# Decision Trees 

6601

CHAPTER 3 DECISION TREE LEARNING


## Classification

| $A_{1}$ | $A_{2}$ | $\ldots$ | $A_{N}$ | $C$ |
| :---: | :---: | :---: | :---: | :---: |
| $v_{11}$ | $v_{12}$ | $\ldots$ | $v_{1 N}$ | $c_{1}$ |
| $v_{21}$ | $v_{22}$ | $\ldots$ | $v_{2 N}$ | $c_{2}$ |

## Decision Tree



## Decision Tree

| $H$ | $W$ | $O$ | $P$ |
| :---: | :---: | :---: | :---: |
| $H$ | $S$ | $S$ | $Y / N ?$ |
| $N$ | $S$ | $O$ | $Y / N ?$ |
| $H$ | $W$ | $R$ | $Y / N ?$ |



## Decision Tree



CHAPTER 3


## Decision Tree

| $H$ | $W$ | $O$ | $P$ |
| :---: | :---: | :---: | :---: |
| $H$ | $S$ | $S$ | $N$ |
| $N$ | $S$ | $O$ | $Y$ |
| $H$ | $W$ | $R$ | $Y / N ?$ |



## Decision Tree




## Continuos Attributes



## Continues Attributes

Guillotine Cut


What is the tree?


## Continues Attributes



## Decision Tree



## Decision Tree



## Which Tree is better if the classification accuracy is the same?

## Decision Tree Learning



Idea: a good attribute splits the examples into subsets that are (ideally) "all positive" or "all negative"


Patrons? is a better choice-gives information about the classification

## Entropy



Claude Shannon

## Measure Of

Uncertainty / Unpredictability in a Random Variable

$$
H(X)=-\sum_{i=1}^{n} p\left(x_{i}\right) \log _{2} p\left(x_{i}\right)
$$

Quantifies Information in a Message

## Entropy

$$
\begin{gathered}
H(X)=-\sum_{i=1}^{n} p\left(x_{i}\right) \log _{b} p\left(x_{i}\right) \\
\left(-\frac{6}{12} * \log _{2} \frac{6}{12}\right)-\left(-\frac{6}{12} * \log _{2} \frac{6}{12}\right)=1
\end{gathered}
$$

Idea: a good attribute splits the examples into subsets that are (ideally) "all positive" or "all negative"


Patrons? is a better choice-gives information about the classification

## Information Gain

$$
\operatorname{Gain}(S, A) \equiv \operatorname{Entropy}(S)-\sum_{v \in \operatorname{Values}(A)} \frac{\left|S_{v}\right|}{|S|} \operatorname{Entropy}\left(S_{v}\right)
$$

Idea: a good attribute splits the examples into subsets that are (ideally) "all positive" or "all negative"


Patrons? is a better choice-gives information about the classification

## Calculate

|  | $H(X)=$ | $\sum_{i=1}^{n} p\left(x_{i}\right) \operatorname{lOg}_{b} p\left(x_{i}\right)$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Gain $(S, A) \equiv$ Entropy $(S)-\sum_{\text {veValues(A) }}$ | $\frac{\left\|S_{v}\right\|}{\|S\|}$ Entropy $\left(S_{v}\right)$ |  |  |  |  |
|  |  |  |  |  |  |
| Day | Outlook | Temperature | Humidity | Wind | PlayTennis |
| D1 | Sunny | Hot | High | Weak | No |
| D2 | Sunny | Hot | High | Strong | No |
| D3 | Overcast | Hot | High | Weak | Yes |
| D4 | Rain | Mild | High | Weak | Yes |
| D5 | Rain | Cool | Normal | Weak | Yes |
| D6 | Rain | Cool | Normal | Strong | No |
| D7 | Overcast | Cool | Normal | Strong | Yes |
| D8 | Sunny | Mild | High | Weak | No |
| D9 | Sunny | Cool | Normal | Weak | Yes |
| D10 | Rain | Mild | Normal | Weak | Yes |
| D11 | Sunny | Mild | Normal | Strong | Yes |
| D12 | Overcast | Mild | High | Strong | Yes |
| D13 | Overcast | Hot | Normal | Weak | Yes |
| D14 | Rain | Mild | High | Strong | No |

## Result

$$
\begin{aligned}
\operatorname{Gain}(S, \text { Outlook }) & =0.246 \\
\operatorname{Gain}(S, \text { Humidity }) & =0.151 \\
\text { Gain }(S, \text { Wind }) & =0.048 \\
\text { Gain }(S, \text { Temperature }) & =0.029
\end{aligned}
$$

## Decision Tree Learning

function DTL(examples, attributes, default) returns a decision tree
if examples is empty then return default else if all examples have the same classification then return the classification else if attributes is empty then return $\operatorname{MODE}$ (examples) else
best $\leftarrow$ Choose-Attribute(attributes, examples) tree $\leftarrow$ a new decision tree with root test best for each value $v_{i}$ of best do examples $_{i} \leftarrow\left\{\right.$ elements of examples with best $\left.=v_{i}\right\}$ subtree $\leftarrow \mathrm{DTL}\left(\right.$ examples $_{i}$, attributes - best, $\operatorname{MODE}($ examples $)$ ) add a branch to tree with label $v_{i}$ and subtree subtree
return tree

## Greedy Algorithm

- In search terms:A greedy algorithm with the Information Gain as a Heuristic
- Could we do better?


## Example

@relation 'gatech_admission'
@attribute 'recommendation' \{'strong', 'weak'\}
@attribute 'gpa' real
@attribute 'gre_math' real
@attribute 'gre_verbal' real
@attribute 'admitted' \{'yes','no'\}
@data
'strong', 4, 800, 800, 'yes' 'weak', 3.4, 600, 500, 'yes' 'strong', 3.6, 800, 550, 'yes' 'strong', 3, 700, 650, 'yes' 'weak', 3.2, 800, 800, 'yes' 'strong', 4, 550, 500, 'yes' 'strong', 3.7, 700, 750, 'yes' 'weak', 4, 800, 200, 'yes' 'strong', 4, 200, 800, 'yes' 'strong', 3.4, 600, 500, 'yes' 'strong', 3.6, 800, 550, 'yes' 'weak', 3, 700, 650, 'yes' 'strong', 4, 550, 500, 'yes' 'strong', 3.7, 700, 750, 'yes' 'weak', 2.8, 800, 800, 'no' 'weak', 4, 200, 200, 'no'
'strong', 2, 500, 200, 'no' 'strong', 3.5, 200, 800, 'no' 'weak', 2, 800, 800, 'no' 'weak', 1.7, I00, 100, 'no' 'weak', 3.7, 50, 0, 'no' 'weak', 2.8, 100, 100, 'no' 'weak', 4, 200, 200, 'no' 'strong', 2, 100,100, 'no' 'weak', 1.7, I00, 100, 'no' 'weak', 3.7, 50, 0, 'no' 'weak', 2.8, I00, 800, 'no' 'weak', 4, 200, 200, 'no' 'strong', 2, 500, 200, 'no' 'strong', 3.5, 200, 800, 'no'

## What is the best?

'weak', 2, 800, 800, 'no' 'weak', 1.7, 100, 100, 'no' 'strong', 3.7,50, 0, 'no' 'weak', 2.8, I00, 800, 'no' 'weak', 4, 200, 200, 'no' 'strong', 2, 500, 200, 'no' 'strong', 3.5, 200, 800, 'no' 'weak', 2, 800, 800, 'no' 'weak', I.7, 100, 100, 'no'<br>'weak', 3.7,50, 0, 'no'

## Weka Demo



## Use Case: Mobile Text

## Entry





## Use Case: Fat Thumbs

prob $\leq 0$
| prevcuradjacent_nom = False
| | curfutadjacent_nom = False
| | | dropprobdiffiabs $\leq 0.000934$ : repeat ( $104.0 / 45.0$ )
| | | dropprobdiff1abs > 0.000934
| | | futneighborprob $\leq 0$ : nonobo ( $552.0 / 11.0$ )
| | | futneighborprob $>0$
| | | | | neighborprob $\leq 0.011765$; nonobo ( $143.0 / 21.0$ )
| | | | | neighborprob > 0.011765
| | | | | ud_sub1 $\leq-140$ : nonobo (172.0/80.0)
| | | | | | ud.sub1 > -140: obosubstitute (527.0/109.0)
| | curfutadjacent_nom $=$ True: rollon (323.0/73.0)
| prevcuradjacent_nom $=$ True: rolloff (555.0/46.0)
prob $>0$
| dt.ud.0.p1 $\leq 124$
| | dropprobdiffsign $\leq-1$ : rolloff (155.0/36.0)
| | dropprobdiffaign >-1: nonerror (111.0/49.0)
| | | | | neighborprob $\leq 0.003185$ : nonerror ( $142.0 / 35.0$ )
| | | | | neighborprob $>0.003185$
| | | | | | average.dd. $2 \leq-88$
| | | | | | | letterfreq $\leq 0.04853$ : obosubstitute (106.0/35.0)
| | | | | | | | letterfreq > 0.04853 : nonobo (142.0/69.0)
| | | | | | average.dd. $2>-88$
| | | | | | | neighborprobdiff $\leq 0.009091$ : nonerror (181.0/65.0)
| | | | | | | neighborprobdiff > 0.009091: obosubstitute (229.0/82.0)
| | | | curfutadjacent_non = True: nonerror (129.0/62.0)
| | | futprob > 0; nonerror (94376.0/1715.0)

## Use Case: Fat Thumbs



No Dictionaries were harmed in the making of this Decision Tree.

# What are the drawbacks? 

Decision Trees are known to over fit the data

## Ensemble Learning



Mixture of Experts: Have we seen one before?

## Random Forests

Winning!


## Random Forests

## Bagging: Bootstrap AGGregation

- INPUT: Data Set of size $N$ with $M$ dimensions
- I) SAMPLE n times from Data
- 2) SAMPLE $m$ times from Attributes
- 3) LEARN TREE on sampled Data and Attributes
- REPEAT UNTIL $k$ trees


## Use Case: Kinect



## Use Case: Kinect



## Use Case: Kinect



Pixel to classify

> (a)

(b)


## Use Case: Kinect

Training 3 trees to depth 20 from I million images takes about a day on a 1000 core cluster

