

Posterior Gaussian Process

Hong Ge

Oct, 2025

Key concepts

- we are not interested in random functions
- we want to *condition* on the training data
- when both prior and likelihood are Gaussian, then
 - posterior is a Gaussian process
 - predictive distributions are Gaussian
- pictorial representation of prior and posterior
- interpretation of predictive equations

Gaussian Process Inference

Recall Bayesian inference in a parametric model.

The posterior is proportional to the prior times the likelihood.

The predictive distribution is the predictions marginalized over the parameters.

How does this work in a Gaussian Process model?

Answer: in our non-parametric model, the “parameters” are the function itself!

Non-parametric Gaussian process models

In our non-parametric model, the “parameters” are the function itself!
The joint distribution

$$\begin{aligned} p(\mathbf{f}, \mathbf{y}) &= p(\mathbf{f}) p(\mathbf{y}|\mathbf{f}) = p(\mathbf{y}) p(\mathbf{f}|\mathbf{y}) \\ \implies \mathcal{N}(\mathbf{f}|\mathbf{m}, \mathbf{k}) \mathcal{N}(\mathbf{y}|\mathbf{f}) &= Z_{|\mathbf{y}} \mathcal{N}(\mathbf{f}|\mathbf{m}_{|\mathbf{y}}, \mathbf{k}_{|\mathbf{y}}). \end{aligned}$$

Gaussian process prior with zero mean and covariance function k

$$p(\mathbf{f}|\mathcal{M}_i) \sim \mathcal{N}(\mathbf{f}|\mathbf{m} \equiv \mathbf{0}, \mathbf{k}),$$

Gaussian likelihood, with noise variance σ_{noise}^2



$$p(\mathbf{y}|\mathbf{f}, \mathcal{M}_i) \sim \mathcal{N}(\mathbf{f}, \sigma_{\text{noise}}^2 \mathbf{I}),$$

leads to a Gaussian process posterior

$$p(\mathbf{f}|\mathbf{y}, \mathcal{M}_i) \sim \mathcal{N}(\mathbf{f}|\mathbf{m}_{|\mathbf{y}}, \mathbf{k}_{|\mathbf{y}}),$$

$$\text{where } \begin{cases} \mathbf{m}_{|\mathbf{y}}(\mathbf{x}) = \mathbf{k}(\mathbf{x}, \mathbf{x}) [\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{y}, \\ \mathbf{k}_{|\mathbf{y}}(\mathbf{x}, \mathbf{x}') = \mathbf{k}(\mathbf{x}, \mathbf{x}') - \mathbf{k}(\mathbf{x}, \mathbf{x}) [\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{k}(\mathbf{x}, \mathbf{x}'). \end{cases}$$

Prior and Posterior



Predictive distribution:

$$p(y_* | x_*, \mathbf{x}, \mathbf{y}) \sim \mathcal{N}(\mathbf{k}(x_*, \mathbf{x})^\top [\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{y}, \\ \mathbf{k}(x_*, x_*) + \sigma_{\text{noise}}^2 - \mathbf{k}(x_*, \mathbf{x})^\top [\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{k}(x_*, \mathbf{x}))$$

Some interpretation

Recall our main result:

$$f_* | \mathbf{x}_*, \mathbf{x}, \mathbf{y} \sim \mathcal{N}(\underbrace{\mathbf{K}(\mathbf{x}_*, \mathbf{x})[\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{y}}_{\text{mean}}, \underbrace{\mathbf{K}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{K}(\mathbf{x}_*, \mathbf{x})[\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{K}(\mathbf{x}, \mathbf{x}_*)}_{\text{variance}}).$$

The **mean** is linear in two ways:

$$\mu(\mathbf{x}_*) = \underbrace{\mathbf{k}(\mathbf{x}_*, \mathbf{x})[\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} \mathbf{y}}_{\text{mean}} = \sum_{n=1}^N \beta_n y_n = \sum_{n=1}^N \alpha_n k(\mathbf{x}_*, \mathbf{x}_n).$$

The last form is most commonly encountered in the kernel literature.

The **variance** is the difference between two terms:

$$V(\mathbf{x}_*) = \underbrace{k(\mathbf{x}_*, \mathbf{x}_*)}_{\text{prior variance}} - \underbrace{k(\mathbf{x}_*, \mathbf{x})[\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_{\text{noise}}^2 \mathbf{I}]^{-1} k(\mathbf{x}, \mathbf{x}_*)}_{\text{explained variance}},$$

the first term is the *prior variance*, from which we subtract a (positive) term, telling how much the data \mathbf{x} has explained.

Note, that the variance is independent of the observed outputs \mathbf{y} .