

Intro. to machine learning (CSI 5325)

Lecture 18: Instance-based learning

Greg Hamerly

Spring 2008

Some content from Tom Mitchell.

1 k -Nearest Neighbor

Instance-Based Learning

Key idea: just store all training examples $\langle x_i, f(x_i) \rangle$

Nearest neighbor algorithm:

- Given query instance x_q , first locate nearest training example x_i , then estimate $\hat{f}(x_q) \leftarrow f(x_i)$

This is also called the '1-nearest neighbor' algorithm.

Using multiple neighbors

Why not use more than just one neighbor?

k -Nearest neighbor:

- Given x_q , take vote among its k nearest neighbors (if discrete-valued target function)
- Let $N(x_q, k)$ be the indexes of the k nearest neighbors of x_q
- Regression: mean of f values over k nearest neighbors:

$$\hat{f}(x_q) \leftarrow \frac{1}{k} \sum_{i \in N(x_q, k)} f(x_i)$$

- Classification: most popular f value over k nearest neighbors:

$$\hat{f}(x_q) \leftarrow \underset{i \in N(x_q, k)}{\text{mode}} f(x_i)$$

When To Consider Nearest Neighbor

- Instances map to points in \mathcal{R}^d
- Less than 20 attributes per instance
- Lots of training data



Advantages:

- Training is very fast
- Learn complex target functions
- Don't lose information

Disadvantages:

- Slow at query time
- Easily fooled by irrelevant attributes

Voronoi Diagram



Behavior in the Limit

Consider $p(x)$ defines probability that instance x will be labeled 1 (positive) versus 0 (negative).

Nearest neighbor:

- As number of training examples $\rightarrow \infty$, approaches Gibbs Algorithm
Gibbs: with probability $p(x)$ predict 1, else 0

k -Nearest neighbor:

- As number of training examples $\rightarrow \infty$ and k gets large, approaches Bayes optimal
Bayes optimal: if $p(x) > .5$ then predict 1, else 0

Note Gibbs has at most twice the expected error of Bayes optimal

Bias and variance in k -NN

What is the inductive bias of k -NN?

What is the *statistical* bias and *statistical* variance of k -NN?

- In particular, how do they vary with k ?
- Think about extreme k values, e.g. $k = 1$ or $k = n$.

Distance-Weighted k -NN

Might want to weight nearer neighbors more heavily...

$$\hat{f}(x_q) \leftarrow \frac{\sum_{i \in N(x_q, k)} w_i f(x_i)}{\sum_{i \in N(x_q, k)} w_i}$$

where

$$w_i \equiv \frac{1}{d(x_q, x_i)^2}$$

and $d(x_q, x_i)$ is distance between x_q and x_i

Note: we could use *all* training examples instead of just k

→ Shepard's method

Curse of Dimensionality

Imagine instances described by 20 attributes, but only 2 are relevant to target function

Curse of dimensionality: nearest neighbor is easily misled when high-dimensional X

One approach:

- Stretch j th axis by weight z_j , where z_1, \dots, z_d chosen to minimize prediction error
- Use cross-validation to automatically choose weights z_1, \dots, z_d
- Note setting z_j to zero eliminates this dimension altogether

see [Moore and Lee, 1994]