

The Kernel Trick

Consider the assumption of linear regression with an explicit feature expansion $\phi: \mathbf{x} \mapsto \phi(\mathbf{x})$:

$$y(\mathbf{x}) = \phi(\mathbf{x})^\top \mathbf{w} + \varepsilon(\mathbf{x}).$$

Given training data $\mathcal{D} = \{(\mathbf{x}, y)\} = (\mathbf{X}, \mathbf{y})$, recall we modeled the residuals $\varepsilon(\mathbf{x})$ as zero-mean independent, identically distributed Gaussians with variance σ^2 :

$$p(\varepsilon) = \mathcal{N}(\varepsilon; \mathbf{0}, \sigma^2 \mathbf{I}),$$

giving rise to the following likelihood:

$$p(\mathbf{y} | \mathbf{X}, \mathbf{w}, \sigma^2) = \mathcal{N}(\mathbf{y}; \Phi^\top \mathbf{w}, \sigma^2 \mathbf{I}),$$

where we have defined $\Phi = \phi(\mathbf{X})$.

In Bayesian linear regression, we further chose a multivariate Gaussian prior for \mathbf{w} :

$$p(\mathbf{w}) = \mathcal{N}(\mathbf{w}; \boldsymbol{\mu}, \boldsymbol{\Sigma}).$$

For simplicity, in the below we assume the prior mean for \mathbf{w} is $\boldsymbol{\mu} = \mathbf{0}$.

Given these assumptions, we were able to derive the posterior distribution of \mathbf{w} given the data \mathcal{D} :

$$p(\mathbf{w} | \mathcal{D}, \sigma^2) = \mathcal{N}(\mathbf{w}; \boldsymbol{\mu}_{\mathbf{w}|\mathcal{D}}, \boldsymbol{\Sigma}_{\mathbf{w}|\mathcal{D}}),$$

where

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{w}|\mathcal{D}} &= \boldsymbol{\Sigma} \Phi^\top (\Phi \boldsymbol{\Sigma} \Phi^\top + \sigma^2 \mathbf{I})^{-1} \mathbf{y}; \\ \boldsymbol{\Sigma}_{\mathbf{w}|\mathcal{D}} &= \boldsymbol{\Sigma} - \boldsymbol{\Sigma} \Phi^\top (\Phi \boldsymbol{\Sigma} \Phi^\top + \sigma^2 \mathbf{I})^{-1} \Phi \boldsymbol{\Sigma}. \end{aligned}$$

If we wish to use our model to predict the outputs \mathbf{y}_* associated with a set of inputs \mathbf{X}_* , we previously derived:

$$p(\mathbf{y}_* | \mathbf{X}_*, \mathcal{D}, \sigma^2) = \mathcal{N}(\mathbf{y}_*; \Phi_* \boldsymbol{\mu}_{\mathbf{w}|\mathcal{D}}, \Phi_* \boldsymbol{\Sigma}_{\mathbf{w}|\mathcal{D}} \Phi_*^\top + \sigma^2 \mathbf{I}).$$

Examining the forms of these expressions, we see that the feature expansion ϕ always appears in one of the following expressions:

$$\Phi \boldsymbol{\Sigma} \Phi^\top \quad \Phi_* \boldsymbol{\Sigma} \Phi^\top \quad \Phi \boldsymbol{\Sigma} \Phi_*^\top \quad \Phi_* \boldsymbol{\Sigma} \Phi_*^\top.$$

The entries of these matrices are always of the form $\phi(\mathbf{x})^\top \boldsymbol{\Sigma} \phi(\mathbf{x}')$, where \mathbf{x} and \mathbf{x}' are two arbitrary inputs. To simplify our expressions, we define a function

$$K(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \boldsymbol{\Sigma} \phi(\mathbf{x}').$$

Because $\boldsymbol{\Sigma}$ is positive definite, it has a “matrix square root,” $\boldsymbol{\Sigma}^{1/2}$ with the property $(\boldsymbol{\Sigma}^{1/2})^2 = \boldsymbol{\Sigma}$.¹ If we define the function $\psi(\mathbf{x}) = \boldsymbol{\Sigma}^{1/2} \phi(\mathbf{x})$, we can see that K is simply an inner product:

$$K(\mathbf{x}, \mathbf{x}') = \psi(\mathbf{x})^\top \psi(\mathbf{x}').$$

¹You can prove this via the singular value decomposition (svd): write $\boldsymbol{\Sigma} = \mathbf{U} \mathbf{D} \mathbf{U}^\top$, where \mathbf{U} is unitary and \mathbf{D} is diagonal with positive entries (because $\boldsymbol{\Sigma}$ is positive definite), then define $\boldsymbol{\Sigma}^{1/2} = \mathbf{U} \mathbf{D}^{1/2} \mathbf{U}^\top$.

Such a function K is guaranteed to always produce positive-definite Gram matrices (a *Gram matrix* is a square matrix of inner products between pairs of elements), and is called a *kernel* or *covariance function*.

Sometimes it is possible to specify a covariance function K directly without ever computing the feature map explicitly. With such a function, we could perform efficient Bayesian linear regression even with a high-dimensional (or even infinite dimensional!) feature expansion ϕ *implicitly*. This idea of computing inner products in a feature space directly is called the *kernel trick* and has been the basis of a large amount of work in the machine-learning community. Effectively, any algorithm that operates purely in terms of inner products between input vectors can be made nonlinear by replacing normal inner products with the evaluation of a kernel.

With this definition, we may rewrite the predictive distribution for \mathbf{y}_* :

$$p(\mathbf{y}_* | \mathbf{X}_*, \mathcal{D}, \sigma^2) = \mathcal{N}(\mathbf{y}_*; \mu_{\mathbf{y}_*|\mathcal{D}}, K_{\mathbf{y}_*|\mathcal{D}}),$$

where

$$\begin{aligned} \mu_{\mathbf{y}_*|\mathcal{D}} &= K(\mathbf{X}_*, \mathbf{X})(K(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I})^{-1}\mathbf{y}; \\ K_{\mathbf{y}_*|\mathcal{D}} &= K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X})(K(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I})^{-1}K(\mathbf{X}, \mathbf{X}_*). \end{aligned}$$

Examples

Perhaps the most-commonly used kernel is the *squared exponential covariance function*:

$$K(\mathbf{x}, \mathbf{x}'; \lambda, \ell) = \lambda^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right),$$

where λ and ℓ are parameters. The former is simply a multiplicative scaling constant (you can think of this as an implicit scalar multiplication in the implicit feature map ϕ). The latter takes the role of a *length scale*; vectors separated by more than a couple length scales will have a kernel value near zero.

An example of Bayesian linear regression using this kernel function is shown in Figure 1. We see that the use of this kernel function allowed us to achieve nice nonlinear regression without computing explicit basis expansions. In fact, you can show that the squared exponential kernel corresponds to an *infinite-dimensional* basis expansion, where we use a Gaussian basis function *centered on every point*. Such a feature expansion would be impossible to use if we attempted to use explicit feature computation.

Sometimes thinking in terms of the kernel can help even when you have an explicit feature expansion on hand. As an example, imagine our inputs are binary vectors of length n (so each input \mathbf{x} is a subset, a member of the power set $\mathcal{P}(n)$). One rather expensive feature expansion we could try would be to enumerate every member of $\mathcal{P}(n)$ and define $\phi(\mathbf{x})_i = \mathbf{s}_i \subset \mathbf{x}$, where \mathbf{s}_i is the i th element of the power set. So we represent our set \mathbf{x} by a feature vector of length 2^n indicating every subset of \mathbf{x} . This is a very expensive feature expansion, requiring exponential space to store for each input. However, if we take $\Sigma = \mathbf{I}$, we can compute the dot product as:

$$K(\mathbf{x}, \mathbf{x}') = 2^{|\mathbf{x} \cap \mathbf{x}'|}, \tag{1}$$

which only requires time and space linear in n !




Figure 1: Example of Bayesian linear regression using the squared exponential covariance function. The true function is $f = \sin(x)$. The kernel parameters are $\lambda = \ell = 1$, and the noise variance was set to $\sigma^2 = 0.1^2$.