# 6.867 Section 6: Nonparametric models

# Contents

1	Intro	2
2	Trees	2
	2.1 Regression	3
	2.1.1 Building a tree	3
	2.1.2 Pruning	4
	2.2 Classification	4
3	Bagging	5
	3.1 Random Forests	6
4	Boosting	7
	4.1 Adaboost Algorithm	7
	4.2 Additive models	8
	4.3 Margin	10
5	Nearest-neighbor methods	10
	5.1 Regression	11
	5.2 Kernel density estimation	12
	5.3 Classification	13
	5.4 Curse of dimensionality	13

#### Fall 2012

# 1 Intro

We will continue to broaden the class of models that we can fit to our data. Graphical models allowed us to fit joint distributions that are compact to represent, and efficient for learning and inference. The complexity was adaptable, in the sense that structure search could try different dependence models and select one that was expected to be a good model of future data.

We now turn to models that are not directly interpretable as fitting a fixed-dimension distribution to data. The name *non-parametric methods* is misleading: it is really a class of methods that does not have a fixed parameterization in advance. Some non-parametric models, such as trees and boosting, which we might call *semi-parametric methods*, can be seen as dynamically constructing something that ends up looking like a more traditional parametric model, but where the actual training data affects exactly what the form of the model will be. Other non-parametric methods, such as nearest-neighbor and Bayesian non-parametric methods, rely directly on the data to make predictions and do not compute a model that summarizes the data.

The semi-parametric methods tend to have the form of a composition of simple models. We'll look at:

- *Tree models*: partition the input space and use different simple predictions on different regions of the space; this increases the hypothesis space.
- *Additive models*: train several different classifiers on the whole space and average the answers; this decreases the variance.

*Boosting* is a way to construct an additive model that both increases hypothesis space and decreases variance.

# 2 Trees

The idea here is that we would like to find a partition of the input space and then fit very simple models to predict the output in each piece. The partition is described using a (typically binary) "decision tree," which recursively splits the space.

These methods differ by:

- The class of possible ways to split the space at each node; these are generally linear splits, either aligned with the axes of the space, or more general.
- The class of predictors within the partitions; these are often simply constants, but may be probability distribution or more general classification or regression models.
- The way in which we control the complexity of the hypothesis: it would be within the capacity of these methods to have a separate partition for each individual training example.
- The algorithm for making the partitions and fitting the models.

The primary advantage of tree models is that they are understandable by humans. This is important in application domains, such as medicine, where there are human experts who think they know what they're doing and where decisions are critically important.

We'll concentrate on the CART/ID3 family of algorithms, which were invented independently in the statistics and the artificial intelligence communities. They work by greedily constructing a partition, where the splits are *axis aligned* and by fitting a *constant* model in the leaves. The interesting questions are how to select the splits and and how to control capacity. The regression and classification versions are very similar.

#### 2.1 Regression

	No model	Prediction rule	Prob model	Dist over models
Regression		*		

The classifier is made up of

- A partition function, π, mapping elements of the input space into exactly one of M regions, R<sub>1</sub>,..., R<sub>M</sub>.
- A collection of M output values, O<sub>m</sub>, one for each region.

If we already knew a division of the space into regions, we would set  $\hat{y}_m$ , the constant output for region  $R_m$  to be the average of the output values in that region; that is:

$$O_{\mathfrak{m}} = \operatorname{average}_{\{\mathfrak{i}|\mathfrak{x}^{(\mathfrak{i})} \in \mathfrak{R}_{\mathfrak{m}}\}} \mathfrak{y}^{(\mathfrak{i})}$$

Define the error in a region as

$$\mathsf{E}_{\mathfrak{m}} = \sum_{\{\mathfrak{i} \mid \mathbf{x}^{(\mathfrak{i})} \in \mathsf{R}_{\mathfrak{m}}\}} (\mathfrak{y}^{(\mathfrak{i})} - \mathsf{O}_{\mathfrak{m}})^2 \ .$$

Ideally, we would select the partition to minimize

$$cM+\sum_{\mathfrak{m}=1}^{M} E_{\mathfrak{m}}$$
 ,

for some regularization constant c. It is enough to search over all partitions of the training data (not all partitions of the input space) to optimize this, but the problem is NP-complete.

#### 2.1.1 Building a tree

So, we'll be greedy. We establish a criterion, given a set of data, for finding the best single split of that data, and then apply it recursively to partition the space. We will select the partition of the data that *minimizes the sum of the mean squared errors of each partition*.

Given a data set D, let

- $R_{j,s}^+(D) = \{x \in D \mid x_j \ge s\}$
- $R_{j,s}^{-}(D) = \{x \in D \mid x_j < s\}$
- $\hat{y}_{j,s}^+ = average_{\{i|x^{(i)} \in R_{i,s}^+(D)\}} y^{(i)}$
- $\hat{y}_{j,s}^- = \operatorname{average}_{\{i|x^{(i)} \in R_{i,s}^-(D)\}} y^{(i)}$

**BuildTree**(D):

- If  $|D| \leq k$ : return Leaf(D)
- Find the variable j and split point s that minimizes:

$$E_{R_{j,s}^+(D)} + E_{R_{j,s}^+(D)}$$
.

• Return Node(j, s, BuildTree(R<sup>+</sup><sub>i,s</sub>(D)), BuildTree(R<sup>-</sup><sub>i,s</sub>(D))

Each call to **BuildTree** considers O(dn) splits (only need to split between each data point in each dimension); each requires O(n) work.

#### Fall 2012

#### 2.1.2 Pruning

It might be tempting to regularize by stopping for a somewhat large k, or by stopping when splitting a node does not significantly decrease the error. One problem with short-sighted stopping criteria is that they might not see the value of a split that is, essentially, two-dimensional. So, we will tend to build a tree that is much too large, and then prune it back.

Define cost complexity of a tree T, where m ranges over its leaves as

$$C_{\alpha}(T) = \sum_{m=1}^{|T|} E_m(T) + \alpha |T| \ . \label{eq:calculation}$$

For a fixed  $\alpha$ , find a T that (approximately) minimizes  $C_{\alpha}(T)$  by "weakest-link" pruning. Create a sequence of trees by successively removing the bottom-level split that minimizes the increase in overall error, until the root is reached. Return the T in the sequence that minimizes the criterion.

Pick  $\alpha$  using cross validation.

#### 2.2 Classification

	No model	Prediction rule	Prob model	Dist over models
Classification		*		

The strategy for building and pruning classification trees is very similar to the one for regression trees.

The output is now the majority of the output values in the leaf:

$$O_{\mathfrak{m}} = \operatorname{majority}_{\{\mathfrak{i} | \mathfrak{x}^{(\mathfrak{i})} \in \mathfrak{R}_{\mathfrak{m}}\}} \mathfrak{y}^{(\mathfrak{i})}$$

Define the error in a region as the number of data points that do not have the value O<sub>m</sub>:

$$\mathsf{E}_{\mathfrak{m}} = \left| \{ \mathfrak{i} \mid \mathfrak{x}^{(\mathfrak{i})} \in \mathsf{R}_{\mathfrak{m}} \text{ and } \mathfrak{y}^{(\mathfrak{i})} \neq \mathsf{O}_{\mathfrak{m}} \} \right|$$

Define the *empirical probability* of an item from class k occurring in region m as:

$$\hat{P}_{mk} = \hat{P}(R_m)(k) = \frac{|\{i \mid x^{(i)} \in R_m \text{ and } y^{(i)} = k\}}{N_m}$$

We'll define the empirical probabilities of feature values, as well, for later use.

$$\hat{P}_{mj\nu} = \hat{P}(R_{mj})(\nu) = \frac{\left| \{i \mid x^{(i)} \in R_m \text{ and } x_j^{(i)} = \nu \} \right|}{N_m}$$

Splitting criteria Minimize "impurity" in child nodes. Some measures include:

• *Misclassification error*:

$$Q_{m}(T) = \frac{E_{m}}{N_{m}} = 1 - \hat{P}_{mO_{m}}$$

• Gini index:

$$Q_{\mathfrak{m}}(\mathsf{T}) = \sum_{k} \hat{\mathsf{P}}_{\mathfrak{m}k}(1 - \hat{\mathsf{P}}_{\mathfrak{m}k})$$

• Entropy:

$$Q_{\mathfrak{m}}(T) = H(R_{\mathfrak{m}}) = -\sum_{k} \hat{P}_{\mathfrak{m}k} \log \hat{P}_{\mathfrak{m}k}$$

Choosing the split that minimizes the entropy of the children is equivalent to maximize the *information gain* of the test  $X_j = v$ , defined by

infoGain
$$(X_j = \nu, R_m) = H(R_m) - \left(\hat{P}_{mj\nu}H(R_{j,\nu}^+) + (1 - \hat{P}_{mj\nu})H(R_{j,\nu}^-)\right)$$

Consider the two-class case. All have value

 $\begin{cases} 0.0 & \text{when } \hat{P}_{m0} = 0.0 \\ 0.5 & \text{when } \hat{P}_{m0} = 0.5 \\ 0.0 & \text{when } \hat{P}_{m0} = 1.0 \end{cases}$ 

There used to be endless haggling about which one to use. It seems to be traditional to use:

- Entropy to select which node to split while growing the tree
- Misclassification error in the pruning criterion

Points about trees There are many variations on this theme:

- Linear regression or other regression or classification method in each leaf
- Non-axis-parallel splits: e.g., run a perceptron for a while to get a split.

What's good about trees:

- People understand them
- Easy to handle multi-class classification
- Easy to handle different loss functions (just change predictor in the leaves)

What's bad about trees:

- High variance: small changes in the data can result in very big changes in the hypothesis.
- Usually not the best predictions

**Hierarchical mixture of experts** Make a "soft" version of trees, in which the splits are probabilistic (so every point has some degree of membership in every leaf). Can be trained with a form of EM.

# 3 Bagging

*Bootstrap aggregation* is a technique for reducing the variance of a non-linear predictor, or one that is adaptive to the data.

- Construct B new data sets of size n by sampling them with replacement from  $\mathcal D$
- Train a predictor on each one: f<sup>b</sup>

• Regression case: bagged predictor is

$$\hat{f}_{bag}(x) = \frac{1}{B} \sum_{b=1}^{B} \hat{f}^{b}(x)$$

• Classification case: majority bagged predictor: let  $\hat{f}^{b}(x)$  be a vector with a single 1 and K - 1 zeros, so that  $\hat{y}^{b}(x) = \arg \max_{k} \hat{f}^{b}(x)_{k}$ . Then

$$\hat{f}_{bag}(x) = \frac{1}{B} \sum_{b=1}^{B} \hat{f}^{b}(x),$$

which is a vector containing the proportion of classifiers that predicted each class k for input x; and the predicted output is

$$\hat{y}_{bag}(x) = \arg \max_{k} \hat{f}_{bag}(x)_{k}$$
 .

Alternatively, we can average the class probabilities from the individual classifier, which gives us an even lower-variance estimate.

In the case of regression and *squared error*, we can show that expected squared error of a classifier trained on a single data set of size n is bounded below by the squared error of a classifier that is an average over classifiers trained on infinitely many data sets of size n *drawn from the population*. This suggests that bagging (in which new training sets are drawn from the data set, not population) will also decrease error.

For classification under 0-1 loss, bagging can improve a good classifier, but it can also make a bad classifier worse. But, here's a way to understand its advantage:

- Let the Bayes optimal decision at x be Y(x) = 1 in a two-class example.
- Suppose each "committee member"  $Y_b$  has an error rate  $e_b < e < 0.5$
- Let  $S_1(x) = \sum_{b=1}^{B} I(G_b(x) = 1)$  be the number of votes for class 1 given input x
- If the committee members are independent,  $S_1(x) \sim Bin(B, 1 e)$  and so  $Pr(S_1 > B/2) \rightarrow 1$  as B gets large.

The main issue with this analysis is that it assumes the committee members are independent, which they definitely are not in this case.

Also, when we bag a model, any simple predictability is lost.

#### 3.1 Random Forests

Random forests are collections of trees that are constructed to be de-correlated, so that using them to vote gives maximal advantage. For b = 1..B

- Draw a bootstrap sample  $\mathcal{D}_b$  of size n from  $\mathcal{D}$
- Grow a tree on data  $\mathcal{D}_b$  by recursively repeating these steps:
  - Select m variables at random from the d variables
  - Pick the best variable and split point among them
  - Split the node
- return tree T<sub>b</sub>

Given the ensemble of trees, vote to make a prediction on a new x.

From Dietterich, via Hastie, Tibshirani, and Friedman.

This is the "Wisdom of Crowds."

#### Fall 2012

## 4 Boosting

We will explore a method for making an additive model, but where we explicitly construct new data sets for training each new member of the ensemble, so that new classifiers attempt to "make up for" weaknesses of the current committee.

Assume a two-class problem with  $y \in \{-1, 1\}$ . The *training error rate* of a predictor G is

$$\hat{\mathsf{E}}(\mathsf{G}) = \frac{1}{n} \sum_{i=1}^{N} \mathsf{I}(\mathsf{y}^{(i)} \neq \mathsf{G}(\mathsf{x}^{(i)}))$$

*Expected error rate* is

$$E_{X,Y}I(Y \neq G(X))$$

G is a *weak classifier* if its expected error is less than 0.5.

Assume we have an algorithm WL that takes in (weighted) data sets  $(x^{(i)}, y^{(i)}, w^{(i)})$ and produces a weak classifier that attempts to minimize *weighted training error rate*:

$$\hat{E}_{W}(G) = \frac{\sum_{i=1}^{n} w^{(i)} I(y^{(i)} \neq G(x^{(i)}))}{\sum_{i=1}^{n} w^{(i)}} = \frac{ww}{W}$$

I just took this opportunity to define ww as *weight on wrong predictions* and W as the total weight. We will also define *weight on correct predictions* as wc = W - ww.

#### 4.1 Adaboost Algorithm

The boosting algorithm works in a loop: feeding the training data into WL, evaluating the resulting classifier  $G_1$  on the data, reweighting it to place more emphasis that were classified incorrectly, feeding the reweighted data into WL to get  $G_2$ , etc. The final classifier has the form

$$G(x) = sign\left(\sum_{m=1}^{M} \alpha_m G_m(x)\right) \ .$$

Here is the *Adaboost.M1* algorithm:

- 1. Initialize data weights  $w_1^{(i)} = 1/n$
- 2. For m = 1 ... M
  - (a)  $G_m = \mathcal{WL}(\mathcal{D}, w_m)$
  - (b) Compute  $E_{w_m}(G_m)$
  - (c) Compute

$$\alpha_{\rm m} = \log \frac{1 - \hat{\mathsf{E}}_{w_{\rm m}}(\mathsf{G}_{\rm m})}{\hat{\mathsf{E}}_{w_{\rm m}}(\mathsf{G}_{\rm m})}$$

(d) For all i,

$$w_{m+1}^{(i)} = w_m^{(i)} \cdot \exp\left(\alpha_m I(y^{(i)} \neq G_m(x^{(i)}))\right)$$

Some versions of boosting normalize the weights after this step.

3. Output

$$G(x) = \text{sign} \left( \sum_{m=1}^M \alpha_m G_m(x) \right) \ .$$

Points that are misclassified in round m have their weights increased by  $exp(\alpha_m)$ , increasing their relative influence in the next round. Another way to see the update is

$$w_{m+1}^{(i)} = w_m^{(i)} \cdot \begin{cases} \frac{1 - \hat{E}_{w_m}(G_m)}{\hat{E}_{w_m}(G_m)} & \text{if } I(y^{(i)} \neq G_m(x^{(i)})) \\ 1 & \text{otherwise} \end{cases}$$

Our premise is that although  $W\mathcal{L}$  may be stupid, it will be able to find a classifier with training error rate less than 0.5.

This view emphasizes that the  $w^{(i)}$  remain positive and that, if the current classifier  $G_m$  is working very poorly (that is, that  $\hat{E}_{w_m}(G_m)$  is near 0.5 ), then the weights of the points we got wrong is not increased by much. If, on the other hand,  $G_m$  is working very well, so  $\alpha_m$  is high, then the examples it got wrong will have their weights increased very substantially.

**Weak learners** You can use any non-ridiculous classifier as a weak learner. A very popular choice is the class of "decision stumps": these are decision trees that just make a single split. It's a terrible classifier all by itself, but can be boosted into generating very good performance.

#### 4.2 Additive models

We can see that boosting ends up fitting a classifier of the form

$$f(x) = \sum_{m=1}^M \beta_m b(x \mbox{ ; } \gamma_m)$$
 ,

where  $\beta_m$  are expansion coefficients and b represents a class of basis functions, parameterized by  $\gamma_m$ . We could try to fit this form directly to data but it is a very difficult optimization problem.

**Forward stagewise additive modeling** Instead, we will fit the model incrementally, or "stagewise". This will be computationally simpler and also end up helping guard against overfitting. The idea is to iteratively add new classifiers that greedily improve the loss of the overall classifier, but not to go back revisit the parameters or expansion coefficients of previous stages.

FSAM algorithm

- 1. Initialize  $f_0(x) = 0$
- 2. For m = 1 ... M
  - (a) Compute new component

$$(\beta_{\mathfrak{m}}, \gamma_{\mathfrak{m}}) = \arg\min_{\beta, \gamma} \sum_{i=1}^{n} L(y^{(i)}, f_{\mathfrak{m}-1}(x^{(i)}) + \beta b(x^{(i)}; \gamma))$$

(b) Add it to the ensemble predictor

$$f_{\mathfrak{m}}(x)=f_{\mathfrak{m}-1}(x)+\beta_{\mathfrak{m}}\mathfrak{b}(x\mbox{ ; }\gamma_{\mathfrak{m}})$$
 .

Return  $f_M$ 

Adaboost as FSAM We can see Adaboost.M1 as an instance of FSAM with the exponential loss function

$$L(\textbf{y},f(\textbf{x}))=exp(-\textbf{y}f(\textbf{x}))$$
 .

At each stage of FSAM, it is necessary to solve the problem

$$\begin{aligned} (\beta_{\mathfrak{m}}, G_{\mathfrak{m}}) &= & \arg\min_{\beta, G} \sum_{i=1}^{n} \exp\left(-y^{(i)} \left(f_{\mathfrak{m}-1}(x^{(i)}) + \beta G(x^{(i)})\right)\right) \\ &= & \arg\min_{\beta, G} \sum_{i=1}^{n} w_{\mathfrak{m}}^{(i)} \exp\left(-y^{(i)} \beta G(x^{(i)})\right) \end{aligned}$$

where

$$w_{\mathfrak{m}}^{(\mathfrak{i})} = \exp(-y^{(\mathfrak{i})}f_{\mathfrak{m}-1}(x^{(\mathfrak{i})}))$$
 .

For any positive  $\beta$ ,

$$\begin{split} G_{\mathfrak{m}} &= & \arg\min_{G} \quad \exp(-\beta) \sum_{y^{(i)} = G(x^{(1)})} w_{\mathfrak{m}}^{(i)} \; + \; \exp(\beta) \sum_{y^{(i)} \neq G(x^{(1)})} w_{\mathfrak{m}}^{(i)} \\ &= & \arg\min_{G} \quad (\exp(\beta) - \exp(-\beta)) \sum_{i=1}^{n} w_{\mathfrak{m}}^{(i)} I(y^{(i)} \neq G(x^{(i)})) \; + \; \exp(-\beta) \sum_{i=1}^{n} w_{\mathfrak{m}}^{(i)} \\ &= & \arg\min_{G} \sum_{i=1}^{n} w_{\mathfrak{m}}^{(i)} I(y^{(i)} \neq G(x^{(i)})) \\ &= & \mathcal{WL}(\mathcal{D}, w_{\mathfrak{m}}) \end{split}$$

Using a negative  $\beta$  is like flipping the signs on the  $y^{(i)}$ , and just ask for the classifier to generate the negation of the correct answers. It doesn't add any value, so we will not consider doing so.

Now, given  $G_m$ , we can solve for  $\beta$ 

$$\beta_{\mathfrak{m}} = \arg \min_{\beta} \exp(-\beta) \sum_{\mathbf{y}^{(i)} = G_{\mathfrak{m}}(\mathbf{x}^{(1)})} w_{\mathfrak{m}}^{(i)} + \exp(\beta) \sum_{\mathbf{y}^{(i)} \neq G_{\mathfrak{m}}(\mathbf{x}^{(1)})} w_{\mathfrak{m}}^{(i)}$$
  
$$\beta_{\mathfrak{m}} = \arg \min_{\beta} \exp(-\beta) w_{\mathfrak{m}} + \exp(\beta) w_{\mathfrak{m}}$$
  
$$= \arg \min_{\beta} \exp(-\beta) W_{\mathfrak{m}} (1 - \hat{\mathsf{E}}_{w_{\mathfrak{m}}}(G_{\mathfrak{m}})) + \exp(\beta) W_{\mathfrak{m}} \hat{\mathsf{E}}_{w_{\mathfrak{m}}}(G_{\mathfrak{m}})$$

where  $W_m = \sum_{i=1}^n w_m^{(i)}$ . Taking the derivative with respect to  $\beta$  and setting to 0, we find that

$$\beta_{\mathfrak{m}} = \frac{1}{2} \log \frac{1 - \hat{E}_{w_{\mathfrak{m}}}(G_{\mathfrak{m}})}{\hat{E}_{w_{\mathfrak{m}}}(G_{\mathfrak{m}})} \ .$$

Now

$$f_{\mathfrak{m}}(x)=f_{\mathfrak{m}-1}(x)+\beta_{\mathfrak{m}}G_{\mathfrak{m}}(x)$$
 ,

so the next weights are

$$\begin{split} w_{m+1}^{(i)} &= \exp(-y^{(i)}f_m(x^{(i)})) \\ &= \exp(-y^{(i)}(f_{m-1}(x^{(i)}) + \beta_m G_m(x^{(i)}))) \\ &= \exp(-y^{(i)}f_{m-1}(x^{(i)})) \exp(-y^{(i)}\beta_m G_m(x^{(i)}))) \\ &= w_m^{(i)}\exp(-y^{(i)}\beta_m G_m(x^{(i)})) \\ &= w_m^{(i)}\exp(\beta_m(2I(G_m(x^{(i)}) \neq y^{(i)}) - 1)) \\ &= w_m^{(i)}\exp(2\beta_m I(G_m(x^{(i)}) \neq y^{(i)}))\exp(-\beta_m) \\ &= w_m^{(i)}\exp(\alpha_m I(G_m(x^{(i)}) \neq y^{(i)}))\exp(-\beta_m) \end{split}$$

This has the same effect as the adaboost reweighting, even though it multiplies by a fixed factor of  $exp(-\beta_m)$ .

So! Adaboost minimizes exponential loss by doing stagewise minimization of weighted misclassification error. This behavior is very interesting. If we plot error vs iteration, we find that:

- Misclassification rate on the training set falls to zero and stays there
- Even after training error is 0, the exponential loss continues to decrease.

If we knew the true data distribution, exponential loss would be minimized by picking

$$f^*(x) = arg\min_{f(x)} \mathsf{E}_{Y|x}(exp(-Yf(x))) = \frac{1}{2}\log\frac{Pr(Y=1\mid x)}{Pr(Y=-1\mid x)} \ .$$

Somewhat sensitive to outliers (mis-labeled points). Can make it more robust using a different loss function (HTF likes binomial deviance), but then there isn't a closed form for the optimization so finding each  $G_m$  will require a gradient descent.

Better than bagging.

#### 4.3 Margin

Define the *voting margin* of a point to be

$$margin_{f}(x,y) = \frac{yf(x)}{\sum_{m} |\alpha_{m}|} = \frac{y\sum_{m} \alpha_{m}G_{m}(x)}{\sum_{m} |\alpha_{m}|}$$

It is a number in [-1, +1] and is positive iff f correctly classifies the example. It can be interpreted as a measure of confidence in the prediction and it continues to decrease with rounds of boosting.

Schapire et al showed that larger margins on the training set result in superior upper bounds on the generalization error. For any  $\theta > 0$ , with high probability,

$$\mathsf{E}_{gen} \leqslant \hat{\mathsf{Pr}}(\mathsf{margin}_{\mathsf{f}}(\mathsf{x}, \mathsf{y}) \leqslant \theta) + O\left(\sqrt{\frac{d}{\mathfrak{n}\theta^2}}\right) \ ,$$

where Pr is the empirical probability of the event in the data set. Note that the bound is independent of M.

### 5 Nearest-neighbor methods

	No model	Prediction rule	Prob model	Dist over models
Regression	*			
Classification	*			
Density estimation	*			

Nearest neighbor methods are some of the simplest and are often very effective. They make few model assumptions and require little or no "training time". They can be very expensive in space and/or time to make predictions.

The idea is just that we remember all the data we have ever seen, and don't attempt to try to compute a model of it. Given a query, we answer it "directly" using a computation on the data.

These methods are often, confusingly, called "kernel methods." The notion of kernel as a distance function is used here as well as in other "kernelized" learning algorithms, but the details of the algorithms are very different.

#### 5.1 Regression

Remember all your data!  $\mathcal{D} = \{(x^{(i)}, y^{(i)})\}.$ 

k-nearest-neighbor (KNN) regression:

$$\hat{f}(x) = \frac{1}{k} \sum_{\{i \mid x^{(i)} \in N_k(x)\}} y^{(i)} \text{ ,}$$

where *neighbors*  $N_k(x)$  are the k points in  $\mathcal{D}$  minimizing (typically) Euclidean distance  $||x - x^{(i)}||$  from x.

**Nadaraya-Watson Kernel-weighted average (KWA)** : **KNN** is straightforward, but gives discontinuous (piecewise constant) regression. If we would prefer for it to be smooth, we can use

$$\label{eq:f(x)} \hat{f}(x) = \frac{\sum_{i=1}^{n} K_{\lambda}(x,x^{(i)}) y^{(i)}}{\sum_{i=1}^{n} K_{\lambda}(x,x^{(i)})} \ ,$$

where the kernel is defined as

$$\mathsf{K}_\lambda(x,x') = \mathsf{D}\left(\frac{\|x-x'\|}{\lambda}\right) \ .$$

The distance function D can be

• Gaussian:

$$\mathsf{D}(\mathsf{t}) = \frac{1}{\sqrt{2\pi}} \exp(-\mathsf{t}^2/2)$$

• Epanechnikov:

$$\mathsf{D}(\mathsf{t}) = \begin{cases} \frac{3}{4}(1-\mathsf{t}^2) & \text{if } |\mathsf{t}| \leqslant 1\\ 0 & \text{otherwise} \end{cases}$$

Many other choices of kernels. They have to be symmetric (K(u) - K(-u)) and integrate to 1.

Some properties of KWA:

- Fitted function is continuous
- λ is a *bandwidth* or *smoothing* parameter; larger λ implies lower variance (it is more smoothed) but higher bias
- **KWA** with a fixed window width keeps bias constant but variance depends on local data density.
- KNN keeps variance constant but bias depends on data density.

KNN adapts the neighborhood size to the data density, which is useful in places where, for example, data is sparse. We can unify KNN and KWA by seeing them as both using a kernel

$$K_{\lambda}(x, x') = D\left(\frac{\|x - x'\|}{h_{\lambda}(x)}\right)$$

- For **KWA**,  $h_{\lambda}(x) = \lambda$
- For **KNN**,  $h_{\lambda}(x) = ||x x_{[k]}||$  where  $x_{[k]}$  is the kth closest  $x^{(i)}$  to x.

**Locally weighted regression (LWR)** : Locally weighted averages are biased at the boundaries. We can improve them by doing regression. Our predictor will use a different  $\hat{\alpha}$  and  $\hat{\beta}$  for each predicted point x:

$$\hat{f}(x) = \hat{\alpha}(x) + \hat{\beta}(x) \cdot x$$
 ,

where

$$\hat{\alpha}(x), \hat{\beta}(x) = arg\min_{\alpha,\beta} \sum_{i=1}^n K_\lambda(x,x^{(i)}) \left[y^{(i)} - \alpha - \beta \cdot x^{(i)}\right]^2 \ .$$

Inside the sum, the K term is a local weight and the other term is a squared error. To minimize this, we let:

- X be the  $n \times d + 1$  data matrix with a column of 1's added
- W(x) be an  $n \times n$  diagonal weight matrix, with  $W_{ii}(x) = K_{\lambda}(x^{(i)}, x)$ .
- Y be the  $n \times 1$  vector of  $y^{(i)}$  values.

Then

$$\hat{f}(x) = \begin{bmatrix} 1 \\ x \end{bmatrix} (X^{T}W(x)X)^{-1} X^{T}W(x)Y$$
$$= L(x)Y$$

Think of L as a weight vector combining the least-squares with the kernel weights. We call this the *equivalent kernel*.

We could try to remove bias in curved regions of f by using local polynomial regression...at risk of increased variance.

#### **General points**

- Select kernel width with cross validation. Leave-one-out cross validation is *really easy* for NN methods.
- Raw features may need to be *standardized* (scaled to have a fixed variance), or the Euclidean metric might not make any sense.
- As number of dimensions increases it is hard to keep both bias and variance low (need a lot of data in a small neighborhood).

#### 5.2 Kernel density estimation

To estimate the density at a query point x, compute

$$\hat{P}(x) = \frac{num\{x^{(i)} \mid ||x^{(i)} - x||^2 < \lambda\}}{n\lambda}$$

Think of this as putting a rectangular bump of width  $\lambda$  and height  $1/\lambda$  down, centered at each  $x^{(i)}$ . To get a density, divide by the number of bumps.

To make a smoother estimate, use a different kernel. This is known as a *Parzen density estimate*:

$$\hat{P}(x) = \frac{1}{n\lambda} \sum_{i=1}^{n} K_{\lambda}(x, x^{(i)}) \ . \label{eq:prod}$$

Cylinders in a car compared to number of pounds weight.

More on this later.

### 5.3 Classification

Easy to do multiple classes:

 $\hat{f}(x) = arg\max$  number of times y occurs in  $\{y^{(\mathfrak{i})} \mid x^{(\mathfrak{i})} \in N_k(x)\}$  .

Good for irregular decision boundaries. Pick k by LOOCV.

**Theorem:** The asymptotic error rate of 1NN is less than or equal to 2 times the Bayes error rate.

• Assume query point is equal to some training point (this will be true asymptotically).

- Bayes rule tells us to predict  $y^* = \arg \max_y \Pr(y \mid x)$ , which has error  $1 \Pr(y^* \mid x)$
- 1NN error is

$$\sum_{y=1}^{k} \Pr(y = k \mid x) (1 - \Pr(y = k \mid x)) .$$

The first term inside the sum is the probability we got label k for input x; the second is the probability of getting an error if we use it to make a prediction.

For two classes,

$$\begin{aligned} 1\text{NN error} &= & \Pr(Y = 1 \mid x)(1 - \Pr(Y = 1 \mid x)) \cdot 2 \\ &= & (1 - \Pr(Y = 0 \mid x))(1 - \Pr(Y = 1 \mid x)) \cdot 2 \\ &\leqslant & 2(1 - \Pr(Y = y^* \mid x)) \\ &\leqslant & 2 \cdot \text{Bayes error rate} \end{aligned}$$

Useful because it can accommodate distance metrics that respect the underlying invariance properties in the data (e.g., rotation and translation of images).

#### 5.4 Curse of dimensionality

In high dimensions, all the  $x^{(i)}$  will be far from each other and x. Lots of cool examples.

If we want to find a sub-hypercube of a unit-dimensional hypercube in d dimensions that contains 10% of the volume, what would the length of the sides of the subcube be? If n = 1, then 0.1. If n = 10, then 0.8.

All sample points are close to the boundary:

- Let M data points be uniformly distributed in a d-dimensional unit hypersphere
- Make a query at the origin
- The median distance to the closest data point is

$$\left(1-\frac{1}{2}^{1/M}\right)^{1/d} \ .$$

- For M = 5000, d = 10, this value is approximately 0.52!!
- So, if we're doing regression, we're mostly extrapolating.

**Moral of the story:** if your data are *uniformly* distributed in high dimensions, these methods are hopeless.

**But!** If they live in a low-dimensional subspace of that high-dimensional space, it might turn out okay.

Bayes error rate is the error rate of the best possible predictor.