

Warm-up: 20 questions

What animal am I thinking of?

- You can ask one Yes/No question at a time
- Try to get to the answer in as few questions as possible

15-388/688 - Practical Data Science: Decision trees and interpretable models

Pat Virtue
Carnegie Mellon University
Spring 2022

Outline

Decision trees

Training (classification) decision trees

Interpreting predictions

Boosting

Examples

Outline

Decision trees

Training (classification) decision trees

Interpreting predictions

Boosting

Examples

Overview

Decision trees and boosted decision trees are some of the most ubiquitous algorithms in data science

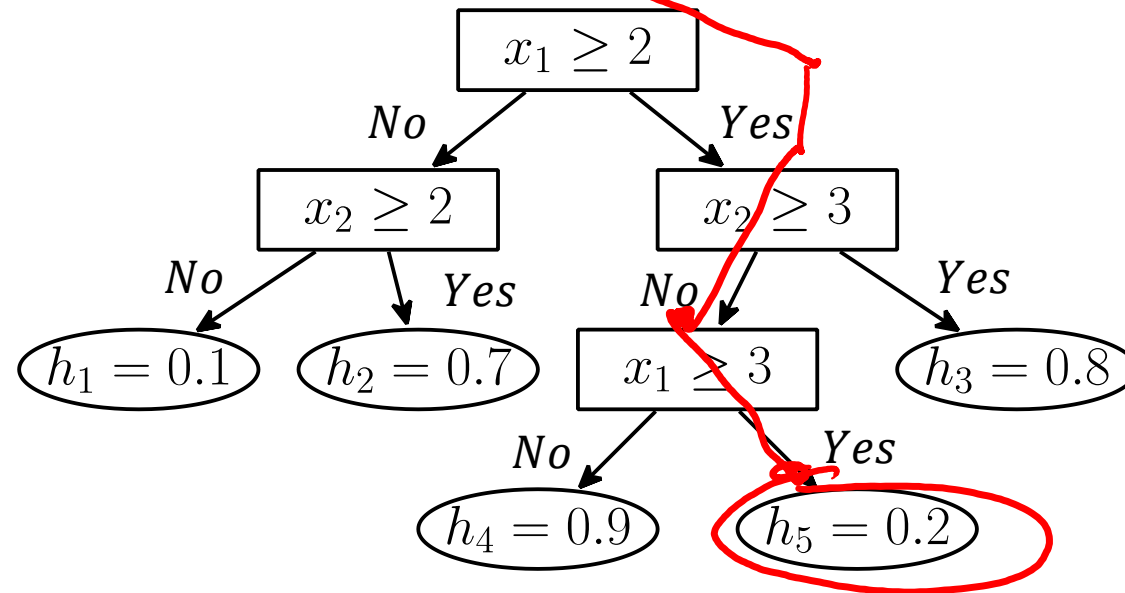
Boosted decision trees typically perform very well without much tuning (the majority of Kaggle contests, for instance, are won with boosting methods)

Decision trees, while not as powerful from a pure ML standpoint, are still one of the canonical examples of an “understandable” ML algorithm

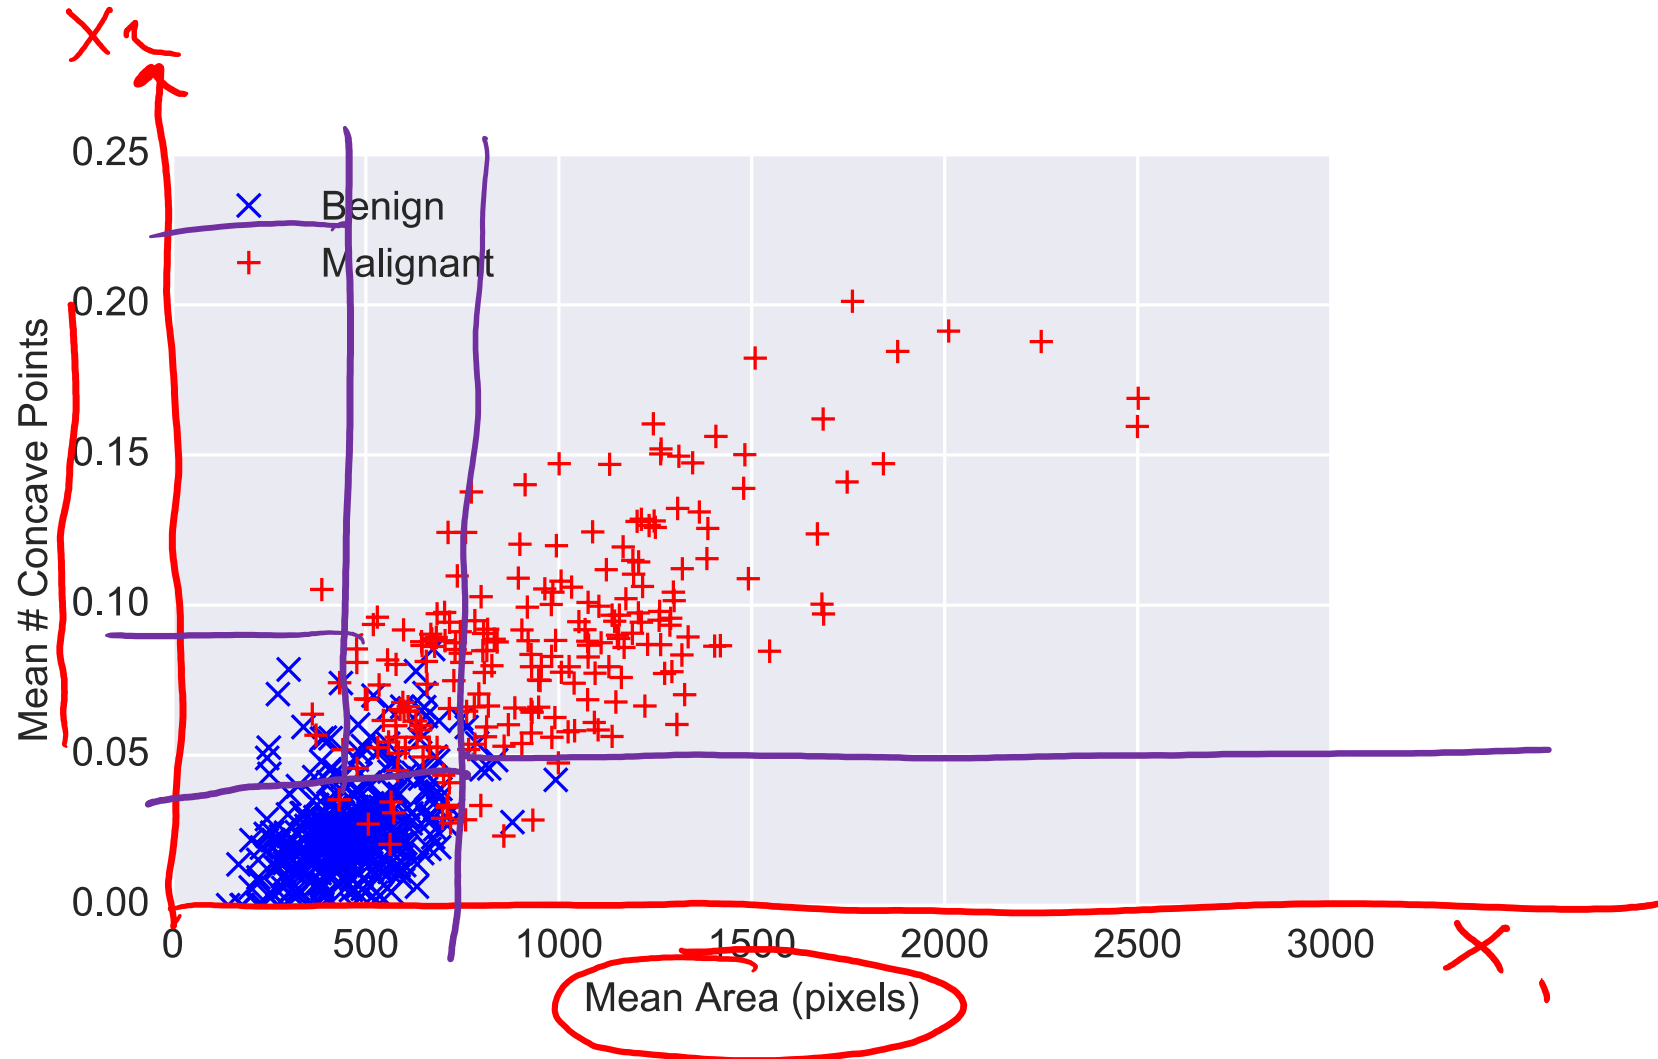
Poll 1

Consider the decision tree below. Given a new point $x = \begin{bmatrix} 3.5 \\ 2.5 \end{bmatrix}$, what should the decision tree predict?

- A. 0.1
- B. 0.2**
- C. 0.7
- D. 0.8
- E. 0.9




Decision tree for cancer prediction



Decision trees

Decision trees were one of the first machine learning algorithms


Basic idea: make classification/regression predictions by tracing through rules in a tree, with a constant prediction at each leaf node



Partitioning the input space

You can think of the hypothesis function of decision trees as partitioning the input space with axis-aligned boundaries

In each partition, predict a constant value



Outline

Decision trees

Training (classification) decision trees


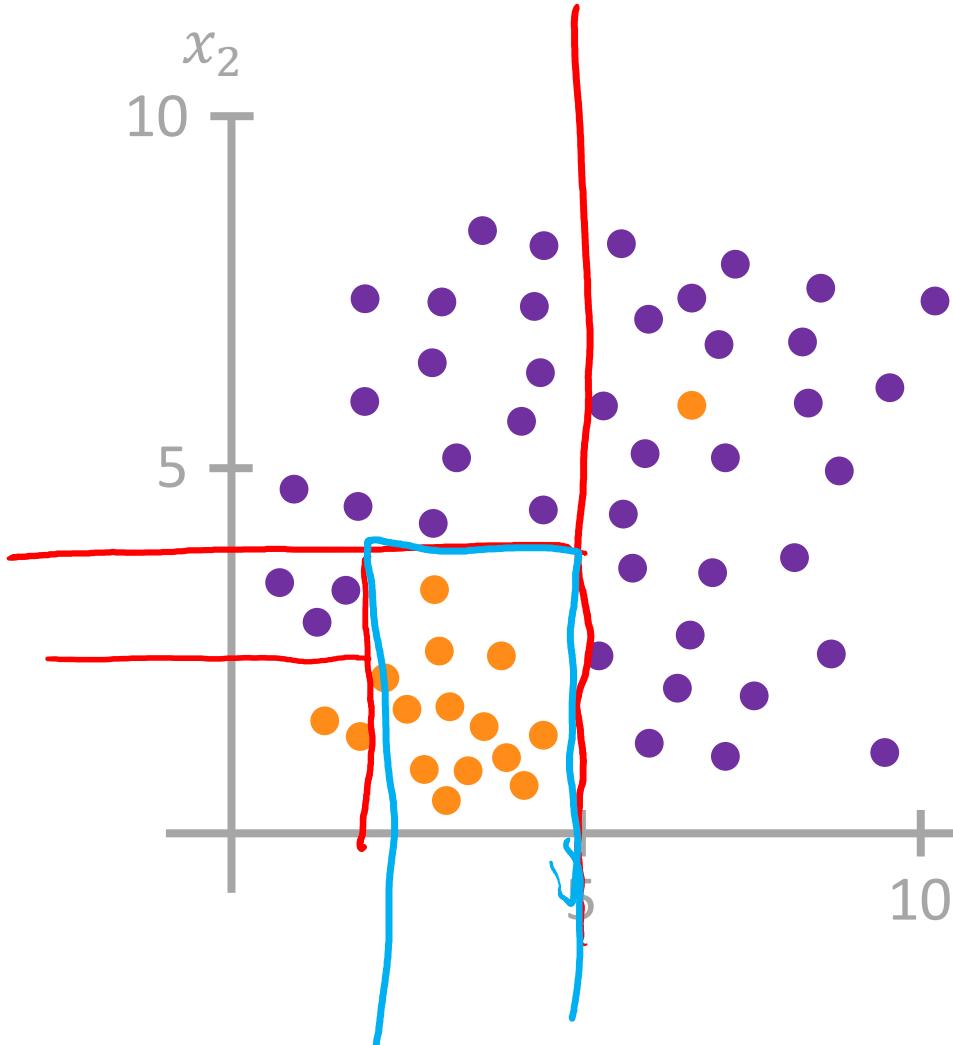
Interpreting predictions

Boosting

Examples

Use training data to learn decision tree

Consider input features $x \in \mathbb{R}^2$



Decision trees as ML algorithms

To specify the decision trees from a machine learning standpoint, we need to specify

1. What is the hypothesis function $h_{\theta}(x)$?
2. What is the loss function $\ell(h_{\theta}(x), y)$?
3. How do we minimize the loss function?

$$\text{minimize}_{\theta} \frac{1}{m} \sum_{i=1}^m \ell(h_{\theta}(x^{(i)}), y^{(i)})$$

Decision trees as ML algorithms

To specify the decision trees from a machine learning standpoint, we need to specify

1. **What is the hypothesis function $h_{\theta}(x)$?**

2. What is the loss function $\ell(h_{\theta}(x), y)$?

3. How do we minimize the loss function?

$$\text{minimize}_{\theta} \frac{1}{m} \sum_{i=1}^m \ell(h_{\theta}(x^{(i)}), y^{(i)})$$

...a decision tree (θ is shorthand for all the parameters that define the tree: tree structure, values to split on, leaf predictions, etc)

Decision trees as ML algorithms


To specify the decision trees from a machine learning standpoint, we need to specify

1. What is the hypothesis function $h_{\theta}(x)$?
- 2. What is the loss function $\ell(h_{\theta}(x), y)$?**
3. How do we minimize the loss function?

$$\text{minimize}_{\theta} \frac{1}{m} \sum_{i=1}^m \ell(h_{\theta}(x^{(i)}), y^{(i)})$$

Reminder: MLE for Bernoulli

Maximum likelihood solution for Bernoulli given by

$$\text{maximize}_{\phi} \prod_{i=1}^m p(y^{(i)}; \phi) = \text{maximize}_{\phi} \prod_{i=1}^m \phi^{y^{(i)}} (1 - \phi)^{1-y^{(i)}}$$


Taking the negative log of the optimization objective (just to be consistent with our usual notation of optimization as minimization)

$$\text{maximize}_{\phi} \ell(\phi) = \sum_{i=1}^m (y^{(i)} \log \phi + (1 - y^{(i)}) \log(1 - \phi))$$

$$\Rightarrow \phi = \frac{\sum_{i=1}^m y^{(i)}}{m}$$

Loss functions in decision trees

Let's assume the output is binary for now (classification task, we will deal with regression shortly), and assume $y \in \{0,1\}$

The typical decision tree algorithm using a probabilistic loss function that considers y to be a Bernoulli random variable with probability $h_\theta(x)$

$$p(y | h_\theta(x)) = h_\theta(x)^y (1 - h_\theta(x))^{1-y}$$

The loss function is just the negative log probability of the output (like in maximum likelihood estimation)


$$\begin{aligned} \ell(h_\theta(x), y) &= -\log p(y | h_\theta(x)) \\ &= -y \log h_\theta(x) - (1 - y) \log(1 - h_\theta(x)) \end{aligned}$$

Decision trees as ML algorithms

To specify the decision trees from a machine learning standpoint, we need to specify

1. What is the hypothesis function $h_{\theta}(x)$?
2. What is the loss function $\ell(h_{\theta}(x), y)$?


3. How do we minimize the loss function?

$$\underset{\theta}{\text{minimize}} \frac{1}{m} \sum_{i=1}^m \ell(h_{\theta}(x^{(i)}), y^{(i)})$$


Optimizing decision trees

Key challenge: unlike models we have considered previously, discrete tree structure means there are no gradients


Additionally, even if we assume binary inputs i.e., $x \in \{0,1\}^n$, there are 2^{2^n} possible decision trees: $n = 7$ means 3.4×10^{38} possible trees



Instead, we're going to use *greedy* methods to incrementally build the tree (i.e., minimize the loss function) one node at a time

Optimizing a single leaf

Consider input features $x \in \mathbb{R}^2$



No $x_1 \geq 5$ Yes

x

$$\frac{|x^+|}{|x|}$$

$$x_2 \geq 4$$

$$\frac{|x_0^+|}{|x_0|}$$

$$\frac{|x_1^+|}{|x_1|}$$

$\frac{3}{14}$

1.0

Optimizing a single leaf

Consider a single leaf in a decision tree (could be root of initial tree)

Let \mathcal{X} denote the examples at this leaf (e.g., in this partition), where \mathcal{X}^+ denotes the positive examples and \mathcal{X}^- denotes negative (zero) examples

What should we choose as the (constant) prediction h at this leaf?

$$\begin{aligned} \underset{h}{\text{minimize}} \sum_{x,y \in \mathcal{X}} \ell(h, y) &= -\frac{|\mathcal{X}^+|}{|\mathcal{X}|} \log h - \frac{|\mathcal{X}^-|}{|\mathcal{X}|} \log(1 - h) \\ \Rightarrow h &= \frac{|\mathcal{X}^+|}{|\mathcal{X}|}, \end{aligned}$$

Which achieves loss:


$$\ell = -h \log h - (1 - h) \log(1 - h)$$

Optimizing splits

Now suppose we want to split this leaf into two leaves, assuming for the time being that $x \in \{0,1\}^n$ is binary

If we split on a given feature j , this will separate \mathcal{X} into two sets: \mathcal{X}_0 and \mathcal{X}_1 (with $\mathcal{X}_0^{+/-}$ and defined similarly to before), and we would choose optimal prediction

$$h_0 = \frac{|\mathcal{X}_0^+|}{|\mathcal{X}_0|}, h_1 = \frac{|\mathcal{X}_1^+|}{|\mathcal{X}_1|}$$



Loss of split

The new leafs will each now suffer loss

$$\ell_0 = -h_0 \log h_0 - (1 - h_0) \log(1 - h_0)$$

$$\ell_1 = -h_1 \log h_1 - (1 - h_1) \log(1 - h_1)$$

Thus, if we split the original leaf on feature j , we no longer suffer our original loss ℓ , but we do suffer losses ℓ_0 and ℓ_1 , i.e., we have decreased the overall loss function

by $\ell - \frac{|x^-|}{|x|} \ell_0 - \frac{|x^+|}{|x|} \ell_1$ (this quantity is called **information gain**)

Greedy decision tree learning – repeat:

- For all leaf nodes, evaluate information gain (i.e., decrease in loss) when splitting on each feature j
- Split the node/feature that minimizes loss the most
- (Run cross-validation to determine when to stop, or after N nodes)

Poll 2


Consider a decision tree training algorithm that stops after creating N nodes. If we increase N , the resulting tree is more likely to:

- A. overfit
- B. underfit

Continuous features

What if x_j 's are continuous?

Solution: sort the examples by their x_j values, compute information gain at each possible split point



Regression trees

Regression trees are the same, except that the hypothesis h are real-valued instead of probabilities, and we use squared loss

$$\ell(h, y) = (h - y)^2$$

This means that the loss a node is given by

$$\underset{h}{\text{minimize}} \frac{1}{|\mathcal{X}|} \sum_{x,y \in \mathcal{X}} (h - y)^2 \implies h = \frac{1}{|\mathcal{X}|} \sum_{x,y \in \mathcal{X}} y \text{ (i. e. mean)}$$

and suffers loss

$$\ell = \frac{1}{|\mathcal{X}|} \sum_{x,y \in \mathcal{X}} (y - h)^2 \text{ (i. e. variance)}$$

Outline

Decision trees

Training (classification) decision trees

Interpreting predictions

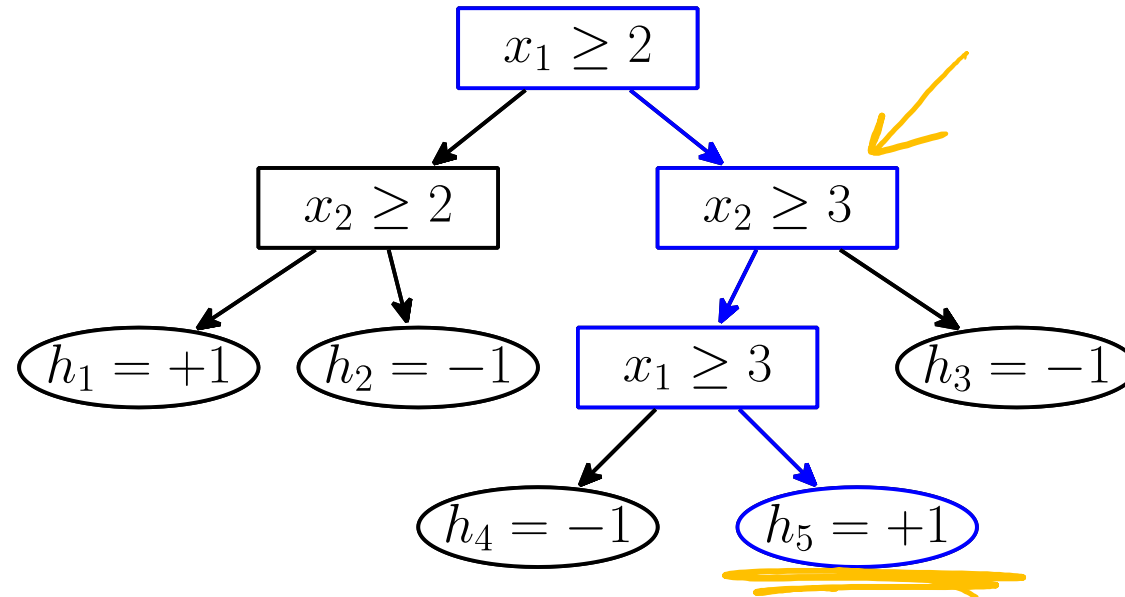
Boosting

Examples

Interpretable models

Decision trees are the canonical example of an interpretable model


Why did we we predict +1?



...because $x_1 \geq 2, x_2 \geq 3, x_1 \geq 3$

...because $x_1 \geq 3, x_2 \geq 3$

Decision tree surface for cancer prediction



...because mean concave points > 0.05, mean area > 791

Explanations in higher dimensions

Explanatory power works “well” even for data with high dimension

Example from full breast cancer dataset with 30 features, “example classified as positive because”

$$\text{max_perimeter} > 117.41, \text{max_concave_points} > 0.087$$

Compare to linear classifier, “example classified positive because”


$$2.142 * \text{mean_radius} + 0.119 * \text{mean_texture} + \dots > 0$$

Do we care about interpretability?

A philosophical question underlying a lot of work in machine learning: should we seek to make predictions that can be “explained”

My own take:

- **Data scientist 1:** I'd rather have a doctor that was 90% accurate than one who was 80% accurate, even if the 80% one gave me explanations.
- **Data scientist 2:** But if the classifier isn't interpretable it may often make “foolish” mistakes because of idiosyncrasies in data.
- **Data scientist 1:** That is not a question of interpretability, that is a question of training set / testing set mismatch.
- **Data scientist 2:** Want to know a good way of telling if you have training set / testing set mismatch? Train an interpretable classifier.

Outline

Decision trees

Training (classification) decision trees

Interpreting predictions

Boosting

Examples

Ensembles of trees

Decision trees have notable advantages (they are relatively easy to interpret, usually fast to train, insensitive to scale of input features)

But, they are also quite limited in their representation power (require axis-aligned splits, don't model probabilities very smoothly)

Basic idea of tree *ensemble* methods is to combine *multiple* tree models together to form a better predictor

Two of the most popular ensemble methods: random forests and boosting

(These models no longer interpretable in the sense we have discussed)

Boosting

Boosting originated as an idea in theoretical machine learning, for “boosting” the performance of weak classifiers (i.e., combining many classifiers that each had modest accuracy to one that had high accuracy)

After some initial success in theory, during the 90s there came to be several practical applications of boosting methods

There are *many* interpretations of boosting (experts still disagree on the “right” one!), and I’m going to highlight one:

We focus on the Gradient Boosted Regression Trees (GBRT) algorithm

Machine learning with general predictions

Let's consider the basic machine learning optimization problem (could be any loss function, classification or regression)

$$\text{minimize}_{\theta} \sum_{i=1}^m \ell(\hat{y}^{(i)}, y^{(i)})$$

where $\hat{y}^{(i)}$ denotes our prediction for the i th example

The gradient with respect to these *predictions* themselves to determine the best way to adjust our predictions (ignoring whether we have any hypothesis function that could actually change the predictions in this way)

$$\frac{\partial}{\partial \hat{y}^{(i)}} \ell(\hat{y}^{(i)}, y^{(i)})$$

Basic idea of GBRTs

Basic idea: GBRTs are effectively performing *gradient descent* on our loss function, using *regression trees* to approximate the gradient

Given: Data set $(x^{(i)}, y^{(i)})_{i=1, \dots, m}$, # trees T , loss ℓ , step size α

Initialize:

$$\hat{y}^{(i)} := 0, \quad \forall i = 1, \dots, m$$

For $t = 1, \dots, T$:

$$h_t := \text{Train_Regression_Tree} \left(\left\{ x^{(i)}, \frac{\partial}{\partial \hat{y}} \ell(\hat{y}^{(i)}, y^{(i)}) \right\}_{i=1}^m \right)$$

Update predictions:

$$\hat{y}^{(i)} := \hat{y}^{(i)} - \alpha h_t(x^{(i)})$$

For new data point x :

$$\text{Predict: } \hat{y} = -\alpha \sum_{i=1}^T h_t(x^{(i)})$$

GBRTs a bit more practically

In practice, fitting the trees is slow, so we actually do a line search to determine how large of a gradient step to take

Can take different gradient steps at different leaves in tree

Here you probably want to use a library (you *could* write your own implementation, but it would be about ~100 lines of code, not ~5 like for an SVM)

Outline

Decision trees

Training (classification) decision trees

Interpreting predictions

Boosting

Examples

Decision trees and GBRTs in scikit-learn

Interface for decision trees and GBRTs in scikit-learn is just like any other classifier


```
from sklearn.tree import DecisionTreeClassifier

clf = DecisionTreeClassifier(criterion='entropy', max_depth=5)
clf.fit(X, y)
```

```
from sklearn.ensemble import GradientBoostingClassifier

clf = GradientBoostingClassifier(loss='deviance', max_depth=3, n_estimators=100)
clf.fit(X, y)
```

Decision tree surface for cancer prediction



GBRT surface for cancer prediction

