

Kernel functions

Linear predictors have a large bias error. Indeed, it takes only d parameters to specify a linear predictor in \mathbb{R}^d . A popular technique to reduce this bias is feature expansion, which adds new parameters by constructing new features through nonlinear combinations of the base features. Formally, this can be viewed in terms of a function $\phi : \mathbb{R}^d \rightarrow \mathcal{V}$ mapping data points $\mathbf{x} \in \mathbb{R}^d$ to a higher-dimensional space \mathcal{V} . By training a linear predictor on a feature-expanded training set, one actually learns a more complex nonlinear predictor in the original space.

For example, consider the quadratic feature-expansion map $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^6$ defined by $\phi(x_1, x_2) = (1, x_1^2, x_2^2, x_1, x_2, x_1x_2)$. Recall that a homogeneous hyperplane in \mathbb{R}^6 with coefficients given by $\mathbf{w} = (w_1, \dots, w_6)$ is the set of points $\{\mathbf{z} \in \mathbb{R}^6 : \mathbf{w}^\top \mathbf{z} = 0\}$. Now note that $\mathbf{w}^\top \phi(\mathbf{x}) = w_1 + w_2x_1^2 + w_3x_2^2 + w_4x_1 + w_5x_2 + w_6x_1x_2$. The sets $\{\mathbf{x} \in \mathbb{R}^2 : w_1 + w_2x_1^2 + w_3x_2^2 + w_4x_1 + w_5x_2 + w_6x_1x_2 = 0\}$ describe second-degree surfaces in the Euclidean bidimensional space. These surfaces include ellipses, parabolas, and hyperbolas.

In general, we may consider polynomial feature expansion maps $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^N$ that use features of the form $\prod_{s=1}^k x_{v_s}$ for all $\mathbf{v} \in \{1, \dots, d\}^k$ and for all $k = 0, 1, \dots, n$ (the previous example is a special case for $d = 2$ and $n = 2$). Fix such a ϕ and consider the classifier $h : \mathbb{R}^d \rightarrow \{-1, 1\}$ defined by

$$h(\mathbf{x}) = \text{sgn}(\mathbf{w}^\top \phi(\mathbf{x})) \quad \text{where} \quad \mathbf{w}^\top \phi(\mathbf{x}) = \sum_{i=1}^N w_i \phi(\mathbf{x})_i$$

This nonlinear classifier in \mathbb{R}^d corresponds to a linear classifier in the space \mathbb{R}^N . Note, however, that

$$N = \sum_{k=0}^n |\{1, \dots, d\}^k| = \sum_{k=0}^n d^k = \frac{d^{n+1} - 1}{d - 1}$$

This implies that $N = \Theta(d^n)$ is exponential in the degree n , and computing ϕ becomes infeasible even for moderately large n .

This computational barrier can be fully sidestepped using the so-called **kernel trick**, which can be applied to many algorithms for learning linear predictors. For example, recall the Perceptron update rule $\mathbf{w}_{t+1} = \mathbf{w}_t + y_t \mathbf{x}_t$ where $\mathbf{w}_1 = \mathbf{0}$. Then, the linear classifiers learned through the Perceptron are of the form

$$h(\mathbf{x}) = \text{sgn} \left(\sum_{s \in S} y_s \mathbf{x}_s^\top \mathbf{x} \right)$$

where S is the set of indices s of training examples (\mathbf{x}_s, y_s) on which the Perceptron made an update. If we run the Perceptron in the space \mathbb{R}^N , the linear classifier h becomes

$$h_\phi(\mathbf{x}) = \text{sgn} \left(\sum_{s \in S} y_s \phi(\mathbf{x}_s)^\top \phi(\mathbf{x}) \right).$$

Hence, in order to compute $h_{\phi}(\mathbf{x})$ quickly we need a way of computing quickly each term $\phi(\mathbf{x}_s)^\top \phi(\mathbf{x})$. The kernel trick helps us find an efficiently computable kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$K(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \phi(\mathbf{x}') \quad \text{for all } \mathbf{x}, \mathbf{x}' \in \mathbb{R}^d. \quad (1)$$

For example, the quadratic kernel, corresponding to the quadratic feature-expansion map

$$\phi(x_1, x_2) = (1, x_1^2, x_2^2, \sqrt{2}x_1, \sqrt{2}x_2, \sqrt{2}x_1x_2)$$

is $K(\mathbf{x}, \mathbf{x}') = (1 + \mathbf{x}^\top \mathbf{x}')^2$, as one can easily verify (the presence of the $\sqrt{2}$ coefficients, which is needed for the math, does not change the class of linear predictors that are representable using the mapping).

Given a kernel K , we can then write the linear classifier generated by the Perceptron as

$$h_K(\mathbf{x}) = \text{sgn} \left(\sum_{s \in S} y_s K(\mathbf{x}_s, \mathbf{x}) \right). \quad (2)$$

Below here, we give the pseudo-code of the Kernel Perceptron algorithm.

Algorithm: Kernel Perceptron

Let S be the empty set.

For all $t = 1, 2, \dots$

1. Get next example (\mathbf{x}_t, y_t)
2. Compute $\hat{y}_t = \text{sgn} \left(\sum_{s \in S} y_s K(\mathbf{x}_s, \mathbf{x}_t) \right)$
3. If $\hat{y}_t \neq y_t$ add t to S

The polynomial kernel $K_n(\mathbf{x}, \mathbf{x}') = (1 + \mathbf{x}^\top \mathbf{x}')^n$ for all $n \in \mathbb{N}$ generalizes the quadratic kernel defined earlier. Using Newton's Binomial Theorem, we can explicitly compute the map ϕ_n such that $K_n(\mathbf{x}, \mathbf{x}') = \phi_n(\mathbf{x})^\top \phi_n(\mathbf{x}')$,

$$(1 + \mathbf{x}^\top \mathbf{x}')^n = \sum_{k=0}^n \binom{n}{k} (\mathbf{x}^\top \mathbf{x}')^k. \quad (3)$$

Now observe that

$$(\mathbf{x}^\top \mathbf{x}')^k = \left(\sum_{i=1}^d x_i x'_i \right)^k = \sum_{\mathbf{v} \in \{1, \dots, d\}^k} \left(\prod_{s=1}^k x_{v_s} x'_{v_s} \right).$$

Therefore,

$$\phi_n(\mathbf{x}) = \left(\sqrt{\binom{n}{k}} \prod_{s=1}^k x_{v_s} \right)_{k=0, \dots, n, \mathbf{v} \in \{1, \dots, d\}^k}. \quad (4)$$

In other words, the feature map $\phi_n : \mathbb{R}^d \rightarrow \mathbb{R}^N$ associated with the polynomial kernel K_n sends each $\mathbf{x} \in \mathbb{R}^d$ to a vector whose components are indexed by all monomials $\prod_{s=1}^k x_{v_s}$ of degree at most n and weighted by the roots of the binomial coefficients.

Another type of kernel is the Gaussian kernel,

$$K_\gamma(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{1}{2\gamma} \|\mathbf{x} - \mathbf{x}'\|^2\right) \quad \gamma > 0.$$

In order to derive the map ϕ_γ associated with K_γ we proceed as follows,

$$\begin{aligned} \exp\left(-\frac{1}{2\gamma} \|\mathbf{x} - \mathbf{x}'\|^2\right) &= \exp\left(-\frac{1}{2\gamma} (\|\mathbf{x}\|^2 + \|\mathbf{x}'\|^2)\right) \exp\left(\frac{1}{\gamma} (\mathbf{x}^\top \mathbf{x}')\right) \\ &= \exp\left(-\frac{\|\mathbf{x}\|^2}{2\gamma}\right) \exp\left(-\frac{\|\mathbf{x}'\|^2}{2\gamma}\right) \sum_{n=0}^{\infty} \frac{1}{n!} \frac{(\mathbf{x}^\top \mathbf{x}')^n}{\gamma^n} \end{aligned} \quad (5)$$

where we used the Taylor series expansion $e^x = 1 + x + \frac{x^2}{2!} + \dots$. A closer look at (5) reveals that the Gaussian kernel is a linear combination of infinitely many polynomial kernels (3) of increasing degree, each weighted by the reciprocal of the factorial of its degree. The parameter γ is a scaling factor for the products $\mathbf{x}^\top \mathbf{x}'$. Finally, the factors $e^{-\|\mathbf{x}\|^2/(2\gamma)} e^{-\|\mathbf{x}'\|^2/(2\gamma)}$ normalize with respect to \mathbf{x} and \mathbf{x}' giving $K_\gamma(\mathbf{x}, \mathbf{x}) = 1$ for each $\mathbf{x} \in \mathbb{R}^d$.

Note that predictors of the form (2) that utilize Gaussian kernels predict any point \mathbf{x} using a linear combination (with coefficients y_t) of Gaussians $e^{-z^2/(2\gamma)}$ centered on \mathbf{x}_s for $s \in S$ and evaluated at \mathbf{x} .

The Gaussian kernel is *universal* in the following sense: for each $\gamma > 0$, for each continuous function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, and for each $\varepsilon > 0$, there exists $g \in \mathcal{H}_\gamma$, with

$$\mathcal{H}_\gamma = \left\{ \sum_{i=1}^N \alpha_i K_\gamma(\mathbf{x}_i, \cdot) : \mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^d, \alpha_1, \dots, \alpha_N \in \mathbb{R}, N \in \mathbb{N} \right\} \quad (6)$$

such that h approximates f with error bounded by ε . An important consequence of this fact is that learning algorithms that use Gaussian kernels can be consistent (that is, the expected risk of their predictors converges to the Bayes risk as the sample size grows to infinity).

Given a data space \mathcal{X} (not necessarily \mathbb{R}^d) and a symmetric function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ how can we check whether K is a kernel? In other words, we want to know whether there exist a feature map $\phi_K : \mathcal{X} \rightarrow \mathcal{H}_K$ and an inner product $\langle \cdot, \cdot \rangle_K$ in \mathcal{H}_K such that $\langle \phi_K(\mathbf{x}), \phi_K(\mathbf{x}') \rangle_K = K(\mathbf{x}, \mathbf{x}')$. Luckily, there is a very simple way of checking that: K is a kernel if and only if for all $m \in \mathbb{N}$ and for all $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$, the $m \times m$ matrix \mathbf{K} such that $\mathbf{K}_{i,j} = K(\mathbf{x}_i, \mathbf{x}_j)$ is positive semidefinite. That is, $\mathbf{z}^\top \mathbf{K} \mathbf{z} \geq 0$ for all $\mathbf{z} \in \mathbb{R}^m$.

The above result tells us the conditions under which ϕ_K exists for a given K . On the other hand, it does not tell us how ϕ_K looks like. In fact, there is no unique representation of the pair ϕ_K, \mathcal{H}_K for a given kernel K . However, it can be shown that if K is a kernel, then we can always represent ϕ_K as $\phi_K(\mathbf{x}) = K(\mathbf{x}, \cdot)$ and \mathcal{H}_K as

$$\mathcal{H}_K \equiv \left\{ \sum_{i=1}^N \alpha_i K(\mathbf{x}_i, \cdot) : \mathbf{x}_1, \dots, \mathbf{x}_N \in \mathcal{X}, \alpha_1, \dots, \alpha_N \in \mathbb{R}, N \in \mathbb{N} \right\}$$

which is a generalization of the space (6) for Gaussian kernels.

A learning algorithm producing linear predictors $g \in \mathcal{H}_K$ becomes nonparametric when K is the Gaussian kernel and N (the number of terms in the sum defining any $g \in \mathcal{H}_K$) is free to grow unbounded as the sample size increases. If, instead, K is a kernel such that ϕ_K maps \mathcal{X} to a finite dimensional space, then any $g \in \mathcal{H}_K$ can be always represented with a fixed number of parameters, and therefore any algorithm choosing predictors from \mathcal{H}_K is parametric.

Since ϕ_K is the feature map for K , we must have $\langle \phi_K(\mathbf{x}), \phi_K(\mathbf{x}') \rangle_K = K(\mathbf{x}, \mathbf{x}')$. Now, recalling that the inner product is a bilinear operator, we can compute the inner product between $\phi_K(\mathbf{x}) = K(\mathbf{x}, \cdot)$ and any $f \in \mathcal{H}_K$,

$$\langle f, K(\mathbf{x}, \cdot) \rangle_K = \sum_{i=1}^N \alpha_i \langle K(\mathbf{x}_i, \cdot), K(\mathbf{x}, \cdot) \rangle_K = \sum_{i=1}^N \alpha_i K(\mathbf{x}_i, \mathbf{x}) = f(\mathbf{x})$$

The equality $\langle f, K(\mathbf{x}, \cdot) \rangle_K = f(\mathbf{x})$ is important and is known as *reproducing property*. For this reason, \mathcal{H}_K is also known as *reproducing kernel Hilbert space* (RKHS). Note that for linear predictors $\mathbf{w} \in \mathbb{R}^d$, we compute $\mathbf{w}^\top \mathbf{x}$ using the standard notion of Euclidean inner product. For linear predictors $f \in \mathcal{H}_K$, instead, we compute $\langle f, \phi_K(\mathbf{x}) \rangle_K = f(\mathbf{x})$ using the inner product for \mathcal{H}_K .

Next, we see how the inner product $\langle f, g \rangle_K$ between two arbitrary $f, g \in \mathcal{H}_K$ is computed, where

$$f = \sum_{i=1}^N \alpha_i K(\mathbf{x}_i, \cdot) \quad \text{and} \quad g = \sum_{j=1}^M \beta_j K(\mathbf{x}'_j, \cdot)$$

Using once more the bilinearity of inner products,

$$\begin{aligned} \langle f, g \rangle_K &= \left\langle \sum_{i=1}^N \alpha_i K(\mathbf{x}_i, \cdot), \sum_{j=1}^M \beta_j K(\mathbf{x}'_j, \cdot) \right\rangle_K \\ &= \sum_{i=1}^N \alpha_i \left\langle K(\mathbf{x}_i, \cdot), \sum_{j=1}^M \beta_j K(\mathbf{x}'_j, \cdot) \right\rangle_K \\ &= \sum_{i=1}^N \sum_{j=1}^M \alpha_i \beta_j \langle K(\mathbf{x}_i, \cdot), K(\mathbf{x}'_j, \cdot) \rangle_K \\ &= \sum_{i=1}^N \sum_{j=1}^M \alpha_i \beta_j K(\mathbf{x}_i, \mathbf{x}'_j) \end{aligned}$$

We can lift to any RKHS the bounds we derived for learning algorithms that can be run with kernels. For instance recall the bound on the number of mistakes provided by the Perceptron convergence theorem,

$$\|\mathbf{u}\|^2 \left(\max_t \|\mathbf{x}_t\|^2 \right)$$

which holds for any $\mathbf{u} \in \mathbb{R}^d$ such that $y_t \mathbf{u}^\top \mathbf{x}_t \geq 1$ for $t = 1, \dots, m$.

In a generic RKHS \mathcal{H}_K , the linear separator \mathbf{u} is some $g \in \mathcal{H}_K$ such that $y_t g(\mathbf{x}_t) \geq 1$ for $t = 1, \dots, m$. The squared norm $\|\mathbf{x}_t\|^2 = \mathbf{x}_t^\top \mathbf{x}_t$ becomes $\|\phi_K(\mathbf{x})\|_K^2 = \langle K(\mathbf{x}, \cdot), K(\mathbf{x}, \cdot) \rangle_K = K(\mathbf{x}, \mathbf{x})$.

Finally $\|\mathbf{u}\|^2$ is replaced by

$$\|f\|_K^2 = \left\| \sum_{i=1}^N \alpha_i K(\mathbf{x}_i, \cdot) \right\|_K^2 = \left\langle \sum_{i=1}^N \alpha_i K(\mathbf{x}_i, \cdot), \sum_{j=1}^N \alpha_j K(\mathbf{x}_j, \cdot) \right\rangle_K = \sum_{i,j=1}^N \alpha_i \alpha_j K(\mathbf{x}_i, \mathbf{x}_j).$$

A more complex linear predictor that can be kernelized is ridge regression, $\mathbf{w} = (\alpha I + S^\top S)^{-1} S^\top \mathbf{y}$, where S is the $m \times d$ matrix whose rows are the training points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathbb{R}^d$ and $\mathbf{y} = (y_1, \dots, y_m)$ is the vector of training labels $y_t \in \mathbb{R}$. Using the identity

$$(\alpha I + S^\top S)^{-1} S^\top = S^\top (\alpha I + S S^\top)^{-1}$$

we can represent the ridge regression predictor in a generic RKHS by $\mathbf{k}(\cdot)^\top (\alpha I + \mathbf{K})^{-1} \mathbf{y}$, where \mathbf{K} is the $m \times m$ matrix with entries $\mathbf{K}_{i,j} = K(\mathbf{x}_i, \mathbf{x}_j)$ and $\mathbf{k}(\cdot)$ is the vector $(K(\mathbf{x}_1, \cdot), \dots, K(\mathbf{x}_m, \cdot))$ of functions $K(\mathbf{x}_t, \cdot)$. Similarly to the non-kernel case, where the prediction on \mathbf{x} is $\mathbf{w}^\top \mathbf{x}$, the prediction on $\phi_K(\mathbf{x})$ is $\langle g, \phi_K(\mathbf{x}) \rangle_K$ which evaluates to $g(\mathbf{x}) = \mathbf{y}^\top (\alpha I + \mathbf{K})^{-1} \mathbf{k}(\mathbf{x})$ where $\mathbf{k}(\mathbf{x}) = (K(\mathbf{x}_1, \mathbf{x}), \dots, K(\mathbf{x}_m, \mathbf{x}))$.

Linear predictors in RKHS can incur overfitting. For example, by increasing the degree n of a polynomial kernel K_n we reduce the training error because higher-degree curves can be used to separate the training points. If the degree is too high, the predictor will overfit. A similar reasoning applies to Gaussian kernels K_γ . The γ parameter corresponds to the width of the Gaussians centered on training points \mathbf{x}_s . If γ is small relative to the typical squared distances $\|\mathbf{x}_s - \mathbf{x}\|^2$ between training and test points, then the classification of a test point \mathbf{x} is essentially determined by the training point closest to it. This implies a training error equal or close to zero, because—similarly to 1-NN classifiers—the training points are never misclassified. Once again, the resulting predictor is likely to overfit. On the other hand, for values of γ that are large with respect to the squared distances $\|\mathbf{x}_s - \mathbf{x}\|^2$, the Gaussians centered on training points are very wide, and the resulting predictors are similar to k -NN classifiers when k is chosen close to the training set size, which is likely to cause underfitting.

We established that any symmetric function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel if and only if the kernel matrix \mathbf{K} is positive semidefinite. This is an important result, because it holds irrespective to the choice of the data space \mathcal{X} . We can therefore define kernels on any set \mathcal{X} : be it a set of matrices, sequences, trees, graphs, and so on. Kernels can be viewed as a way to encapsulate the data space, offering a uniform interface to a learning algorithm that can be efficiently run in the corresponding RKHS. In order to guide the intuition when designing a kernel function on a given data space \mathcal{X} , one should recall that $K(x, x')$ implements an inner product between $K(x, \cdot)$ and $K(x', \cdot)$ in some RKHS. We can then start by defining a notion of similarity on \mathcal{X} , and then adjust it so that arbitrary kernel matrices are always positive semidefinite.