

Title: Exploring Uncertainty with Gaussian Processes: An Interactive Tutorial

Brief description:

Uncertainty is central to trustworthy AI. Gaussian Processes (GPs) are a key probabilistic model in machine learning, providing not just predictions but also confidence. This interactive tutorial makes GPs accessible through a weather forecasting example. Learners progress from Gaussian distributions → regression → Gaussian Processes, exploring how priors encode expectations, Bayes' rule updates beliefs with data, and marginalisation yields predictions at new points. The focus is on intuition and interactivity, with formulas minimal and optional.

Target audience:

- High school learners, early undergraduates, and professionals new to GPs
- Assumed knowledge: mean, variance, Gaussian distribution, simple regression

Learning objectives:

1. Recall the Gaussian distribution as a foundation for uncertainty
2. Understand how GPs extend linear regression by modelling many possible functions
3. See how prior beliefs and data can be combined through Bayes' rule (conditioning)
4. Explore the trade-off between simple and complex models (via hyperparameters)
5. Learn how GPs make predictions at unseen data points (marginalisation)

Duration

- Main tutorial: 10-12 min
- Optional deep-dive: 3-5min
- Total: ~12-15 min

Materials

- *Intuitive_GaussianProcesses_interactive.html* : Interactive tutorial with plots, sliders, and animations
- *marginalisation.html*: Optional deep-dive animation + explanation
- *Intuitive_GaussianProcesses_static.pdf*: Static fallback version

How to Use: Open the HTML file in any modern browser (tested in Chrome/Firefox).