

Epistemic and Aleatoric Uncertainty Quantification and Surrogate Modelling in High-Performance Multiscale Plasma Physics Simulations*

Yehor Yudin^{1,2}[0000-0001-8867-32073], David Coster¹[0000-0002-2470-9706], Udo von Toussaint¹[0000-0002-8867-1014], and Frank Jenko^{1,2}[0000-0001-6686-1469]

¹ Max-Planck-Institute for Plasma Physics, Garching 85748, Germany

² CS Department, Technical University of Munich, Garching 85748, Germany
{yehor.yudin, david.coster, udo.v.toussaint, frank.jenko}@ipp.mpg.de

Abstract. This work suggests several methods of uncertainty treatment in multiscale modelling and describes their application to a system of coupled turbulent transport simulations of a tokamak plasma. We propose a method to quantify the usually aleatoric uncertainty of a system in a quasistationary state, estimating the mean values and their errors for quantities of interest, which is average heat fluxes in the case of turbulence simulations. The method defines the stationarity of the system and suggests a way to balance the computational cost of simulation and the accuracy of estimation. This allows, contrary to many approaches, to incorporate aleatoric uncertainties in the analysis of the model and to have a quantifiable decision for simulation runtime. Furthermore, the paper describes methods for quantifying the epistemic uncertainty of a model and the results of such a procedure for turbulence simulations, identifying the model's sensitivity to particular input parameters and sensitivity to uncertainties in total. Finally, we introduce a surrogate model approach based on Gaussian Process Regression and present a preliminary result of training and analysing the performance of such a model based on turbulence simulation data. Such an approach shows a potential to significantly decrease the computational cost of the uncertainty propagation for the given model, making it feasible on current HPC systems.

Keywords: Uncertainty Quantification · Sensitivity Analysis · Plasma Physics · Multiscale Modelling · High Performance Computing · Surrogate Modelling · Gaussian Process Regression

1 Introduction

Thermonuclear fusion is a prospective source of clean and abundant renewable energy with a potential for worldwide deployment, independent of geography

* Supported by organisation “Munich School of Data Science”

and politics. To achieve fusion conditions, one needs to bring plasma temperature and density to sufficiently high values and isolate the plasma from the external matter for a sufficiently long time. One of the most studied methods to confine a plasma is to do it in a toroidal magnetic trap. The most developed type of device for confinement is a tokamak, an axisymmetric toroidal chamber with magnetic coils. One of the critical phenomena interesting for designing and analysing such confinement devices is the distribution of heat and density of the plasma over the minor radius of the torus. In the latest decades, it was discovered that transport properties of plasmas are dominated by microscopic dynamics of turbulent nature[7].

One of the challenges of such a study is the multiscale nature of the occurring phenomena, with properties of the processes happening across the size of an entire device, of the order of meters, during the confinement relevant times, of the order of seconds, are influenced and driven by microscopic processes, of the order of millimetres and microseconds.

One of the techniques in computational fields applied to tackle such challenges is coupling single-scale models, each solving a subset of equations describing processes happening on a particular scale and being implemented as a separate computer code. Such codes can, in principle, be run as a standalone program, accepting inputs and producing outputs, adhering to defined data structures.

In this work, we study the simulation of temperature profiles in a tokamak plasma, where plasma properties on the smallest scales define heat transport, and the equations for transport, equilibrium and microscopic turbulence are solved via separate codes, similar to the approach described by Falchetto et al. [7].

The following practical step interesting for scientific purposes is estimating uncertainties in the quantities of interest (QoIs) studied on the largest scale, considering that those uncertainties come from different sources.

To handle uncertainties, a quantity of interest can be treated as a random function of its parameters, also understood as random variables. In order to estimate uncertainties of function values using only a finite number of evaluations, assumptions can be made on the function, and an intermediate model or meta-model can be created based on a given sample of evaluations. This surrogate model describes the function's dependency on a selected subset of parameters and serves as a replacement for evaluations of the original function for practical purposes.

In this work, we apply surrogate modelling to the turbulent transport plasma simulation, replacing the solution of the microscopic turbulence equation with a fast model able to infer mean values and uncertainties of turbulent heat flux.

2 Methodology

2.1 Epistemic Uncertainty

One of the most common concerns in science is the uncertainty of inference or predictions due to a lack of knowledge of actual values of specific parameters

or lack of knowledge incorporated into the model, in other words, epistemic uncertainty.

In theoretical fields, particularly in computational areas, mathematical models often involve unknown parameters that can only be estimated probabilistically based on prior knowledge, observations or experiments. The uncertainty propagation problem means finding an approximation of the probability density function of quantities of interest in the model, given a set of uncertain parameters that are understood as random variables with known probability density functions.

As it is often difficult to retrieve a particular parametric form of QoI PDFs, the useful statistics describing the properties of output random variables are the statistical moments of the density.

Polynomial Chaos Expansion One of the most commonly used methods for uncertainty quantification is Polynomial Chaos Expansion (PCE)[18]. The core idea of the method is to expand the function of output Y of inputs X using the finite basis of orthonormal polynomials $P_i(x)$ up to the order of p , each depending on a subset of input components and being orthogonal to their PDFs:

$$Y \approx \hat{Y}(X) = \sum_{i=0}^N c_i P_i(X) \quad (1)$$

For the function $Y = f_X(x)$, the expansion coefficient ω_k and the ordinate values x_k could be determined via a quadrature scheme based on a Spectral Projection method exploiting the orthogonality of the basis polynomials and using normalisation factors H_i for polynomials P_i :

$$c_i = \frac{1}{H_i} \int_{\Omega} Y(x) P_i(x) f_X(x) dx \approx \frac{1}{H_i} \sum_{k=1}^N Y(x_k) P_i(x_k) \omega_k \quad (2)$$

The benefit of such an orthonormal approach is that the expansion coefficients are sufficient to calculate the estimates of the moments of the output PDF.

$$\hat{\mathbb{E}}[Y] = c_1, \quad \hat{\mathbb{V}}[Y] = \sum_{i=2}^N \mathbb{E}[P_i^2(X)] c_i \quad (3)$$

However, this method suffers from the curse of dimensionality as the number of required samples grows exponentially, like a binomial coefficient, with the number of dimensions d : $N = (d + p)!/p! \cdot d!$

Sensitivity Analysis A particular type of statistics over a sample of code's QoI responses aims to define the contribution of a single parameter, or a subset of parameters, to the uncertainty of the output. In this work, we study Sobol indices[17], which is a global variance-based sensitivity metric. The quantity we

estimate is the variance of the mathematical expectations of the function output, conditioned on a subset of input components $\{i\}$ and normalised by the total variance of the output. Here we use the first-order Sobol indices, identifying the influence of a single parameter i ; higher-order Sobol indices, determining the influence of a subset of interacting parameters $\{i\}$; and total Sobol indices, identifying the total effect of each parameter, using a set of all components $-i$, excluding a single i component:

$$S_{\{i\}} = \frac{\mathbb{V}[\mathbb{E}[Y|X_{\{i\}}]]}{\mathbb{V}[Y]}, \quad S_i^{tot} = 1 - \frac{\mathbb{V}[\mathbb{E}[Y|X_{-i}]]}{\mathbb{V}[Y]} \quad (4)$$

Using such an indicator, one could judge which parameters are most important for the behaviour of the computational model and decide whether a particular parameter's value can be fixed in subsequent studies or whether more values should be tried out.

2.2 Aleatoric Uncertainty

One of the types of uncertainties, different from the epistemic, is the aleatoric uncertainty coming from the inherent properties of the model and its solver. In the first place, it is related to the dynamic stochastic behaviour of the quantities of interests, for which the model is solved.

Several methods for analysing aleatoric uncertainties commonly involve two principles: separating epistemic uncertainties and performing statistical analyses of output quantities. One approach uses replica sampling, in which a certain number of model replicas are solved for each set of input parameters, followed by analysing the distributions of quantities of interest values[19]. Other methods are suitable for analysing the quasistationary nonlinear behaviour of the model in the vicinity of a certain attractor over a considerable time of solution[3].

This work utilises a practical measure for an aleatoric uncertainty based on analysis of the standard error of the mean values of quantities of interest $E = \text{SEM}[y] = \sigma[y]/\sqrt{n_{\text{eff}}}$ for a sample of observations \mathbf{Y}^{eff} , each taken in an autocorrelation time window of length t_{Δ} apart from the previous one. Such an approach allows judging the requirements of a single solution for a given set of parameters and controlling the level of aleatoric uncertainty. However, such a measure only provides information on the second moments of the quantity of interest distribution. Furthermore, it is not always well incorporable into a single framework of analysis of both epistemic and aleatoric uncertainties.

2.3 Surrogate Modelling

This work primarily focuses on applying Gaussian Process Regression (GPR), a non-parametric probabilistic machine learning model fully defined by the covariance matrix of observed data samples and capable of providing a posterior PDF of output given any input values from the considered support[15].

When the Matern-3/2 kernel k describes covariance between any two parametric points x_i, x_j , determining a PDF $p(f(x_*))$ of a posterior for new data x_* using a standard GPR based on a sample of observed inputs X and outputs y with a Normal likelihood looks like following:

$$p(f(x_*)|\mathbf{X}, \mathbf{y}, x_*) \sim \mathcal{N}(\mu(x_*), \sigma^2(x_*)) \quad (5)$$

$$\sigma^2(x_*) = K(x_*, x_*) - K(x_*, \mathbf{X})^\top (K(\mathbf{X}, \mathbf{X}) + \sigma_e^2(I))^{-1} K(x_*, \mathbf{X}) \quad (6)$$

$$K = (k(x_i, x_j))_{i,j=1..N} \quad (7)$$

$$k(x_i, x_j) = \sigma^2 \left(1 + \sqrt{3} \frac{|x_i - x_j|}{l} \right) \exp \left(-\sqrt{3} \frac{|x_i - x_j|}{l} \right) \quad (8)$$

The method allows for capturing uncertainties in the underlying data that come from the lack of information on the parametric dependencies and internal noise. Furthermore, this method has a high prediction uncertainty in the region of the parametric space for which there was not enough training data, as the posterior would have a higher variance for points decorrelated with the observed sample. Such a high uncertainty indicates the untrustworthiness of the regression results and the model's limitations.

One of the ideas behind the surrogate modelling approach is to create flexible and adaptive surrogates based on new data. Online Learning allows adding new data into the training sample to improve the surrogate gradually used as a proxy solver[10]. Active Learning uses the regression model's prediction uncertainty and information about sought dependency to determine new points in the parameter space for the evaluation of the function of interest[14].

In this work, we apply Gaussian Process Regression trained on a data set describing performed turbulence simulations. We trained such a surrogate model for the dependency of heat fluxes on the temperature values and its gradient for a particular configurational location in a tokamak plasma.

3 Numerical Results

3.1 Computational Model

The object of the study is the effects of turbulent processes happening at microscopic time and space scales on the temperature profiles of a nuclear fusion device across the confinement scales. For that, we use the Multiscale Fusion Workflow (MFW)[11] consisting of three main models, each implemented as a separate computer code, serving as an independent workflow component capable of exchanging its solution with others in a black-box fashion. The models are the equilibrium code that describes plasma geometry, the transport code that evolves the plasma profiles and the turbulence code used to compute effective transport coefficients (Fig 1). The code that solves equilibrium Grad-Shafranov equations for plasma density and the magnetic field is CHEASE[12], a high-resolution fixed-boundary solver. Evolution of temperature and density is performed with ETS[4], a 1D code for energy and particle transport.




Fig. 1: Main components of the Multiscale Fusion Workflow. The turbulence model is the one that requires the most computational resources to solve, and that exposes the highest sensitivity to its inputs.

In this work, we analyse the turbulence component of the workflow, for which we use the 3D nonlinear gyrofluid turbulence code GEM[16] in its local flux tube approximation. Such turbulence models are very sensitive to input parameters and computationally expensive to solve. Thus, it is a prime point for performing a study of uncertainties in the joint multiscale model and the component for which a fast and easy-to-evaluate surrogate model would be of interest to create. As discussed in [5], performing UQ on the MFW with GEM is currently computationally unreasonable since 100 000 core hours are required for a single run, whereas a surrogate model should be at least 10 000 times cheaper (1 core instead of 1024, and a factor of 10 per cycle around the MFW loop).

Codes store their respective solution in standardised data structures of Consistent Physical Objects (CPO)[8], and the turbulence quantities of interests are stored in a structure describing core transport properties. In principle, the workflow implemented to analyse the flux time traces and prepare a dataset for surrogate training and utilisation can support any code capable of interfacing via CPOs.

The resulting data structures are passed from a code to a code instance in a point-to-point fashion using the MUSCLE2 coupling library[2]. Currently, a version supporting various turbulence models is being implemented using the MUSCLE3 library[20].

3.2 Simulations and resulting time traces

A number of simulations of turbulence code for different input parameter values were performed to study the behaviour of quantities of interest in turbulence. Furthermore, this data was used to analyse what information can be included in a surrogate model and as an initial data set for the model training.

Transport equations describe temperature and density behaviour using a 1D approach, where quantities are functions of radial coordinate and time but must

be solved globally for the entire domain. Turbulence occurs on a smaller scale, and behaviour around a single magnetic flux tube or a small range of the radial domain typically defines it very well. In order to interpolate the turbulence solution for the entire device, the tokamak turbulence model for the ASDEX Upgrade[7] must be solved for at least eight different radial positions[5].

The turbulent model solution's quantity of interest is the heat flux averaged across a magnetic field surface at a specific radial position in the toroidal device. The computed result is nonlinear on microsecond timescales, with several features and stages typical for all relevant parameter space. Initially non-turbulent, the effective QoI value grows exponentially as turbulence develops until nonlinear effects dominate, causing turbulence to saturate and QoI to behave quasistationary. Although mean values remain around a particular mean value, the QoI is also locally chaotic and fluctuates substantially over time.

Defining whether the system is saturated, accurately representing the mean output values for a particular case, and decoupling fluctuating behaviour are crucial to creating an effective surrogate model parametrising the transport properties.

Macroscopic kinetic profiles behave similarly to micro-level turbulent fluxes over confinement-relevant times. Starting from initial conditions, they converge towards a stationary level self-consistent with the coupled model. The profiles dynamically and nonlinearly interact with local turbulent behaviour, resulting in a quasistationary solution. In order to capture the dependency of turbulent fluxes in relevant parameter space, a surrogate model must be based on a large sample of code solutions covering the relevant regions.

Creating a sample for surrogate training involves running the same code for multiple input parameter values and managing the simulation and training data, which pose a significant challenge. Using the resulting surrogate as a microscopic code replacement requires ensuring it captures QoI dependency well for all practically possible parameter values. This requires a training dataset covering enough parameter space information; the surrogate should indicate its applicability limits and epistemic uncertainty due to a lack of training data.

Given the practical challenges, this work suggests another novel computational workflow besides one of the multiscale simulations to solve the turbulent transport problem. The new workflow analyses turbulent fluxes time traces on the microscopic temporal scales and defines the stationarity of the QoI as well as its mean values and estimation error.

3.3 Aleatoric uncertainty

One of the work's challenges was to separate fundamental irreducible uncertainty from epistemic parametric uncertainty in a fluctuating system. It was important to establish the stationary mean values for analysing key parametric dependencies and to use fluctuations only for statistical error estimation. However, solving the model for a long enough time to gather sufficient data is computationally demanding and gathering a large sample would be prohibitively expensive. Es-

timating error levels during the course of a simulation run can help decide when to stop the simulation and significantly save computational costs.

The procedure for analysing the QoI time traces $y(t)$, representing a model solution for a single parametric point, constitutes of the following steps:

1. For the time-traces $y(t)$ of length t_n representing a model solution in a scalar quantity of interest at time steps $\{t\}$, we select a part in the saturated phase:
 - (a) Define the $y(t)$ ramp-up phase duration $t_{r.u.}$: here, in practice, for the long term, we chose an initial 15% of readings
 - (b) Discard the readings from the ramp-up phase
2. Downsample the readings:
 - (a) Calculate the Auto-Correlation Time: $t_A = \text{ACT}[y(t)] = \min t^* :$

$$\frac{1}{(t_n - t^*)} \sum_{t=t^*}^{t_n} (y(t) - \bar{y}) \cdot (y(t - t^*) - \bar{y}) < \frac{1}{t_n} \sum_{t=t_1}^{t_n} y^2(t) \cdot e^{-1}$$
 - (b) Split time series in saturated phase into $n_{\text{eff}} = \lfloor \frac{t_n}{t_A} \rfloor$ windows
 - (c) For every autocorrelation time window of length t_A choose a downsampled reading as a mean value $y_i^{\text{eff}} = \frac{1}{t_A} \sum_{j=i \cdot t_A}^{(i+1) \cdot t_A} y_j$ for an effective time step $t_i^{\text{eff}} = \frac{1}{t_A} \sum_{j=i \cdot t_A}^{(i+1) \cdot t_A} j$
 - (d) Collect downsampled readings into a new set $\mathbf{Y}^{\text{eff}} = \{y_i^{\text{eff}}\}$
3. Test the stationarity of the resulting time series:
 - (a) Here: compute an ordinary least-squares linear multivariate regression model of downsampled QoI readings over the time
 - (b) Apply Normal Equations to find coefficients: $\hat{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$ where \mathbf{X} consists of effective time steps t_i^{eff}
 - (c) Test if the linear regression coefficients are below a chosen relative tolerance: $\hat{\theta} < \epsilon_{\text{tol}}$
 - i. If the coefficients are too large, continue running the simulation for another t_{run} time steps
 - ii. If the coefficients are small enough, stop the simulation and proceed to the statistics calculation
4. Compute the essential statistics estimates:
 - (a) Mean: $\mu[y] = \frac{1}{n_{\text{eff}}} \sum_{y \in \mathbf{Y}^{\text{eff}}} y$
 - (b) Standard deviation: $\sigma[y] = \left(\frac{1}{n_{\text{eff}} - 1} \sum_{y \in \mathbf{Y}^{\text{eff}}} (y - \mu[y])^2 \right)^{-\frac{1}{2}}$
 - (c) Standard Error of the Mean: $\text{SEM}[y] = \sigma[y] / \sqrt{n_{\text{eff}}}$

We applied this procedure for the time traces produced by turbulence code iteratively, in particular to ion heat flux values over time, sampling after a number of solution time steps, and analysed the behaviour of the mean estimate and its error.

Figure 2 describes mean values calculated with the sequential procedure based on the data of time traces of heat fluxes in a quasistationary phase.

After the relative change of the mean and the relative change of the standard error of the mean converge below a certain threshold, as shown in Figure 3, we considered the duration of a simulation sufficient enough to have a reasonable estimate of the average flux and its error and added this data to a dataset describing the computational turbulence model behaviour. Subsequently, this dataset is used to draw conclusions on the model's parametric uncertainties and train a surrogate model.

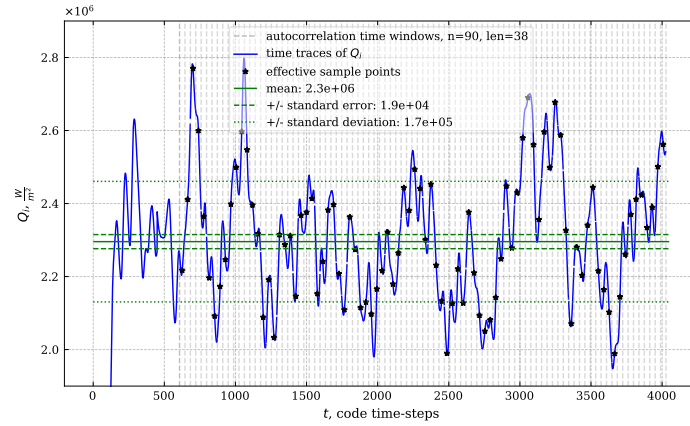


Fig. 2: The standard error of the mean is calculated accounting to discarding 15% of reading that do not belong to stationary phase, calculating the autocorrelation time of the remaining time traces, taken a sub-sample constituent of a single mean reading per autocorrelation window and renormalising mean error using the effective size of the sample.

3.4 Epistemic Uncertainty

To explore the epistemic uncertainties of the properties of the turbulence model, we performed an uncertainty propagation, taking a particular input of macroscopic parameter with a value taken from experimental measurement and, assuming that on top of this value, there is uncertainty, described as a normally distributed random variable with a Coefficient of Variation, a ratio of the variance of a quantity to its mean value, $CV(x) = \sqrt{V[x]}/E[x]$ equal to 0.1. The statistics for average ion and electron heat flux $\mathbf{y} = (Q_i(\rho), Q_e(\rho))$ were estimated using a PCE method with Hermite polynomials of order $p = 2$ applied for four input parameters of ion and electron temperatures profile values and gradients: $\mathbf{x} = T_i(\rho), T_e(\rho), \nabla T_i(\rho), \nabla T_e(\rho)$.

The uncertainty propagation was performed using EasyVVUQ[9] library, which manages the definition of input uncertainties, sampling scheme, generation of input data structures and collection of the output ones, as well as calculating statistics. The batches of the turbulence code instances were run on the MPCDF COBRA supercomputer using QCG-PilotJob[13] middle-ware library, which manages HPC resources of a single computational allocation and distributes it to run multiple instances of the code with a particular parallelisation set-up.

One of the issues of uncertainty propagation through the model that exposes chaotic behaviour on some scales is to untangle the epistemic and aleatoric uncertainties. Before analysing the parametric uncertainty propagation results, the time traces produced by the simulations are processed as described in the previous subsection.




Fig. 3: Sequential analysis of time traces with an interval of 150 time-steps per window, defined by 6 hours of computation on a single MPCDF’s COBRA node. After each step of time traces computation, we recalculate the autocorrelation time, repopulate an effective data sample, and estimate new values for the mean, standard deviation and standard error of the mean. With the growth of the sample, the mean estimation converges to the true value, and the standard error of the mean decreases to around 1% of the respective mean value after 2300 time steps out of 4000 in total. Here, one can decide to stop the computation using a tolerance threshold on the relative change of the mean estimate of 10^{-3} or by the error of the mean dropping below 1.1% Coefficient of Variation. The plots of the standard error of the mean and autocorrelation time dependency on number simulation steps indicate a convergence of the time traces analysis with simulation time.

The result of this study with an error of $CV(\mathbf{x}) = 0.1$ was inconclusive, as there was a significant overlap of the error bounds for mean values of quantities of interest for neighbouring points in parametric space (Fig. 4). In this case, it was impossible to conclude with certainty whether the ion heat flux trend was rising or falling with respect to the growth of ion temperature and its gradient. As




Fig. 4: Dependency of the mean ion heat flux on the ion temperature. For the small variation of input parameters, the mean estimation could not allow for drawing good conclusions on the gradient of this dependency, and the means are within uncertainty intervals of neighbouring points. For the larger variation of input parameters, the relative errors of the mean, also stated next to the points, are small enough to reconstruct the dependency with sufficient certainty.

the next step, we have taken an input parameter with normal distribution with $CV(\mathbf{x}) = 0.25$ and propagated this uncertainty through the model. Enlarging the input variation allowed us to study the behaviour and uncertainties of the model over the broader range of input parameter values and recover the character of the dependency of ion heat flux on ion temperature and its gradient (Fig. 4). Furthermore, longer simulation runs allowed reaching lower levels of the aleatoric uncertainty estimates in the quantities of interest,

The uncertainty analysis results per input component show a dominant parameter for which the model is most sensitive (Fig. 5). For this particular prior PDF over parametric space, the behaviour of the electron heat flux is most sensitive to variation in electron temperature. Still, other parameters and their interactions also play a significant role. Other metrics that were chosen to analyse the model's epistemic uncertainty are the Coefficient of Variation in the quantity of interest and the Uncertainty Amplification Factor[6]. In this work, the coefficient of variation of the model's output ion heat flux is about $CV(Q_i) \approx 2.96$, showing a significant range of possible values that QoI can take in the study.

The Uncertainty Amplification Factor is the ratio of the Coefficient of Variation of a QoI to the coefficient of variation of the input parameter $UAF(y|x) = CV(y)/CV(x)$. In the case of the studied model and chosen input parameters prior distribution, the Uncertainty Amplification Factor of ion heat flux is $UAF(Q_i|\mathbf{x}) \approx 11.9$ showing a significant sensitivity of the model to the parametric uncertainty.

3.5 Surrogate model

Having performed the original epistemic uncertainty propagation, we have collected a substantial sample of the physical model's input-output pairs. This sample was used to fit a surrogate model able to capture the behaviour of the physi-




Fig. 5: Sobol indices for electron and ion heat fluxes. Most of the variation in electron heat flux is produced by the variation in the electron temperature. Interactions of the input parameters explain some of the non-negligible variation. For the ion heat flux, the output variation is dominated by a single parameter in this regime. However, a fraction is still explained by the nonlinear interaction of parameters.

cal model QoI means values and uncertainties over the selected input parametric subspace. Here, the QoIs and the input parameters are the same as described in the epistemic uncertainty propagation results section.

The GPR surrogate provides an intermediate model based on the assumptions on the regularity of the actual dependency, primarily encoded in the covariance kernel and the likelihood of the regression model. Unlike the PCE surrogate, which depends on the predefined polynomial order and is difficult to adapt given new information about the sought dependency, the uncertainties of the GPR one produce strong indications that the information on the actual dependency is sufficient for the regression and is easily adaptable through sequential design and active learning algorithms.

Here we present the results of the Gaussian Process Regression surrogate trained on 50% of the collected simulation dataset (Fig. 6). We developed functionality for the training and analysis of surrogates based on EasySurrogate[1] library. After performing a hyper-parameter optimisation, a particular set of categorical parameter values was chosen to judge surrogates by the highest coefficient of determination of QoI values predicted by a surrogate y_i^{sur} compared to the ground truth simulation values $y_i^{\text{g.t.}}$ for the model testing: $R^2 = 1 - \frac{\sum_{i=1}^{n_{\text{test}}} (y_i^{\text{g.t.}} - y_i^{\text{sur}})^2}{\sum_{i=1}^{n_{\text{test}}} (y_i^{\text{g.t.}} - \bar{y}^{\text{g.t.}})^2}$

The future plan is to apply a validation test for the applicability of such a surrogate as a replacement for the solution of the turbulence model for given inputs describing core profiles and inferring the average heat fluxes and effective transport coefficients. The original simulation workflow utilising the turbulence code GEM required around 27.0 s on 1024 cores of Cineca MARCONI supercomputer[11] to predict fluxes for the next iteration of core profiles advancement, whether the EasySurrogate GPR model requires a fraction of a second on a single




Fig. 6: GPR surrogate model test results with the coefficient of determination $R^2 = 0.988$. Inferred values of ion heat flux compared to the observed values produced by the physical model. Vertical bars denote the ± 1.96 of the standard deviation of the predictive posterior distribution, grey bounds denote 15% relative error, and dotted lines denote the region of interest. For most test samples, the true value is within 95% prediction interval, and the predicted mean is within 15% relative error.

core for the same procedure. Furthermore, applying a GPR model would allow the detection of points in parameter space where the surrogate uncertainty is too high and should not be applied as a proxy for a high-fidelity solver. When it is understood that a GPR surrogate performs sub-par for a region of interest in parametric space, some work was done on designing an Active Learning scheme based on Bayesian Optimisation and Sequential Design of Experiments.

4 Discussion

In this paper, we presented a number of methods to treat uncertainties in multi-scale simulations. This includes ways to measure aleatoric and epistemic uncertainties, to control and reduce the amount of computational resources required to quantification of the uncertainties, as well as to identify situations when uncertainties are hindering analysis of the physical model behaviour. The suggested method for aleatoric uncertainty processing is applicable when some model dependency could be decomposed into the mean and fluctuation behaviour and allows for a trade-off between the computational effort for the model solution and the accuracy of the mean behaviour estimation. The surrogate modelling method proposes a way to make epistemic uncertainty propagation cheaper by constructing a data-based machine learning GPR proxy model for an expensive physical model solver that could be used as a replacement in multiscale coupled simulations and for sampling in uncertainty quantification methods.

We applied the described methods for the dependency of average heat fluxes produced by plasma turbulence code GEM based on tokamak core temperature profiles and their gradients. While studying such fluxes' time traces, we identified cases when the simulations converged early enough to allow reducing the duration of runs and saving computational resources. The QoI mean levels were estimated with sufficient certainty to draw conclusions about sought dependency and use solution data as a basis for further analysis and surrogate model training. The subsequently performed uncertainty propagation and sensitivity analysis showed situations where aleatoric uncertainties could not allow concluding on parametric dependencies. Also, we defined cases with a strong indication that a particular set of variables influences output variation dominantly. The surrogate modelling approach of fitting a Gaussian Process Regression model to the dataset of flux computing turbulent simulations showed that it is possible to train a metamodel that allows capturing the mean behaviour of the fluxes well within the posterior predictive bounds for a region of interest in parameter space of the model.

The development of a surrogate for the turbulence code opens up the much-needed possibility of doing detailed UQ for the whole MFW workflow with the accuracy of the underlying turbulence code but without the expense of the turbulence code in the workflow.

Future work would include testing the surrogate as a proxy of the expensive turbulence code, capable of inferring fluxes and subsequently defining the effective transport coefficients. This includes implementing an expanded version of a Multiscale Fusion Workflow, capable of switching between different implementations of the turbulence model and validating the results of the coupled turbulent transport simulations using a turbulence surrogate against a version using the gyrofluid code. Subsequently, such a modified workflow should include capacity for adaptive surrogate training through Online Learning and Active Learning algorithms, as well as methods to select turbulence model implementation in the course of a single simulation.

5 Acknowledgements

The authors of this paper would like to acknowledge the support of the Poznan Supercomputer and Networking Centre (PSNC) and Max-Planck Computational Data Facility (MPCDF). Research by Yehor Yudin is funded by the Helmholtz Association under the "Munich School for Data Science - MuDS".

References

1. Easysurrogate github page, <https://github.com/wedeling/EasySurrogate>, last accessed 2 Mar 2023
2. Borgdorff, J., et al.: Distributed multiscale computing with MUSCLE 2, the Multi-scale Coupling Library and Environment. *Journal of Computational Science* **5**(5), 719–731 (Sep 2014). <https://doi.org/10.1016/j.jocs.2014.04.004>

3. Brajard, J., et al.: Combining data assimilation and machine learning to infer unresolved scale parametrization. *Phil. Trans. of the Roy. Soc. A: Math., Phys. and Eng. Sc.* **379**(2194) (2021). <https://doi.org/10.1098/rsta.2020.0086>
4. Coster, D., et al.: The European Transport Solver. *IEEE Transactions on Plasma Science* **38**(9), 2085–2092 (Sep 2010). <https://doi.org/10.1109/TPS.2010.2056707>
5. Coster, D., et al.: Building a turbulence-transport workflow incorporating uncertainty quantification for predicting core profiles in a tokamak plasma. *Nuclear Fusion* (2021)
6. Edeling, W., et al.: The impact of uncertainty on predictions of the CovidSim epidemiological code. *Nature Computational Science* **1**(2). <https://doi.org/10.1038/s43588-021-00028-9>
7. Falchetto, G., et al.: The European Integrated Tokamak Modelling (ITM) effort: achievements and first physics results. *Nuclear Fusion* **54**(4), 043018 (Apr 2014). <https://doi.org/10.1088/0029-5515/54/4/043018,00005>
8. Imbeaux, F., et al.: A generic data structure for integrated modelling of tokamak physics and subsystems. *Computer Physics Communications* **181**(6), 987–998 (Jun 2010). <https://doi.org/10.1016/j.cpc.2010.02.001>
9. Jancauskas, V., Lakhilili, J., Richardson, R., Wright, D.: EasyVVUQ: Verification, validation and uncertainty quantification for HPC simulations (2021), <https://github.com/UCL-CCS/EasyVVUQ>
10. Leiter, K., Barnes, B., Becker, R., Knap, J.: Accelerated scale-bridging through adaptive surrogate model evaluation. *Journal of Computational Science* **27**, 91–106 (2018). <https://doi.org/10.1016/j.jocs.2018.04.010>
11. Luk, O., et al.: ComPat framework for multiscale simulations applied to fusion plasmas. *Computer Physics Communications* **239**, 126–133 (Jun 2019). <https://doi.org/https://doi.org/10.1016/j.cpc.2018.12.021>
12. Lütjens, H., Bondeson, A., Sauter, O.: The cheese code for toroidal mhd equilibria. *Computer Physics Communications* **97**(3), 219–260 (1996)
13. Piontek, T., et al.: Development of science gateways using qcg – lessons learned from the deployment on large scale distributed and hpc infrastructures. *J. of Grid Comp.* **14**(4), 559–573 (2016). <https://doi.org/10.1007/s10723-016-9384-9>
14. Preuss, R., von Toussaint, U.: Global optimization employing gaussian process-based bayesian surrogates. *Entropy* **20** (2018). <https://doi.org/10.3390/e20030201>
15. Rasmussen, C.E., W.K.: *Gaussian Processes for Machine Learning.* the MIT Press (2006). <https://doi.org/10.7551/mitpress/3206.001.0001>
16. Scott, B.D.: Free-energy conservation in local gyrofluid models. *Physics of Plasmas* **12**(10), 102307 (Oct 2005). <https://doi.org/10.1063/1.2064968>
17. Sobol, I.: On sensitivity estimation for nonlinear mathematical models. *Mathematical Modeling* **2**(1), 112–118 (1990)
18. Sullivan, T.: *Introduction to Uncertainty Quantification, TAM*, vol. 63. Springer Cham (2015). <https://doi.org/10.1007/978-3-319-23395-6>
19. Vassaux, M., et al.: Ensembles are required to handle aleatoric and parametric uncertainty in molecular dynamics simulation. *Journ. of Chem. Th. and Comp.* **17**(8) (2021). <https://doi.org/10.1021/acs.jctc.1c00526>
20. Veen, L., Hoekstra, A.: Easing multiscale model design and coupling with MUSCLE 3. In: *Computational Science – ICCS 2020*. pp. 425–438. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-50433-5_33