

Gaussian processes


Gaussian processes (GPs) are distributions over functions from an input x , which could be high-dimensional and could be a complex object like a string or graph, to a scalar or 1-dimensional output $f(x)$. We will use a Gaussian process prior over functions in a Bayesian approach to regression. Gaussian processes can also be used as a part of larger models.

Goals: 1) model complex functions: not just straight lines/planes, or a combination of a fixed number of basis functions. 2) Represent our uncertainty about the function, for example with error bars. 3) Do as much of the mathematics as possible analytically.

The value of uncertainty

Using Gaussian processes, like Bayesian linear regression, we can compute our uncertainty about a function given data. Error bars are useful when making decisions, or optimizing expensive functions. Examples of expensive functions include ‘efficacy of a drug’, or ‘efficacy of a training procedure for a large-scale neural network’.

In the cartoon below, a black dashed line shows an estimate of a function, with a gray region indicating uncertainty. The function could be performance of a system as a function of its settings. Often the performance of a system doesn’t vary up and down rapidly as a function of one setting, but in high-dimensions there could easily be multiple modes — different combinations of inputs that work quite well.



The largest value appears to be around $x=0.4$. If we can afford to run a couple of experiments, we might also test the system for $x=0.8$, which might actually perform better. According to our point estimates, the performance at $x=0$ is similar to $x=0.8$. However, we aren’t uncertain at $x=0$, so it isn’t worth running experiments there.

Using uncertainty in a probabilistic model to guide search is called *Bayesian optimization*. We could do Bayesian optimization with Bayesian linear regression. Although we would have to be careful: Bayesian linear regression is often too certain if the model is too simple, or doesn’t include enough basis functions.

Bayesian optimization can also be based on neural networks. However, representing the posterior distribution over parameters, and integrating to give a predictive distribution, is challenging for neural networks.

The multivariate Gaussian distribution

We’ll see that Gaussian processes are really just high-dimensional multivariate Gaussian distributions. This section has a quick review of some of the things we can do with Gaussians. It’s perhaps surprising that these manipulations can answer interesting statistical and machine learning questions!

A Gaussian distribution is completely described by its parameters μ and Σ :

$$p(\mathbf{f} | \Sigma, \mu) = |2\pi\Sigma|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\mathbf{f} - \mu)^\top \Sigma^{-1}(\mathbf{f} - \mu)\right),$$

where μ and Σ are the mean and covariance:

$$\mu_i = \mathbb{E}[f_i]$$

$$\Sigma_{ij} = \mathbb{E}[f_i f_j] - \mu_i \mu_j.$$

The covariance matrix Σ must be positive definite. (Or positive semi-definite if we allow some of the elements of \mathbf{f} to depend deterministically on each other.) If we know a distribution is Gaussian and know its mean and covariances, we know its density function.

Any marginal distribution of a Gaussian is also Gaussian. So given a joint distribution:

$$p(\mathbf{f}, \mathbf{g}) = \mathcal{N}\left(\begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}; \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}, \begin{bmatrix} A & C \\ C^\top & B \end{bmatrix}\right),$$

then as soon as you convince yourself that the marginal

$$p(\mathbf{f}) = \int p(\mathbf{f}, \mathbf{g}) \, d\mathbf{g}$$

is Gaussian, you already know the means and covariances:

$$p(\mathbf{f}) = \mathcal{N}(\mathbf{f}; \mathbf{a}, A).$$

Any conditional distribution of a Gaussian is also Gaussian:

$$p(\mathbf{f} | \mathbf{g}) = \mathcal{N}(\mathbf{f}; \mathbf{a} + C B^{-1}(\mathbf{g} - \mathbf{b}), A - C B^{-1} C^\top)$$

Showing this result from scratch requires quite a few lines of linear algebra. However, this is a standard result that is easily looked up (and I wouldn't make you show it in an exam!).

Representing function values with a Gaussian

We can think of functions as infinite-dimensional vectors. Dealing with the mathematics of infinities and real numbers rigorously is possible but involved. I will side-step these difficulties with a low-tech description.

We could imagine a huge finite discretization of our input space. For example, we could consider only the input vectors \mathbf{x} that can be realized with IEEE floating point numbers. After all, those are the only values we'll ever consider in our code. In theory (but not in practice) we could evaluate a function f at every discrete location, $f_i = f(\mathbf{x}^{(i)})$. If we put all the f_i 's into a vector \mathbf{f} , that vector specifies the whole function (up to an arbitrarily fine discretization).

While we can't store a vector containing the function values for every possible input, it is possible to define a multivariate Gaussian prior on it. Given noisy observations of some of the elements of this vector, it will turn out that we can infer other elements of the vector without explicitly representing the whole object.

If our model is going to be useful, the prior needs to be sensible. If we don't have any specific domain knowledge, we'll set the mean vector to zero. If we drew "function vectors" from our prior, an element $f_i = f(\mathbf{x}^{(i)})$, will be positive just as often as it is negative. Defining the covariance is harder. We certainly can't use a diagonal covariance matrix: if our beliefs about the function values are independent, then observing the function in one location will tell us nothing about its values in other locations. We usually want to model continuous functions: so function values for nearby inputs should have large covariances.

We will define the covariance between two function values using a covariance or *kernel* function:

$$\text{cov}[f_i, f_j] = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}).$$

There are families of positive definite kernel functions ("Mercer kernels"), which always produce a positive definite matrix K if each element is set to $K_{ij} = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$. If the kernel wasn't positive definite, our prior wouldn't define a valid Gaussian distribution. One kernel that is positive definite is proportional to a Gaussian:

$$k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|^2),$$


but there are many other choices. For example, another valid kernel is:

$$k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = (1 + \|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|) \exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|).$$

Gaussian processes get their name because they define a Gaussian distribution over a vector of function values. *Not* because they sometimes use a Gaussian kernel to set the covariances.

Using the Gaussian process for regression

The plot below shows three functions drawn from a Gaussian process prior with zero mean and a Gaussian kernel function.




The 'functions' plotted here are actually samples from a 100-dimensional Gaussian distribution. Each sample gave the height of a curve at 100 locations along the x -axis, which were joined up with straight lines. If we chose a different kernel we could get rough functions, straighter functions, or periodic functions.

Our prior is that the function comes from a Gaussian process, $f \sim \mathcal{GP}$, where a vector of function values has prior $p(\mathbf{f}) = \mathcal{N}(\mathbf{f}; \mathbf{0}, K)$, where $f_i = f(\mathbf{x}^{(i)})$ and $K_{ij} = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$.

Given noisy observations of some of the function values:

$$y_i \sim \mathcal{N}(f_i, \sigma_n^2),$$

we can update our beliefs about the function. Here the n in σ_n^2 stands for *noise*. According to the model, the observations are centred around one underlying true function as follows:



But we don't know where the function is, we only see the y observations.

In 1D we can represent the whole function fairly accurately (as in the plot above) with ~ 100 values. In high-dimensions however, we can't grid up the input space and explicitly estimate the function everywhere. Instead, we only consider the function at places we have observed, and at test locations $X_* = \{\mathbf{x}^{(*,i)}\}$ where we will make predictions. We call the vector of function values at the test locations \mathbf{f}_* .

We can write down the model's joint distribution of the observations and the \mathbf{f}_* function values that we're interested in. This joint distribution is Gaussian, as it's simply a marginal of our prior over the whole function, with noise added to some of the values:

$$P\left(\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix}\right) = \mathcal{N}\left(\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix}; \mathbf{0}, \begin{bmatrix} K(X, X) + \sigma_n^2 \mathbf{I} & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix}\right)$$

The compact notation $K(X, Y)$ follows the Rasmussen and Williams GP textbook (Murphy uses an even more compact notation). Given an $N \times D$ matrix¹ X and an $M \times D$ matrix Y , $K(X, Y)$ is an $N \times M$ matrix of kernel values:

$$K(X, Y)_{ij} = k(\mathbf{x}^{(i)}, \mathbf{y}^{(j)}),$$

where $\mathbf{x}^{(i)}$ is the i th row of X and $\mathbf{y}^{(j)}$ is the j th row of Y .

The marginal covariance of the function values \mathbf{f}_* at the test locations, $K(X_*, X_*)$ is given directly by the prior. The marginal covariance of the observations, $K(X, X) + \sigma_n^2 \mathbf{I}$, is given by the prior, plus the independent noise variance. The cross covariances, $K(X, X_*)$ come from the prior on functions, the noise has no effect:

$$\begin{aligned} \text{cov}(y_i, f_{*,j}) &= \mathbb{E}[y_i f_{*,j}] - \mathbb{E}[y_i] \mathbb{E}[f_{*,j}] \\ &= \mathbb{E}[(f_i + v_i) f_{*,j}], \quad v_i \text{ is noise from } N(0, \sigma_n^2) \\ &= \mathbb{E}[f_i f_{*,j}] + \mathbb{E}[v_i] \mathbb{E}[f_{*,j}], \\ &= \mathbb{E}[f_i f_{*,j}] = k(\mathbf{x}^{(i)}, \mathbf{x}^{(*j)}). \end{aligned}$$

Because Gaussians marginalize so easily, we can ignore the enormous (or infinite) prior covariance matrix over all of the function values that we're not interested in.


Using the rule for conditional Gaussians, we can immediately identify that the posterior over function values $p(\mathbf{f}_* | \mathbf{y})$ is Gaussian with:

$$\begin{aligned} \text{mean}, \bar{\mathbf{f}}_* &= K(X_*, X)(K(X, X) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \\ \text{cov}(\mathbf{f}_*) &= K(X_*, X_*) - K(X_*, X)(K(X, X) + \sigma_n^2 \mathbf{I})^{-1} K(X, X_*) \end{aligned}$$

This posterior distribution over a few function values is a marginal distribution of the posterior distribution over the whole function. The posterior distribution over possible functions is itself a Gaussian Process (a large or infinite Gaussian distribution).

Visualizing the posterior


One way to visualize what we believe about the function is to sample plausible realizations from the posterior distribution. The figure below shows three functions sampled from the posterior, two of them are very close to each other.



Really I plotted three samples from a 100-dimensional Gaussian, $p(\mathbf{f}_* | \text{data})$, where 100 test locations were chosen on a grid. We can see from these samples that we're fairly sure about the function for $x \in [0, 0.4]$, but less sure of what it's doing for $x \in [0.5, 0.8]$.

We can summarize the uncertainty at each input location by plotting the mean of the posterior distribution and error bars. The figure below shows the mean of the posterior plotted as a black line, with a grey band indicating ± 2 standard deviations.

1. Actually the inputs don't have to be D -dimensional vectors. There are kernel functions for graphs and strings. Once we have computed the covariances, none of the remaining mathematics looks at the feature vectors themselves.



This figure doesn't show how the functions might vary within that band, for example whether they are smooth or rough.² However, it might be easier to read, and is cheaper to compute. We don't need to compute the posterior covariances between test locations to plot the mean and error band. We just need the 1-dimensional posterior at each test point:

$$p(f(\mathbf{x}^{(*)}) | \text{data}) = \mathcal{N}(f; m, s^2).$$

To identify the mean m and variance s^2 , we need the covariances:

$$K_{ij} = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}), \quad (\mathbf{k}^{(*)})_i = k(\mathbf{x}^{(*)}, \mathbf{x}^{(i)}).$$

Then the computation is a special case of the posterior expression already provided, but with only one test point:

$$\begin{aligned} M &= K + \sigma_n^2 \mathbf{I} \\ m &= \mathbf{k}^{(*)\top} M^{-1} \mathbf{y} \\ s^2 &= k(\mathbf{x}^{(*)}, \mathbf{x}^{(*)}) - \underbrace{\mathbf{k}^{(*)\top} M^{-1} \mathbf{k}^{(*)}}_{\text{positive}}. \end{aligned}$$

The mean prediction m is just a linear combination of observed function values \mathbf{y} .

After observing data, our beliefs are always more confident than under the prior. The amount that the variance improves, $\mathbf{k}^{(*)\top} M^{-1} \mathbf{k}^{(*)}$, doesn't actually depend on the \mathbf{y} values we observe! These properties of inference with Gaussians are computationally convenient: we can pick experiments, knowing how certain we will be after getting its result. However, these properties are somewhat unrealistic. When we see really surprising results, we usually become less certain, because we realize our model is probably wrong. In the Gaussian process framework we can adjust our kernel function in response to the observations, and then our uncertainties will depend on the observations again.

What you should know

- The marginal and conditional of a joint Gaussian are Gaussian. Be able to write down a marginal distribution given a joint.
- GP idea: explain how a smooth curve or surface relates to a draw from a multivariate Gaussian distribution. How would you plot one?

Check your understanding

- Given a GP model and N observations in a regression task, what do you need to compute to predict the function value (with an error bar) at a test location? What is the computational cost of these operations? Would you actually be able to do it? Or is there anything else you need to know?

² Sometimes we care about dependencies in predictions: for example, we might care whether a model says that several stock-prices are likely to fall together when they fall.

Reading

The core recommended reading for Gaussian processes is Murphy, Chapter 15, to section 15.2.4 inclusive.

Alternatives are: Barber Chapter 19 to section 19.3 inclusive, or the dedicated Rasmussen and Williams book³ up to section 2.5. There is also a chapter on GPs in MacKay's book.

Gaussian processes are Bayesian kernel methods. Murphy Chapter 14 has more information about kernels in general if you want to read about the wider picture.

The idea of *Bayesian optimization* is old. However, a relatively recent paper rekindled interest in the idea as a way to tune machine learning methods: <http://papers.nips.cc/paper/4522-practical-bayesian-optimization-of-machine-learning-algorithms> These authors then formed a startup around the idea, which was acquired by Twitter. Other companies like Netflix have also used their software.

3. <http://gaussianprocess.org/gpml/>