## Decision Tree Learning

## Decision Trees

- Examples of systems that learn decision trees: c4.5, CLS, IDR, ASSISTANT, ID5, CART, ID3.
- Suitable problems:
- Instances are described by attribute-value couples
- The target function has discrete values
- Disjunctive descriptions of concepts may be required
- The training set may contain errors (noise)
- The training set may contain incomplete data


## c4.5

- c4.5 [Qui93b, Qui96]: evolution of ID, also by J. R. Quinlan
- Inspired to one of the first decision tree learning system, CLS (Concept Learning Systems) by E.B. Hunt
- Benchmark for many learning systems


## Example

- Instances: Saturday mornings
- Classes:
- Good day for playing tennis
- Bad day for playing tennis
- Attributes
- outlook, discrete, values=\{sunny,overcast,rain\}
- temperature, continuous
- humidity, continuous
- windy, discrete, values=\{true, false\}


## Training set

| No | Outlook | Temp $\left({ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D1 | sunny | 75 | 70 | T | P |
| D2 | sunny | 80 | 90 | T | N |
| D3 | sunny | 85 | 85 | F | N |
| D4 | sunny | 72 | 95 | F | N |
| D5 | sunny | 69 | 70 | F | P |
| D6 | overcast | 72 | 90 | T | P |
| D7 | overcast | 83 | 78 | F | P |
| D8 | overcast | 64 | 65 | T | P |
| D9 | overcast | 81 | 75 | F | P |
| D10 | rain | 71 | 80 | T | N |
| D11 | rain | 65 | 70 | T | N |
| D12 | rain | 75 | 80 | F | P |
| D13 | rain | 68 | 80 | F | P |
| D14 | rain | 70 | 96 | F | P |

## Decision Tree



## Decision Tree

## Outlook=sunny

\| Humidity $\leq 75: P$
| Humidity > 75: N
Outlook=overcast: P
Outlook=rain
Windy=True: $N$
|
Windy=False: P

## Notation

- Let T be the training set,
- Let $\left\{\mathrm{c}_{1}, \mathrm{c}_{2}, \ldots, \mathrm{c}_{\mathrm{k}}\right\}$ be the set of classes;


## Tree Buildling Algorithm

- build_tree(T) returns a tree:
- T contains examples from the same class
- Return a leaf with label the class
- T contains examples from more than one class
- T is partitioned into subsets $\mathrm{T}_{1}, \mathrm{~T}_{2}, \ldots, \mathrm{~T}_{\mathrm{n}}$ according to a test on an attribute
- Call the algorithm recursively on the subsets:
- child $=$ build_tree( $T_{i}$ ) for $\mathrm{i}=1, \ldots \mathrm{n}$
- Return a subtree with the root associated to the test and children child ${ }_{1}, \ldots$, child $_{n}$.


## Example of build_tree

| T= <br> No | Outlook | Temp ( ${ }^{\circ} \mathrm{F}$ ) | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D1 | sunny | 75 | 70 | T | P |
| D2 | sunny | 80 | 90 | T | N |
| D3 | sunny | 85 | 85 | F | N |
| D4 | sunny | 72 | 95 | F | N |
| D5 | sunny | 69 | 70 | F | P |
| D6 | overcast | 72 | 90 | T | P |
| D7 | overcast | 83 | 78 | F | P |
| D8 | overcast | 64 | 65 | T | P |
| D9 | overcast | 81 | 75 | F | P |
| D10 | rain | 71 | 80 | T | N |
| D11 | rain | 65 | 70 | T | N |
| D12 | rain | 75 | 80 | F | P |
| D13 | rain | 68 | 80 | F | P |
| D14 | rain | 70 | 96 | F | P |

## Test on Outlook

- $\mathrm{T}_{\text {sunny }}=$

| No | Outlook | Temp $\left({ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D1 | sunny | 75 | 70 | T | P |
| D2 | sunny | 80 | 90 | T | N |
| D3 | sunny | 85 | 85 | F | N |
| D4 | sunny | 72 | 95 | F | N |
| D5 | sunny | 69 | 70 | F | P |

- $\mathrm{T}_{\text {overcast }}=$

| No | Outlook | Temp $\left({ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D7 | overcast | 83 | 78 | F | P |
| D8 | overcast | 64 | 65 | T | P |
| D9 | overcast | 81 | 75 | F | P |

## Test on Outlook

- $\mathrm{T}_{\text {rain }}=$

| No | Outlook | Temp $\left({ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D10 | rain | 71 | 80 | T | N |
| D11 | rain | 65 | 70 | T | N |
| D12 | rain | 75 | 80 | F | P |
| D13 | rain | 68 | 80 | F | P |
| D14 | rain | 70 | 96 | F | P |

## build_tree( $T_{\text {sunny }}$ )

- Test: Humidity $\leq 75$
- $\mathrm{T}_{\text {sunny, }}$ Humidity $\leq 75$

| No | Outlook | Temp ( ${ }^{\circ} \mathrm{F}$ ) | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D1 | sunny | 75 | 70 | T | P |
| D5 | sunny | 69 | 70 | F | P |

- Leaf, P label
- $\mathrm{T}_{\text {sunny, }}$ Humidity $>75$

| No | Outlook | Temp ( $\left.{ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D2 | sunny | 80 | 90 | T | N |
| D3 | sunny | 85 | 85 | F | N |
| D4 | sunny | 72 | 95 | F | N |

- Leaf, N label


## build_tree(T $T_{\text {overcast }}$ )

- $\mathrm{T}_{\text {overcast }}=$

| No | Outlook | Temp ( ${ }^{\circ}$ F) | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D7 | overcast | 83 | 78 | F | P |
| D8 | overcast | 64 | 65 | T | P |
| D9 | overcast | 81 | 75 | F | P |

- Leaf, P label


## builld_tree( $T_{\text {rain }}$ )

- Test: Windy=?
- $\mathrm{T}_{\text {rain,true }}=$

| No | Outlook | Temp $\left({ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D10 | rain | 71 | 80 | T | N |
| D11 | rain | 65 | 70 | T | N |

- Leaf, N label
- $\mathrm{T}_{\text {rain,false }}=$

| No | Outlook | Temp $\left({ }^{\circ} \mathrm{F}\right)$ | Humid (\%) | Windy | Class |
| :--- | :--- | :--- | :--- | :--- | :--- |
| D12 | rain | 75 | 80 | F | P |
| D13 | rain | 68 | 80 | F | P |
| D14 | rain | 70 | 96 | F | P |

- Leaf, P label


## Decision Tree



## Tests on Attributes

- Discrete attribute X with n possible values $\mathrm{x}_{1}, \ldots, \mathrm{x}_{\mathrm{n}}$ :
- Equality with a constant: $\mathrm{X}=$ cost, 2 possible outcomes: yes, no
- Equality test: $\mathrm{X}=$ ?, n possible outcomes
- Membership in a set: $X \in S, 2$ possible outcomes: yes, no
- Membership in a set of a partition of $\left\{x_{1}, \ldots, x_{n}\right\}$ : one outcome per set


## Tests on Attributes

- Example of membership in a set of a partition:
- Attribute Outlook, partition of the set of values \{\{sunny\},\{rain,overcast\}\}

- Continuous attribute $X$
- Comparison with a threshold $\mathrm{X} \leq$ cost, 2 possible outcomes: yes, no


## Termination Condition

- c4.5 stops
- When an uniform set is found
- When an empty set is found
- A leaf is returned with label the most frequent class in the father
- When no test is such that at least two subsets contain a minimum number of cases.
- The minimum number of cases is a userdefined parameter (e.g. 2)


## Building the Tree

- Search in space of all possible trees
- Once a test is assigned to a node it is possible to backtrack
- Infeasible
- Greedy search
- Tests on nodes chosen irrevocably: once a test is assigned to a node it is not possible to backtrack
- Choice on the basis of a heuristic
- Most used heuristics
- Entropy
- Gini index


## Choice of the Test

- Choice of the attribute
- Discrete attributes:
- Choice of the type of test
- Possibly choice of the constant or partition
- Continuous attributes
- Choice of the threshold
- Usually only the equality test $\mathrm{X}=$ ? is used for discrete attributes
- Only the attribute must be chosen
- Constraints that the test must satisfy: at least two among $\mathrm{T}_{1}, \mathrm{~T}_{2}, \ldots, \mathrm{~T}_{\mathrm{n}}$ must contain a minimum number of examples


## Entropy

- In information theory, entropy is a measure of the uncertainty associated with a discrete random variable.
- Random variable C with k possible values $\mathrm{c}_{1}, \ldots, \mathrm{c}_{\mathrm{k}}$, entropy $\mathrm{H}(\mathrm{C})$ is given by

$$
H(C)=-\sum_{j=1}^{k} P\left(c_{j}\right) \log _{2} P\left(c_{j}\right)=E\left[-\log _{2} P(C)\right]
$$

- Also known as Shannon entropy


## Information Theory Interpretation

- Suppose you want to transmit a series of messages made of the values of N independent and identically distributed (i.i.d.) random variables
- What is the minimum number of bits necessary to transmit these messages?
- Source coding theorem (Shannon 1948)
- "The minimum number of bits necessary to encode the values of N i.i.d. random variables with negligible risk of information loss is $\mathrm{N} \times \mathrm{H}(\mathrm{X})$ as N tends to infinity, where $\mathrm{H}(\mathrm{X})$ is the entropy of the random variables"
$-H(X)$ is the minimum number of bits to encode the value of a random variable $X$


## Application to Decision Trees

- Random variable C: class of an instance randomly selected from a set T
- In this case

$$
P\left(c_{j}\right)=\frac{\left|T_{j}\right|}{|T|}
$$

- where $T_{j}$ is the set of examples from $T$ that belong to class $\mathrm{c}_{\mathrm{j}}$
- We can define the entropy of $T$ as the entropy of the random variable C

$$
H(T)=H(C)
$$

## Entropy

- $H(T)$ measures the minimum number of bits necessary for encoding, without loss, a message of the form
- "The present example, randomly selected from the training set, belongs to class $\mathrm{c}_{j}$ "
- $H(T)$ is also called info(T)


## Entropy for Two Classes

- In the case of two classes, + and -, with probabilities $p_{+}=P(+)$ and $p_{-}=P(-)$

$$
H(T)=-p_{+} \times \log _{2} p_{+}-p_{-} \times \log _{2} p_{-}
$$

- Only one variable is independent: $p=1-p_{+}$

$$
H(T)=-p_{+} \times \log _{2} p_{+}-\left(1-p_{+}\right) \times \log _{2}\left(1-p_{+}\right)
$$

## Entropy for Two Classes

- For $\mathrm{p}_{+}=0.5$ : $\mathrm{p}_{-}=0.5$

$$
\begin{array}{r}
H(T)=-0.5 \times \log _{2} 0.5-0.5 \times \log _{2} 0.5= \\
-0.5 \times(-1)-0.5 \times(-1)=1
\end{array}
$$

- For $p_{+}=0$ : $p_{=}=1.0$

$$
H(T)=-0 \times \log _{2} 0-1 \times \log _{2} 1=(-0 \times-\infty)-0
$$

- In this case, we define

$$
p_{+} \times\left.\log _{2} p_{+}\right|_{p_{+}=0} \stackrel{\text { def }}{=} \lim _{p_{+} \rightarrow 0} p_{+} \times \log _{2} p_{+}=0
$$

- So $H(T)=0-0=0$
- Similarly, if $p_{+}=1 H(T)=0$


## Entropy for Two Classes



## Entropy

- Entropy measures also the non-uniformity or impurity of a set:
- It is minimal when the set is most pure, i.e., when it contains examples from only one class
- It is maximal when the set is most impure, i.e., when it contains an equal number of examples of the two classes


## Intuition for Two Classes

- If all the examples belongs to the same class, I do not need to communicate the class of a randomly drawn example
- If the examples are uniformly distributed over classes in the training set, I need to use 1 bit on average: message 0 encodes one class, message 1 the other


## Example

- Play tennis problem
- Training set T: 14 examples, 9 positive, 5 negative
- Entropy of T
- Remember: $\log _{a} x=\log _{10} x / \log _{10} a=\ln x / \ln a$
- $\log _{10} 2=0.301$

$$
\begin{aligned}
& H(T)=-(9 / 14) \log _{2}(9 / 14)-(5 / 14) \log _{2}(5 / 14)= \\
& =-(9 / 14) \log _{10}(9 / 14) / 0.301-(5 / 14) \log _{10}(5 / 14) / 0.301= \\
& =-0.643 *(-0.192) / 0.301-0.357 *(-0.447) / 0.301= \\
& =0.410+0.530=0.940
\end{aligned}
$$

## Entropy for Three Classes

- Three classes $\left\{\mathrm{c}_{1}, \mathrm{c}_{2}, \mathrm{c}_{3}\right\}$ :
$-p_{1}=P\left(c_{1}\right)$
$-p_{2}=P\left(c_{2}\right)$
$-p_{3}=P\left(c_{3}\right)$
- Only two independent variables: $p_{3}=1-p_{1}-p_{2}$

$$
\begin{aligned}
& H(T)=-p_{1} \times \log _{2} p_{1}-p_{2} \times \log _{2} p_{2}- \\
& -\left(1-p_{1}-p_{2}\right) \times \log _{2}\left(1-p_{1}-p_{2}\right)
\end{aligned}
$$

## Entropy for Three Classes



## Maximum of the Entropy

- The maximum of the entropy is obtained for a uniform distribution of the examples over classes
- For three classes: $p_{1}=p_{2}=p_{2}=1 / 3=0,333$
- $\mathrm{H}(\mathrm{T})=-3^{*} 1 / 3^{*} \log _{2} 1 / 3=\log _{2} 3=1,585$
- For $k$ classes, we get the maximum for $P\left(c_{j}\right)=1 / k$

$$
H_{\max }(T)=-\sum_{j=1}^{k} \frac{1}{k} \times \log _{2}\left(\frac{1}{k}\right)=k \times \frac{1}{k} \times \log _{2} k=\log _{2} k
$$

## c4.5 Heuristic

- Information gain of a test: the value given by the decrease of entropy due to the partitioning of the set of examples according to the test
- Test $X$ with $n$ possible outcomes
- Set of examples $T$ is partitioned into $n$ subsets $T_{1}$,
$\ldots, \mathrm{T}_{\mathrm{n}}$
- Entropy after the partitioning: weighted average entropy in the subsets

$$
H_{X}(T)=\sum_{i=1}^{n} \frac{\left|T_{i}\right|}{|T|} \times H\left(T_{i}\right)
$$

## Information Gain

$$
\operatorname{gain}(X)=H(T)-H_{X}(T)
$$

- It is equivalent to the number of bits necessary for encoding the class of examples that we save when we also know the result of the test X on the examples
- Entropy decreases as more and more uniform subsets are obtained


## Attribute Choice

- The attribute with the highest information gain is selected for splitting the current set of examples


## Example

- Test on the attribute Windy:
- Windy=F => $T_{F}$ contains 6 positive and 2 negative examples
- Windy=T => $T_{T}$ contains 3 positive and 3 negative examples
- $\mathrm{T}=[9,5]$
- $\mathrm{T}_{\mathrm{F}}=[6,2]$
- $\mathrm{T}_{\mathrm{T}}=[3,3]$


## Example

| Windy | Play | Don't Play | Total |
| :--- | :--- | :--- | :--- |
| F | 6 | 2 | 8 |
| T | 3 | 3 | 6 |
| Total | 9 | 5 | 14 |

- $H\left(T_{F}\right)=-6 / 8^{*} \log _{2}(6 / 8)-2 / 8^{*} \log _{2}(2 / 8)=0.811$
- $H\left(T_{T}\right)=-3 / 6 * \log _{2}(3 / 6)-3 / 6 * \log _{2}(3 / 6)=1$ gain $($ Windy $)=H(T)-\left(8 / 14^{*} H\left(T_{F}\right)+6 / 14 * H\left(T_{T}\right)\right)=$
$0.940-(8 / 14)^{*} 0.811-(6 / 14)^{*} 1=0.048$


## Example

- Test on the attribute outlook:
- Outlook=sunny => $\mathrm{T}_{\text {sunny }}$ contains 2 pos and 3 neg
- Outlook=overcast => $T_{\text {overcast }}$ contains 4 pos and 0 neg
- Outlook=rain => $\mathrm{T}_{\text {rain }}$ contains 3 pos and 2 neg
- $\mathrm{T}=[9,5], \mathrm{T}_{\text {rain }}=[3,2], \mathrm{T}_{\text {sunny }}=[2,3], \mathrm{T}_{\text {overcast }}=[4,0]$


## Example

| Outlook | Play | Don't Play | Total |
| :--- | :--- | :--- | :--- |
| Sunny | 2 | 3 | 5 |
| Overcast | 4 | 0 | 4 |
| Rain | 3 | 2 | 5 |
| Total | 9 | 5 | 14 |

$H\left(\mathrm{~T}_{\text {sunny }}\right)=-2 / 5^{*} \log _{2}(2 / 5)-3 / 5^{*} \log _{2}(3 / 5)=0.971$
$H\left(T_{\text {overcast }}\right)=-4 / 4^{*} \log _{2}(4 / 4)-0 / 4^{*} \log _{2}(0 / 4)=0$
$H\left(T_{\text {rain }}\right)=-3 / 5^{*} \log _{2}(3 / 5)-2 / 5^{*} \log _{2}(2 / 5)=0.971$
gain(Outlook) $=\mathrm{H}(\mathrm{T})-\left((5 / 14)^{*} \mathrm{H}\left(\mathrm{T}_{\text {sunny }}\right)+(4 / 14)^{*} \mathrm{H}\left(\mathrm{T}_{\text {overcast }}\right)\right.$
$\left.+(5 / 14)^{*} \mathrm{H}\left(\mathrm{T}_{\text {rain }}\right)\right)=$
$0.940-0.357^{*} 0.971-0.286 * 0-0.357^{*} 0.971=0.246$

## Problems of Information Gain

- High tendency to favor tests with many results
- Example: test on an attribute that is a key: $\left|\mathrm{T}_{\mathrm{i}}\right|=1, \mathrm{n}=$ |T|

$$
H\left(T_{i}\right)=0 \forall i \in[1,2, \ldots, n] \Rightarrow H_{X}(T)=0
$$

- So the gain is maximum:

$$
\operatorname{gain}(X)=H(T)
$$

- Example: diagnosis problem, examples=patients, attribute="Patient's name"
- Useless attribute but has the highest gain


## Normallized Gain

- The gain is normalized with respect to the entropy of the test itself
- The random variable in this case is the value of the attribute X

$$
H(X)=-\sum_{i=1}^{n} \frac{\left|T_{i}\right|}{|T|} \times \log _{2}\left(\frac{\left|T_{i}\right|}{|T|}\right)
$$

- Gain ratio

$$
\operatorname{gainratio}(X)=\frac{\operatorname{gain}(X)}{H(X)}
$$

## Gain Ratio in the Diagnosis example

- X=patient name

$$
\begin{gathered}
\operatorname{gain}(X)=H(T) \leq \log _{2} k \\
H(X)=-\sum_{i=1}^{n} \frac{1}{n} \log _{2} \frac{1}{n}=n \frac{1}{n} \log _{2} n=\log _{2} n \\
\operatorname{gainratio}(X) \leq \frac{\log _{2} k}{\log _{2} n}
\end{gathered}
$$

- If $n \gg k \Rightarrow$ gain ratio( $X$ ) is small
- $H(X)$ is also called splitinfo


## Example

gainratio(Windy)=gain(Windy)/H(Windy)
$\mathrm{H}($ Windy $)=-8 / 14^{*} \log _{2}(8 / 14)-6 / 14^{*} \log _{2}(6 / 14)=0.985$
gainratio(Windy) $=0.048 / 0.985=0.048$
gainratio(Outlook)=gain(Outlook)/H(Outlook)
H(Outlook) $=-5 / 14^{*} \log _{2}(5 / 14)-5 / 14^{*} \log _{2}(5 / 14)$
$-4 / 14 * \log _{2}(4 / 14)=1.577$
gainratio(Outlook) $=0.246 / 1.577=0.156$

## Gini Index

- Other impurity measure

$$
\operatorname{gini}(T)=1-\sum_{j=1}^{k} P\left(C_{j}\right)^{2}
$$

- For two classes

$$
\operatorname{gini}(T)=1-p_{+}^{2}-\left(1-p_{+}\right)^{2}=2 p_{+}-2 p_{+}^{2}
$$

## Gini Index

- The max of Gini index is achieved for $p_{+}=p_{-}=0,5$ : $\operatorname{gini}(T)=0,5$
- The min is achieved for $p_{+}=0$ or $p_{-}=0$ : $\operatorname{gini}(T)=0$
- For three classes

$$
\operatorname{gini}(T)=1-p_{1}^{2}-p_{2}^{2}-\left(1-p_{1}-p_{2}\right)^{2}
$$

## Gini index



## Gini Index for Three Classes



## Ginỉ Index

- In general, the Gini index has a maximum for $\mathrm{p}_{1}=\mathrm{p}_{2}=\ldots \mathrm{p}_{\mathrm{k}}=1 / \mathrm{k}$

$$
\operatorname{gini}(T)=1-\sum_{i=1}^{k}\left(\frac{1}{k}\right)^{2}=1-k \times \frac{1}{k^{2}}=1-\frac{1}{k}
$$

- and a minimum for $p_{1}=1, p_{2}=\ldots p_{k}=0$ : gini $(T)=0$


## Gini index

- Gini index after the split on attribute $X$

$$
\operatorname{gini}_{X}(T)=\sum_{j=1}^{n} \frac{\left|T_{i}\right|}{|T|} \times \operatorname{gini}\left(T_{i}\right)
$$

- The attribute $X$ that gives the lowest $\operatorname{gini}_{x}(T)$ is chosen for the split
- Gini index is used by CART


## Tests on Continuous Attributes

- Suppose continuous attribute $X$ assumes $m$ values in T
- Let $\left\langle\mathrm{V}_{1}, \mathrm{~V}_{2}, \ldots, \mathrm{~V}_{\mathrm{m}}\right\rangle$ be the values ordered from the smallest to the largest
- The test $\mathrm{X} \leq \mathrm{V}_{\mathrm{i}}$ divides the m values into two groups:
- $\mathrm{X} \leq \mathrm{V}_{\mathrm{i}}:\left\{\mathrm{V}_{1}, \ldots, \mathrm{~V}_{\mathrm{i}}\right\}$
- $X>V_{i}:\left\{V_{i+1}, \ldots, V_{m}\right\}$
- So we have m-1 candidates for the threshold: $\mathrm{V}_{1}, \ldots$, $V_{m-1}$
- Each candidate must be evaluated


## Tests on Continuous Attributes

- Evaluation of a test:
- Sort the examples on the basis of the attribute X
- Count the examples
- Let $\mathrm{e}_{\mathrm{j}, \mathrm{i}}$ be the number of examples that have value $\mathrm{V}_{\mathrm{i}}$ for X and belong to class $\mathrm{c}_{\mathrm{j}}$ for $\mathrm{j}=1, \ldots, \mathrm{k}$, $\mathrm{i}=1, \ldots, \mathrm{~m}-1$
- Let $\mathrm{e}_{\mathrm{j}}$ be the number of examples that belong to class $\mathrm{c}_{\mathrm{j}}$ for $\mathrm{j}=1, \ldots, \mathrm{k}$
- Each test has two possible outcomes: yes or no


## Tests on Continuous Attributes

- $\mathrm{d}_{\mathrm{j}, \text { yes }}=$ examples of class j in branch $\mathrm{X} \leq \mathrm{V}_{\mathrm{i}}$
- $d_{j, \text { no }}=$ examples of class $j$ in branch $X>V_{i}$
- For $\mathrm{j}=1$ to $k \mathrm{~d}_{\mathrm{j}, \mathrm{yes}}=0 \mathrm{~d}_{\mathrm{j}, \mathrm{no}}=\mathrm{e}_{\mathrm{j}}$
- For $\mathrm{i}=1$ to $\mathrm{m}-1$
- Let the test be $\mathrm{X} \leq \mathrm{V}_{\mathrm{i}}$
- For $\mathrm{j}=1$ to k
- $\mathrm{d}_{\mathrm{j}, \text { yes }}:=\mathrm{d}_{\mathrm{j}, \text { yes }}+\mathrm{e}_{\mathrm{j}, \mathrm{i}}$
- $\mathrm{d}_{\mathrm{j}, \mathrm{no}}:=\mathrm{d}_{\mathrm{j}, \mathrm{no}}-\mathrm{e}_{\mathrm{j}, \mathrm{i}}$
- The heuristic for attribute $X$, threshold $V_{i}$ can be computed from the values $\mathrm{d}_{1, \text { yes }}, \ldots, \mathrm{d}_{\mathrm{k}, \text { yes }}, \mathrm{d}_{1, \text { no }}, \ldots, \mathrm{d}_{\mathrm{k}, \text { no }}$


## Attributes with Unknown Value

- If the training set contains examples with one or more attributes unspecified
- For example
<?,72,90,T,P>
- How do we take into account this example when computing the heuristic and when splitting the example set?


## Attributes with Unknown Value

- Test evaluation:
- Consider a discrete attribute $X$ with $n$ values $\left\{x_{1}, \ldots, x_{n}\right\}$
- Let $F$ be the set of examples of $T$ that have $X$ known
- $F$ provides us with a distribution of examples into values and class $\mathrm{P}\left(\mathrm{x}_{\mathrm{i}}, \mathrm{c}_{\mathrm{j}}\right)$
- We assume that the unknown values have the same distribution $\mathrm{P}\left(\mathrm{x}_{\mathrm{i}}, \mathrm{c}_{\mathrm{j}}\right)$ with respect to the attribute and the class
- Therefore the examples with unknown values do not alter the value of the entropy or of the Gini index


## Attributes with Unknown Value

- Entropy of the partitioning: the unknown value is considered as a value of its own
- Let $T_{n+1}$ be the set of examples of $T$ with unknown value for $X$

$$
H(X)=-\sum_{i=1}^{n+1} \frac{\left|T_{i}\right|}{|T|} \times \log _{2}\left(\frac{\left|T_{i}\right|}{|T|}\right)
$$

## Attributes with Unknown Value

- c4.5 further penalizes attributes with unknown values by multiplying the gain by the probability that the attribute is known
- If $F$ is the set of examples of $T$ with $X$ known

$$
\operatorname{gain}(X)=\frac{|F|}{|T|} \times\left(H(F)-H_{X}(F)\right)
$$

## Example

- Suppose that in the example database the case D6 : Outlook=overcast, Temperature=72, Humidity=90, Windy=T
- is replaced by

Outlook=?, Temperature=72, Humidity=90, Windy=T

- The frequencies of Outlook over F are

| Outlook | Play | Don't Play | Total |
| :--- | :--- | :--- | :--- |
| Sunny | 2 | 3 | 5 |
| Overcast | 3 | 0 | 3 |
| Rain | 3 | 2 | 5 |
| Total | 8 | 5 | 13 |

## Example

| Outlook | Play | Don't Play | Total |
| :--- | :--- | :--- | :--- |
| Sunny | 2 | 3 | 5 |
| Overcast | 3 | 0 | 3 |
| Rain | 3 | 2 | 5 |
| Total | 8 | 5 | 13 |

- $\mathrm{H}(\mathrm{F})=-8 / 13^{*} \log _{2}(8 / 13)-5 / 13^{*} \log _{2}(5 / 13)=0.961$
- $\mathrm{H}_{\text {outlook }}(\mathrm{F})=5 / 13^{*}\left(-2 / 5^{*} \log _{2}(2 / 5)-3 / 5^{*} \log _{2}(3 / 5)\right)$
$+3 / 13^{*}\left(-3 / 3^{*} \log _{2}(3 / 3)-0 / 3^{*} \log _{2}(0 / 3)\right)$
$+5 / 13^{*}\left(-3 / 5^{*} \log 2(3 / 5)-2 / 5^{*} \log 2(2 / 5)\right)$
$=0.747$
- gain(Outlook)=13/14*(0.961-0.747)=0.199 (slightly less than before)


## Example

- $\mathrm{H}($ Outlook $)=-5 / 14 * \log _{2}(5 / 14)-3 / 14 * \log _{2}(3 / 14)-$ $5 / 14 * \log _{2}(5 / 14)-1 / 14^{*} \log _{2}(1 / 14)=1.809$
- gainratio(Outlook)=0.199/1.809=0.110


## Partitioning

- In which subset of a node with test on $X$ do we put an example with X unknown?
- Partitioning is generalized in a probabilistic sense:
- Each example of $T$ is assigned a weight, initially 1
- Each example is inserted in every subset $T_{i}$ with a weight
- If we partition on discrete attribute $X$
- If example $e$, with weight $w$ in $T$, has $X=x_{i}$, $e$ is put in $T_{i}$ with weight $w$ and in $T_{j}$ with $j \neq i$ with weight 0
- If example $e$, with weight $w$ in $T$, has $X=$ ?, $e$ is put in $T_{i}$ with weight $w \times P\left(x_{i}\right)$
$-P\left(x_{i}\right)$ can be estimated by relative frequency in $F$


## Example

- Partitioning according to Outlook
- No problem for the 13 cases of $F$
- Example D6 is assigned to sets $\mathrm{T}_{\text {sunny, }}, \mathrm{T}_{\text {overcast }}$ and $\mathrm{T}_{\text {rain }}$ with weights $5 / 13,3 / 13$ e $5 / 13$ respectively


## Example

- Consider $\mathrm{T}_{\text {sunny }}$

| No Outlook | Temp | Humidity |  | Windy |  |
| :--- | :---: | :---: | :---: | :---: | :---: |
| Class | Weight |  |  |  |  |
| D1 sunny | 75 | 70 | T | P | 1 |
| D2 sunny | 80 | 90 | T | N | 1 |
| D3 sunny | 85 | 85 | F | N | 1 |
| D4 sunny | 72 | 95 | F | N | 1 |
| D5 sunny | 69 | 70 | F | P | 1 |
| D6 ? | 72 | 90 | T | P | $5 / 13$ |

- If $\mathrm{T}_{\text {sunny }}$ is partitioned on Humidity with threshold 75 , we get the following distribution:
- Humidity $\leq 752$ class P, 0 class $N$
- Humidity > 75 5/13 class P, 3 class N


## Example

- The second subset still contains examples from two classes but no test produces two subsets with at least two examples=> stop
Outlook=sunny


Humidity $\leq 75: \mathrm{P}$ (2.0)
Humidity > 75: N (3.4/0.4)
Outlook=overcast: P (3.2)
Outlook=rain
Windy=True: N (2.4/0.4)
Wndy=False: P (3.0)

## Output Interpretation

Numbers (A/B) associated to leaves

- A=total weight of examples associated to the leaf
- $B=t o t a l$ weight of examples associated to the leaf that are misclassified
- For example
N (3.4/0.4)
- Means that 3.4 cases belong to the leaf of which 0.4 do not belong to class N


## Classification of Unseen Cases with All Values Known

- An unseen (new) case e with all attributes known is classified by
- Traversing the tree from the root by following the branches that correspond to the values of the case
- Suppose that the leaf

> Pos (A/B)
is reached, then $e$ is classified as Pos with probability (A-B)/A and Neg with probability B/A

- In the case of more than two classes, the distribution of examples into classes at the leaf is used


## Classification of Unseen Cases with Unkown Values

- T: node with a test on X , e has X unknown
- $e$ is associated with a weight $w$, initially set to 1
- When traversing the tree, if the attribute at the current node is unknown in e we explore all subtrees.
- Subtree $T_{i}$ is explored by assigning e to $T_{i}$ with weight $w \times P\left(x_{j}\right)$
- $P\left(x_{j}\right)$ is estimated by considering by relative frequency over $T$
- The information about relative frequencies are stored in the leaves that are descendants of $T$


## Classification of Unseen Cases with Unkown Values

- In the end more than one leaf will be reached
- Let L be the set of leaves that are reached
- Let $w_{1}=P(\| e)$ be the weight of $e$ that reaches I
- Let $P\left(c_{j} \mid l\right)$ be the probability that an example in I belongs to class $\mathrm{c}_{\mathrm{j}}$
(estimated by relative frequency)

$$
\begin{aligned}
& P\left(c_{j} \mid e\right)=\sum_{l \in L} P\left(c_{j}, l \mid e\right)=\sum_{l \in L} P\left(c_{j} \mid l, e\right) \times P(l \mid e)= \\
& \sum_{l \in L} P\left(c_{j} \mid l\right) \times P(l \mid e)
\end{aligned}
$$

## Example

- Unseen example e

Outlook=sunny, Temperature=70, Humidity=?, Windy=F

- Outlook=sunny => first subtree
- Humidity=?=> we cannot determine whether Humidity $\leq 75$
- We follow both branches with weights
- Branch Humidity $\leq 75$ : w=2.0/5.4=0.370
- Branch Humidity > 75: w=3.4/5.4=0.630


## Example

- Leaf humidity $\leq 75$ :
$-P(P \mid I)=100 \%$
- $\mathrm{P}(\mathrm{N} \mid \mathrm{I})=0 \%$
- Leaf humidity $>75$
$-P(N \mid I)=3 / 3.4=88 \%$
$-P(P \|)=0.4 / 3.4=12 \%$
- Overall we get
- $\mathrm{P}(\mathrm{P} \mid \mathrm{e})=0.370 * 100 \%+0.630 * 12 \%=44 \%$
- $\mathrm{P}(\mathrm{N} \mid \mathrm{e})=0.630 * 88 \%=56 \%$


## Observations on c4.5 Search (1)

- c4.5 performs a hill-climbing search in the space of all the possible decision trees starting from the simplest hypothesis and going to more complex hypotheses
- Observations:
- The space of possible decision trees is equivalent to the powerset of $X$ so for c4.5 the hypothesis space surely includes the target concept
- c4.5 maintains a single hypothesis during the search. This contrasts with Candidate-Elimination that maintains the set of all the hypotheses coherent with the examples.


## Observations on c4.5 Search (2)

- c4.5 does not performs backtracking during the search. So it runs the risk of incurring in a solution only locally optimal
- c4.5 uses all the training examples at every step for deciding how to refine the tree, differently from the methods such as Find-S or CandidateElimination that take decisions in an incremental way on the basis of individual examples. The result is that the search of c4.5 is less sensible to errors in the single examples.


## c4.5 Inductive Bias (1)

- c4.5 inductive bias is [Mit97]:
- Shortest trees are preferred to the longer ones. The trees that put closer to the root the attributes with the highest information gain are preferred.
- Even if c 4.5 hypothesis space is the powerset of $X$, thanks to this inductive bias, c4.5 is able to generalize anyway
- Differences with Candidate-Elimination: CandidateElimination searches in a complete way an incomplete hypothesis space while c4.5 searches in an incomplete way a complete hypothesis space


## c4.5 Inductive Bias (2)

- c4.5: preference bias (or search bias)
- Candidate-Elimination: restriction bias (or language bias)
- In general, a preference bias is better because it allows to work with a complete hypothesis space
- Some algorithms use both


## Why preferring shorter hypotheses?

- Question largely debated by philosophers
- William of Occam in 1320 was the first to tackle the problem [Mit97]:
- Occam's razor: prefer the simplest hypothesis that describes the data
- Is this true?
- In favor: simplest hypothesis are fewer than more complex ones, so it is more difficult to find a short hypothesis that describes the data by chance.
- Against: if we define a small set of hypotheses (e. g. the trees with 17 nodes), then we should prefer those.


## Overfitting

- A hypothesis overfits the data when there exists a simpler hypothesis less accurate on the training set but more accurate on the universe
- Overfitting may happen when the training set contains errors or when it is small


## Overfitting



## Overfitting

- Example: consider the previous training set plus
<sunny,80,70,strong,N>
- which should be positive but it is erroneously classified as negative
- The tree h' learned from the new training set is more complex than the tree h learned from the original training set
- However, since the new example is incorrect, h' will make more mistakes than h on U


## Approaches to Avoid Overfitting

- Stop the growth before reaching a tree that perfectly classifies the training examples (prepruning)
- Grow the tree and prune it afterwards (postpruning)


## How to Determine the Optimal Dimension of the Tree

- Using a separate example set for evaluating the tree ("training and validation set" approach)
- Using all the available examples for training but using a statistical test to decide whether to keep a node or not
- Using an explicit measure of the complexity of encoding the tree and the examples that are exceptions and choose a tree that (locally) minimizes this measure (minimum description length principle)


## Reduced error pruning

- "Reduced error pruning" is a postpruning and "training and validation set" approach.
- Divide the data into a training and a validation set
- While a further pruning reduces the accuracy on the validation set
- Evaluate the impact on the validation set of pruning every node
- Pruning a node means replacing it with a leaf with label the most frequent class in the node
- Prune the node that gives the best accuracy on the validation set
- Problems when the data are limited


## Reduced error pruning



## Simplification of the Tree in c4.5

- "Pessimistic pruning": it is a postpruning that uses a statistical test rather than a testing set
- The tree that is obtained is simpler and is often more accurate when evaluating new cases because it reduces the overfitting
- Pruning by substitution: a single leaf in place of a tree or a sub-tree in place of a tree
- It is based on an estimate of the error rate


## Simplification of the Tree in c4.5

- Consider a leaf that covers N training cases of which E erroneously
- The error on the training set is $\mathrm{E} / \mathrm{N}$
- Consider E as events in N trials
- We ask what is the probability of the error event on the entire population of cases covered by the leaf
- The probability of the error has a binomial distribution


## Binomial Distribution



- The binomial distribution describes the probability that $r$ heads are obtained when throwing $n$ times a coin for which the probability of landing head is $p$


## Binomial Distribution

$$
P(r)=\frac{n!}{r!(n-r)!} p^{r}(1-p)^{n-r}=\binom{n}{r} p^{r}(1-p)^{n-r}
$$

$$
E[r]=n p
$$

$$
\operatorname{Var}[r]=n p(1-p)
$$

## Estimate of p

- Suppose you observed $r$ heads in $n$ throws, you want to estimate p
- In particular, you want

1. an interval $[a, b]$ in which you are sure that $p$ falls with a certain probability $\mathrm{N} \%$, i.e. $\mathrm{P}(\mathrm{a}<\mathrm{p}<\mathrm{b})=\mathrm{N} \%$

- [a,b] is called a (two-tailed) confidence interval,
- $\mathrm{N} \%$ is called the confidence level, or

1. a number a such that $P(a<p)=L \%$

- [a,1] is a one-tailed confidence interval
- L\% is called the probability of the upper tail, or

1. a number $b$ such that $P(p<b)=U \%$

- $[0, b]$ is a one-tailed confidence interval
- $\mathrm{U} \%$ is called the probability of the lower tail


## One-tailled Confidence Interval

- We consider one-tailed confidence of the form [0,b]



## Estimate of p

- Given r, n, and N\%, L\% or U\% there are algorithms that give you ( $a, b$ ), $a$ and $b$
- Either directly using the binomial distribution or
- By approximating the binomial distribution with a normal distribution
- This approximation is good when $\mathrm{n}>30$ or when $\mathrm{np}(1-\mathrm{p})>5$


## Confidence Interval

- Let's go back to learning
- Suppose you have E error over N examples
- E corresponds to r, N to n and the true error probability to $p$
- We want to estimate the true error probability
- We use one-tailed confidence intervals [0,b]
- Let CF be the probability of the confidence interval (U \%)
- Let $U_{C F}(E, N)$ be the upper bound of the confidence interval, i.e. let $U_{C F}(E, N)=b$


## c4.5 Tree Pruining

- c4.5 considers $\mathrm{N}^{*} \mathrm{U}_{\mathrm{CF}}(\mathrm{E}, \mathrm{N})$ instead of E as an estimate of the number of errors in a leaf
- c4.5 uses CF=75\% by default
- The estimate of the number of errors of a tree is the sum of the estimate of the number of errors in each leaf
- c4.5 replaces a tree with a leaf if the estimated number of errors in the leaf is smaller than those of the tree
- The same for replacing trees with sub-trees


## Example: vote Dataset

- Instances: members of the US congress
- Attributes: votes expressed by the member regarding 16 questions
- Classes: democratic, republican


## Example: vote Dataset

```
physician fee freeze = n:
    adoption of the budget resolution = y: democrat (151.0)
    adoption of the budget resolution = u: democrat (1.0)
    adoption of the budget resolution = n:
    education spending = n: democrat (6.0)
        education spending = y: democrat (9.0)
        education spending = u: republican (1.0)
physician fee freeze = y:
    synfuels corporation cutback = n: republican (97.0/3.0)
    synfuels corporation cutback = u: republican (4.0)
    synfuels corporation cutback = y:
    duty free exports = y: democrat (2.0)
    duty free exports = u: republican (1.0)
    duty free exports = n:
            education spending = n: democrat (5.0/2.0)
            education spending = y: republican (13.0/2.0)
            education spending =u: democrat (1.0)
physician fee freeze = u:
    water project cost sharing = n: democrat (0.0)
    water project cost sharing = y: democrat (4.0)
    water project cost sharing = u:
    mx missile = n: republican (0.0)
    mx missile = y: democrat (3.0/1.0)
    mx missile = u: republican (2.0)
```


## Example: vote Dataset

Simplified tree:
physician fee freeze $=\mathrm{n}$ : democrat (168.0/2.6)
physician fee freeze = y: republican (123.0/13.9) physician fee freeze $=\mathrm{u}$ :
mx missile $=\mathrm{n}$ : democrat (3.0/1.1)
mx missile $=\mathrm{y}$ : democrat (4.0/2.2)
mx missile $=\mathrm{u}$ : republican (2.0/1.0)

## Example: vote Dataset

- The sub-tree education spending $=\mathrm{n}$ : democrat (6.0) education spending $=\mathrm{y}$ : democrat (9.0) education spending $=u$ : republican (1.0)
- was replaced by the leaf democrat because
- First leaf $E=0 N=6 U_{75 \%}(0,6)=0.206$
- Second leaf $E=0 N=9 U_{75 \%}(0,9)=0.143$
- Third leaf $\mathrm{E}=0 \mathrm{~N}=1 \mathrm{U}_{75 \%}(0,1)=0.750$
- So the estimate of the number of errors for the subtree is

$$
6 * 0.206+9 * 0.143+1 * 0.750=3.273
$$

## Example: vote Dataset

- If the sub-tree
education spending $=\mathrm{n}$ : democrat (6.0)
education spending $=\mathrm{y}$ : democrat (9.0)
education spending $=u$ : republican (1.0)
- was replaced by the leaf democrat, it would cover 16 cases with 1 error so the estimate of the number of errors would be:

$$
16 * U_{75 \%}(1,16)=16 * 0.157=2.512
$$

- $\quad$ Since the estimate is lower for the leaf than for the sub-tree, this is replaced by the leaf


## Example: vote Dataset

- The sub-tree immediately above takes the form adoption of the budget resolution $=y$ : democrat (151.0) adoption of the budget resolution $=\mathrm{u}$ : democrat (1.0) adoption of the budget resolution $=n$ : democrat (16/1)
- The estimate of the number of errors for the sub-tree is

$$
151 * \mathrm{U}_{75 \%}(0,151)+1^{*} \mathrm{U}_{75 \%}(0,1)+2.512=4.642
$$

- If the sub-tree was replaced by the leaf democrat the estimate of the number of errors would be

$$
168 * U_{75 \%}(1,168)=2.610
$$

- Therefore it is convenient to perform the substitution


## Pruned Trees

- The numbers (N/P) that appear close to the leaves mean:
- $N$ number of training example covered by the leaf
- P number of predicted errors if N new cases would be classified by the tree ( $\mathrm{N}^{*} \mathrm{U}_{\mathrm{CF}}(\mathrm{E}, \mathrm{N})$ )
- Estimate of the error on new cases of the pruned trees=errors predicted in the leaves / number of cases in the training set
- In the vote dataset the sum of errors predicted in the leaves is 20.8 over a total of 300 cases in the training set => 6.9\% error


## Example: vote Dataset

Evaluation on training data (300 items):

$$
\text { Before Pruning } \quad \text { After Pruning }
$$

$$
\begin{array}{ccccc}
\text { Size } & \text { Errors } & \text { Size } & \text { Errors Estimate } \\
25 & 8(2.7 \%) & 7 & 13(4.3 \%) & (6.9 \%)
\end{array} \lll
$$

Evaluation on test data ( 135 items):

## Before Pruning After Pruning

Size Errors Size Errors Estimate
25 7(5.2\%) 7 4(3.0\%) (6.9\%) <<

## Example: vote Dataset

- The error on the training set is higher for the pruned tree than for the unpruned tree
- but, as expected, the error on the unseen cases is lower


## Computational Complexity

- Worst case analysis
- A: number of attributes
- N : number of examples
- Hypothesis: an attribute can be used only once

- The average case can be very different


## Examples of Usage

- Results of the labor negotiations in Canada in the years 1987--88.
- Concepts: acceptable or unacceptable contract.
- Contract: 16 attributes, not all applicable to all cases.
- Classes: acceptable (good) or not acceptable (bad).
- Aim: obtaining a system able to judge the acceptability of a new contract.


## Data Divided in Some File

- labor-neg.names: description of the problem in term of classes and attributes

good, bad.

| Attributes
| ----------
duration:
wage increase first year:
continuous attributes
wage increase second year:
wage increase third year:
cost of living adjustment:
working hours:
continuous
none, tcf, tc
continuous
values
pension:
standby pay:
none, ret_allw, empl_contr
shift differential:
education allowance:
statutory holidays:
vacation:
continuous
continuous
yes, no
continuous
below average, average, generous
longterm disability assistance: yes, no
contribution to dental plan: none, half, full
bereavement assistance:
yes, no
contribution to health plan: none, half, full


## Training set

- Labor-neg.data: attributes not applicable or unknown: "?"

```
1,5.0,?,?,?,40,?,?,2,?,11,average,?,?,yes,?,good
3,3.7,4.0,5.0,tc,?,?,?,?,yes,?,?,?,?,yes,?,good
3,4.5,4.5,5.0,?,40,?,?,?,?,12,average,?,half,yes,half,good
2,2.0,2.5,?,?,35,?,?,6,yes,12,average,?,?,?,?,good
3,6.9,4.8,2.3,?,40,?,?,3,?,12,below average,?,?,?,?,good
3,3.5,4.0,4.6,none,36,?,?,3,?,13,generous,?,?,yes,full,good
1,3.0,?,?, none, 36,?,?,10,no,11,generous,?,?,?,?,good
2,4.0,5.0,?,tcf,35,?,13,5,?,15,generous,?,?,?,?,good
2,2.5,3.0,?,?,40, none,?,?,?,11,below average,?,?,?,?,bad
3,3.5,4.0,4.6,tcf,27,?,?,?,?,?,?,?,?,?,?,good
2,4.5,4.0,?,?,40,?,?,4,?,10,generous,?,half,?,full,good
1,6.0,?,?,?,38,?,8,3,?,9,generous,?, ?,?,?,good
2,3.0,3.0,?,none,33,?,?,?,yes,12,generous,?,?,yes,full,good
3,2.0,2.5,?,?, 35,none,?,?,?,10, average,?,?,yes,full,bad
3,4.5,4.5,5.0,none,40,?,?,?,no,11,average,?,half,?,?,good
2,2.5,2.5,?,?,38,empl_contr, ?, ?,?,10, average, ?, ?, ?, ?, bad
1,2.0,?,?,tc,40,ret_allw, 4,0,no,11,generous,no, none, no, none,bad
3,2.0,2.5,2.0,?,37,empl_contr,?,?,?,10,average,?,?, yes, none,bad
2,4.5,4.0,?, none, 40, ?, ?, 4, ?,12, average, yes, full,yes,half,good
```

Labor-neg.test: set of new cases, to be used to evaluate
the learned tree (same formats of labor-neg.data)

## Example

- Command to run c4.5 on the data:
c4.5-f labor-neg -u
where the -f option is used to indicate the name of the set of files and -u means that the learned tree must be evaluated on the cases in labor-neg.test


## Output

- Two decision trees:
- complete tree
- pruned tree
- Evaluation of the trees


## Output: Complete Tree

```
C4.5 [release 8] decision tree generator
    Options:
        File stem <labor-neg>
        Trees evaluated on unseen cases
Read 40 cases (16 attributes) from labor-neg.data
Decision Tree:
wage increase first year <= 2.5 :
| statutory holidays <= 10 : bad (6.0)
| statutory holidays > 10 :
| | working hours <= 38 : good (2.3/1.0)
| | working hours > 38 : bad (3.0)
wage increase first year > 2.5 :
| statutory holidays > 10 : good (21.2)
| statutory holidays <= 10 :
| | wage increase first year <= 4 : bad (4.5/0.5)
| | wage increase first year > 4 : good (3.0)
```

Sat May 18 17:05:35 1996

## Output: Pruned Tree

```
Simplified Decision Tree:
wage increase first year <= 2.5 : bad (11.3/2.8)
wage increase first year > 2.5 :
| statutory holidays > 10 : good (21.2/1.3)
| statutory holidays <= 10 :
| | wage increase first year <= 4 : bad (4.5/1.7)
| | wage increase first year > 4 : good (3.0/1.1)
```

Tree saved

## Tree Evaluation

```
Evaluation on training data (40 items):
\begin{tabular}{|c|c|c|c|c|}
\hline Bef & ning & \multicolumn{3}{|r|}{After Pruning} \\
\hline Size & Errors & Size & Errors & \\
\hline
\end{tabular}
\(111(2.5 \%) \quad 712.5 \%) \quad(17.4 \%) \ll\)
```

Evaluation on test data (17 items):

| Befo | Pruning | After Pruning |  |  |
| :---: | :---: | :---: | :---: | :---: |
| Size | Errors | Size | Errors | Estimate |
| 11 | $4(23.5 \%)$ | 7 | $3(17.6 \%)$ | (17.4\%) |


| (a) | (b) |  |
| ---: | ---: | :--- |
| ---- | - classified as |  |
| 10 | 1 | (a) : class good |
| 2 | 4 | (b) : class bad |

- Weka is a machine learning suite
- It contains tools for:
- Learning decision trees
- Learning production rules
- Discovering association rules
- Clustering
- Open source, in Java
- Downloadable from
- http://www.cs.waikato.ac.nz/ml/weka/
- Connected to the book: [Wit05]


## Weka

- Contains an implementation of c4.5 Release 8: j4.8
- Uses a format called ARFF (Attribute Relation File Format) for the input files


## ARFF

- A single file
- Two sections:
- Header
- Data
- Header: attribute definition
- @relation <dataset name>
- @attribute <attribute name> \{<val1>, <val2>, ..., <valn>\}
- Or
- @attribute <attribute name> real
- The last attribute is usually used as the class


## File Iabor.arff: Header

```
% labor.arff

\section*{comment}
```

@relation 'labor-neg-data'

```
@relation 'labor-neg-data'
@attribute 'duration' real
@attribute 'duration' real
@attribute 'wage-increase-first-year' real
@attribute 'wage-increase-first-year' real
@attribute 'wage-increase-second-year' real
@attribute 'wage-increase-second-year' real
@attribute 'wage-increase-third-year' real
@attribute 'wage-increase-third-year' real
@attribute 'cost-of-living-adjustment' {'none','tcf','tc'}
@attribute 'cost-of-living-adjustment' {'none','tcf','tc'}
@attribute 'working-hours' real
@attribute 'working-hours' real
@attribute 'pension' {'none','ret_allw','empl_contr'}
@attribute 'pension' {'none','ret_allw','empl_contr'}
@attribute 'standby-pay' real
@attribute 'standby-pay' real
@attribute 'shift-differential' real
@attribute 'shift-differential' real
@attribute 'education-allowance' {'yes','no'}
@attribute 'education-allowance' {'yes','no'}
@attribute 'statutory-holidays' real
@attribute 'statutory-holidays' real
@attribute 'vacation' {'below_average','average','generous'}
@attribute 'vacation' {'below_average','average','generous'}
@attribute 'longterm-disability-assistance' {'yes','no'}
@attribute 'longterm-disability-assistance' {'yes','no'}
@attribute 'contribution-to-dental-plan' {'none','half','full'}
@attribute 'contribution-to-dental-plan' {'none','half','full'}
@attribute 'bereavement-assistance' {'yes','no'}
@attribute 'bereavement-assistance' {'yes','no'}
@attribute 'contribution-to-health-plan' {'none','half','full'}
@attribute 'contribution-to-health-plan' {'none','half','full'}
@attribute 'class' {'bad','good'}
```

```
@attribute 'class' {'bad','good'}
```

```

\section*{Data Section}
- @data tag followed by the description of the examples, one description per row (ending with a line feed) as in c4.5
- Every example is described by the list of values for every attribute separated by , (comma)
- Each value corresponds to the attribute in the same position in the header
- The declarations @relation, @attribute and @data are case insensitive.

\section*{Labor.arffi: Data}
```

@data
1,5,?,?,?,40,?,?,2,?,11,'average',?,?,'yes',?,'good'
3,3.7,4,5,'tc',?,?,?,?,'yes',?,?,?,?,'yes',?,'good'
2,2,2.5,?,?,35,?,?,6,'yes',12,'average',?,?,?,?,'good'
3,6.9,4.8,2.3,?,40,?,?,3,?,12,'below_average',?,?,?,?,'good'

```
...

The single quotes (') are necessary only if the values contain spaces

\section*{ARFF}
- Other formats for the attributes
- integer, string, date
- integer: the attribute can take only integer values
- string: the attribute can take arbitrary string values
- String attributes are treated as nominal attributes

\section*{ARFF}
- date: date attributes. Format:
@attribute <name> date [<date-format>]
- where <date-format> is an optional string that specifies how the values are to be interpreted and printed. The default format is the ISO-8601 format "yyyy-MM-dd'T'HH:mm:ss"
- Example
@RELATION Timestamps
@ATTRIBUTE timestamp DATE "yyyy-MM-dd HH:mm:ss"
@DATA
"2001-04-03 12:12:12"
"2001-05-03 12:59:55"

\section*{ARFF}
- date: date attributes. Format:
@attribute <name> date [<date-format>]
- where <date-format> is an optional string that specifies how the values are to be interpreted and printed. The default format is the ISO-8601 format "yyyy-MM-dd'T'HH:mm:ss"
- Example
@RELATION Timestamps
@ATTRIBUTE timestamp DATE "yyyy-MM-dd HH:mm:ss"
@DATA
"2001-04-03 12:12:12"
"2001-05-03 12:59:55"

\section*{Weka Execution}
- Double click on Weka (weka.jar)
- Click on Explorer
- Click on Open file... : select the file to be opened, for example labor.arff
- Click on Classify
- Choice of the algorithm: j48
- Setting the test options: for example, select Percentage Split to select a fraction of the examples to be used for testing
- Click on Start

\section*{Output}

J48 pruned tree
wage-increase-first-year <= 2.5: bad (15.27/2.27)
wage-increase-first-year > 2.5
| statutory-holidays <= 10: bad (10.77/4.77)
| statutory-holidays > 10: good (30.96/1.0)
Number of Leaves : 3

Size of the tree :
5

Time taken to build model: 0.08 seconds

\section*{Output}
=== Evaluation on test split ===
=== Summary ===

Correctly Classified Instances
Incorrectly Classified Instances
Kappa statistic
Mean absolute error
Root mean squared error
Relative absolute error
Root relative squared error
Total Number of Instances

17
3
0.6875
0.2806
0.3729
58.3712 \%
72.0645 \%

20

\section*{Output}
=== Detailed Accuracy By Class ===
\begin{tabular}{cllcccc} 
TP Rate & FP Rate & Precision & Recall & F-Measure & Class & \\
0.667 & 0 & 1 & & 0.667 & 0.8 & \\
1 & 0.333 & 0.786 & 1 & 0.88 & good &
\end{tabular}
=== Confusion Matrix ===
a b <-- classified as
6 3| a = bad
0 11| b = good

\section*{Output Explanation}
- Confusion matrix:
pos neg <-- classification
TP FN
P pos
FP TN
N neg
- TP=true positive, \(\mathrm{FN}=\) false negative, \(\mathrm{FP}=\) false positive, \(\mathrm{TN}=\) true negative, \(\mathrm{P}=\mathrm{TP}+\mathrm{FN}=\) positive, \(\mathrm{N}=\mathrm{FP}+\mathrm{TN}=\) negative
- Accuracy \(=(T P+T N) /(T P+T N+F N+F P)\)
- Error rate \(=(F P+F N) /(T P+T N+F N+F P)=1-A c c u r a t e z z a\)
- TP Rate(TPR)=Sensitivity=Recall=TP/(TP+FN)=TP/P
- FP Rate(FPR)=FP/(FP+TN)=FP/N
- FN Rate(FNR)=FP/(TP + FN) \(=F P / P=1-T P R\)
- TN Rate(TNR)=Specificity=TN/(FP+TN)=TN/N=1-FPR
- Precision=TP/(TP+FP)
- Couples: TPR-FPR, Sensitivity-Specificity, Recall-Precision
- F-meaure=2*Precision*Recall/(Precision+Recall)

\section*{F-measure}


\section*{Useful Links}
- C4.5
http://www.rulequest.com/Personal/
C source code
- Weka: open source suite of machine learning algorithms written in Java http://www.cs.waikato.ac.nz/ml/weka/

\section*{References}
[Mit97] T. M. Mitchell, Machine Learning, McGraw-Hill, 1997
[Qui93b] J. R. Quinlan, C4.5: Programs for machine learning, Morgan Kaufmann Publishers, San Mateo, California, 1993
[Qui96] J. R. Quinlan, Improved Use of Continuous Attributes in C4.5, Journal of Artificial Intelligence Research, 4, pag. 77--90, 1996. ftp://ftp.cs.cmu.edu/ project/jair/volume4/quinlan96a.ps
[Wit05] I.H. Witten, E. Frank, Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations, Second Edition, Morgan Kaufmann, 2005.```

