f (e.g. $f = 1/\sqrt{\|\vec{p}_1 - \vec{p}_j\|_2 + \sigma^2}$);
4. initial hyperparameter guess $\sigma_0 = \frac{0.815}{n} \sum_{i=1}^n \|\vec{p}_i - \vec{p}_\ell\|_2$, where ℓ is the closest point to i in the parameter space [Hardy \(1971\)](#).

Output

1. Optimal hyperparameter σ_{opt}

Optimise σ minimising RMSE

```
Select  $\sigma$  # e.g. via conjugate gradient method
Compute  $\hat{D}$  and  $\vec{w}_i$  as in algorithm 3
# use Rippa's algorithm Rippa \(2011\) and LOOCV
 $[\hat{L}, \hat{U}] = \text{lu}(\hat{D})$  # Compute LU decomposition
for  $j=1, \dots, n$  do
   $\vec{x} = \hat{U}^{-1} \hat{L}^{-1} \vec{e}_j$  #  $\vec{e}_j$  is the  $j$ -th column of the identity matrix
   $\delta_j = \|\frac{w_{i,j}}{x_j}\|$  # compute interpolation error with  $j$ -th components of  $\vec{x}$  and  $\vec{w}_i$ 
end
RMSE =  $\sqrt{\frac{1}{n} \sum_{j=1}^n \delta_j^2}$ 
```

end

Algorithm 5: Bootstrap method

Input

1. number of reboots N_r ;
2. training parameter values $\hat{P} = [\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n]$;
3. training snapshots $\hat{Y} = [\vec{y}_1, \vec{y}_2, \dots, \vec{y}_n]$;

OutputEnsemble of ROMs, $\mathcal{R}_1, \dots, \mathcal{R}_{N_r}$

this can be done in parallel

for $i=1, \dots, N_r$ **do****for** $j=1, \dots, N$ **do**Sample one parameter point \vec{p}_j from $[\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n]$ with replacement;**if** \vec{p}_j exists **then**| Discard \vec{p}_j to avoid singular RBF kernel;**end**Reduce dataset with [Algorithm 2](#);Train and tune the RBF network with [Algorithms 3, 4](#);**end**
