Statistical test to detect covariance misspecification in Gaussian process regression: New insights into non-stationarity detection

Théo Sylvestre[†,1,3], Guillaume Damblin[§,1], Raksmy Nop[§,1], Amandine Marrel[¶,2,3]

[†] PhD student [¶] PhD supervisor [§] PhD co-supervisor

PhD expected duration: Nov. 2024 – Oct. 2027

[1] CEA Saclay, 91190 Gif-sur-Yvette, France,
{theo.sylvestre,guillaume.damblin,raksmy.nop}@cea.fr

[2] CEA, DES, IRESNE, DER, Cadarache, France,
amandine.marrel@cea.fr

[3] LMA Université d'Avignon, EA 2151, 84029 Avignon,

Abstract

In computer experiments, numerical codes aim to simulate physical phenomena as accurately as possible. However, those high-fidelity simulation tools are time-consuming, preventing their intensive use for an in-depth understanding of the simulated phenomena. To overcome this, the simulator is commonly replaced by a surrogate model. One popular choice is the Gaussian Process Regression ($\mathcal{GPR}$) model, a non-parametric approach that naturally quantifies uncertainty in its predictions. A $\mathcal{GPR}$ model is fully specified by its mean function and covariance kernel $k(\cdot,\cdot)$. In most applications, $k(\cdot,\cdot)$ is assumed to be stationary, meaning that it only depends on the difference between two input locations, not on their absolute positions. Assuming such stationary covariance may be inappropriate when the numerical model exhibits distinct behaviours across different regions of the input space. This situation arises in some thermal-hydraulic transients studied in reactor safety. One notable example is reactivity insertion accidents occurring in pool-type research reactors, where non-stationarity is suspected by experts. Appropriate covariance kernels, whose structure depends on absolute input locations rather than solely on relative differences, should then be used, leading to complex non-stationary $\mathcal{GPR}$ models. However, before resorting to such advanced methods, it is important to assess whether non-stationarity is actually present. In practice, this is commonly inferred from diagnostics applied to an initially fitted stationary $\mathcal{GPR}$ model. In this framework, our contribution is a new statistical test designed to assess whether the assumed covariance $k(\cdot,\cdot)$ is adequate for a given $\mathcal{GPR}$ problem. This approach provides an alternative way to validate $\mathcal{GPR}$ models and offers new insights into the detection of non-stationarity. In particular, if a broad class of stationary kernels is rejected, this strongly suggests the presence of non-stationarity in the data. For each tested covariance kernel, the idea is to compare some observed validation criteria with the values that these criteria would have taken if the data had truly been generated from the corresponding $\mathcal{GPR}$ model. To construct the test, we specifically focus on two complementary criteria commonly used for $\mathcal{GPR}$ model validation, with one based on predictions and the other on predictive variances. After showing encouraging results for detecting covariance misspecification on several numerical examples, the test is applied to the reactivity insertion use case. In this setting, the approach supports the experts' prior assessment of the presence of non-stationarity and further justifies the need for a non-stationary $\mathcal{GPR}$ model.