## Classification Trees

## William Long

## MIT Lab for Computer Science

## Data Mining

- Prediction vs Knowledge Discovery
- Statistics vs Machine Learning
- Phases:
- Problem selection
- Data preparation
- Data reduction
- Method application
- Evaluation of results


## Machine Learning



## Classification Tree



## Classification Trees

- Data consisting of learning set of cases
- Each case consists of a set of attributes with values and has a known class
$\bullet$ Classes are one of a small number of possible values, usually binary
- Attributes may be binary, multivalued, or continuous


## Background

- Classification trees were invented twice
- Statistical community: CART
- Brieman 1984
- Machine Learning community
- Quinlan and others
- Originally called "decision trees"


## Example

| Outlook | Temp | Humidity | Windy? | Class |
| :---: | :---: | :---: | :---: | :---: |
| sunny | 75 | 70 | yes | play |
| sunny | 80 | 90 | yes | dont play |
| sunny | 85 | 85 | no | dont play |
| sunny | 72 | 95 | no | dont play |
| sunny | 69 | 70 | no | play |
| cloudy | 72 | 90 | yes | play |
| cloudy | 83 | 78 | no | play |
| cloudy | 64 | 65 | yes | play |
| cloudy | 81 | 75 | no | play |
| rain | 71 | 80 | yes | dont play |
| rain | 65 | 70 | yes | dont play |
| rain | 75 | 80 | no | play |
| rain | 68 | 80 | no | play |
| rain | 70 | 96 | no | play |

## Example: classified

| Outlook | Temp | Humidity | Windy? | Class |
| :---: | :---: | :---: | :---: | :---: |
| sunny | 75 | 70 | yes | play |
| sunny | 80 | 90 | yes | dont play |
| sunny | 85 | 85 | no | dont play |
| sunny | 72 | 95 | no | dont play |
| sunny | 69 | 70 | no | play |
| cloudy | 72 | 90 | yes | play |
| cloudy | 83 | 78 | no | play |
| cloudy | 64 | 65 | yes | play |
| cloudy | 81 | 75 | no | play |
| rain | 71 | 80 | yes | dont play |
| rain | 65 | 70 | yes | dont play |
| rain | 75 | 80 | no | play |
| rain | 68 | 80 | no | play |
| rain | 70 | 96 | no | play |

## Tree

- Outlook=sunny
- Humidity <=75: play
- Humidity > 75: don't play
- Outlook=cloudy: play

Outlook=rain

- Windy=yes: don't play
- Windy=no: play


## Assumptions

- Independence of partitions
- Branching on individual variables captures behavior
- No linearity assumption Classification
- Although probabilities possible


## Data Types

- Binary
- Multiple valued
- N branches
- Select subsets of values
- Continuous
- Find cut point


## Divide and Conquer

- 9/14: play



## Splitting Criteria

- Information gain
- gain $=-\Sigma \mathrm{p}^{*} \log _{2} \mathrm{p}$

Gini statistic (weighted average impurity)

- Gini $=1-\Sigma \mathrm{p}^{2}$
- Information gain ratio
- Others


## Information Gain

- gain $=-\Sigma \mathrm{p} * \log _{2} \mathrm{p}$
$\Delta \operatorname{info}()=-9 / 14 * \log _{2}(9 / 14)-5 / 14 * \log _{2}(5 / 14)=.940$ bits
$-\operatorname{info}($ outlk $)=5 / 14 *\left(-2 / 5 * \log _{2}(2 / 5)-3 / 5 * \log _{2}(3 / 5)\right)$
$+4 / 14 *\left(-4 / 4 * \log _{2}(4 / 4)-0 / 4 * \log _{2}(0 / 4)\right)$
$5 / 14 *\left(-3 / 5 * \log _{2}(3 / 5)-2 / 5 * \log _{2}(2 / 5)\right)$
$=.694$ bits
$\checkmark$ gain $=.246$ bits
- vs info(windy) $=.892$ bits


## Divide and Conquer

- 9/14: play



## Continuous Variable

| Temp | Class | Ratio | Gain | Gini |
| :---: | :---: | :---: | ---: | ---: |
| 64 | play | $1 / 1+8 / 13$ | 0.048 | 0.577 |
| 65 | dont play | $1 / 2+8 / 12$ | 0.010 | 0.583 |
| 68 | play | $2 / 3+7 / 11$ | 0.000 | 0.587 |
| 69 | play | $3 / 4+6 / 10$ | 0.015 | 0.582 |
| 70 | play | $4 / 5+5 / 9$ | 0.045 | 0.573 |
| 71 | dont play | $4 / 6+5 / 8$ | 0.001 | 0.586 |
| 72 | dont play | $4 / 7+5 / 7$ | 0.016 | 0.582 |
| 72 | play | $5 / 8+4 / 6$ | 0.001 | 0.586 |
| 75 | play | $6 / 9+3 / 5$ | 0.003 | 0.586 |
| 75 | play | $7 / 10+2 / 4$ | 0.025 | 0.579 |
| 80 | dont play | $7 / 11+2 / 3$ | 0.000 | 0.587 |
| 81 | play | $8 / 12+1 / 2$ | 0.010 | 0.583 |
| 83 | play | $9 / 13+0 / 1$ | 0.113 | 0.555 |
| 85 | dont play |  |  |  |

## Information Gain Ratio

- Attributes with multiple values favored by information gain
Correction provided by analogous split info split info $=-\Sigma \mathrm{T}^{*} \log _{2} \mathrm{~T}$ split info $=-5 / 14 * \log _{2}(5 / 14)-4 / 14 * \log _{2}(4 / 14)-$ $5 / 14 * \log _{2}(5 / 14)=1.577$ bits gain ratio $=.246 / 1.577=.156$


## Missing Values

- Adjust gain ratio
$-\operatorname{Gain}(x)=$ prob A is known $* \operatorname{info}(\mathrm{x})$
$-\operatorname{Split}(x)=-u * \log _{2} u-\Sigma T^{*} \log _{2} t$
- Partitioning of training set cases
- Use weights based on prevalence of values
- Classification
- Use weights and sum the weighted leaves


## Example with missing value

| Outlook | Temp | Humidity | W indy? | Class |
| :---: | :---: | :---: | :---: | :---: |
| sunny | 75 | 70 | yes | play |
| sunny | 80 | 90 | yes | dont play |
| sunny | 85 | 85 | no | dont play |
| sunny | 72 | 95 | no | dont play |
| sunny | 69 | 70 | no | play |
| $?$ | 72 | 90 | yes | play |
| cloudy | 83 | 78 | no | play |
| cloudy | 64 | 65 | yes | play |
| cloudy | 81 | 75 | no | play |
| rain | 71 | 80 | yes | dont play |
| rain | 65 | 70 | yes | dont play |
| rain | 75 | 80 | no | play |
| rain | 68 | 80 | no | play |
| rain | 70 | 96 | no | play |

## Frequencies for Outlook

|  | play | don't play | total |
| :--- | :---: | :---: | :---: |
| sunny | 2 | 3 | 5 |
| cloudy | 3 | 0 | 3 |
| rain | 3 | 2 | 5 |
| total | 8 | 5 | 13 |

## Information Gain With Missing

$\operatorname{info}()=-8 / 13 * \log _{2}(8 / 13)-5 / 13 * \log _{2}(5 / 13)=.961$ bits
$-\operatorname{info}($ outlk $)=5 / 13 *\left(-2 / 5 * \log _{2}(2 / 5)-3 / 5 * \log _{2}(3 / 5)\right)$

$$
+3 / 13 *\left(-3 / 3 * \log _{2}(3 / 3)-0 / 3 * \log _{2}(0 / 3)\right)
$$

$5 / 13 *\left(-3 / 5 * \log _{2}(3 / 5)-2 / 5 * \log _{2}(2 / 5)\right)$
$=.747$ bits
gain $=13 / 14 *(.961-.747)=.199$ bits split $=-5 / 14 * \log _{2}(5 / 14)-3 / 14 * \log _{2}(3 / 14)-$ $5 / 14 * \log _{2}(5 / 14)-1 / 14 * \log _{2}(1 / 14)=1.809$ gain ratio $=.199 / 1.809=.110$

## Dividing Sunny

| Outlook | Temp | Humidity | Windy? | Class | W eight |
| :--- | :---: | :---: | :---: | :--- | :---: |
| sunny | 75 | 70 | yes | play | 1 |
| sunny | 80 | 90 | yes | dont play | 1 |
| sunny | 85 | 85 | no | dont play | 1 |
| sunny | 72 | 95 | no | dont play | 1 |
| sunny | 69 | 70 | no | play | 1 |
| $?$ | 72 | 90 | yes | play | $5 / 13$ |

## What Next?

- Most trees are less than perfect
- Variables don't completely predict the outcome
- Data is noisy
- Data is incomplete (not all cases covered)
- Determine the best tree without overfitting or underfitting the data
- Stop generating branches appropriately
- Prune back the branches that aren't justified


## Pruning

$\Delta$ Use a test set for pruning

- Cost complexity: (CART)
» $\mathrm{E} / \mathrm{N}+\alpha^{*} \mathrm{~L}($ tree $)$
- Reduced error
» $\mathrm{E}^{\prime}=\Sigma \mathrm{J}+\mathrm{l}(\mathrm{s}) / 2$
» $\mathrm{E}+1 / 2<\mathrm{e}^{\prime}+\mathrm{se}\left(\mathrm{e}^{\prime}\right)$
- Cross validation
- Split training set into N parts
- Generate N trees, each leaving 1 part for validation


## Pessimistic Pruning: (C4.5)

$\rightarrow$ Estimate errors: $\sum \mathrm{N}^{*} \mathrm{U}_{\mathrm{CF}}(\mathrm{E}, \mathrm{N})$ Example:

$$
-\mathrm{v}=\mathrm{a}: \mathrm{T}(6) \mathrm{U}_{25 \%}(0,6)=.206
$$

$$
-\mathrm{v}=\mathrm{b}: \mathrm{T}(9) \mathrm{U}_{25 \%}(0,9)=.143
$$


$-\mathrm{v}=\mathrm{c}: \mathrm{F}(1) \mathrm{U}_{25 \%}(0,1)=.750$
$-6^{*} .206+9^{*} .143+1^{*} .750=3.273$

- vs $16^{*} U_{25 \%}(1,16)=16^{*} .157=2.512$
$-=>$ eliminate subtree


## Developing a Tree for Ischemia

- Data:
- learning set 3453 cases
- test set 2320 cases
- Attributes: 52
- Types: dichotomous (chest pain), multiple (primary symptom), continuous (heart rate)
- Related attributes
- Missing values


## Concerns

- Probability rather than classification
- Compare to other methods (LR, NN)
- Clinical usefulness


## Probability of Disease

- Fraction at leaf estimates probability
- Small leaves give poor estimates
$\rightarrow$ Correction: $\quad \frac{i\left(n^{\prime}-i^{\prime}\right)+i^{\prime}}{n\left(n^{\prime}-i^{\prime}\right)+n^{\prime}}$


## Tree for Ischemia

```
STCHANGE = 1: ISCHEMIA (166.0/57.3)
STCHANGE = 6: ISCHEMIA (273.0/43.2)
STCHANGE = 0:
| NCPNITRO = 2: NO-ISCHEMIA (1613.0/219.1)
NCPNITRO = 1:
| SYMPTOM1 = 2: NO-ISCHEMIA (6.1/4.8)
    SYMPTOM1 = 4: NO-ISCHEMIA (6.1/4.0)
    SYMPTOM1 = 7: ISCHEMIA (3.0/2.4)
| SYMPTOM1 = 8: ISCHEMIA (17.2/9.3)
| SYMPTOM1 = 9: NO-ISCHEMIA (52.5/16.8)
| | SYMPTOM1 = 1:
| | | SEX = 1: NO-ISCHEMIA (10.1/3.4)
| | SEX = 2: ISCHEMIA (8.1/4.4)
| | SYMPTOM1 = 3:
| | | AGE <= 63 : ISCHEMIA (7.0/4.2)
| | | AGE > 63 : NO-ISCHEMIA (7.1/3.2)
| | SYMPTOM1 = 10:
| | | SEX = 2: NO-ISCHEMIA (135.5/55.8)
| | | SEX = 1:
| | | | TWAVES = 1: NO-ISCHEMIA (1.0/0.9)
| | | TWAVES = 2: ISCHEMIA (46.0/15.6)
| | | TWAVES = 4: ISCHEMIA (10.0/6.4)
| | | | TWAVES = 0:
| | | | AGE > 76 : NO-ISCHEMIA (12.7/4.7)
| | | | | AGE <= 76 :
| | | | | | SYSBP > 178 : ISCHEMIA (10.2/4.7)
```


## Tree for Ischemia: Results

Evaluation on training data (3453 items):
Before Pruning After Pruning
Size Errors Size Errors Estimate
462 494(14.3\%) 74 668(19.3\%) (24.5\%) $\ll$

Evaluation on test data (2320 items):
Before Pruning After Pruning
Size Errors Size Errors Estimate 462 502(21.6\%) 74 426(18.4\%) (24.5\%) $\ll$
(a) (b) <-classified as

490223 (a): class ISCHEMIA
2031404 (b): class NO-ISCHEMIA

## Issues

- Using related attributes in different parts of the tree
- Use a subset of variables in final tree

Overfitting: need more severe pruning

- Adjust confidence level
- Small leaves
- Set a large minimum leaf size
- Need relative balance of outcomes
- Enrich outcomes of training set


## Treatment of Variables

- Continuous => Ranges
- When fine distinctions are inappropriate
- Avoids overfitting
- Age: 20,30,40,50,60,70,80,90
- Categorical $=>$ Continuous
- When there is some order to the categories
- Natural subsetting
- Smoking: never $=>0$, quit $>5 \mathrm{yr}=>1$, quit $1-5 \mathrm{yr}=>2$, quit $<1 \mathrm{yr}$ (or unk) $=>3$, current $=>4$


## Treatment of Variables

-Specify a value for unknown

- Stroke: unknown => false

Combining variables

- "Or" across drugs by class or implications
- Picking variables on pragmatic grounds
- Start with many variables and narrow to ones most clinically relevant


## Variables Cont'd

- Missing values
- Force, if likely value different from average of knowns
- Derived values
- E.g., pulse pressure or product values
- Combine related variables


## Combinations of Variables



## Comparison with Logistic Regression

- Trees:
- Automatic selection
- Classification
- Assumes independence of subgroups
- Handles interactions automatically
- Handles missing values
- Linear relationships chopped into categories
- Handles outliers
- LR:
- Manual selection
- Probability
- Assumes same behavior over all cases
- Requires interaction variables
- Requires complete data
- Handles linear relationships
- Sensitive to outliers


## Multiple Trees

Weakness: Limited number of categories (leaf nodes) in optimal tree - there is only one way to categorize a case
Strategy: Generate several different trees and use them to vote on a classification
Advantage: Allows multiple ways of categorizing a case
Disadvantage: Makes it much harder to explain the classification of a case

## Generating Multiple Trees

- Use different subsets of the learning set
- Bagging: uniformly sampling $m$ cases with replacement for each tree
- Divide set into 10 parts and use each 9 to generate a tree
- Adapt the learning set
- Boosting: after generating each tree, increase the weight of cases misclassified by the tree


## Voting on a Classification

- Equal votes
- Votes in proportion to the size of the leaves
- Votes weighted by the $\alpha$ used to reweigh the cases (standard for boosting)


## Boosting

$\bullet \mathrm{C}_{1}$ constructed from training \& $\mathrm{e}_{1}=$ error rate
$-W(c)=w(c) /\left\{\begin{array}{l}2 \mathrm{e} \text { if case misclassified } \\ 2(1-\mathrm{e}) \text { otherwise }\end{array}\right.$

- Composite classifier obtained by voting
- Weight $\left(\mathrm{C}_{\mathrm{i}}\right)=\log \left(\left(1-\mathrm{e}_{\mathrm{i}}\right) / \mathrm{e}_{\mathrm{i}}\right)$


## Boosting

- Adaboost: Freund \& Schapire, 1997
- many classifiers: 25, 100, 1000
- Miniboost: Quinlan 1998
- 3 classifiers and take majority vote
- allows simplifications
- computationally efficient


## MiniBoosting

$\bullet$ Performance is improved

- Combined trees are possible but very complex
Even the leafless branches of combined trees contribute to the performance improvement


## Empirical Comparison

- Bauer \& Kohavi, Mach Learn 36:105, '99
- Bagging, AdaBoost, Arc (bag+reweigh)
- AdaBoost \& Arc better than Bagging on avg
- AdaBoost had problems with noisy datasets
- Reweighing can be unstable when error rates are small
- Not pruning decreased errors for bagging and increased them for AdaBoost


## Literature

Breiman et al., Classification and Regression Trees
Quinlan, C4.5 Programs for Machine Learning

- Resources: http://www.kdnuggets.com/

