

CS-GY 6923: Lecture 8

k-Nearest Neighbors, Kernel Methods

NYU Tandon School of Engineering, Prof. Christopher Musco

- Previous methods studied (regression, logistic regression) are considered linear methods. They make predictions based on $\langle \mathbf{x}, \underline{\beta} \rangle$ – i.e. based on weighted sums of features.
- In the next part of the course we move on to non-linear methods. Specifically, **kernel methods** and **neural networks**.
- Both are very closely related to feature transformations!

WARM UP: k -NEAREST NEIGHBOR METHOD

k -NN algorithm: a simple but powerful baseline for classification.

Training data: $(\underline{x_1}, y_1), \dots, (\underline{x_n}, y_n)$ where $y_1, \dots, y_n \in \{1, \dots, q\}$.

Classification algorithm:

Given new input $\underline{x_{new}}$,

- Compute $\underline{sim}(\underline{x_{new}}, \underline{x_1}), \dots, \underline{sim}(\underline{x_{new}}, \underline{x_n})$.¹
- Let $\underline{x_{j_1}} \dots \underline{x_{j_k}}$ be the training data vectors with highest similarity to $\underline{x_{new}}$.
- Predict $\underline{y_{new}}$ as $\underline{majority}(\underline{y_{j_1}}, \dots, \underline{y_{j_k}})$.

¹ $\underline{sim}(\underline{x_{new}}, \underline{x_i})$ is any chosen similarity function, like $\underline{1 - \|\underline{x_{new}} - \underline{x_i}\|_2}$.

k -NEAREST NEIGHBOR METHOD




Fig. 1. The dataset.

$k=1$




Fig. 2. The 1NN classification map.

$n = \text{size of training set}$





Fig. 3. The 5NN classification map.


- Smaller k , more complex classification function.
- Larger k , more robust to noisy labels.

Works remarkably well for many datasets.



MNIST IMAGE DATA

Especially good for large datasets with lots of repetition. Works well on MNIST for example:




$\approx 95\%$ Accuracy out-of-the-box.²

Let's look into this example a bit more...

²Can be improved to $\approx 99.5\%$ with a fancy similarity function!

MNIST IMAGE DATA

Each pixel is number from [0, 1]. 0 is black, 1 is white.
Represent 28×28 matrix of pixel values as a flattened vector.



```
xmat = np.array([[1,2,3],[4,5,6],[7,8,9]])
```

```
array([[1, 2, 3],  
       [4, 5, 6],  
       [7, 8, 9]])
```


```
xvec = xmat.ravel()
```

```
array([1, 2, 3, 4, 5, 6, 7, 8, 9])
```

INNER PRODUCT SIMILARITY

Given data vectors $\mathbf{x}, \mathbf{w} \in \mathbb{R}^d$, the inner product $\langle \mathbf{x}, \mathbf{w} \rangle$ is a natural similarity measure.

$$\langle \mathbf{x}, \mathbf{w} \rangle = \sum_{i=1}^d x_i w_i = \cos(\theta) \|\mathbf{x}\|_2 \|\mathbf{w}\|_2.$$



Also called "cosine similarity".

INNER PRODUCT SIMILARITY

$$\|a\|_r^2 = a^T a \quad (x - w)^T (x - w) = \underbrace{x^T x}_{\|x\|_2^2} - \underbrace{2x^T w} + \underbrace{w^T w}_{\|w\|_2^2}$$

Connection to Euclidean (ℓ_2) Distance:

$$\|x - w\|_2^2 = \|x\|_2^2 + \|w\|_2^2 - 2\langle x, w \rangle$$

If all data vectors has the same norm, the pair of vectors with largest inner product is the pair with smallest Euclidean distance.

INNER PRODUCT FOR MNIST

Inner product between MNIST digits:

$$\langle \mathbf{x}, \mathbf{w} \rangle$$




$$\langle \mathbf{x}, \mathbf{w} \rangle = \sum_{i=1}^{28} \sum_{j=1}^{28} \underline{\text{matx}[i,j]} \cdot \underline{\text{matw}[i,j]}.$$

Inner product similarity is higher when the images have large pixel values (close to 1) in the same locations. I.e. when they have a lot of overlapping white/light gray pixels.

INNER PRODUCT FOR MNIST

Visualizing the inner product between two images:



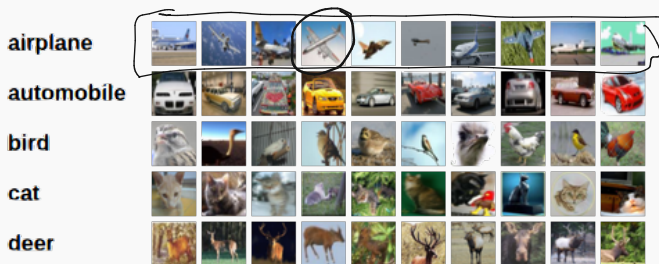
Images with high inner product have a lot of overlap.

Most similar images during k -nn search, $k = 9$:



K-NN FOR OTHER IMAGES

Does not work as well for less standardized classes of images:




CIFAR 10 Images

Even after scaling to have same size, converting to separate RGB channels, etc. something as simple as k -nn won't work.

One-vs.-all or Multiclass Cross-entropy Classification with Logistic Regression:

- Learn q classifiers with parameters $\underline{\beta}^{(1)}, \underline{\beta}^{(2)}, \dots, \underline{\beta}^{(q)}$.
- Given \underline{x}_{new} compute $\langle \underline{x}_{new}, \underline{\beta}^{(1)} \rangle, \dots, \langle \underline{x}_{new}, \underline{\beta}^{(q)} \rangle$
- Predict class $\underline{y}_{new} = \arg \max_i \langle \underline{x}_{new}, \underline{\beta}^{(i)} \rangle$.


If each \mathbf{x} is a vector with $28 \times 28 = \underline{784}$ entries than each $\underline{\beta}^{(i)}$ also has $\underline{784}$ entries. Each parameter vector can be viewed as a 28×28 image.



MATCHED FILTER

Visualizing $\beta_m^{(r)}, \dots, \beta_q^{(s)}$

log loss



For an input image X_{new} , compute inner product similarity with all weight matrices and choose most similar one.

In contrast to k -NN, only need to compute similarity with q items instead of n .

Logistic Regression Model:

Given data matrix $X \in \mathbb{R}^{n \times d}$ (here $d = 784$) and binary label vector $\mathbf{y} \in \{0, 1\}^n$ for class i (1 if in class i , 0 if not), find $\underline{\beta} \in \mathbb{R}^d$ to minimize the log loss between:

\mathbf{y}

and

$h(\underline{X}\underline{\beta})$

where $\underline{h}(z) = \frac{1}{1+e^{-z}}$ applies the logistic function entrywise to $\underline{X}\underline{\beta}$.

$$\text{Loss} = - \sum_{j=1}^n y_j \log(h(\underline{X}\underline{\beta})_j) + (1 - y_j) \log(1 - h(\underline{X}\underline{\beta})_j)$$

$h(x_i^T \beta)$


ALTERNATIVE VIEW

Reminder from linear algebra: Without loss of generality, can assume that β lies in the row span of X . $X\beta$

So for any $\beta \in \mathbb{R}^d$, there exists a vector $\alpha \in \mathbb{R}^n$ such that:

$$X\beta = \underline{X^T\alpha}$$

$X^T\alpha \rightarrow$ linear combination of rows of X



$$b = b_r + b^+ \quad Xb^+$$

Logistic Regression Equivalent Formulation:

Given data matrix $X \in \mathbb{R}^{n \times d}$ (here $d = 784$) and binary label vector $\mathbf{y} \in \{0, 1\}^n$ for class i (1 if in class i , 0 if not), find $\alpha \in \mathbb{R}^n$ to minimize the log loss between:

$$\underline{\mathbf{y}} \quad \text{and} \quad \underline{h(\mathbf{X}\mathbf{X}^T\boldsymbol{\alpha})}.$$

$$L(\underline{\boldsymbol{\alpha}}) = \sum_{j=1}^n y_j \log(h(\underline{\mathbf{X}\mathbf{X}^T\boldsymbol{\alpha}})_j) + (1-y_j) \log(1-h(\underline{\mathbf{X}\mathbf{X}^T\boldsymbol{\alpha}})_j)$$

Can still be minimized via gradient descent:

$$\nabla L(\boldsymbol{\alpha}) = \mathbf{X}\mathbf{X}^T(h(\mathbf{X}\mathbf{X}^T\boldsymbol{\alpha}) - \mathbf{y}).$$

What does classification for a new point \mathbf{x}_{new} look like? Recall that for a given one-vs-all classification for class i , the original parameter vector $\underline{\beta}^{(i)} = \underline{\mathbf{X}}^T \underline{\alpha}^{(i)}$.

- Learn q classifiers with parameters $\underline{\alpha}^{(1)}, \underline{\alpha}^{(2)}, \dots, \underline{\alpha}^{(q)}$.
- Given \mathbf{x}_{new} compute $\langle \mathbf{x}_{new}, \mathbf{X}^T \alpha^{(1)} \rangle, \dots, \langle \mathbf{x}_{new}, \mathbf{X}^T \alpha^{(q)} \rangle$
- Predict class $y_{new} = \arg \max_i \langle \mathbf{x}_{new}, \mathbf{X}^T \alpha^{(i)} \rangle$.


REFORMULATED VIEW

$$\langle x_{\text{new}}, \beta^{(i)} \rangle$$


$$\langle a, b \rangle = a^T b$$

Score for class i :

$$\begin{aligned} \langle x_{\text{new}}, X^T \alpha_j \rangle &= x_{\text{new}}^T X^T \alpha_j^{(i)} = \langle a, b \rangle \\ &= \langle X x_{\text{new}}, \alpha_j^{(i)} \rangle = \sum_{j=1}^n \alpha_j^{(i)} \cdot [X x_{\text{new}}]_j \\ &= \sum_{j=1}^n \alpha_j^{(i)} \langle x_{\text{new}}, x_j \rangle. \end{aligned}$$



ORIGINAL VIEW OF LOGISTIC REGRESSION



$$\langle \vec{\beta}^{(0)}, \vec{x}_{new} \rangle = 45$$

The equation shows the inner product of the initial weight vector $\vec{\beta}^{(0)}$ and a new input vector \vec{x}_{new} . The weight vector is represented by a blurry image of a '5', and the input vector is a clear image of a '5'. The result is 45.

$$\langle \vec{\beta}^{(5)}, \vec{x}_{new} \rangle = 212$$

The equation shows the inner product of the weight vector $\vec{\beta}^{(5)}$ and the same input vector \vec{x}_{new} . The weight vector is a more distinct image of a '5'. The result is 212, which is highlighted in green.

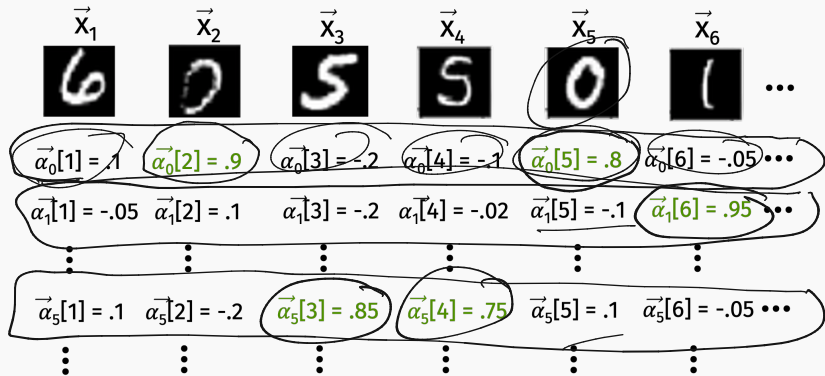
$$\langle \vec{\beta}^{(6)}, \vec{x}_{new} \rangle = 84$$

The equation shows the inner product of the weight vector $\vec{\beta}^{(6)}$ and the same input vector \vec{x}_{new} . The weight vector is a blurry image of a '6'. The result is 84.

$X^T \varphi$



NEW VIEW OF LOGISTIC REGRESSION



Learn n length parameter vectors $\alpha^{(0)}, \dots, \alpha^{(9)}$, one for each class.

$$\alpha_j^{(b)} \rightarrow \alpha_b[j]$$

NEW VIEW OF LOGISTIC REGRESSION

$$(1 \times d)(d \times 1)(n \times 1) = 1$$

$$\alpha^{(0)_1} \times \langle \vec{x}_1, \vec{x}_{new} \rangle + \alpha^{(0)_2} \times \langle \vec{x}_2, \vec{x}_{new} \rangle + \alpha^{(0)_3} \times \langle \vec{x}_3, \vec{x}_{new} \rangle + \dots = 45$$

$\chi_{new}^T \chi^T q$

$$\alpha^{(5)_1} \times \langle \vec{x}_1, \vec{x}_{new} \rangle + \alpha^{(5)_2} \times \langle \vec{x}_2, \vec{x}_{new} \rangle + \alpha^{(5)_3} \times \langle \vec{x}_3, \vec{x}_{new} \rangle + \dots = 212$$

$$\alpha^{(6)_1} \times \langle \vec{x}_1, \vec{x}_{new} \rangle + \alpha^{(6)_2} \times \langle \vec{x}_2, \vec{x}_{new} \rangle + \alpha^{(6)_3} \times \langle \vec{x}_3, \vec{x}_{new} \rangle + \dots = 84$$


Classification looks similar to k -NN: we compute the similarity between x_{new} and every other vector in our training data set. A weighted sum of the similarities leads to scores for each class.

Assign x_{new} to the class with highest score.


q } n → size of training set.

DIVING INTO SIMILARITY

Often the inner product **does not make sense** as a similarity measure between data vectors. Here's an example (recall that smaller inner product means less similar):



But clearly the first image is more similar.



Here's a more realistic scenario.

KERNEL FUNCTIONS: PERSPECTIVE ONE

A kernel function $k(\mathbf{x}, \mathbf{y})$ is simply a similarity measure between data points.

$$k(\mathbf{x}, \mathbf{y}) = \begin{cases} \text{large if } \mathbf{x} \text{ and } \mathbf{y} \text{ are similar.} \\ \text{close to 0 if } \mathbf{x} \text{ and } \mathbf{y} \text{ are different.} \end{cases}$$

Example: The Radial Basis Function (RBF) kernel, aka the Gaussian kernel:

$$\|x\|_2^2 + \|y\|_2^2 - 2x \cdot y = \|x - y\|_2^2$$

$$k(\mathbf{x}, \mathbf{y}) = e^{-\|\mathbf{x} - \mathbf{y}\|_2^2 / \sigma^2}$$



for some scaling factor σ .

$$k(\overset{\vec{z}}{\text{5}}, \overset{\vec{x}}{\text{5}}) > k(\overset{\vec{y}}{\square}, \overset{\vec{x}}{\text{5}})$$

KERNEL FUNCTIONS: PERSPECTIVE ONE



Lots of kernel functions involve transformations of $\langle \mathbf{x}, \mathbf{y} \rangle$ or $\|\mathbf{x} - \mathbf{y}\|_2$:

- Gaussian RBF Kernel: $k(\mathbf{x}, \mathbf{y}) = \underline{e^{-\|\mathbf{x}-\mathbf{y}\|_2^2/\sigma^2}}$
- Laplace Kernel: $k(\mathbf{x}, \mathbf{y}) = \underline{e^{-\|\mathbf{x}-\mathbf{y}\|_2/\sigma}}$
- Polynomial Kernel: $k(\mathbf{x}, \mathbf{y}) = (\underline{\langle \mathbf{x}, \mathbf{y} \rangle} + \underline{1})^q$

But you can imagine much more complex similarity metrics.

KERNEL FUNCTIONS: PERSPECTIVE TWO

For a simple algorithm like k -NN you can swap out the inner product similarity with any similarity function you could possibly imagine.

For a methods like logistic regression, this is not the case...

Recall: We learned a parameter vector α to minimize $LL(\mathbf{y}, \mathbf{X}^T \alpha)$ where $LL()$ denotes the logistic loss. Then we classified via:

$$\langle \mathbf{x}_{new}, \mathbf{X}^T \alpha \rangle = \mathbf{x}_{new}^T \mathbf{X}^T \alpha = \sum_{j=1}^n \alpha_j \langle \mathbf{x}_{new}, \mathbf{x}_j \rangle.$$

The inner product similarity came from the fact that our predictions were based on the linear function $\langle \mathbf{x}_{new}, \mathbf{X}^T \alpha \rangle$.

KERNEL FUNCTIONS AS FEATURE TRANSFORMATION


A positive semidefinite (PSD) kernel is any similarity function with the following form:

$$k(\mathbf{x}, \mathbf{w}) = \phi(\mathbf{x})^T \phi(\mathbf{w})$$

where ϕ : $\mathbb{R}^d \rightarrow \mathbb{R}^m$ is a some feature transformation function.

KERNEL FUNCTIONS AND FEATURE TRANSFORMATION

Example: Degree 2 polynomial kernel, $k(\mathbf{x}, \mathbf{w}) = (\mathbf{x}^T \mathbf{w} + 1)^2$



$$\begin{aligned}
 \underline{(\mathbf{x}^T \mathbf{w} + 1)^2} &= (x_1 w_1 + x_2 w_2 + x_3 w_3 + 1)^2 \\
 &= \underline{1 + 2x_1 w_1 + 2x_2 w_2 + 2x_3 w_3 + x_1^2 w_1^2 + x_2^2 w_2^2 + x_3^2 w_3^2} \\
 &\quad + \underline{2x_1 w_1 x_2 w_2 + 2x_1 w_1 x_3 w_3 + 2x_2 w_2 x_3 w_3} \\
 &= \underline{\phi(\mathbf{x})^T \phi(\mathbf{w})}
 \end{aligned}$$

KERNEL FUNCTIONS AND FEATURE TRANSFORMATION

$$\langle \underline{\phi(x)}, \phi(y) \rangle = e^{-\|x-y\|_2^2 / \sigma^2}$$

Not all similarity metrics are positive semidefinite (PSD), but all of the ones we saw earlier are:

• Gaussian RBF Kernel: $k(\mathbf{x}, \mathbf{y}) = \underline{e^{-\|\mathbf{x}-\mathbf{y}\|_2^2 / \sigma^2}}$

• Laplace Kernel: $k(\mathbf{x}, \mathbf{y}) = \underline{e^{-\|\mathbf{x}-\mathbf{y}\|_2 / \sigma}}$

• Polynomial Kernel: $k(\mathbf{x}, \mathbf{y}) = \underline{(\langle \mathbf{x}, \mathbf{y} \rangle + 1)^q}$.

And there are many more...


Sometimes $\phi(\vec{x})$ is simple and explicit. **More often, it is not.**

As we will discuss shortly, it doesn't necessarily matter – we often don't even need to know ϕ .

KERNEL FUNCTIONS AND FEATURE TRANSFORMATION

Feature transformations \iff new similarity metrics.


Using $k(\cdot, \cdot)$ in place of the inner product $\langle \cdot, \cdot \rangle$ is **equivalent** to replacing every data point $\underline{x}_1, \dots, \underline{x}_n$ by $\phi(\underline{x}_1), \dots, \phi(\underline{x}_n)$.³



³Transform dimension m is often very large: e.g. $m = O(d^q)$ for a degree q polynomial kernel. For many kernels (e.g. the Gaussian kernel) m is actually infinite. Typically you need to use regularization.

TAKEAWAY ONE

We can improve performance by replacing the inner product with another kernel $k(\cdot, \cdot)$ for the same reason that feature transformations improved performance.



When you add features, it becomes possible to learn more complex decision boundaries (in this case a circle) with a linear classifier.

PSD kernel functions give a principled way of “swapping out” the inner product with a new similarity metric for linear algorithms like multiple linear regression or logistic regression. For non-PSD kernels it is not clear how to do this.

KERNEL LOGISTIC REGRESSION

Standard logistic regression

Loss function: \rightarrow log loss

$$L(\alpha) = \sum_{j=1}^n LL(y_j, X_j^T \alpha)$$

Gradient:

$$\nabla L(\alpha) = \underbrace{XX^T}_{n \times n} (h(\underbrace{XX^T \alpha}_{\rightarrow y}) - \underline{y})$$

Prediction:

$$z = \sum_{j=1}^n \alpha_j \langle \underline{x_{new}}, \underline{x_j} \rangle$$

$$y_{new} = \mathbb{1}[z > 0]$$

Kernel logistic regression

Loss function:

$$L(\alpha) = LL(y, \phi(X)^T \alpha)$$

Gradient:

$$\nabla L(\alpha) = \phi(X) \phi(X)^T (h(\phi(X) \phi(X)^T \alpha) - y)$$

Prediction:

$$z = \sum_{j=1}^n \alpha_j \langle \underbrace{\phi(x_{new}), \phi(x_j)}_{\leftarrow (x_{new}, x_j)} \rangle$$

$$y_{new} = \mathbb{1}[z > 0]$$

Standard linear regression

Loss function:

$$L(\boldsymbol{\alpha}) = \|\mathbf{y} - \mathbf{X}\mathbf{X}^T\boldsymbol{\alpha}\|_2$$

Gradient:

$$\nabla L(\boldsymbol{\alpha}) = 2\mathbf{X}\mathbf{X}^T(\mathbf{X}\mathbf{X}^T\boldsymbol{\alpha} - \mathbf{y}).$$

Prediction:

$$y_{new} = \sum_{j=1}^n \alpha_j \cdot \langle \mathbf{x}_{new}, \mathbf{x}_j \rangle.$$

Kernel linear regression

Loss function:

$$L(\boldsymbol{\alpha}) = \|\mathbf{y} - \phi(\mathbf{X})\phi(\mathbf{X})^T\boldsymbol{\alpha}\|_2$$

Gradient:

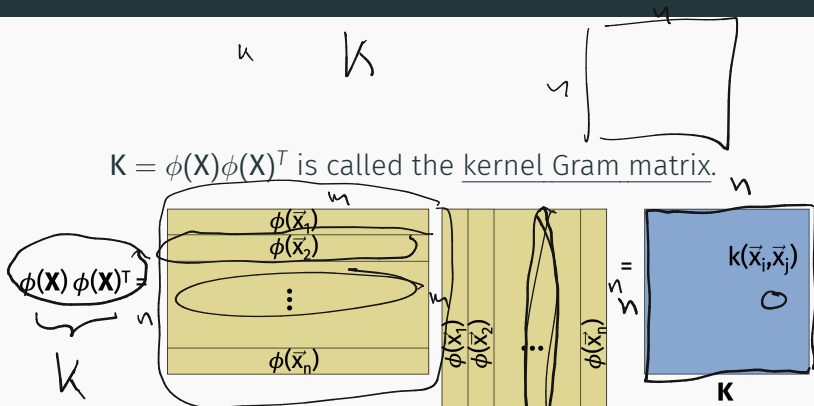
$$\nabla L(\boldsymbol{\alpha}) = 2\phi(\mathbf{X})\phi(\mathbf{X})^T(\phi(\mathbf{X})\phi(\mathbf{X})^T\boldsymbol{\alpha} - \mathbf{y}).$$

Prediction:

$$y_{new} = \sum_{j=1}^n \alpha_j \cdot \langle \phi(\mathbf{x}_{new}), \phi(\mathbf{x}_j) \rangle.$$

$\frac{2k(\mathbf{x}_{new}, \mathbf{x}_j)}{2} = k(\mathbf{x}_{new}, \mathbf{x}_j)$

KERNEL MATRIX



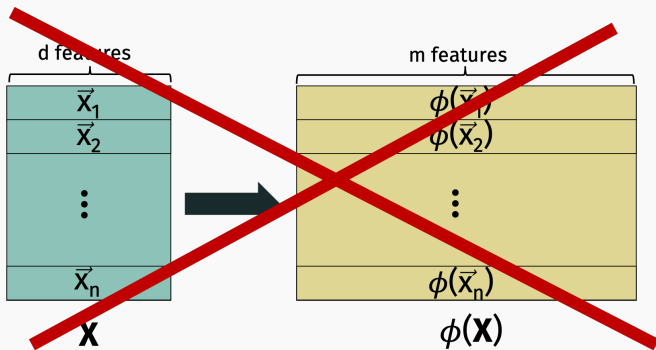
$$\phi(x_i)^T \phi(x_j) = k(x_i, x_j)$$

$$e^{-\|x_i - x_j\|^2}$$

KERNEL TRICK

We never need to actually compute $\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_n)$ explicitly!

- For training we just need the kernel matrix \mathbf{K} , which requires computing $k(\mathbf{x}_i, \mathbf{x}_j)$ for all i, j .
- For testing we just need to compute $k(\mathbf{x}_{new}, \mathbf{x}_i)$ for all i .



$$\left(\langle x, w \rangle + 1 \right)^q$$


$\hookrightarrow O(d)$

This can lead to significant computational savings!

- Transform dimension m is often very large: e.g. $m = \underline{O(d^q)}$ for a degree q polynomial kernel.
- For many kernels (e.g. the Gaussian kernel) m is actually *infinite*. So kernel trick is your only option.

BEYOND THE KERNEL TRICK

The kernel matrix \mathbf{K} is still $n \times n$ though which is huge when the size of the training set n is large. Has made the kernel trick less appealing in some modern ML applications.




There is an inherent quadratic dependence on n in the computational and space complexity of kernel methods.

- 10,000 data points \rightarrow runtime scales as \sim 100,000,000, \mathbf{K} takes 800MB of space.
- 1,000,000 data points \rightarrow runtime scales as \sim 10^{12} , \mathbf{K} takes 8TB of space.

BEYOND THE KERNEL TRICK

Tensor Sketch

Many algorithmic advances in recent years partially address this computational challenge (random Fourier features methods, Nystrom methods, etc.)





Often based on “reversing” the kernel trick to find a compact feature set that well approximates the kernel.⁴

⁴This was a major topic of my research 3-5 years ago.

KERNEL REGRESSION

We won't study kernel regression in detail, but it's a very important statistical tool, especially when dealing with spatial or temporal data.



Also known as Gaussian Process (GP) Regression or Kriging.
Can be justified from a Bayesian modeling perspective.

KERNEL REGRESSION

Reformulation of linear regression:

$$\beta = X^T \alpha$$

$$\min_{\beta} \|X\beta - y\|_2^2 + \lambda \|\beta\|_2^2 \rightarrow \min_{\alpha} \|XX^T \alpha - y\|_2^2 + \lambda \|X^T \alpha\|_2^2$$

Replace XX^T by kernel matrix K during training.

$$= (X^T v)^T (X^T \alpha)$$

$$= \alpha^T X X^T \alpha$$

Prediction:

$$y_{new} = \sum_{i=1}^n \alpha_i \cdot \underline{k(x_{new}, x_i)}.$$

Added benefit: Relatively numerically stable. E.g. is a much better option for performing multivariate or even single variate polynomial regression than direct feature expansion.

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_{10} \end{bmatrix}$$



$$[1 \ x_1 \ \dots \ x_1^8] \quad x_1 = .5$$

$$[1 \ .5 \ .25 \ \dots \ 0 \ 0 \ 0 \ 0]$$

$$x_1 = 2$$

$$[1 \ x_{10} \ \dots \ x_{10}^8]$$

$$1 \ 2 \ 4 \ \dots \ 10$$

$$\left[(x_i x_{i+1})^8 \right]$$