

CSC 411 Lecture 20: Gaussian Processes


Roger Grosse, Amir-massoud Farahmand, and Juan Carrasquilla

University of Toronto

- Last lecture: Bayesian linear regression, a parametric model
- This lecture: Gaussian processes
 - Derive as a generalization of Bayesian linear regression, with possibly infinitely many basis functions
 - Define a distribution directly over functions (i.e., a [stochastic process](#))
 - Based on the Kernel Trick, one of the most important ideas in machine learning
 - Conceptually cleaner, since we can specify priors directly over functions. This lets us easily incorporate assumptions like smoothness, periodicity, etc., which are hard to encode as priors over regression weights.

Towards Gaussian Processes

- **Gaussian Processes** are distributions over functions.
- They're actually a simpler and more intuitive way to think about regression, once you're used to them.



— GPML

Towards Gaussian Processes

- A Bayesian linear regression model defines a distribution over functions:

$$f(\mathbf{x}) = \mathbf{w}^\top \boldsymbol{\psi}(\mathbf{x})$$

Here, \mathbf{w} is sampled from the prior $\mathcal{N}(\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)$.

- Let $\mathbf{f} = (f_1, \dots, f_N)$ denote the vector of function values at $(\mathbf{x}_1, \dots, \mathbf{x}_N)$.
- By the linear transformation rules for Gaussian random variables, the distribution of \mathbf{f} is a Gaussian with

$$\mathbb{E}[f_i] = \boldsymbol{\mu}_w^\top \boldsymbol{\psi}(\mathbf{x}_i)$$

$$\text{Cov}(f_i, f_j) = \boldsymbol{\psi}(\mathbf{x}_i)^\top \boldsymbol{\Sigma}_w \boldsymbol{\psi}(\mathbf{x}_j)$$

- In vectorized form, $\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}_f, \boldsymbol{\Sigma}_f)$ with

$$\boldsymbol{\mu}_f = \mathbb{E}[\mathbf{f}] = \boldsymbol{\Psi} \boldsymbol{\mu}_w$$

$$\boldsymbol{\Sigma}_f = \text{Cov}(\mathbf{f}) = \boldsymbol{\Psi} \boldsymbol{\Sigma}_w \boldsymbol{\Psi}^\top$$

Towards Gaussian Processes

- Recall that in Bayesian linear regression, we assume noisy Gaussian observations of the underlying function.

$$y_i \sim \mathcal{N}(f_i, \sigma^2) = \mathcal{N}(\mathbf{w}^\top \boldsymbol{\psi}(\mathbf{x}_i), \sigma^2).$$

- The observations \mathbf{y} are jointly Gaussian, just like \mathbf{f} .

$$\begin{aligned} \mathbb{E}[y_i] &= \mathbb{E}[f(\mathbf{x}_i)] \\ \text{Cov}(y_i, y_j) &= \begin{cases} \text{Var}(f(\mathbf{x}_i)) + \sigma^2 & \text{if } i = j \\ \text{Cov}(f(\mathbf{x}_i), f(\mathbf{x}_j)) & \text{if } i \neq j \end{cases} \end{aligned}$$

- In vectorized form, $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$, with

$$\begin{aligned} \boldsymbol{\mu}_y &= \boldsymbol{\mu}_f \\ \boldsymbol{\Sigma}_y &= \boldsymbol{\Sigma}_f + \sigma^2 \mathbf{I} \end{aligned}$$

Towards Gaussian Processes

- Bayesian linear regression is just computing the conditional distribution in a multivariate Gaussian!
- Let \mathbf{y} and \mathbf{y}' denote the observables at the training and test data.
- They are jointly Gaussian:

$$\begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \boldsymbol{\mu}_{\mathbf{y}} \\ \boldsymbol{\mu}_{\mathbf{y}'} \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}} & \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}'} \\ \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}} & \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}'} \end{pmatrix} \right).$$

- The predictive distribution is a special case of the conditioning formula for a multivariate Gaussian:

$$\begin{aligned} \mathbf{y}' | \mathbf{y} &\sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{y}'|\mathbf{y}}, \boldsymbol{\Sigma}_{\mathbf{y}'|\mathbf{y}}) \\ \boldsymbol{\mu}_{\mathbf{y}'|\mathbf{y}} &= \boldsymbol{\mu}_{\mathbf{y}'} + \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}}^{-1} (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}) \\ \boldsymbol{\Sigma}_{\mathbf{y}'|\mathbf{y}} &= \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}'} - \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}}^{-1} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}'} \end{aligned}$$

- We're implicitly marginalizing out \mathbf{w} !

- The marginal likelihood is just the PDF of a multivariate Gaussian:

$$\begin{aligned} p(\mathbf{y} | \mathbf{X}) &= \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y) \\ &= \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}_y|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_y)^\top \boldsymbol{\Sigma}_y^{-1}(\mathbf{y} - \boldsymbol{\mu}_y)\right) \end{aligned}$$

- To summarize:

$$\boldsymbol{\mu}_f = \boldsymbol{\Psi} \boldsymbol{\mu}_w$$

$$\boldsymbol{\Sigma}_f = \boldsymbol{\Psi} \boldsymbol{\Sigma}_w \boldsymbol{\Psi}^\top$$

$$\boldsymbol{\mu}_y = \boldsymbol{\mu}_f$$

$$\boldsymbol{\Sigma}_y = \boldsymbol{\Sigma}_f + \sigma^2 \mathbf{I}$$

$$\boldsymbol{\mu}_{y'|y} = \boldsymbol{\mu}_{y'} + \boldsymbol{\Sigma}_{y'y} \boldsymbol{\Sigma}_{yy}^{-1} (\mathbf{y} - \boldsymbol{\mu}_y)$$

$$\boldsymbol{\Sigma}_{y'|y} = \boldsymbol{\Sigma}_{y'y'} - \boldsymbol{\Sigma}_{y'y} \boldsymbol{\Sigma}_{yy}^{-1} \boldsymbol{\Sigma}_{yy'}$$


$$p(\mathbf{y} | \mathbf{X}) = \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$$

- After defining $\boldsymbol{\mu}_f$ and $\boldsymbol{\Sigma}_f$, we can forget about \mathbf{w} !
- What if we just let $\boldsymbol{\mu}_f$ and $\boldsymbol{\Sigma}_f$ be anything?


- When I say let μ_f and Σ_f be anything, I mean let them have an arbitrary functional dependence on the inputs.
- We need to specify
 - a mean function $\mathbb{E}[f(\mathbf{x}_i)] = \mu(\mathbf{x}_i)$
 - a covariance function called a **kernel function**:
$$\text{Cov}(f(\mathbf{x}_i), f(\mathbf{x}_j)) = k(\mathbf{x}_i, \mathbf{x}_j)$$
- Let $\mathbf{K}_\mathbf{X}$ denote the kernel matrix for points \mathbf{X} . This is a matrix whose (i, j) entry is $k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$, and is called the **Gram matrix**.
- We require that $\mathbf{K}_\mathbf{X}$ be positive semidefinite for *any* \mathbf{X} . Other than that, μ and k can be arbitrary.

Gaussian Processes

- We've just defined a distribution over *function values* at an arbitrary finite set of points.
- This can be extended to a distribution over *functions* using a kind of black magic called the Kolmogorov Extension Theorem. This distribution over functions is called a **Gaussian process (GP)**.
- We only ever need to compute with distributions over function values. The formulas from a few slides ago are all you need to do regression with GPs.
- But distributions over functions are conceptually cleaner.



(a), prior



(b), posterior

- How do you think these plots were generated?

Kernel Trick

- This is an instance of a more general trick called the **Kernel Trick**.
- Many algorithms (e.g. linear regression, logistic regression, SVMs) can be written in terms of dot products between feature vectors, $\langle \mathbf{x}, \mathbf{x}' \rangle = \boldsymbol{\psi}(\mathbf{x})^\top \boldsymbol{\psi}(\mathbf{x}')$.
- A **kernel** implements an inner product between feature vectors, typically implicitly, and often much more efficiently than the explicit dot product.
- For instance, the following feature vector is quadratic in size:

$$\boldsymbol{\phi}(\mathbf{x}) = (1, \sqrt{2}x_1, \dots, \sqrt{2}x_d, \sqrt{2}x_1x_2, \sqrt{2}x_1x_3, \dots, \sqrt{2}x_{d-1}x_d, x_1^2, \dots, x_d^2)$$

- But the **quadratic kernel** can compute the inner product in linear time:

$$k(\mathbf{x}, \mathbf{x}') = \langle \boldsymbol{\phi}(\mathbf{x}), \boldsymbol{\phi}(\mathbf{x}') \rangle = 1 + \sum_{i=1}^d 2x_i x'_i + \sum_{i,j=1}^d x_i x_j x'_i x'_j = (1 + \langle \mathbf{x}, \mathbf{x}' \rangle)^2$$

- Many algorithms can be **kernelized**, i.e. written in terms of kernels, rather than explicit feature representations.
- We rarely think about the underlying feature space explicitly. Instead, we build kernels directly.
- Useful composition rules for kernels (to be proved in Homework 7):
 - A constant function $k(\mathbf{x}, \mathbf{x}') = \alpha$ is a kernel.
 - If k_1 and k_2 are kernels and $a, b \geq 0$, then $ak_1 + bk_2$ is a kernel.
 - If k_1 and k_2 are kernels, then the product $k(\mathbf{x}, \mathbf{x}') = k_1(\mathbf{x}, \mathbf{x}')k_2(\mathbf{x}, \mathbf{x}')$ is a kernel. (Interesting and surprising fact!)
- Before neural nets took over, kernel SVMs were probably the best-performing general-purpose classification algorithm.

Kernel Trick: Computational Cost

- The kernel trick lets us implicitly use very high-dimensional (even infinite-dimensional) feature spaces, but this comes at a cost.
- **Bayesian linear regression:**

$$\begin{aligned}\boldsymbol{\mu} &= \sigma^{-2} \boldsymbol{\Sigma} \boldsymbol{\Psi}^T \mathbf{t} \\ \boldsymbol{\Sigma}^{-1} &= \sigma^{-2} \boldsymbol{\Psi}^T \boldsymbol{\Psi} + \mathbf{S}^{-1}\end{aligned}$$

- Need to compute the inverse of a $D \times D$ matrix, which is an $\mathcal{O}(D^3)$ operation. (D is the number of features.)
- **GP regression:**

$$\begin{aligned}\boldsymbol{\mu}_{\mathbf{y}'|\mathbf{y}} &= \boldsymbol{\mu}_{\mathbf{y}'} + \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}}^{-1} (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}) \\ \boldsymbol{\Sigma}_{\mathbf{y}'|\mathbf{y}} &= \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}'} - \boldsymbol{\Sigma}_{\mathbf{y}'\mathbf{y}} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}}^{-1} \boldsymbol{\Sigma}_{\mathbf{y}\mathbf{y}'}\end{aligned}$$

- Need to invert an $N \times N$ matrix! (N is the number of training examples.)


Kernel Trick: Computational Cost

- This $\mathcal{O}(N^3)$ cost is typical of kernel methods. Most exact kernel methods don't scale to more than a few thousand data points.
- Kernel SVMs can be scaled further, since you can show you only need to consider the kernel over the support vectors, not the entire training set. (This is part of why they were so useful.)
- Scaling GP methods to large datasets is an active (and fascinating) research area.

- One way to define a kernel function is to give a set of basis functions and put a Gaussian prior on \mathbf{w} .
- But we have lots of other options. Here's a useful one, called the **squared-exp**, or **Gaussian**, or **radial basis function (RBF)** kernel:




$$k_{\text{SE}}(\mathbf{x}_i, \mathbf{x}_j) = \sigma^2 \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\ell^2}\right)$$

- More accurately, this is a **kernel family** with **hyperparameters** σ and ℓ .
- It gives a distribution over smooth functions:






$$k_{\text{SE}}(x_i, x_j) = \sigma^2 \exp\left(-\frac{(x_i - x_j)^2}{2\ell^2}\right)$$


- The hyperparameters determine key properties of the function.
- Varying the **output variance** σ^2 :




- Varying the **lengthscale** ℓ :




- The choice of hyperparameters heavily influences the predictions:



(b), $\ell = 0.3$



(a), $\ell = 1$




(c), $\ell = 3$

- In practice, it's very important to tune the hyperparameters (e.g. by maximizing the marginal likelihood).



$$k_{\text{SE}}(x_i, x_j) = \sigma^2 \exp\left(-\frac{(x_i - x_j)^2}{2\ell^2}\right)$$

- The squared-exp kernel is **stationary** because it only depends on $x_i - x_j$. Most kernels we use in practice are stationary.
- We can visualize the function $k(0, x)$:




GP Kernels (optional)

- The periodic kernel encodes for a probability distribution over periodic functions
- The linear kernel results in a probability distribution over linear functions




GP Kernels (optional)

- The Matern kernel is similar to the squared-exp kernel, but less smooth.
- See Chapter 4 of GPML for an explanation (advanced).
- Imagine trying to get this behavior by designing basis functions!



(a)



(b)


GP Kernels (optional)

- We get exponentially more flexibility by combining kernels.
- The sum of two kernels is a kernel.
 - This is because valid covariance matrices (i.e. PSD matrices) are closed under addition.
- The sum of two kernels corresponds to the sum of functions.



Additive kernel

Linear + Periodic

e.g. seasonal pattern w/ trend




$$k(x, y, x', y') = k_1(x, x') + k_2(y, y')$$




GP Kernels (optional)

- A kernel is like a similarity function on the input space. The sum of two kernels is like the OR of their similarity.
- Amazingly, the product of two kernels is a kernel. (Follows from the Schur Product Theorem.)
- The product of two kernels is like the AND of their similarity functions.
- Example: the product of a squared-exp kernel (spatial similarity) and a periodic kernel (similar location within cycle) gives a locally periodic function.



GP Kernels (optional)

- Modeling CO₂ concentrations:
trend + (changing) seasonal pattern + short-term variability + noise
- Encoding the structure allows sensible extrapolation.



- Bayesian linear regression lets us determine uncertainty in our predictions.
- Bayesian Occam's Razor is a sophisticated way of penalizing the complexity of a distribution over functions.
- Gaussian processes are an elegant framework for doing Bayesian inference directly over functions.
- The choice of kernels gives us much more control over what sort of functions our prior would allow or favor.