Gaussian Process and Derivative Matching

Reduced Order Modeling Project

Problem Statement

Gaussian Processes (GP) are non-parametric statistical models used in regression, classification, and other tasks. They are particularly useful for making predictions with uncertainty quantification. However, certain applications, such as solving differential equations or gradient-based optimization, require not only accurate predictions but also accurate derivative estimations.
The primary challenge addressed by this project is:

- How to train a Gaussian Process (GP) model to ensure the derivative of the GP's posterior mean, $\mu'(x)$, closely aligns with the derivative of the true function, $f'(x)$, while also providing accurate predictions?

Project Contributions

- Implement a Gaussian Process (GP) model using a Radial Basis Function (RBF) kernel:

$$k(x, x') = \sigma^2 \exp\left(-\frac{||x - x'||^2}{2l^2}\right)$$

- Introduce a noise model to account for observation noise, adding $\sigma_n^2 I$ to the covariance matrix.
- Develop a training loop incorporating the Negative Log Marginal Likelihood (NLML), defined as:

$$\text{NLML} = \frac{1}{2} y^T K^{-1} y + \frac{1}{2} \log |K| + \frac{n}{2} \log(2\pi)$$

- Create a method to match the GP's derivative with the true function's derivative by computing the gradient of the GP's predictive mean
- Visualize the GP's predictions, including plotting the GP's posterior mean and examining the impact of derivative matching on training loss and prediction accuracy.