

UNCERTAINTY QUANTIFICATION IN COUPLED MULTIPHYSICS SYSTEMS VIA GAUSSIAN PROCESS SURROGATES: APPLICATION TO FUEL ASSEMBLY BOW

Ali Abboud,^{1,*} Stanislas de Lambert,² Josselin Garnier,¹ & Bertrand Leturcq²

¹Centre de Mathématiques Appliquées, Ecole Polytechnique, Institut Polytechnique de Paris, 91120, Palaiseau, France

²Université Paris-Saclay, CEA, Service d'Études Mécaniques et Thermiques, 91191, Gif-sur-Yvette, France

*Address all correspondence to: Ali Abboud, Centre de Mathématiques Appliquées, Ecole Polytechnique, Institut Polytechnique de Paris, 91120, Palaiseau, France, E-mail: ali.abboud@polytechnique.edu

Predicting fuel assembly bow in pressurized water reactors requires solving tightly coupled fluid-structure interaction problems, whose direct simulations can be computationally prohibitive, making large-scale uncertainty quantification very challenging. This work introduces a general mathematical framework for coupling Gaussian process surrogate models representing distinct physical solvers, aimed at enabling rigorous uncertainty quantification in coupled multiphysics systems. A theoretical analysis establishes that the predictive variance of the coupled GP system remains bounded under mild regularity and stability assumptions, ensuring that uncertainty does not grow uncontrollably through the iterative coupling process. The methodology is then applied to the coupled hydraulic-structural simulation of the fuel assembly bow, enabling global sensitivity analysis and full uncertainty quantification at a fraction of the computational cost of direct code coupling. The results demonstrate accurate uncertainty propagation and stable predictions, establishing a solid mathematical basis for surrogate-based coupling in large-scale multiphysics simulations.

KEY WORDS: *uncertainty quantification, Gaussian process regression, fuel assembly bow, fluid-structure interaction*

1. INTRODUCTION

Many multiphysics problems require coupling distinct computational codes, each solving a specific physical subsystem. Such partitioned code-coupling schemes iterate between the solvers, exchanging boundary data until a global equilibrium is reached. While conceptually flexible, this approach becomes computationally demanding when high-fidelity solvers are involved and many iterations are needed. This challenge is magnified in uncertainty quantification and global sensitivity analysis, where thousands of coupled evaluations are typically required. The goal is no longer just to compute a coupled state but to quantify how input uncertainties propagate through an iterative multicode coupling. To make this tractable, high-fidelity solvers can be replaced by probabilistic surrogate models, notably Gaussian process surrogates [1], which provide both predictions and associated uncertainty estimates. Integrating these surrogates in a coupled system yields a Gaussian process-based coupled model, where the outputs of one Gaussian process feed the inputs of another until convergence. This raises fundamental mathematical questions: how to propagate uncertainty consistently through the coupled system, under what conditions the predictive variance remains bounded, and how to ensure stability and consistency of the coupled predictions. This work addresses these issues by formulating Gaussian process coupling as a fixed-point iteration and by proving variance-control results under mild regularity assumptions. The methodology is first illustrated on an analytical test case with a known reference solution and is then applied to a realistic fluid–structure interaction (FSI) problem involving the deformation of fuel assemblies in pressurized water reactors [2,3].

The main contributions of this work are as follows:

- A Monte Carlo strategy tailored to coupled fixed-point solvers, based on a deterministic mean-path run and constant-offset perturbations of Gaussian process posterior means, enabling scalable uncertainty quantification without sampling full Gaussian process realizations over continuous input spaces.
- A theoretical analysis that links Gaussian processes posterior variances to design fill distances [including a lift to multioutput Gaussian process through a linear model of coregionalization (LMC) kernel representation] and combines this with Lipschitz stability of the coupled solution map to obtain finite-sample, high-probability bounds on the coupled outputs.
- An analytical benchmark validation comparing the proposed scheme against a rigorous but computationally heavy trajectory-conditioned sampling procedure.
- An application to fuel-assembly bow in pressurized-water-reactor simulations through a coupled fluid–structure interaction workflow (hydraulic and mechanical surrogates), illustrating the practical impact of the method on predictive uncertainty assessment and sensitivity analysis.

2. GENERAL METHODOLOGY

This section investigates the propagation of uncertainty arising from the coupling of two independently constructed surrogate models, each representing a distinct physical code. Building on the Bayesian framework of Gaussian process, which inherently provides uncertainty estimates, we develop and analyze a strategy to quantify how these uncertainties interact through the coupling. A combination of theoretical analysis and Monte Carlo simulations is used to validate the approach, culminating in an application to the modeling of fuel assembly bow phenomena in the next section.

2.1 Coupling of Two Computational Codes as a Fixed-Point Problem

Before addressing uncertainty quantification in the context of coupled surrogate models, it is necessary to give a precise mathematical formulation of the deterministic coupling problem. We present here the abstract framework, introduce the operators involved, and establish that the coupled solution is well defined and that its approximation can be characterized through finitely many evaluations of the solvers.

We consider the coupling of a finite number $C \geq 2$ of computational solvers. For readability, Sections 2.1 and 2.2 focus on the practical case $C = 2$. The general case C is treated in Section 2.3 for the theoretical analysis.

Interface space. The interaction between the two solvers is mediated through a set of interface variables. These variables are collected in a vector u belonging to the interface space $\mathcal{U}$. In theory, $\mathcal{U}$ may be viewed as a Banach space, but in practice we only work with finite-dimensional subspaces $\mathcal{U} \subset \mathbb{R}^d$. By construction, $\mathcal{U}$ is distinct from the input spaces $\mathcal{X}_c$ ($c = 1, 2$) of the individual solvers: it contains only the quantities that are communicated between subsystems. Similarly, the input and output spaces $\mathcal{X}_c, \mathcal{Y}_c$ are finite-dimensional subsets of $\mathbb{R}^{d_c}$ and $\mathbb{R}^{D_c}$, respectively.

Definition 1 (Computational solvers). For $c = 1, 2$, let $\mathcal{S}_c : \mathcal{X}_c \rightarrow \mathcal{Y}_c$ be a deterministic solver associated with subsystem c . The space $\mathcal{X}_c$ denotes admissible inputs, and $\mathcal{Y}_c$ the corresponding outputs.

In practice, each solver can only be accessed through a single evaluation at each iteration.

Definition 2 (Transfer operators). The solvers exchange data through transfer operators,

$$\Gamma_1 : \mathcal{U} \rightarrow \mathcal{X}_1, \quad \Gamma_{12} : \mathcal{Y}_1 \rightarrow \mathcal{X}_2, \quad \Gamma_2 : \mathcal{Y}_2 \rightarrow \mathcal{U}. \quad (1)$$

The operator Γ_1 maps the interface state into input data for solver $\mathcal{S}_1$, Γ_{12} transfers the output of $\mathcal{S}_1$ to an admissible input for solver $\mathcal{S}_2$, and Γ_2 updates the interface state from the output of $\mathcal{S}_2$.

Definition 3 (Coupling operator). The global coupling operator is defined as

$$\mathcal{T} := \Gamma_2 \circ \mathcal{S}_2 \circ \Gamma_{12} \circ \mathcal{S}_1 \circ \Gamma_1 : \mathcal{U} \rightarrow \mathcal{U}. \quad (2)$$

Given an interface state $u \in \mathcal{U}$, one application of $\mathcal{T}$ consists of: mapping u into an input for solver $\mathcal{S}_1$ via Γ_1 , running $\mathcal{S}_1$ once, transferring its output to $\mathcal{X}_2$ using Γ_{12} , running $\mathcal{S}_2$ once, and mapping the result back into $\mathcal{U}$ with Γ_2 . A coupled solution is a fixed point $u^* = \mathcal{T}(u^*)$, i.e., an interface state consistent with both solvers.

Proposition 1 (Banach fixed-point theorem [4]). Assume $\mathcal{T}$ is a contraction on a Banach space $(\mathcal{U}, \|\cdot\|_{\mathcal{U}})$. Then, for any initial $u^{(0)} \in \mathcal{U}$, the sequence

$$u^{(m+1)} := \mathcal{T}(u^{(m)}), \quad m \geq 0, \quad (3)$$

converges to the unique fixed point u^* .

Definition 4 (Iterative termination). In computations, the iteration is stopped at index $M < \infty$ when the criterion

$$\|u^{(M)} - u^{(M-1)}\|_{\mathcal{U}} \leq \varepsilon \quad (4)$$

is satisfied for a prescribed tolerance $\varepsilon > 0$.

Finite evaluations. Since each solver $\mathcal{S}_c$ is called exactly once per iteration, we denote by $x_m^{(c)} \in \mathcal{X}_c$ the input passed to solver $\mathcal{S}_c$ at iteration m , for $c = 1, 2$. After M iterations, the total number of evaluations of each solver is therefore M . Thus the approximation $u^{(M)}$ is entirely determined by the finite sets of evaluations,

$$\left\{ \mathcal{S}_c(x_m^{(c)}) : m = 1, \dots, M, \ c = 1, 2 \right\}. \quad (5)$$

Finite characterization of the coupled approximation. The computed approximation $u^{(M)}$ depends only on these finitely many evaluations. Moreover, by Proposition 1, $u^{(M)} \rightarrow u^*$ as $M \rightarrow \infty$, so the fixed point u^* is approximated by iterates $u^{(M)}$, each of which involves only finitely many evaluations.

Proof of Eq. (5). We proceed by induction on m . At $m = 0$, $u^{(0)}$ is prescribed. The first update is

$$u^{(1)} = \Gamma_2 \left(\mathcal{S}_2 \left(\Gamma_{12} \left(\mathcal{S}_1 \left(\Gamma_1 \left(u^{(0)} \right) \right) \right) \right) \right), \quad (6)$$

which requires one evaluation of $\mathcal{S}_1$ at $x_1^{(1)} = \Gamma_1(u^{(0)}) \in \mathcal{X}_1$ and one evaluation of $\mathcal{S}_2$ at $x_1^{(2)} = \Gamma_{12}(\mathcal{S}_1(x_1^{(1)})) \in \mathcal{X}_2$.

Suppose now that $u^{(m)}$ depends only on the finite collection of evaluations

$$\left\{ \mathcal{S}_1(x_{m'}^{(1)}) : m' = 1, \dots, m \right\}, \quad \left\{ \mathcal{S}_2(x_{m'}^{(2)}) : m' = 1, \dots, m \right\}. \quad (7)$$

Then computing $u^{(m+1)} = \mathcal{T}(u^{(m)})$ requires evaluating $\mathcal{S}_1$ at a new input $x_{m+1}^{(1)}$ determined from $u^{(m)}$, and subsequently $\mathcal{S}_2$ at a new input $x_{m+1}^{(2)} = \Gamma_{12}(\mathcal{S}_1(x_{m+1}^{(1)}))$. Thus $u^{(m+1)}$ depends only on the previous evaluations together with these two new ones.

By induction, for each m the interface variable $u^{(m)}$ is determined by finitely many evaluations of the form $\mathcal{S}_c(x_m^{(c)})$. Since the algorithm terminates after $M < \infty$ iterations, $u^{(M)}$ depends only on finitely many such evaluations.

The coupled solution is well defined as a fixed point of the operator $\mathcal{T}$, and its approximation reduces to a finite set of evaluations of the solvers at points $x_m^{(c)}$. This finite characterization is the foundation for introducing surrogate models in place of $\mathcal{S}_c$, and for developing uncertainty quantification in the next subsection.

2.2 Uncertainty Quantification in Coupled Gaussian Process Surrogates

In the context of coupling physical codes, replacing computationally expensive numerical solvers with surrogate models—typically Gaussian processes—offers significant performance gains. In general cases, if the surrogates show strong predictive capabilities, it can justify the use of their mean predictions to replace the original codes for uncertainty quantification methodology. However, these Gaussian process surrogates provide a model-based representation of surrogate uncertainty through the posterior variance, conditional on the chosen prior, kernel family, and hyperparameter estimation. Assuming that the mean prediction alone is sufficient in coupled multiphysics systems via Gaussian process neglects this uncertainty, which may propagate through the coupled simulation workflow and impact the accuracy and robustness of the final results. To study the propagation of this uncertainty in this context we describe a Monte Carlo strategy below.

Notation. For $c = 1, 2$, $f^{(c)}$ denotes the surrogate of solver $\mathcal{S}_c$ with input space $\mathcal{X}^{(c)} \subset \mathbb{R}^{d_c}$. The outputs of $f^{(c)}$ may be vector-valued in $\mathbb{R}^{D_c}$, in which case the covariance kernel is matrix-valued and $f^{(c)}$ is a (vector-valued) multioutput Gaussian process [5,6]. Let $\mathbf{y}$ denote the vector of coupled outputs returned by the fixed-point algorithm (it may be equal to the interface variable u or be a function of it; we keep the notation $\mathbf{y}$ to stress vector-valued quantities).

Throughout, we assume an isotopic design for each surrogate c (the same input set is used for all output coordinates).

Correlation structures for multioutput Gaussian processes. A multioutput (vector-valued) Gaussian process (MOGP) is a random field $x \in \mathcal{X}^{(c)} \mapsto f^{(c)}(x) \in \mathbb{R}^{D_c}$ with mean $\mu^{(c)}(x) \in \mathbb{R}^{D_c}$. Writing

$$f^{(c)}(x) = \left(f_1^{(c)}(x), \dots, f_{D_c}^{(c)}(x) \right), \quad (8)$$

the process $\left\{ f_\ell^{(c)}(x) \right\}_{\ell, x}$ is Gaussian if and only if, for any $n \geq 1$, any inputs $x_1, \dots, x_n \in \mathcal{X}^{(c)}$, and any indices $\ell_1, \dots, \ell_n \in \{1, \dots, D_c\}$, the linear combination $\sum_{j=1}^n \lambda_j f_{\ell_j}^{(c)}(x_j)$ has a univariate normal distribution. Defining the cross-covariance functions

$$\kappa_{\ell, \ell'}^{(c)}(x, x') := \text{Cov}\left(f_\ell^{(c)}(x), f_{\ell'}^{(c)}(x')\right), \quad (9)$$

we can collect them into a matrix-valued covariance kernel:

$$\kappa^{(c)} : \mathcal{X}^{(c)} \times \mathcal{X}^{(c)} \rightarrow \mathbb{R}^{D_c \times D_c}, \quad \kappa^{(c)}(x, x') = \left[\kappa_{\ell, \ell'}^{(c)}(x, x') \right]_{\ell, \ell'=1}^{D_c}. \quad (10)$$

For any finite design set $X = \{x_i\}_{i=1}^n$, stacking $\left(f^{(c)}(x_1)^\top, \dots, f^{(c)}(x_n)^\top \right)^\top \in \mathbb{R}^{nD_c}$ yields a block covariance matrix $K^{(c)}(X, X)$ with entries

$$K_{(i, \ell), (j, \ell')}^{(c)} = \kappa_{\ell, \ell'}^{(c)}(x_i, x_j). \quad (11)$$

The matrix $K^{(c)}(X, X)$ is symmetric and positive semidefinite (PSD), i.e.,

$$v^\top K^{(c)}(X, X) v \geq 0 \quad \text{for all } v \in \mathbb{R}^{nD_c}. \quad (12)$$

Thus, specifying a MOGP amounts to choosing cross-covariance functions $\kappa_{\ell, \ell'}^{(c)}$ such that the associated block covariance matrices are PSD for all finite sets of inputs. Two standard choices are

- *Independent outputs.* Take a diagonal kernel,

$$\kappa^{(c)}(x, x') = \text{diag}\left(\kappa_1^{(c)}(x, x'), \dots, \kappa_{D_c}^{(c)}(x, x')\right), \quad (13)$$

so the D_c coordinates are a priori uncorrelated, i.e., $\kappa_{\ell, \ell'}^{(c)}(x, x') = 0$ for $\ell \neq \ell'$. For an isotopic design (same inputs for all outputs), the joint covariance is block diagonal:

$$K^{(c)}(X, X) = \text{diag}\left(K_1^{(c)}, \dots, K_{D_c}^{(c)}\right) \quad \text{with} \quad \left[K_\ell^{(c)} \right]_{ij} = \kappa_\ell^{(c)}(x_i, x_j). \quad (14)$$

- *Structured correlations [linear model of coregionalization (LMC)]*. Let $Q \geq 1$ and set

$$\kappa^{(c)}(x, x') = \sum_{q=1}^Q B_q^{(c)} \kappa_q^{(c)}(x, x'), \quad (15)$$

where each $B_q^{(c)} \in \mathbb{R}^{D_c \times D_c}$ is a PSD coregionalization matrix and $\kappa_q^{(c)}$ is a scalar covariance on $\mathcal{X}^{(c)}$ [5,7]. Equivalently, writing $B_q^{(c)} = A_q^{(c)} \left(A_q^{(c)} \right)^\top$ shows that outputs are linear mixtures of Q independent latent Gaussian processes with kernels $\kappa_q^{(c)}$, which guarantees PSD of the full block covariance. In the isotopic case, the joint covariance over all outputs and all inputs reads

$$K^{(c)}(X, X) = \sum_{q=1}^Q B_q^{(c)} \otimes K_q^{(c)}, \quad \left[K_q^{(c)} \right]_{ij} = \kappa_q^{(c)}(x_i, x_j), \quad (16)$$

making explicit the separation between output-output dependence (via $B_q^{(c)}$) and input-input dependence (via $K_q^{(c)}$). In our application we adopt independent outputs (diagonal $B_q^{(c)}$), but we emphasize that the methodology accommodates LMC without further changes.

Surrogate models. Each surrogate $f^{(c)}$ is trained on evaluations of $\mathcal{S}_c$ from a space-filling design (e.g., Latin hypercube or low-discrepancy sampling [8]). *A priori* we have

$$f^{(c)}(x) \sim \mathcal{GP}\left(\mu^{(c)}(x), \kappa^{(c)}(x, x')\right),$$

where $\mu^{(c)}(x) \in \mathbb{R}^{D_c}$ and, in the multioutput case, $\kappa^{(c)}(x, x') \in \mathbb{R}^{D_c \times D_c}$ is matrix-valued. If $\bar{\mu}^{(c)}$ denotes the posterior mean and $\bar{\kappa}^{(c)}$ denotes the posterior kernel, then for each x we have $\mathbf{0} \preceq \bar{\kappa}^{(c)}(x, x) \preceq \kappa^{(c)}(x, x)$, with $\preceq$ the Loewner order on symmetric PSD matrices. In the independent-output case, this reduces componentwise to $0 \leq \bar{\kappa}_\ell^{(c)}(x, x) \leq \kappa_\ell^{(c)}(x, x)$.

Monte Carlo framework for uncertainty quantification of the coupled Gaussian process. Uncertainty quantification is performed by a standard Monte Carlo procedure. An outer loop generates i.i.d. realizations of the coupled solution: for each iteration $j = 1, \dots, N$, independent realizations $\{f^{(c,j)}\}_c$ of the Gaussian process surrogates are drawn, the coupling algorithm is executed with these surrogates starting from a prescribed initial interface state $u^{(0)}$, and the corresponding coupled solution $\mathbf{y}^{(j)}$ is obtained. We denote by $M^{(j)}$ the number of fixed-point iterations performed for simulation j before the stopping criterion $\|u^{(m,j)} - u^{(m-1,j)}\|_{\mathcal{U}} \leq \varepsilon$ is met. In general, $M^{(j)}$ depends on the Monte Carlo realization j . The sample $\{\mathbf{y}^{(j)}\}_{j=1}^N$ is then used to compute empirical statistics such as mean, variance, and empirical uncertainty intervals (from Monte Carlo samples). Unless stated otherwise, any 95% interval reported from Monte Carlo outputs refers to a central empirical quantile interval, and should not be interpreted as a frequentist confidence interval.

Three approaches can in principle be used to generate the surrogate realizations inside the coupling loop:

- *Ideal but infeasible.* For each Monte Carlo replication j , a full realization $f^{(c,j)}$ of the Gaussian process surrogate $f^{(c)}$ over the entire input space $\mathcal{X}_c$ is drawn. The coupling algorithm is then run with $f^{(c,j)}$. This would provide exact i.i.d. samples of the coupled solution, but it is computationally impossible since $\mathcal{X}_c$ is infinite.
- *Rigorous but computationally heavy.* For each Monte Carlo replication j , the coupled solver is run with a fully stochastic surrogate family $\{f^{(c,j)}\}_c$ that is sampled consistently with the posterior Gaussian process laws. More precisely, at iteration m of replication j , when the coupling algorithm queries surrogate c at a new input $x_m^{(c,j)} \in \mathcal{X}^{(c)}$, let

$$X_{m-1}^{(c,j)} := \left(x_1^{(c,j)}, \dots, x_{m-1}^{(c,j)} \right), \quad Z_{m-1}^{(c,j)} := \left(f^{(c,j)}(x_1^{(c,j)}), \dots, f^{(c,j)}(x_{m-1}^{(c,j)}) \right),$$

denote the locations and values already sampled of the surrogate c along the trajectory of replication j . By consistency of the Gaussian process posterior, the conditional distribution of the vector $f^{(c)}(x_m^{(c,j)})$ given $f^{(c)}(X_{m-1}^{(c,j)}) = Z_{m-1}^{(c,j)}$ is multivariate Gaussian with explicitly computable mean and variance. The value $f^{(c,j)}(x_m^{(c,j)})$ is then drawn from this conditional law and appended to $Z_{m-1}^{(c,j)}$, and the coupled fixed-point iteration is updated using these sampled values. This construction ensures that, for each j , the random trajectory $\left\{f^{(c,j)}(x_m^{(c,j)})\right\}_{m,c}$ is exactly distributed according to the joint Gaussian process posterior along the (random) coupling path, so that the outputs $\mathbf{y}^{(j)}$ are rigorous samples of the coupled Gaussian process-based model. However, at each new query the conditional Gaussian formulas involve covariance matrices of size $m \times m$ built at locations $\{x_1^{(c,j)}, \dots, x_m^{(c,j)}\}$ which depend on both m and j , so that repeated inversions of growing matrices are required. The resulting cost is very high, but the procedure is fully coherent with the Gaussian process prior and posterior, and can serve, at least on moderate grids, as a reference method for the example of Section 2.4.

- *Proposed method.* The proposed method is a practical compromise described below, which fixes the input sets after one deterministic coupling run and reuses them for all MC iterations, thereby greatly reducing the cost while retaining a nontrivial description of uncertainty.

We first fix an initial interface state $u^{(0)}$ and run the deterministic coupling using the Gaussian process posterior means $\bar{\mu}^{(c)}$. This fixed-point iteration produces a sequence $(u^{(m)})_{m \geq 0}$ and, by the stopping criterion, converges after M iterations. The points called during this deterministic computation at the m th iteration are denoted

$$x_m^{(c)} \in \mathcal{X}^{(c)}, \quad m = 1, \dots, M, \quad c = 1, 2.$$

These same points $\{x_m^{(c)}\}_{m=1}^M$ are reused for all Monte Carlo replications, which keeps the sampling cost low and fixes a unique reference path starting from $u^{(0)}$.

For each Monte Carlo replication $j = 1, \dots, N$ and for each c , we draw the Gaussian vector

$$\left(\tilde{f}^{(c,j)}(x_m^{(c)})\right)_{m=1}^M \sim \mathcal{N}\left(\bar{\mu}^{(c)}(x^{(c)}), \bar{\mathbf{K}}^{(c)}\right),$$

that is, a realization of the vector $\left(f^{(c)}(x_m^{(c)})\right)_{m=1}^M$. Here $\bar{\mu}^{(c)}(x^{(c)}) = \left(\bar{\mu}^{(c)}(x_m^{(c)})\right)_{m=1}^M$ and $\bar{\mathbf{K}}^{(c)}$ is the block covariance matrix whose (m, m') block is $\bar{\kappa}^{(c)}(x_m^{(c)}, x_{m'}^{(c)}) \in \mathbb{R}^{D_c \times D_c}$.

We then define the errors,

$$\delta_m^{(c,j)} = \tilde{f}^{(c,j)}(x_m^{(c)}) - \bar{\mu}^{(c)}(x_m^{(c)}), \quad m = 1, \dots, M, \quad (17)$$

and, at simulation j and iteration m , we use the perturbed surrogate:

$$\hat{f}^{(c,m,j)}(x) = \bar{\mu}^{(c)}(x) + \delta_m^{(c,j)}. \quad (18)$$

Running the coupling with $\{\hat{f}^{(c,m,j)}\}$, starting from the same initial state $u^{(0)}$, yields a vector output $\mathbf{y}^{(j)}$. We can anticipate that the i.i.d. family $\{\mathbf{y}^{(j)}\}_{j=1}^N$ obtained by this procedure has the same distribution as the one of the rigorous but computationally heavy procedures described above, provided the sampled points $x_m^{(c,j)}$ stay close enough (relative to the correlation radius) to $x_m^{(c)}$. This will be quantified in the next subsection.

From the family $\{\mathbf{y}^{(j)}\}_{j=1}^N$, we compute

$$\hat{\mathbf{y}} = \frac{1}{N} \sum_{j=1}^N \mathbf{y}^{(j)}, \quad \widehat{\text{Cov}}[\mathbf{y}] = \frac{1}{N-1} \sum_{j=1}^N \left(\mathbf{y}^{(j)} - \hat{\mathbf{y}}\right) \left(\mathbf{y}^{(j)} - \hat{\mathbf{y}}\right)^\top. \quad (19)$$

When needed, componentwise summaries [mean, variance, and empirical intervals / Gaussian intervals (depending on the estimator used)] are obtained from $\hat{\mathbf{y}}$ and the diagonal of $\widehat{\text{Cov}}[\mathbf{y}]$; these quantities describe the uncertainty on the coupled outputs $\mathbf{y}$.

2.3 Theoretical Control of Predictive Variance

Objective and strategy. The purpose of this subsection is to provide a theoretical justification for the Monte Carlo perturbation scheme of Section 2.2, which is used to propagate Gaussian process uncertainty through coupled surrogate models. More precisely, the aim is to show that the random coupled outputs produced by the practical algorithm are concentrated around the deterministic fixed point obtained with the Gaussian process mean surrogate, and that their dispersion is explicitly controlled by the Gaussian process posterior variances via finite-sample bounds on the effect of surrogate uncertainty on the coupled solution.

The argument proceeds in three steps: First, relate the posterior variances of scalar and multioutput Gaussian processes to the fill distance of the design points, by exploiting the kernel-based approximation results of Kanagawa et al. [9], which connect Gaussian process posterior variance to worst-case interpolation error in the associated reproducing kernel Hilbert space (RKHS). In particular, Matérn kernels induce Sobolev RKHSs, so that the variance can be controlled through the fill distance of a representative input set. Second, establish a Lipschitz stability result for the coupled fixed point with respect to perturbations of the surrogates. And finally, combine these two ingredients to obtain a finite-sample, high-probability bound (Proposition 4), which constitutes the main theoretical result of this section and quantifies the uncertainty of the Monte Carlo outputs.

Number of coupled solvers. In Sections 2.1 and 2.2, the presentation focused on the case of two coupled solvers ($C = 2$), which corresponds to the application considered in this work. In the present subsection, the notation is slightly generalized and the analysis is carried out for a generic finite number $C \geq 2$ of coupled surrogates; the case $C = 2$ is recovered by specializing the formulas to two components.

Vector-valued convention. In coherence with the previous subsection, outputs (including $\mathbf{y}$) are vectors and Gaussian process kernels may be matrix-valued. Accordingly, covariance inequalities are to be understood in Loewner (PSD) order and pointwise deviations are measured with the Euclidean norm. When we write a sup-norm for vector-valued functions $g : \mathcal{X} \rightarrow \mathbb{R}^D$, we mean the sup-Euclidean norm $\|g\|_{\infty,2} := \sup_{x \in \mathcal{X}} \|g(x)\|_2$.

A first ingredient is a geometric quantity that characterizes how well a design X covers the input domain and will later control the decay of the Gaussian process posterior variance.

Definition 5 (Fill distance). The fill distance of a point set $X \subset \mathcal{X}$ with respect to a domain $\mathcal{X} \subset \mathbb{R}^d$ is defined as

$$h_{X,\mathcal{X}} := \sup_{x \in \mathcal{X}} \min_{x_i \in X} \|x - x_i\|, \quad (20)$$

which measures the largest gap between any point in the domain $\mathcal{X}$ and its nearest representative in X .

The following result recalls the standard multioutput Gaussian process regression formulas and fixes notation for the block covariance matrices used in the subsequent bounds.

Theorem 1 (Posterior mean and covariance for multioutput Gaussian process regression [5,6]). *Let $f : \mathcal{X} \rightarrow \mathbb{R}^D$ have a Gaussian process prior $f \sim \mathcal{GP}(\mu, \kappa)$ with mean function $\mu : \mathcal{X} \rightarrow \mathbb{R}^D$ and matrix-valued kernel $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{D \times D}$ (positive semidefinite in the sense of Loewner order). Given inputs $X = (x_1, \dots, x_n) \in \mathcal{X}^n$ and noisy vector-valued observations $\mathbf{Z} = (z(x_1)^\top, \dots, z(x_n)^\top)^\top \in \mathbb{R}^{nD}$, $\mathbf{Z} = f(X) + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_{nD})$, the posterior is $f | \mathbf{Z} \sim \mathcal{GP}(\bar{\mu}, \bar{\kappa})$ with*

$$\bar{\mu}(x) = \mu(x) + K_{xX} (K_{XX} + \sigma^2 I_{nD})^{-1} (\mathbf{Z} - \mu_X), \quad (21)$$

$$\bar{\kappa}(x, x') = \kappa(x, x') - K_{xX} (K_{XX} + \sigma^2 I_{nD})^{-1} K_{Xx'}, \quad (22)$$

where $\mu_X = (\mu(x_1)^\top, \dots, \mu(x_n)^\top)^\top \in \mathbb{R}^{nD}$, $K_{XX} \in \mathbb{R}^{nD \times nD}$ is the block covariance with (i, j) -block $[K_{XX}]_{ij} = \kappa(x_i, x_j) \in \mathbb{R}^{D \times D}$, $K_{xX} = [\kappa(x, x_1) \ \dots \ \kappa(x, x_n)] \in \mathbb{R}^{D \times nD}$, and $K_{Xx'} = K_{x'X}^\top$. In the noise-free case, replace $K_{XX} + \sigma^2 I_{nD}$ by K_{XX} .

In order to relate Gaussian process posterior variances to approximation properties in Sobolev spaces, a classical characterization of Matérn kernels in terms of their associated RKHS is recalled next.

Proposition 2 (RKHS of Matérn kernels [9]). *Let $\kappa_{\alpha,h}$ be the Matérn kernel on $\mathcal{X} \subset \mathbb{R}^{d_c}$ with smoothness $\alpha > 0$ and scaling $h > 0$. If $s = \alpha + d_c/2$ is an integer, the RKHS $\mathcal{H}_{\kappa_{\alpha,h}}$ is norm-equivalent to the Sobolev space $W_2^s(\mathcal{X})$:*

$$W_2^s(\mathcal{X}) = \left\{ f \in L_2(\mathcal{X}) : \sum_{|\beta| \leq s} \|D^\beta f\|_{L_2(\mathcal{X})}^2 < \infty \right\}.$$

Thus $\mathcal{H}_{\kappa_{\alpha,h}} = W_2^s(\mathcal{X})$ as sets of functions, and there exist constants $c_1, c_2 > 0$ with $c_1 \|f\|_{W_2^s(\mathcal{X})} \leq \|f\|_{\mathcal{H}_{\kappa_{\alpha,h}}} \leq c_2 \|f\|_{W_2^s(\mathcal{X})}, \forall f \in \mathcal{H}_{\kappa_{\alpha,h}}$.

Under these regularity assumptions, the next result of Kanagawa et al. [9] links the posterior variance of a scalar Gaussian process to the local fill distance of the design.

Theorem 2 (Scalar fill-distance bound, Kanagawa et al. [9]). *Let κ be a scalar kernel on $\mathbb{R}^{d_c}$ whose RKHS is norm-equivalent to the Sobolev space $W_2^s(\mathbb{R}^{d_c})$ with $s > d_c/2$. For $\rho > 0$ and any finite set of points $X = \{x_1, \dots, x_n\} \subset \mathbb{R}^{d_c}$, define the (local) fill distance at $x \in \mathbb{R}^{d_c}$ by*

$$h_{\rho,X}(x) := \sup_{\substack{x' \in \mathbb{R}^{d_c} \\ \|x' - x\| \leq \rho}} \min_{x_i \in X} \|x' - x_i\|. \quad (23)$$

Then, for any $\rho > 0$, there exist $h_0 > 0$ and $C > 0$ such that for all $x \in \mathbb{R}^{d_c}$ and all point sets X with $h_{\rho,X}(x) \leq h_0$, the posterior variance satisfies

$$\bar{\kappa}(x, x) \leq C h_{\rho,X}(x)^{2s-d_c}. \quad (24)$$

This scalar variance bound can be extended to the multioutput setting by exploiting the LMC representation of the matrix-valued kernel.

Theorem 3 (Using LMC representation a lift of Kanagawa's scalar bound to multioutput Gaussian processes). *Let the matrix-valued kernel be of LMC form $\kappa(x, x') = \sum_{q=1}^Q B_q \kappa_q(x, x')$, with $B_q \succeq \mathbf{0}$ and each scalar κ_q satisfying Theorem 2 with parameters $(s_q, C_q, h_{0,q})$. Then, for every x with $h_{\rho,X}(x) \leq \min_q h_{0,q}$,*

$$\bar{\kappa}(x, x) \preceq \left(\sum_{q=1}^Q C_q h_{\rho,X}(x)^{2s_q-d_c} \right) \left(\sum_{q=1}^Q B_q \right). \quad (25)$$

Proof. Fix $x \in \mathcal{X}$ and $v \in \mathbb{R}^D$. The scalar projection $g_v(x) := v^\top f(x)$ is a Gaussian process with kernel $\kappa_v(x, x') = v^\top \kappa(x, x') v = \sum_{q=1}^Q a_q \kappa_q(x, x')$, where $a_q := v^\top B_q v \geq 0$ [10,11]. By Gaussian conditioning monotonicity

$$v^\top \bar{\kappa}(x, x) v = \text{Var}(v^\top f(x) \mid f(X)) \leq \text{Var}(v^\top f(x) \mid v^\top f(X)) = \bar{\kappa}_v(x, x). \quad (26)$$

Let φ_q be feature maps for the κ_q so that $\kappa_q(x, x') = \langle \varphi_q(x), \varphi_q(x') \rangle_{\mathcal{H}_q}$. A feature map for κ_v is $\varphi = (\sqrt{a_q} \varphi_q)_q$ into the Hilbert direct sum $\bigoplus_q \mathcal{H}_q$ [11,12], since

$$\langle \varphi(x), \varphi(x') \rangle_{\bigoplus_q \mathcal{H}_q} = \sum_{q=1}^Q \langle \sqrt{a_q} \varphi_q(x), \sqrt{a_q} \varphi_q(x') \rangle_{\mathcal{H}_q} = \sum_{q=1}^Q a_q \kappa_q(x, x') = \kappa_v(x, x').$$

For this scalar Gaussian process with kernel κ_v , the posterior variance admits the usual RKHS interpretation as the squared distance from $\varphi(x)$ to the span $S := \text{span}\{\varphi(x_i)\}$ in feature space [9]:

$$\sqrt{\bar{\kappa}_v(x, x)} = \text{dist}(\varphi(x), S) \leq \sum_{q=1}^Q \sqrt{a_q} \text{dist}(\varphi_q(x), S_q) = \sum_{q=1}^Q \sqrt{a_q} \sqrt{\bar{\kappa}_q(x, x)},$$

where $S_q := \text{span}\{\varphi_q(x_i)\}$ and we used the triangle inequality in the direct sum. Squaring and applying Cauchy-Schwarz yields

$$\overline{\kappa}_v(x, x) \leq \left(\sum_{q=1}^Q \sqrt{a_q} \sqrt{\overline{\kappa}_q(x, x)} \right)^2 \leq \left(\sum_{q=1}^Q a_q \right) \left(\sum_{q=1}^Q \overline{\kappa}_q(x, x) \right).$$

By Theorem 2, $\overline{\kappa}_q(x, x) \leq C_q h_{\rho, X}(x)^{2s_q - d_c}$ whenever $h_{\rho, X}(x) \leq h_{0, q}$. Therefore,

$$\overline{\kappa}_v(x, x) \leq \left(\sum_{q=1}^Q a_q \right) \left(\sum_{q=1}^Q C_q h_{\rho, X}(x)^{2s_q - d_c} \right) = v^\top \left(\sum_{q=1}^Q B_q \right) v \cdot \left(\sum_{q=1}^Q C_q h_{\rho, X}(x)^{2s_q - d_c} \right).$$

Combining with Eq. (26), for all v , gives

$$v^\top \overline{\kappa}(x, x) v \leq v^\top \left(\sum_{q=1}^Q B_q \right) v \cdot \left(\sum_{q=1}^Q C_q h_{\rho, X}(x)^{2s_q - d_c} \right),$$

which is equivalent to the stated Loewner bound. $\square$

Remark 1 (Interpretation of Theorem 3). Theorem 3 states that, for a multioutput Gaussian process with an LMC kernel, the posterior covariance at any point x is bounded above by a scalar factor depending on the fill distance $h_{\rho, X}(x)$, times a fixed output-output matrix $\sum_q B_q$. In other words, as the design becomes denser (the fill distance decreases), the posterior uncertainty on all outputs decays at a rate determined by the smoothness parameters s_q . This result lifts the scalar bound of Kanagawa et al. [9] to the multioutput setting relevant for our coupled surrogates.

After controlling the posterior variances of the individual surrogates, the second ingredient concerns the stability of the coupled fixed point with respect to perturbations of these surrogates.

Coupling operator and solution map. Let $\mathcal{U}$ be the interface space and $H : \mathcal{U} \rightarrow \mathcal{Y}$ the post-processing map to outputs. Fix an integer $C \geq 2$, and let $f = (f^{(1)}, \dots, f^{(C)})$ be a C -tuple of surrogates, with fixed interface operators

$$\Gamma_1 : \mathcal{U} \rightarrow \mathcal{X}_1, \Gamma_{r, r+1} : \mathcal{Y}_r \rightarrow \mathcal{X}_{r+1}, \quad r = 1, \dots, C-1,$$

and $\Gamma_C : \mathcal{Y}_C \rightarrow \mathcal{U}$; define the coupling operator

$$\mathcal{T}_f := \Gamma_C \circ f^{(C)} \circ \Gamma_{C-1, C} \circ \dots \circ f^{(1)} \circ \Gamma_1 : \mathcal{U} \rightarrow \mathcal{U}.$$

We then define the solution map $\Phi : \mathcal{A} \rightarrow \mathcal{Y}$, $\Phi(f) := H(u^*(f))$, $u^*(f) \in \text{Fix}(\mathcal{T}_f)$, where $\mathcal{A}$ is an admissible class of surrogate tuples on which the fixed point exists and is unique.

This stability analysis is based on a uniform contraction property of the coupling operator together with a Lipschitz condition on the post-processing map H .

Assumption 1 (Uniform contraction). There exists $\rho \in (0, 1)$ such that for all $f \in \mathcal{A}$, the operator $\mathcal{T}_f$ is a contraction on a Banach space $(\mathcal{U}, \|\cdot\|_{\mathcal{U}})$ with modulus at most ρ , and H is L_H -Lipschitz. Moreover, there are constants $L_1, \dots, L_C$ (depending only on the interface operators and uniform Lipschitz bounds of the admissible surrogates) such that for any $f, g \in \mathcal{A}$ and any $u \in \mathcal{U}$,

$$\|\mathcal{T}_f(u) - \mathcal{T}_g(u)\|_{\mathcal{U}} \leq \sum_{c=1}^C L_c \|f^{(c)} - g^{(c)}\|_{\infty, 2}. \quad (27)$$

Under the above contraction assumption, the following proposition shows that the coupled solution depends on Lipschitz continuously on the underlying surrogates.

Proposition 3 (Lipschitz stability of the solution map). *Under Assumption 1, the map Φ is Lipschitz on $\mathcal{A}$:*

$$\|\Phi(f) - \Phi(g)\|_{\mathcal{Y}} \leq \frac{L_H}{1 - \rho} \sum_{c=1}^C L_c \|f^{(c)} - g^{(c)}\|_{\infty,2} \quad \text{for all } f, g \in \mathcal{A}. \quad (28)$$

Proof. For any $f, g \in \mathcal{A}$,

$$\|u^*(f) - u^*(g)\|_{\mathcal{U}} = \|\mathcal{T}_f(u^*(f)) - \mathcal{T}_g(u^*(g))\|_{\mathcal{U}} \leq \|\mathcal{T}_f(u^*(f)) - \mathcal{T}_f(u^*(g))\|_{\mathcal{U}} + \|\mathcal{T}_f(u^*(g)) - \mathcal{T}_g(u^*(g))\|_{\mathcal{U}}.$$

By contraction, the first term is $\leq \rho \|u^*(f) - u^*(g)\|_{\mathcal{U}}$; by Eq. (27), the second term is $\leq \sum_{c=1}^C L_c \|f^{(c)} - g^{(c)}\|_{\infty,2}$. Therefore, $\|u^*(f) - u^*(g)\|_{\mathcal{U}} \leq (1/(1 - \rho)) \sum_{c=1}^C L_c \|f^{(c)} - g^{(c)}\|_{\infty,2}$.

Finally, we apply the L_H -Lipschitzness of H and we obtain

$$\|\Phi(f) - \Phi(g)\|_{\mathcal{Y}} = \|H(u^*(f)) - H(u^*(g))\|_{\mathcal{Y}} \leq L_H \|u^*(f) - u^*(g)\|_{\mathcal{U}} \leq \frac{L_H}{1 - \rho} \sum_{c=1}^C L_c \|f^{(c)} - g^{(c)}\|_{\infty,2}.$$

□

Remark 2 (Interpretation of Proposition 3). Proposition 3 shows that the final coupled output $\Phi(f)$ is Lipschitz-continuous with respect to the surrogates. Small uniform perturbations of the surrogates $f^{(c)}$ (measured in $\|\cdot\|_{\infty,2}$) produce controlled changes in the fixed-point solution and thus in the observable output. The prefactor $(L_H/(1 - \rho)) \sum_c L_c$ can be interpreted as a global “sensitivity” constant of the coupled system with respect to surrogate errors.

The next corollary applies this general Lipschitz bound to the specific form of perturbations used in the Monte Carlo scheme, namely, constant (in x) offsets of the Gaussian process means at each iteration.

Corollary 1 (Constant-offset perturbations used in our MC scheme). *In our Monte Carlo replication (m, j) , we use perturbed surrogates of the form $\hat{f}^{(c)}(x) = \bar{\mu}^{(c)}(x) + \delta_{m,j}^{(c)}$ with $\delta_{m,j}^{(c)} \in \mathbb{R}^{D_c}$ independent of x . Then $\|\hat{f}^{(c)} - \bar{\mu}^{(c)}\|_{\infty,2} = \|\delta_{m,j}^{(c)}\|_2$ and the contraction modulus ρ of $\mathcal{T}_f$ is unchanged. Hence*

$$\|\Phi(\hat{f}) - \Phi(\bar{\mu})\|_{\mathcal{Y}} \leq \frac{L_H}{1 - \rho} \sum_{c=1}^C L_c \|\delta_{m,j}^{(c)}\|_2. \quad (29)$$

Note that Φ and $\mathcal{T}_f$ are different. $\mathcal{T}_f : \mathcal{U} \rightarrow \mathcal{U}$ is the one-step iteration map (depends on f and acts on the interface variable u), whereas Φ maps the surrogate tuple f to the final observable $y = H(u^*(f))$ obtained at the fixed point of $\mathcal{T}_f$.

Remark 3 (Connection to the Monte Carlo scheme). Corollary 1 shows that the specific form of perturbation used in our Monte Carlo scheme (constant offsets with respect to x at each iteration) preserves the contraction property of the coupling operator and turns the Lipschitz bound of Proposition 3 into a simple bound in terms of the Euclidean norms $\|\delta_{m,j}^{(c)}\|_2$. This makes it possible to link the Gaussian process-based perturbations to a global control on the deviation of the coupled solution.

Combining the Gaussian process posterior variance bounds from the first part of this subsection with the stability result above leads to the following finite-sample, high-probability control on the Monte Carlo outputs, which forms the main theoretical result of this section.

Proposition 4 (Finite-sample high-probability uncertainty quantification bound). *For each $c \in \{1, \dots, C\}$, let the surrogate be a multioutput Gaussian process with the matrix-valued kernel of LMC form $\kappa^{(c)}(x, x') = \sum_{q=1}^{Q_c} B_q^{(c)} \kappa_q^{(c)}(x, x')$ where $B_q^{(c)} \succeq \mathbf{0}$, and each scalar latent kernel $\kappa_q^{(c)}$ satisfies the scalar fill-distance bound (Theorem 2) with parameters $(s_{c,q}, C_{c,q}, h_{0,c,q})$, with $s_{c,q} > d_c/2$. Let $h^{(c)}$ denote the (local) fill distance of the design used for $f^{(c)}$, and assume $h^{(c)} \leq h_{0,c} := \min_q h_{0,c,q}$.*

In our Monte Carlo scheme, the perturbed surrogates are constant offsets,

$$\hat{f}^{(c)}(x) = \bar{\mu}^{(c)}(x) + \delta^{(c)}, \quad \delta^{(c)} \sim \mathcal{N}(0, \Sigma_c) \text{ in } \mathbb{R}^{D_c}, \text{ independent of } x,$$

with the covariance matrices dominated (in spectral norm) by the posterior marginals: $\lambda_{\max}(\Sigma_c) \leq \sup_{x \in \mathcal{X}} \lambda_{\max}(\bar{\kappa}^{(c)}(x, x))$. Then, for any $\beta \in (0, 1)$, with probability at least $1 - \beta$,

$$\begin{aligned} \left\| \Phi\left(\left\{\bar{\mu}^{(c)} + \delta^{(c)}\right\}_c\right) - \Phi\left(\left\{\bar{\mu}^{(c)}\right\}_c\right) \right\|_Y &\leq \frac{L_H}{1-\rho} \sum_{c=1}^C L_c \sqrt{2 D_c \lambda_{\max}\left(\sum_{q=1}^{Q_c} B_q^{(c)}\right) \sum_{q=1}^{Q_c} \mathcal{C}_{c,q}(h^{(c)})^{2s_{c,q}-d_c}} \\ &\quad \times \sqrt{\log\left(\frac{2CD_c}{\beta}\right)}. \end{aligned}$$

Proof. By Theorem 3, for all x , $\bar{\kappa}^{(c)}(x, x) \preceq \left(\sum_{q=1}^{Q_c} \mathcal{C}_{c,q}(h^{(c)})^{2s_{c,q}-d_c}\right) \left(\sum_{q=1}^{Q_c} B_q^{(c)}\right)$. Taking the largest eigenvalues and using the homogeneity of $\lambda_{\max}$ for non-negative scalars gives

$$\sup_{x \in \mathcal{X}} \lambda_{\max}(\bar{\kappa}^{(c)}(x, x)) \leq \lambda_{\max}\left(\sum_{q=1}^{Q_c} B_q^{(c)}\right) \sum_{q=1}^{Q_c} \mathcal{C}_{c,q}(h^{(c)})^{2s_{c,q}-d_c}.$$

By the assumption on Σ_c , it follows that

$$\lambda_{\max}(\Sigma_c) \leq \sigma_c^2, \quad \text{where} \quad \sigma_c^2 := \lambda_{\max}\left(\sum_{q=1}^{Q_c} B_q^{(c)}\right) \sum_{q=1}^{Q_c} \mathcal{C}_{c,q}(h^{(c)})^{2s_{c,q}-d_c}.$$

Let $\delta^{(c)} \sim \mathcal{N}(0, \Sigma_c)$ in $\mathbb{R}^{D_c}$. For any $t_c > 0$, using $\|\cdot\|_2 \leq \sqrt{D_c} \|\cdot\|_\infty$, a union bound over coordinates, and the one-dimensional Gaussian tail bound,

$$\begin{aligned} \mathbb{P}\left(\left\|\delta^{(c)}\right\|_2 > t_c\right) &\leq \mathbb{P}\left(\left\|\delta^{(c)}\right\|_\infty > t_c/\sqrt{D_c}\right) \\ &\leq \sum_{\ell=1}^{D_c} \mathbb{P}\left(\left|\delta_\ell^{(c)}\right| > t_c/\sqrt{D_c}\right) \leq 2D_c \exp\left(-\frac{t_c^2}{2D_c \lambda_{\max}(\Sigma_c)}\right) \\ &\leq 2D_c \exp\left(-\frac{t_c^2}{2D_c \sigma_c^2}\right). \end{aligned}$$

With the choice

$$t_c := \sqrt{2D_c} \sigma_c \sqrt{\log\left(\frac{2cD_c}{\beta}\right)} = \sqrt{2 D_c \lambda_{\max}\left(\sum_{q=1}^{Q_c} B_q^{(c)}\right) \sum_{q=1}^{Q_c} \mathcal{C}_{c,q}(h^{(c)})^{2s_{c,q}-d_c}} \sqrt{\log\left(\frac{2cD_c}{\beta}\right)},$$

we get $\mathbb{P}(\|\delta^{(c)}\|_2 > t_c) \leq \beta/c$. A union bound over $c = 1, \dots, C$ yields that, with probability at least $1 - \beta$, simultaneously for all c , $\|\delta^{(c)}\|_2 \leq t_c$. Since the offsets are constant in x , $\|\hat{f}^{(c)} - \bar{\mu}^{(c)}\|_{\infty,2} = \|\delta^{(c)}\|_2$.

By Proposition 3,

$$\left\| \Phi\left(\left\{\hat{f}^{(c)}\right\}_c\right) - \Phi\left(\left\{\bar{\mu}^{(c)}\right\}_c\right) \right\|_Y \leq \frac{L_H}{1-\rho} \sum_{c=1}^C L_c \left\| \hat{f}^{(c)} - \bar{\mu}^{(c)} \right\|_{\infty,2} \leq \frac{L_H}{1-\rho} \sum_{c=1}^C L_c t_c,$$

and substituting the expression of t_c gives the asserted bounds. $\square$

Remark 4 (Interpretation and role of Proposition 4). Proposition 4 is the main theoretical result of this section. It provides a finite-sample, high-probability bound on the deviation between the output obtained by running the coupling algorithm with the mean surrogates $\{\bar{\mu}^{(c)}\}_c$ and the random outputs produced by our practical Monte Carlo scheme based on perturbed surrogates $\{\bar{\mu}^{(c)} + \delta^{(c)}\}_c$. In words, it shows that, with probability at least $1 - \beta$, the random outputs $\Phi(\{\bar{\mu}^{(c)} + \delta^{(c)}\}_c)$ lie in a neighborhood of the deterministic fixed-point output $\Phi(\{\bar{\mu}^{(c)}\}_c)$ whose radius is explicitly controlled by the Lipschitz sensitivity constants of the coupled solver (L_H, ρ, L_c), the posterior variances of the Gaussian process surrogates (through the eigenvalues of $\sum_q B_q^{(c)}$ and the constants $C_{c,q}$), and the fill distances $h^{(c)}$ of the designs used to train the surrogates. As the design becomes denser (so that $h^{(c)} \rightarrow 0$ and the posterior variances decrease), the bound shrinks and the Monte Carlo outputs concentrate around the deterministic fixed point.

2.4 Analytical Validation on a Multidimensional Benchmark Example

This subsection considers an explicit multidimensional setting (3D input / 3D coupled output) in order to assess, in a controlled case, the behavior of Method 2 (trajectory-conditioned sampling) and Method 3 (mean-path constant-offset scheme) introduced in Section 2.2.

Benchmark and reference solution. Let $\mathcal{U} = [0, 1]^3$ and consider two smooth analytical functions $g^{(1)}, g^{(2)} : \mathcal{U} \rightarrow \mathbb{R}^3$ defined componentwise by

$$g^{(1)}(x_1, x_2, x_3) = \begin{pmatrix} 0.30 + 0.08 \cos(\pi x_1) + 0.06x_2^2 + 0.04x_3 + 0.03x_1x_2 \\ 0.34 + 0.07 \sin(\pi x_2) + 0.05x_3^2 - 0.04x_1x_3 + 0.02x_1 \\ 0.32 + 0.06 \cos(\pi x_3) + 0.05x_1^2 + 0.03x_2 + 0.02x_2x_3 \end{pmatrix},$$

$$g^{(2)}(x_1, x_2, x_3) = \begin{pmatrix} 0.36 + 0.07 \sin(\pi x_1) - 0.05x_2 + 0.04x_3^2 + 0.02x_1x_3 \\ 0.31 + 0.08 \cos(\pi x_2) + 0.05x_1^2 + 0.04x_3 - 0.02x_1x_2 \\ 0.35 + 0.06 \sin(\pi x_3) + 0.04x_2^2 - 0.03x_1 + 0.02x_1x_2 \end{pmatrix}.$$

We define the deterministic coupling map by relaxation,

$$\mathcal{T}(x) = \frac{1}{2} \left(g^{(1)}(x) + g^{(2)}(x) \right),$$

and the coupled reference solution is the unique fixed point $x^* \in \text{Fix}(\mathcal{T})$, numerically found as

$$x^* = (0.385949, 0.386257, 0.383437)^\top.$$

Clarification. Following the notation of Section 2.1, we rewrite the relaxation coupling as the composition of two maps $\Gamma_1 : \mathbb{R}^3 \rightarrow \mathbb{R}^6$ and $\Gamma_2 : \mathbb{R}^6 \rightarrow \mathbb{R}^3$. We set

$$\Gamma_1(x) = (x, g^{(1)}(x)),$$

and, for $y = (y_1, \dots, y_6) \in \mathbb{R}^6$,

$$\Gamma_{2,j}(y) = \frac{1}{2} y_{3+j} + \frac{1}{2} g_j^{(2)}((y_\ell)_{\ell=1}^3), \quad j = 1, 2, 3.$$

Then $\mathcal{T} = \Gamma_2 \circ \Gamma_1$ and the coupled solution satisfies $x^* = \mathcal{T}(x^*)$. We also verified numerically that $\mathcal{T}(\mathcal{U}) \subset \mathcal{U}$ on $\mathcal{U} = [0, 1]^3$ and that $\sup_{x \in \mathcal{U}} \|D\mathcal{T}(x)\|_2 \leq \rho$ with $\rho \simeq 0.2079 < 1$, which supports the contractive character of $\mathcal{T}$ on the considered domain and ensures uniqueness of the fixed point in $\mathcal{U}$.

Surrogates and coupled fixed-point solver. We build two structured GP surrogates $f^{(1)}$ and $f^{(2)}$ consistent with the above decomposition: $f^{(1)}$ emulates Γ_1 and $f^{(2)}$ emulates Γ_2 . More precisely, $f^{(1)}$ has mean function $m^{(1)}(x) = (x_1, x_2, x_3, 0, 0, 0)$ and diagonal covariance with diagonal entries $(0, 0, 0, k, k, k)$, where k is a Matérn covariance on $\mathbb{R}^3$, while $f^{(2)}$ has the mean function $m^{(2)}(y) = (y_4/2, y_5/2, y_6/2)$ and diagonal covariance (k, k, k) acting on

the first three coordinates (y_1, y_2, y_3) . In practice, each random component is modeled by an independent GP with Matérn-5/2 kernel, *fixed* isotropic length-scale $\ell = 0.80$ (no hyperparameter optimization), and a very small nugget (jitter) $\alpha = 10^{-12}$ added for numerical stability only (it is not interpreted as observation noise).

For each Monte Carlo replication j , the coupled surrogate solution is obtained as the fixed point of

$$x = \mathcal{T}_{f^{(j)}}(x), \quad \mathcal{T}_{f^{(j)}} := f^{(2,j)} \circ f^{(1,j)},$$

starting from $x^{(0)} = (0.5, 0.5, 0.5)$ and stopped when $\|x^{(m+1)} - x^{(m)}\|_2 \leq \varepsilon$ with tolerance $\varepsilon = 10^{-7}$.

Table 1 reports the deterministic prediction accuracy of the uncertain GP blocks associated with the two analytical functions $g^{(1)}$ and $g^{(2)}$. In the structured implementation, the uncertain block of f_1 corresponds to $g^{(1)}$, whereas the uncertain block of f_2 corresponds to $(1/2)g^{(2)}$; the indicators associated with $g^{(2)}$ are therefore reported after rescaling the posterior mean to the original function $g^{(2)}$. The table includes the predictive coefficient Q^2 , the root mean square error (RMSE), and the mean absolute error (MAE), all computed on an independent validation set.

Benchmark results (small DOE versus large DOE). Two DOE sizes are considered: small DOE ($n = 20$) and large DOE ($n = 500$) training points per function. In both cases, $N = 500$ Monte Carlo replications are performed for each method.

For small DOE, the empirical means are $\hat{\mathbf{m}}_{M2} = (0.379372, 0.373306, 0.388248)^\top$, $\hat{\mathbf{m}}_{M3} = (0.386091, 0.385796, 0.380067)^\top$, with componentwise variances $\widehat{\text{Var}}_{M2} = (2.1222 \times 10^{-2}, 1.9478 \times 10^{-2}, 1.8086 \times 10^{-2})^\top$, $\widehat{\text{Var}}_{M3} = (4.680 \times 10^{-3}, 4.754 \times 10^{-3}, 3.930 \times 10^{-3})^\top$. The mean difference is $\hat{\mathbf{m}}_{M2} - \hat{\mathbf{m}}_{M3} = (-0.006719, -0.012490, 0.008181)^\top$. The componentwise Kolmogorov–Smirnov statistics are (0.100, 0.092, 0.142) with p -values (0.0134, 0.0290, 0.0001), while the Welch tests give $t = (-0.933, -1.794, 1.233)$ with p -values (0.3509, 0.0732, 0.2180). For $\|y\|_2$, the Kolmogorov–Smirnov statistic is 0.118 ($p = 0.0019$) and the Welch test gives $t = 3.437$ ($p = 0.0006$).

For large DOE, the empirical means are $\hat{\mathbf{m}}_{M2} = (0.385719, 0.386774, 0.383345)^\top$, $\hat{\mathbf{m}}_{M3} = (0.385863, 0.386027, 0.383166)^\top$, with componentwise variances $\widehat{\text{Var}}_{M2} = (5.8902, 6.5133, 7.7506) \times 10^{-5}$, $\widehat{\text{Var}}_{M3} = (5.5739, 6.4060, 6.8686) \times 10^{-5}$. The mean difference is $\hat{\mathbf{m}}_{M2} - \hat{\mathbf{m}}_{M3} = (-0.000144, 0.000747, 0.000179)^\top$. The componentwise Kolmogorov–Smirnov statistics are (0.034, 0.072, 0.060) with p -values (0.9351, 0.1497, 0.3294), and the Welch tests give $t = (-0.301, 1.469, 0.332)$ with p -values (0.7636, 0.1421, 0.7402). For $\|y\|_2$, we obtain a Kolmogorov–Smirnov statistic 0.058 ($p = 0.3699$) and a Welch test $t = 0.877$ ($p = 0.3806$). Figure 1 reports the corresponding empirical distributions.

TABLE 1: Deterministic prediction accuracy of the multidimensional GP metamodels, reported on the scale of the original analytical functions. Q^2 , RMSE, and MAE are computed on an independent validation set

DOE/Function	Component	Q^2	RMSE	MAE
small DOE $g^{(1)}$	y_1	0.9493	1.31×10^{-2}	8.61×10^{-3}
small DOE $g^{(1)}$	y_2	0.6654	1.38×10^{-2}	9.63×10^{-3}
small DOE $g^{(1)}$	y_3	0.9149	1.27×10^{-2}	8.33×10^{-3}
small DOE $g^{(2)}$	y_1	0.8663	1.11×10^{-2}	6.82×10^{-3}
small DOE $g^{(2)}$	y_2	0.9463	1.44×10^{-2}	9.30×10^{-3}
small DOE $g^{(2)}$	y_3	0.7947	1.12×10^{-2}	7.40×10^{-3}
large DOE $g^{(1)}$	y_1	1.0000	3.69×10^{-4}	9.82×10^{-5}
large DOE $g^{(1)}$	y_2	0.9999	1.88×10^{-4}	6.18×10^{-5}
large DOE $g^{(1)}$	y_3	1.0000	3.02×10^{-4}	8.64×10^{-5}
large DOE $g^{(2)}$	y_1	0.9999	2.37×10^{-4}	6.85×10^{-5}
large DOE $g^{(2)}$	y_2	1.0000	3.13×10^{-4}	8.57×10^{-5}
large DOE $g^{(2)}$	y_3	0.9999	2.15×10^{-4}	6.82×10^{-5}




FIG. 1: Empirical distributions of the coupled output $\mathbf{y} = (y_1, y_2, y_3)$ for Methods 2 and 3. Vertical dashed lines indicate the true reference solution, and dotted lines indicate the deterministic GP-mean fixed point. Histograms are plotted as probability densities using common bin edges within each panel.

Discussion. For small DOE, Method 3 remains reasonably close to Method 2 in terms of empirical means, but the Kolmogorov–Smirnov tests detect statistically significant differences in distribution shape (componentwise and for $\|\mathbf{y}\|_2$), reflecting the limited informativeness of the design. For large DOE, the two methods become statistically consistent in practice: empirical means and variances are very close and both the Kolmogorov–Smirnov and Welch tests no longer indicate significant discrepancies (componentwise and for $\|\mathbf{y}\|_2$). This behavior is coherent with the rationale of Section 2.3: as the design becomes denser (smaller fill distance), posterior variances decrease and the effect of GP perturbations on the coupled fixed point is better controlled, so the fixed-path constant-offset scheme provides a reliable approximation of the trajectory-conditioned reference method.

3. APPLICATION: UNCERTAINTY QUANTIFICATION OF FUEL ASSEMBLY BOW USING THE SURROGATE MODELS

3.1 Physical Context and Coupled Model

Fuel assembly (see Fig. 2) bow is a long-standing safety and operational concern in pressurized water reactor (PWR) cores, first highlighted in the 1990s through cases of incomplete rod cluster insertion [13,14]. The deformation of fuel assemblies results from the combined action of neutronic, mechanical, and hydraulic phenomena [15–17]. In industrial practice, fuel assembly bow is evaluated over successive irradiation cycles, each cycle corresponding to about 12–18 months of operation (see Fig. 3) and represented in computations by a sequence of quasistatic time steps. At each time step, hydraulic and mechanical conditions are updated and coupled through a fluid-structure interaction (FSI) loop: the hydraulic code Phorcys computes grid-level forces from the current deformed geometry [3], while the thermomechanical code DACC updates the assembly deformation using a beam-type model of the rod bundle and skeleton [18], including nonlinear rod-grid contact, thermal expansion, irradiation growth, creep, variable grid clamping, and hold-down forces [19]. This partitioned Phorcys-DACC coupling is uncertain and can be computationally demanding (slower solving, more iterations) for large systems.

To reduce this cost while retaining predictive accuracy, surrogate models based on Gaussian processes have been developed in previous works. In [21], Gaussian process hydraulic surrogate models were constructed to predict




FIG. 2: Schematic view of a fuel assembly (around 4 m high) with spacer grids (based on the results of [20])




FIG. 3: Hydraulic network model for a row of 15 fuel assemblies under reactor operation (inlet and outlet) conditions (based on the results of [3])

spacer-grid forces from uncertain inputs, and in [19], Gaussian process mechanical surrogates were built to predict assembly deformation. In both cases, Matérn kernels were adopted to provide flexible but regularized representations of the underlying physics. Their coupling was then formalized and validated in [22], enabling full-cycle simulations via chained mechanical surrogate models, a nonlinear penalization to prevent interassembly penetration, and a robust dual-loop FSI algorithm. Comparisons with the Phorcys-DACC reference chain showed that the coupled surrogates reproduce fuel assembly bow with sufficient accuracy while reducing computational time by nearly two orders of magnitude.

The present work builds on this surrogate-based coupling and extends it to a fully probabilistic, uncertainty-aware framework. The objective is to quantify the epistemic uncertainty induced by the coupled Gaussian process surrogates and to assess, on the realistic fuel assembly bow application, the practical relevance of the variance-control results established in Section 2.3 and, finally, to illustrate uncertainty quantification on an example of FA simulation using the surrogate models.

3.2 Uncertainty Quantification of the Surrogate Models in the Context of Fuel Assembly Bow

In this application, the two coupled surrogates play the role of $f^{(1)}$ (hydraulic GP surrogate model, one output per spacer grid) and $f^{(2)}$ (mechanical GP surrogate model, one output per deformation principal mode). Both are multioutput Gaussian processes with Matérn-based kernels. At each irradiation time step t , the coupled Phorcys-DACC solver follows the abstract setting of Section 2.3 with $C = 2$: given a surrogate pair $f = (f^{(1)}, f^{(2)})$, the coupling operator $\mathcal{T}_f$ acts on an interface space $\mathcal{U}$, admits a unique fixed point $u_t^*(f)$, and the observable bow field at step t is obtained as $\mathbf{y}_t = \Phi_t(f) := H(u_t^*(f))$, with Φ_t satisfying the Lipschitz stability property of Proposition 3. The full-cycle output is obtained by composing these stepwise solution maps over $t = 1, \dots, T$. For notational simplicity we denote by $\Phi_{\text{cyc}}(f)$ the resulting full-cycle bow field.

Deterministic reference coupled simulation. Throughout this subsection, we take as deterministic reference the coupled solution obtained by using Gaussian process posterior means only: $f^\mu := (\bar{\mu}^{(1)}, \bar{\mu}^{(2)})$, $\mathbf{y}^\mu := \Phi_{\text{cyc}}(f^\mu)$, which corresponds to the operational surrogate-based simulator used in practice.

Deterministic mean-path run. Uncertainty is propagated with the simplified Monte Carlo strategy (Method 3) of Section 2.2, adapted to a multistep irradiation cycle. Starting from the initial core geometry, we first run the coupled simulation with the Gaussian process posterior means $\{\bar{\mu}^{(c)}\}_{c=1}^2$. For each irradiation step $t = 1, \dots, T$, the inner fixed-point solver produces interface iterates $(u^{(t,m)})_{m \geq 0}$ and stops after M_t iterations. The surrogate inputs queried at inner iteration m of step t are denoted $x_{t,m}^{(c)} \in \mathcal{X}^{(c)}$, $t = 1, \dots, T$, $m = 1, \dots, M_t$, $c = 1, 2$. We define the deterministic cycle path

$$\mathcal{X}_{\text{path}}^{(c)} := \left\{ x_{t,m}^{(c)} : t = 1, \dots, T, m = 1, \dots, M_t \right\} \subset \mathcal{X}^{(c)}, \quad c = 1, 2.$$

This set is the cycle-level analog of the points $\{x_m^{(c)}\}_{m=1}^M$ used in Method 3. It is computed once from the mean-only run and then reused, unchanged, for all Monte Carlo replications.

Monte Carlo perturbations along the deterministic path. For each Monte Carlo replication $j = 1, \dots, N$ and each surrogate $c \in \{1, 2\}$, we sample one joint posterior Gaussian process realization restricted to the deterministic cycle path: $\left(\tilde{f}^{(c,j)}(x_{t,m}^{(c)}) \right)_{t,m} \sim \mathcal{N}(\bar{\mu}^{(c)}(\mathcal{X}_{\text{path}}^{(c)}), \bar{\mathbf{K}}^{(c)})$, where $\bar{\mu}^{(c)}(\mathcal{X}_{\text{path}}^{(c)}) = \left(\bar{\mu}^{(c)}(x_{t,m}^{(k)}) \right)_{t,m}$ and $\bar{\mathbf{K}}^{(c)}$ is the block posterior covariance matrix whose $((t, m), (t', m'))$ block is $\bar{\kappa}^{(c)}(x_{t,m}^{(c)}, x_{t',m'}^{(c)}) \in \mathbb{R}^{D_c \times D_c}$. We then define the pathwise offsets,

$$\delta_{t,m}^{(c,j)} := \tilde{f}^{(c,j)}(x_{t,m}^{(c)}) - \bar{\mu}^{(c)}(x_{t,m}^{(c)}), \quad t = 1, \dots, T, \quad m = 1, \dots, M_t, \quad j = 1, \dots, N,$$

and, at irradiation step t and inner iteration m , we use the constant-offset perturbed surrogate $\hat{f}^{(c,t,m,j)}(x) = \bar{\mu}^{(c)}(x) + \delta_{t,m}^{(c,j)}$. For fixed (t, m, j, c) , the quantity $\delta_{t,m}^{(c,j)}$ is constant in x ; i.e., it acts as an additive shift of the mean

predictor. It is updated only when the algorithm advances to the next inner iteration index (t, m) , which matches the constant-offset construction of Method 3 and the setting analyzed in Section 2.3.

Cycle solve under perturbed surrogates. For each replication j , we rerun the full-cycle coupled solver from the same initial state as in the mean-path run, replacing $\bar{\mu}^{(c)}$ by $\hat{f}^{(c,t,m,j)}$ at each step t and inner iteration m . This produces the Monte Carlo output $\mathbf{y}^{(j)} = \Phi_{\text{cyc}}\left(\left\{\hat{f}^{(c,t,m,j)}\right\}_{c,t,m}\right)$. The sample $\{\mathbf{y}^{(j)}\}_{j=1}^N$ then provides a Monte Carlo approximation of the epistemic uncertainty on the coupled cycle output induced by the Gaussian process surrogates, in direct correspondence with Method 3. We estimate the mean and covariance of the cycle output by $\hat{\mathbf{y}}_N = (1/N) \sum_{j=1}^N \mathbf{y}^{(j)}$, $\widehat{\text{Cov}}[\mathbf{y}] = (1/(N-1)) \sum_{j=1}^N (\mathbf{y}^{(j)} - \hat{\mathbf{y}}_N)(\mathbf{y}^{(j)} - \hat{\mathbf{y}}_N)^\top$. When needed, componentwise summaries (variances, Gaussian uncertainty intervals) are obtained from $\hat{\mathbf{y}}_N$ and the diagonal of $\widehat{\text{Cov}}[\mathbf{y}]$.

Results and interpretation in light of the theory. Figure 4 shows the standard deviation of the coupled Gaussian process predictions of fuel assembly bow over a row of 15 assemblies and across the irradiation cycle, obtained by running the coupled solver with the Gaussian process surrogates. The largest epistemic uncertainty induced by the surrogates reaches 0.26 mm. This remains small compared with the displacement scale of about 10 mm. (In the figure, the actual height of the fuel assembly is around 4 m. However, the diagram is intended to represent the spacing between two assemblies, therefore, the vertical dimension is not to scale and does not reflect the real length of the fuel assembly. Each small square in the figure has a width equivalent to 2 mm, i.e., the nominal water gap between two grids for undeformed fuel assemblies.) This behavior is consistent with the theoretical results of Section 2.3: the training designs for the hydraulic and mechanical surrogates are dense (small fill distances $h^{(c)}$), which keeps posterior variances low along the mean path (Theorem 3). In addition, the Lipschitz stability of the coupled solution map Φ (Proposition 3) limits the amplification of these local surrogate errors through the surrogate-based FSI coupling. Overall, the surrogates capture the fuel assembly bow response with sufficient accuracy: epistemic uncertainty remains negligible at the coupled-output level, supporting the use of the posterior mean predictors alone in large-scale reactor simulations.




FIG. 4: Uncertainty of the coupled Gaussian process in a simulation of a row of 15 fuel assemblies

3.3 Sensitivity Analysis for the Fluid-Structure Coupling

Based on the results above, we use only the surrogate-model posterior means as predictors in this section and throughout the rest of the paper. We perform a global sensitivity analysis of the coupled fuel assembly bow simulation using Sobol' indices [23], leveraging the fast evaluation and predictive accuracy of the Gaussian process surrogate models established in [19,21].

Coupled model with uncertain parameters. In addition to the interface variable $u \in \mathcal{U}$ updated by the fixed-point coupling (Section 2.1), the two deterministic solvers depend on a vector of uncertain inputs $\theta \in \Theta \subset \mathbb{R}^p$ (hydraulic boundary conditions, mechanical parameters, and initial deformations). For $c \in \{1, 2\}$ we write $\mathcal{S}_c : \mathcal{X}_c \times \Theta \rightarrow \mathcal{Y}_c$, $(x, \theta) \mapsto \mathcal{S}_c(x, \theta)$, and for a fixed θ we define $\mathcal{S}_c^\theta(x) := \mathcal{S}_c(x, \theta)$. For each sample θ used for Sobol' index estimation, the coupling algorithm is run with θ kept fixed throughout all fixed-point iterations. The coupled solution therefore depends on θ through the fixed point $u^*(\theta) \in \text{Fix}(\mathcal{T}_\theta)$, $\mathcal{T}_\theta := \Gamma_2 \circ \mathcal{S}_2^\theta \circ \Gamma_{12} \circ \mathcal{S}_1^\theta \circ \Gamma_1$, and the quantity of interest is the coupled output $\mathbf{y}^*(\theta) = H(u^*(\theta))$.

Gaussian process surrogates used in the sensitivity study. To make the Sobol' analysis tractable, we replace each solver $\mathcal{S}_c$ by a Gaussian process surrogate model trained on the joint inputs (x, θ) , and we run the coupling using only its posterior mean $\bar{\mu}^{(c)}(x, \theta)$ (Section 2.2). Accordingly, the sensitivity analysis targets the variability of $\mathbf{y}^*(\theta)$ induced by the uncertain inputs θ , which include both aleatory and epistemic physical parameters through their prescribed probability distributions (see Table 2).

High-dimensional input space. A direct Sobol' analysis is challenging because θ is high-dimensional: it includes $45 = 3 \times 15$ modal coefficients describing the three C , S , and W deformation modes for 15 assemblies (see Fig. 5), five hydraulic boundary-condition parameters, and additional uncertain quantities from the hydraulic and mechanical models.

Methodology. To make this analysis computationally feasible and scientifically meaningful, a structured reduction of the input space is applied. This reduction is guided by earlier sensitivity analyses carried out separately for the mechanical and hydraulic models in [19,21]. The input parameters are organized into coherent groups to reflect the underlying physical processes: Initial modal deformations are aggregated into three representative vectors for the C , S , and W modes; hydraulic boundary conditions are grouped into a distinct set of corresponding parameters; the results of the sensitivity analyses on both the hydraulic and mechanical models allowed us to identify and neglect

TABLE 2: Uncertain input parameters, their names, and associated probability distributions

Symbol	Name	Probability distribution
C_j	Modal coefficient C	$\mathcal{N}(0, \sigma_C^2)$ (σ_C depends on irradiation cycle)
S_j	Modal coefficient S	$\mathcal{N}(0, \sigma_S^2)$ (σ_S depends on irradiation cycle)
W_j	Modal coefficient W	$\mathcal{N}(0, \sigma_W^2)$ (σ_W depends on irradiation cycle)
Hydraulic boundary conditions		
V	Average axial velocity	$\mathcal{N}(\mu = 5, \sigma^2 = 0.05^2)$
M_{in}	Maximum deviation (inlet)	$\mathcal{N}(\mu = 0.05, \sigma^2 = 0.005^2)$
L_{offset}^{in}	Inlet lateral offset	$\mathcal{N}(\mu = 0, \sigma^2 = 0.2^2)$
M_{out}	Maximum deviation (outlet)	$\mathcal{N}(\mu = 0.04, \sigma^2 = 0.004^2)$
L_{offset}^{out}	Outlet lateral offset	$\mathcal{N}(\mu = 0, \sigma^2 = 0.1^2)$
Epistemic parameter		
h_l	Epistemic parameter (hydraulic model)	$\mathcal{U}(0.01, 0.03)$
Mechanical parameters		
MSI	Holddown interaction of the system	$\mathcal{U}(0.01, 0.025)$
$grid_clamping$	Grid clamping force	$\mathcal{N}(\mu, \sigma^2)$ (μ, σ^2 depend on irradiation cycle)
C_{creep}	Creep constant	$\mathcal{N}(\mu = 1, \sigma^2 = 0.3^2)$
C_{growth}	Growth constant	$\mathcal{N}(\mu = 1, \sigma^2 = 0.3^2)$




FIG. 5: C , S , and W deformation modes of a fuel assembly

the variability of certain parameters that exhibited negligible influence on the simulation outputs. These include the inlet temperature in both simulations, the axial resistance of spacer grids in the hydraulic model, and the fast neutron flux. It is important to note that these parameters were not entirely removed from the models but rather fixed at their mean values, because these parameters originally served as input variables for the surrogate models. This structured reformulation reduces redundancy, simplifies interpretation, and enhances the clarity of the resulting Sobol' indices.

The choice of the initial configuration and the uncertain input parameters. Based on a prescribed neutronic configuration, each fuel assembly is positioned within the core according to its irradiation cycle (an assembly of cycle j has already experienced $j - 1$ full irradiation cycles). Assemblies of cycle 3 are placed in the central region, assemblies of cycle 2 are placed at the periphery, and assemblies of cycle 1 are placed in intermediate positions. This influences the input parameters of the mechanical surrogate models, especially the modal coefficients which are used to generate the initial deformations as a function of the age of each assembly. Additionally, the grid clamping force is determined individually for each assembly. Table 2 summarizes the uncertain input parameters used in the coupled simulations, along with their associated probability distributions. All input parameters in this study are supposed to be independent.

Quantities of interest and sensitivity analysis. The primary quantities of interest in this study are the modal coefficients C , S , and W , which characterize the mechanical deformation state of each fuel assembly. A global sensitivity analysis has been conducted using Sobol' indices, which quantify the contribution of each uncertain input parameter to the variance of these outputs. In this study, the number of simulations is set to $n_{\text{simulation}} = 11,000$, which corresponds to the number required to compute the Sobol' sensitivity indices using the Saltelli method [24] implemented in the *Uranie* platform [25]. The Saltelli approach requires a total of $n_s(n_x + 2)$ simulations, where $n_s = 1000$ denotes the base sample size, and $n_x = 9$ represents the number of input variables: three vector inputs (C , S , and W), one vector for the hydraulic boundary conditions, and the five scalar inputs. The analysis has been performed at three key stages of the reactor cycle: prior to irradiation, at the end of the irradiation phase, and after irradiation during the opening of the upper core plate (see Fig. 6).

This temporal segmentation allows the evolution of sensitivities to be captured throughout the operational cycle. In addition, a detailed Sobol' analysis has been conducted by grouping fuel assemblies according to their irradiation age (cycle) and spatial location in the core.

Old assemblies of cycle 3 are located at the center of the core, assemblies of cycle 2 are placed near the core shroud, and new assemblies of cycle 1 are positioned in the intermediate regions. This classification enables a more precise interpretation of how age and radial position influence mechanical behavior. To further enhance interpretability, we have computed aggregated Sobol' indices [26], which quantify the relative contribution of groups of input parameters to the total output variance. For clarity of presentation, these aggregated indices are displayed as normalized




FIG. 6: Schematic of the fuel assembly grid's lateral displacement across reactor cycle stages

pie charts, thereby providing a compact visual summary of the dominant sources of uncertainty across assemblies and time points. It is worth emphasizing that the sum of the aggregated first-order Sobol' indices has been found to be close to unity, indicating that higher-order interactions among input parameters contribute negligibly to the variance. This property justifies both the normalization step for visualization and the use of aggregated first-order indices as a reliable measure of parameter importance in this context.

Results before irradiation phase. During the pre-irradiation phase, the simulation considers the insertion force of control rods in the grids, closure of the core, activation of primary circulation pumps, and the thermal ramp-up to operational temperature. This phase corresponds to the initial mechanical stabilization of the reactor core under thermal and hydraulic loads, in the absence of irradiation-driven phenomena such as creep or growth. Figure 7 shows that the initial geometric deformation of fuel assemblies is the dominant factor influencing the final deformation state across all cycles (1, 2, and 3) before irradiation. These initial deformations determine the contact distribution between assemblies and thus govern the mechanical equilibrium at the early stages.

Results at the end of irradiation phase. At the end of irradiation, deformation is governed by both irradiation-driven effects and the legacy of initial conditions (Fig. 8). Overall, hydraulic boundary conditions and the epistemic parameter h_l seem to dominate the variance across cycles, especially regarding modes C and W . For assemblies of cycle 1, where initial deformations are small, the response is mainly controlled by BC and h_l , with creep acting as a secondary contributor, especially for mode S . For assemblies of cycle 2, BC and h_l remain dominant for C and W modes, and creep remains a notable second-order effect; while regarding the S mode, larger initial deformations play a stronger role; influence is shared between initial deformations, BC, h_l , and creep. For assemblies of cycle 3, initial deformations become the leading driver in several modes (notably S), surpassing the hydraulic influence, while creep has only marginal impact. Growth- and hold-down-related parameters remain negligible throughout.

Results after irradiation phase. At the post-irradiation stage—typically following upper core plate lift-off— assemblies undergo elastic relaxation and retain residual deformation shaped by their irradiation history and boundary conditions (Fig. 9). For assemblies of cycle 1 with low burnup in the intermediate regions, modal deformation variance




FIG. 7: Aggregated first-order Sobol' indices before irradiation phase on the C , S , and W deformation modes




FIG. 8: Aggregated first-order Sobol' indices at the end of irradiation phase on the C , S , and W deformation modes




FIG. 9: Aggregated first order Sobol' indices after the irradiation phase on the C , S , and W deformation modes

is governed mainly by hydraulic boundary conditions (BC) and the epistemic parameter h_l , with creep contributing only second-order effects. For assemblies of cycle 2, regarding modes S and W , one can argue initial deformations; BC and h_l remain influent. However, no input really stands out in mode C . For assemblies of cycle 3 at the core center with highest burnup, while BC and h_l remain influent regarding the C and W modes (along with C and W initial deformations), the S mode is almost only driven by legacy geometric deformations (initial S mode), confirming the persistence of accumulated deformations. Overall, results for cycle 3 are similar to the ones in the end-of-irradiation phase, highlighting that phenomena occurring after this phase only have little impact on input for these assemblies. In this regime, creep becomes negligible compared to geometric effects.

3.4 Predictive Accuracy of the Simulation Codes

In simulations involving coupled mechanical and hydraulic models of fuel assemblies, initial deformation data are subject to measurement uncertainty. Though these uncertainties are often overlooked, it has been demonstrated that they play a significant role in the variance of the final outcomes [19]. It is therefore crucial to investigate how such uncertainties propagate through the computational model and influence the resulting physical quantities. This section primarily aims to evaluate the predictive capability of the coupled model for a given configuration—referred to as the “calculation of the day”—by quantifying the confidence that can be placed in its results, assuming the input measurements (including initial deformations) are known up to a certain measurement error. A secondary objective is to analyze the contribution of initial deformation uncertainty to the overall result variability, in order to better understand the dominant sources of uncertainty.

Methodology. Running high-fidelity coupled simulations that account for both the mechanical and hydraulic behavior of nuclear fuel assemblies can become time consuming for large systems, especially when such simulations must be repeated numerous times for uncertainty quantification. To overcome this limitation, we use the surrogate models coupling strategy. These tools offer an efficient and reasonably accurate alternative to direct numerical simulations, significantly reducing computation time while preserving the essential characteristics of the physical response.

It is important to note that, in this simulation, we use the mean values of the Gaussian processes as predictors, and we do not account for the uncertainty associated with the surrogate model itself. The approach adopted in this study relies on a combination of physical assumptions and data inspired by realistic reactor operating conditions.

Hydraulic modeling. The parabolic velocity profile introduced in [22] is assumed. The uncertainty in these boundary conditions is considered to be small. The hydraulic epistemic parameter, denoted h_l , is treated as a uniformly distributed variable within the range [10 mm, 30 mm], with a nominal reference value of 20 mm.

Geometrical Deformation. The initial deformation is applied to a row of 15 fuel assemblies, as described in Section 3.3). In this study, we do not explore the full space of possible deformation patterns; rather, we focus specifically on the measurement uncertainty surrounding the known (measured) bow shapes. To this end, random variations are introduced in the modal coefficients C , S , and W to simulate the expected variability due to measurement errors.

Material and structural uncertainties. Beyond geometric variability, the model incorporates uncertainties in several material and structural properties. Parameters such as creep behavior, fuel rod growth, and the stiffness of grid spacers—particularly those related to clamping forces—are included in the analysis, as their uncertainties cannot be neglected. The probability distributions used for these parameters are those listed in Table 2.

Propagation of measurement uncertainty on modal coefficients. In the fuel assembly bow application, the initial deformation is measured at ten axial levels (typically at the spacer grids) and represented in a reduced modal basis. Let $\mathbf{U} \in \mathbb{R}^{10}$ denote the true displacement vector and write $\mathbf{U} = \sum_{j=1}^3 C_j \mathbf{u}_j = \mathbf{M}\mathbf{C}$, where $\{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$ is an orthonormal family of deformation modes, $\mathbf{M} \in \mathbb{R}^{10 \times 3}$ collects these modes as columns, and $\mathbf{C} = (C_1, C_2, C_3)^\top$ are the modal coefficients. The measured displacement is modeled as $\mathbf{U}_{\text{mes}} = \mathbf{U} + \boldsymbol{\varepsilon}$, $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{10})$, $\sigma = 0.3$ mm. Since the modes are orthonormal, the least-squares estimate of $\mathbf{C}$ is simply the projection of $\mathbf{U}_{\text{mes}}$ onto this basis. The projection is linear and the measurement noise is Gaussian, so the estimated modal coefficients are themselves Gaussian, mutually uncorrelated, with variance σ^2 inherited from the measurement noise. This justifies modeling the uncertainties on C_j , S_j , and W_j as independent Gaussian variables with variance σ^2 .

Representation of the uncertainty in the initial configuration. An assumed measurement uncertainty of approximately 0.3 mm is therefore applied on the modal coefficients to represent variability in the initial deformation of fuel assemblies. Figure 10 illustrates the resulting uncertainty in the initial deformation field of the reference simulation, showing 95% uncertainty intervals (Gaussian bands under the assumed normal model) associated with each assembly




FIG. 10: Uncertainty in the initial deformation field with the 95% uncertainty intervals associated with each assembly [dark blue (far left and far right): assemblies of cycle 2; red (middle left and middle right): assemblies of cycle 1; light blue (center): assemblies of cycle 3]

assembly. In this study, the same uncertainty level (0.3 mm) is applied regardless of the fuel cycle stage. This is likely conservative at the beginning of cycle 1: factory measurements, performed in air on controlled metrology benches with unirradiated assemblies, generally exhibit smaller uncertainties, and assemblies can be straightened if necessary. In contrast, measurements taken later in the cycle (after cycles 1–3) are made under water on irradiated assemblies and are typically less precise. Throughout this section, the nominal water-gap thickness (distance between two grids, corresponding to one square in Fig. 10) is taken as 2 mm at full power.

Uncertainty on the parabolic velocity profile. The axial velocity distribution is represented by a quadratic law $v(x) = ax^2 + bx + c$ where x is the spatial coordinate in the core and the coefficients (a, b, c) are uncertain. They are inferred from three Gaussian random inputs: the mean axial velocity $\bar{v} \sim \mathcal{N}(\mu_v, \sigma_v^2)$, the maximum deviation $M \sim \mathcal{N}(\mu_M, \sigma_M^2)$, and the lateral offset $L_{off} \sim \mathcal{N}(\mu_{L_{off}}, \sigma_{L_{off}}^2)$. Algebraic constraints link these inputs to (a, b, c) , defining the parabolic shape. Uncertainty propagation is performed by Monte Carlo simulation: generate N realizations of $(\bar{v}, M, L_{off})$, compute coefficients (a_i, b_i, c_i) and evaluate $v_i(x) = a_i x^2 + b_i x + c_i$, and estimate the empirical mean and variance at each x , $\hat{v}(x) = (1/N) \sum_{i=1}^N v_i(x)$, $\sigma_{v(x)}^2 = (1/(N-1)) \sum_{i=1}^N (v_i(x) - \hat{v}(x))^2$.

This procedure quantifies the variability of $v(x)$ induced by uncertainty in the input parameters. Figure 11 illustrates the resulting dispersion of velocity profiles consistent with the prescribed probability laws.

Reference simulation and uncertainty quantification. The reference simulation corresponds to a fully coupled thermomechanical and hydraulic computation performed under nominal conditions. In this context, all uncertain input parameters are fixed at their respective mean values. The geometric configuration considered is that of a row of 15 fuel assemblies. The simulation covers a complete irradiation period of 18 months, during which the evolution of mechanical deformation is tracked. The primary quantity of interest is the final deformation of each assembly in the row.

To assess the impact of parameter uncertainties on the simulation output, a forward uncertainty quantification procedure is conducted using a Monte Carlo approach. A total of 2000 simulations are performed, each corresponding to a distinct realization of the input parameters sampled from their associated probability density function. These uncertain inputs include the uncertainty of the measurements of the initial geometry (the measured value, assumed to be read from the instrument, has been fixed), material behavior parameters (e.g., creep and growth coefficients), and




FIG. 11: Uncertainty in the boundary condition profiles with 95% pointwise uncertainty band (injection profile = inlet conditions and withdrawal profile = outlet conditions in Fig. 3)

the hydraulic parameters. The empirical variance and uncertainty intervals of the output deformation fields are then estimated from the simulation results.

In parallel, a modal decomposition is employed to represent the deformation field $Y(z)$ along the axial coordinate z . This representation is expressed as $Y(z) = \sum_{i=1}^3 C_i u_i(z)$, where $\{u_i(z)\}_{i=1}^3$ are deterministic modal shapes, and $\{C_i\}_{i=1}^3$ are the corresponding random modal coefficients. These coefficients encapsulate the uncertainties arising from measurement errors of the initial geometry. Assuming the modal shapes are fixed and known, the variance of the deformation field $Y(z)$ at each axial location can be computed using the following expression: $\text{Var}(Y(z)) = \sum_{i,j=1}^3 u_i(z) u_j(z) \text{Cov}(C_i, C_j)$, where $\text{Cov}(C_i, C_j)$ denotes the covariance between modal coefficients C_i and C_j , which is estimated from the sampled data set. This formulation allows the construction of a spatial uncertainty map of the deformation field. Moreover, it facilitates the derivation of local variances and empirical (Monte Carlo) intervals at each position z .

Such analysis provides a detailed and probabilistically rigorous characterization of the possible variations in the final deformation profiles, enhancing the interpretability and reliability of the simulation results under uncertainty.

Results. Figure 12 shows the reference final deformation after 18 months of irradiation and removal of the upper plate. The results are now presented in three figures to facilitate analysis by grouping them according to the type of fuel assembly and their spatial location within the row. Figure 13 corresponds to assemblies of cycle 2 positioned at the periphery of the reactor vessel. Figure 14 shows the results for assemblies of cycle 1 located in the intermediate region, and Fig. 15 presents the results for assemblies of cycle 3 placed at the center of the core. In each figure, the standard deviation of the grid displacement is computed. This provides an indicator of the maximum uncertainty associated with each type of assembly across the three figures.

After presenting the uncertainty estimates in the previous figures, we now show the full distributions of 2000 calculations for assemblies A_6 and A_{10} in Fig. 16, using the sample mean as a summary statistic. The distributions clearly deviate from Gaussian, which can be attributed to nonlinear effects and interactions between complex physical phenomena.

The uncertainty in the simulation results seems primarily influenced by two key factors: the spatial position of the assemblies within the core and their irradiation age. Our interpretation is that assemblies located near the core center are more free to deform due to fewer mechanical constraints of their neighbors, which leads to higher variability in their displacement responses. In contrast, peripheral assemblies are restricted by the reactor vessel, limiting their




FIG. 12: Reference simulation results after an 18-month irradiation cycle




FIG. 13: Uncertainty in the final deformation of fuel assemblies of cycle 2




FIG. 14: Uncertainty in the final deformation of fuel assemblies of cycle 1

movement and consequently reducing the associated uncertainty. Additionally, the age of the assemblies plays a significant role. Older assemblies deform more under the same loading, which in turn increases the uncertainty linked to their mechanical behavior. With increasing irradiation cycles, fuel assemblies undergo stress relaxation of the tiny springs retaining the rods and, as a consequence, tend to lose stiffness. As a result, older assemblies are more prone to bowing than newly loaded, straight assemblies under identical loading conditions. These two parameters—position and age—must therefore be carefully considered when interpreting results given by the coupled codes.




FIG. 15: Uncertainty in the final deformation of fuel assemblies of cycle 3




FIG. 16: A sample of size 2000 of deformation curves for assemblies A_6 and A_{10} of cycle 3 (see Fig. 15)

Synthesis of the practical impact. To summarize the main outcomes of the fuel-assembly bow application and highlight the practical contribution of the proposed surrogate-based coupling framework, we report in Tables 3 and 4 the key quantitative indicators: computational cost, feasible UQ sample sizes (Sobol' and Monte Carlo), resulting uncertainty levels, and the dominant parameters' contributors across the reactor-cycle stages.

TABLE 3: Summary of the practical contribution of the proposed surrogate-based coupling framework

Quantity	Direct coupled simulation (reference codes)	Surrogate-based coupled simulation
Time to simulate a full irradiation cycle	≈ 4 h (one coupled run)	≈ 2 min (one coupled run)
Sobol sensitivity analysis sample size	11,000 coupled evaluations is impractical	11,000 coupled evaluations performed
MC simulations for final uncertainty quantification	Impractical	2000 simulations performed
Maximum total output uncertainty (from input uncertainties)	Impractical using reference codes	Quantified: 1.03 mm
Maximum epistemic uncertainty induced by surrogate coupling	N/A	0.26 mm

TABLE 4: Main contributors to the variance of fuel-assembly bow across the three reactor-cycle stages (aggregated first-order Sobol' indices)

Cycle stage	Sobol's results: Dominant parameters influencing the final fuel assembly deformation (by assembly irradiation cycle)
Before irradiation phase	All cycles: Initial geometric deformations (modal coefficients C, S, W).
End of irradiation phase	Cycle 1: BC and h_l dominate; creep is secondary (notably for mode S). Cycle 2: Modes C/W : BC and h_l dominate (creep secondary); mode S : shared influence (initial deformations, BC, h_l , creep). Cycle 3: Initial deformations become leading (notably S); BC and h_l remain influential for C/W ; creep is marginal.
After irradiation phase	Cycle 1: BC and h_l dominate; creep remains second-order. Cycle 2: Modes S/W : initial deformations + BC + h_l ; mode C : no single dominant contributor. Cycle 3: Modes C/W : BC and h_l + initial C/W ; mode S : driven almost exclusively by legacy geometry (initial S); creep negligible.

Note: Cycle 1/2/3 denotes the assembly irradiation cycle (i.e., its age/burnup class in the core).

4. CONCLUSION

This work has introduced a general framework for incorporating Gaussian process surrogates into coupled multiphysics simulations and for quantifying the induced epistemic uncertainty. At the theoretical level, the analysis combines fill-distance bounds for scalar Gaussian processes in Sobolev-type RKHSs, lifted to matrix-valued LMC kernels, with a Lipschitz stability result for the coupled fixed point under uniform perturbations of the surrogates. This leads to a finite-sample, high-probability bound on the deviation between the coupled solution obtained with Gaussian process means and the random outputs of a Monte Carlo perturbation scheme, expressed in Loewner order and explicitly controlled by the fill distances of the designs and the kernel hyperparameters. The resulting bounds apply to general vector-valued surrogates and abstract coupling operators satisfying a uniform contraction property, and are therefore not tied to a specific physical model.

An analytical benchmark with a closed-form coupled solution has been used to validate the proposed constant-offset Monte Carlo scheme against a fully coherent, trajectory-conditioned sampling strategy, showing statistical agreement of the two output distributions at practical sample sizes. Finally, the methodology has been applied to a realistic fluid-structure interaction problem for fuel assembly bow over irradiation cycles, where it was shown that the uncertainty due to the coupled Gaussian process surrogates remains small (on the order of a few tenths of a

millimeter) compared with the dominant input-data uncertainties. This supports the use of Gaussian process mean predictions as accurate surrogates within reactor-scale uncertainty quantification, while the theoretical results ensure that analogous conclusions can be drawn for other coupled Gaussian process-based models under the same regularity and stability assumptions.

REFERENCES

1. Rasmussen, C.E. and Williams, C.K., *Gaussian Processes for Machine Learning*, Cambridge, MA: MIT Press, 2006.
2. Wanninger, A., Mechanical Analysis of the Bow Deformation of Fuel Assemblies in a Pressurized Water Reactor Core, PhD, Technische Universität München, 2018.
3. de Lambert, S., Contribution to the Multiphysical Analysis of Fuel Assembly Bow, PhD, Université Paris-Saclay, 2021.
4. Zeidler, E., *Nonlinear Functional Analysis and Its Applications I*, Berlin: Springer, 1986.
5. Álvarez, M.A., Rosasco, L., and Lawrence, N.D., Kernels for Vector-Valued Functions: A Review, *Found. Trends Mach. Learn.*, **4**(3):195–266, 2012.
6. Micchelli, C.A. and Pontil, M., On Learning Vector-Valued Functions, *Neural Comput.*, **17**(1):177–204, 2005.
7. Bonilla, E.V., Chai, K.M.A., and Williams, C.K.I., Multi-Task Gaussian Process Prediction, in *Advances in Neural Information Processing Systems (NeurIPS)*, La Jolla, CA: NIPS Foundation, pp. 153–160, 2008.
8. Damblin, G., Couplet, M., and Iooss, B., Numerical Studies of Space-Filling Designs: Optimization of Latin Hypercube Samples and Subprojection Properties, *J. Simul.*, **7**:276–289, 2013.
9. Kanagawa, M., Hennig, P., Sejdinovic, D., and Sriperumbudur, B.K., Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences, arXiv:1807.02582, 2018.
10. Micchelli, C.A. and Pontil, M., On Learning Vector-Valued Functions, *Neural Comput.*, **17**(1):177–204, 2005.
11. Álvarez, M.A., Rosasco, L., and Lawrence, N.D., Kernels for Vector-Valued Functions: A Review, *Found. Trends Mach. Learn.*, **4**(3):195–266, 2012.
12. Aronszajn, N., Theory of Reproducing Kernels, *Trans. Am. Math. Soc.*, **68**(3):337–404, 1950.
13. Andersson, T., Almberger, J., and Björnkvist, L., A Decade of Assembly Bow Management at Ringhals, in *IAEA-TECDOC-1454 Structural Behaviour of Fuel Assemblies for Water Cooled Reactors*, Vienna, Austria: IAEA, pp. 129–136, 2005.
14. Kerkar, N. and Paulin, P., *Exploitation Des Coeurs REP*, Les Ulis, France: EDP Sciences, 2008.
15. Karlsson, L. and Manngard, T., Modelling of PWR Fuel Assembly Deformations during Irradiation, *Trans. of the 15th Int. Conf. on Structural Mechanics in Reactor Technology (SMiRT-15)*, Seoul, Korea, 1999.
16. de Lambert, S., Campioni, G., Faucher, V., Leturcq, B., and Cardolaccia, J., Modeling the Consequences of Fuel Assembly Bowing on PWR Core Neutronics Using a Monte-Carlo Code, *Ann. Nucl. Energy*, **134**:330–341, 2019.
17. Wanninger, A., Mechanical Analysis of the Bow Deformation of a Row of Fuel Assemblies in a PWR Core, *Nucl. Eng. Technol.*, **50**:297–305, 2018.
18. de Lambert, S., Cardolaccia, J., Faucher, V., Leturcq, B., and Campioni, G., Semi-Analytical Fluid-Structure Model for the Analysis of Fuel Assembly Bow in Full PWR Cores, *Prog. Nucl. Energy*, **160**:104668, 2023.
19. Abboud, A., Garnier, J., Leturcq, B., Pacull, J., Fandeur, O., and de Lambert, S., Uncertainty Quantification of Fuel Assembly Bow in Pressurized Water Reactor Through a Thermomechanical Simulation, *Proc. of the Int. Conf. on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2025)*, Denver, CO, USA, 2025.
20. Abboud, A., Uncertainty Quantification Applied to Fuel Assembly Bow in a Pressurized Water Reactor, PhD, Institut Polytechnique de Paris, 2025.
21. Abboud, A., de Lambert, S., Garnier, J., Leturcq, B., and Lamorte, N., Sensitivity Analysis of a Flow Redistribution Model for a Multidimensional and Multifidelity Simulation of Fuel Assembly Bow in a Pressurized Water Reactor, *Nucl. Eng. Des.*, **443**:114259, 2025.
22. Abboud, A., Garnier, J., Leturcq, B., Lamorte, N., and de Lambert, S., Towards Uncertainty Quantification: Efficient Surrogate Models in Coupled Fluid-Structure Interaction for Fuel Assembly Bow, *Proc. of the 19th Int. Conf. on Probabilistic Safety Assessment and Analysis (PSA 2025)*, 2025.
23. Sobol', I.M., Global Sensitivity Indices for Nonlinear Mathematical Models and Their Monte Carlo Estimates, *Math. Comput. Simul.*, **55**:271–280, 2001.

24. Saltelli, A., Chan, K., and Scott, E.M., *Sensitivity Analysis*, Hoboken, NJ: Wiley, 2000.
25. Blanchard, J.B., Damblin, G., Martinez, J.M., Arnaud, G., and Gaudier, F., The Uranie Platform: An Open-Source Software for Optimisation, Meta-Modelling and Uncertainty Analysis, *EPJ Nucl. Sci. Technol.*, **5**:4, 2019.
26. Iooss, B. and Le Gratiet, L., Uncertainty and Sensitivity Analysis of Functional Risk Curves Based on Gaussian Processes, *Reliab. Eng. Syst. Saf.*, **187**:58–66, 2019.