S J P N Trust's



HIRASUGAR INSTITUTE OF TECHNOLOGY, NIDASOSHI.

Inculcating Values, Promoting Prosperity

Approved by AICTE, Recognized by Govt. of Karnataka and Affiliated to VTU Belagavi.

Accredited at 'A' Grade by NAAC Programmes Accredited by NBA: CSE, ECE, EEE & ME

CSDecision Treesshi

Prof. Mahesh G Huddar

Dept. of Computer Science and Engineering

3.1 Give decision trees to represent the following boolean functions: (a) $A \land \neg B$ (b) $A \lor [B \land C]$ (c) A XOR B

 $(d) \left[A \wedge B\right] \vee \left[C \wedge D\right]$

(a) $A \land \neg B$ (b) $A \lor [B \land C]$ (c) A X O R B , Nidasoshi(d) $[A \land B] \lor [C \land D]$

- Every Variable in Boolean function such as A, B, C etc. has two possibilities that is True and False
- Every Boolean function is either True or False shi
- If the Boolean function is true we write YES (Y)
- If the Boolean function is False we write NO (N)

(c) A XOR B

(d) $[A \land B] \lor [C \land D]$



3.1 Give decision trees to represent the following boolean functions: (a) $A \land \neg B$ (b) $A \lor [B \land C]$ (c) A XOR B

 $(d) \left[A \wedge B\right] \vee \left[C \wedge D\right]$











Decision Trees

- Decision Trees is one of the most widely used Classification Algorithm
- Features
 - Method for approximating *discrete-valued* functions (including boolean)
 - Learned functions are represented as *decision trees* (or *if-then-else* rules)
 - Expressive hypotheses space, including disjunction
 - Robust to noisy data

Example

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

Decision Tree Representation (PlayTennis)



(Outlook=Sunny, Temp=Hot, Humidity=High, Wind=Strong) No

Decision trees expressivity

• Decision trees represent a disjunction of conjunctions on constraints on the value of attributes:



Decision tree representation (PlayTennis)

- Decision trees classify instances by sorting them down the tree from the root to some leaf node, which provides the classification of the instance.
- Each node in the tree specifies a test of some attribute of the instance, and each branch descending from that node corresponds to one of the possible values for this attribute.
- An instance is classified by starting at the root node of the tree, testing the attribute specified by this node, then moving down the tree branch corresponding to the value of the attribute in the given example.
- This process is then repeated for the subtree rooted at the new node.

Decision tree representation (PlayTennis)

- In general, decision trees represent a disjunction of conjunctions of constraints on the attribute values of instances.
- Each path from the tree root to a leaf corresponds to a conjunction of attribute tests, and the tree itself to a disjunction of these conjunctions.

 $(Outlook = Sunny \land Humidity = Normal)$

- \vee (Outlook = Overcast)
- \vee (Outlook = Rain \wedge Wind = Weak)

Example

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

Although a variety of decision tree learning methods have been developed with somewhat differing capabilities and requirements, decision tree learning is generally best suited to problems with the following characteristics:

1. Instances are represented by attribute-value pairs. Instances are described by a fixed set of attributes (e.g., *Temperature*) and their values (e.g., *Hot*). The easiest situation for decision tree learning is when each attribute takes on a small number of disjoint possible values (e.g., *Hot*, *Mild*, *Cold*). However, extensions to the basic algorithm allow handling real-valued attributes as well (e.g., representing *Temperature* numerically).

- 2. The target function has discrete output values. The decision tree is usually used for Boolean classification (e.g., yes or no) kind of example. Decision tree methods easily extend to learning functions with more than two possible output values. A more substantial extension allows learning target functions with real-valued outputs, though the application of decision trees in this setting is less common.
- **3. Disjunctive descriptions may be required.** Decision trees naturally represent disjunctive expressions.

- **4.** The training data may contain errors. Decision tree learning methods are robust to errors, both errors in classifications of the training examples and errors in the attribute values that describe these examples.
- 5. The training data may contain missing attribute values. Decision tree methods can be used even when some training examples have unknown values (e.g., if the *Humidity* of the day is known for only some of the training examples).

- Many practical problems have been found to fit these characteristics.
- Decision tree learning has therefore been applied to problems such as learning to classify medical patients by their disease, equipment malfunctions by their cause, and loan applicants by their likelihood of defaulting on payments.
- Such problems, in which the task is to classify examples into one of a discrete set of possible categories, are often referred to as *classification problems*.

THE BASIC DECISION TREE LEARNING ALGORITHM

- Most algorithms that have been developed for learning decision trees are variations on a core algorithm that employs a top-down, greedy search through the space of possible decision trees.
- This approach is exemplified by the ID3 algorithm (Quinlan 1986) and its successor C4.5 (Quinlan 1993), which form the primary focus of our discussion here.
- The basic algorithm for decision tree learning, corresponding approximately to the ID3 algorithm.
- Next, we consider a number of extensions to this basic algorithm, including extensions incorporated into C4.5 and other more recent algorithms for decision tree learning.

Example

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

Which Attribute Is the Best Classifier?

- The central choice in the ID3 algorithm is selecting which attribute to test at each node in the tree.
- We would like to select the attribute that is most useful for classifying examples.
- What is a good quantitative measure of the worth of an attribute? We will define a statistical property, called *information gain*, that measures how well a given attribute separates the training examples according to their target classification.
- ID3 uses this information gain measure to select among the candidate attributes at each step while growing the tree.

ENTROPY MEASURES HOMOGENEITY OF EXAMPLES

- *Entropy,* characterizes the (im)purity of an arbitrary collection of examples.
- Given a collection S, containing positive and negative examples of some target concept, the entropy of S relative to this boolean classification is $Entropy(S) \equiv -p_{\oplus} \log_2 p_{\oplus} - p_{\ominus} \log_2 p_{\ominus}$
- where p+, is the proportion of positive examples in S and p-, is the proportion of negative examples in S.
- In all calculations involving entropy we define 0 log 0 to be 0.

ENTROPY MEASURES HOMOGENEITY OF EXAMPLES

- Entropy measures the (*im*)*purity* of a collection of examples. It depends from the distribution of the random variable *p*.
 - *S* is a collection of training examples
 - *p*₊ the proportion of positive examples in *S*

$$Entropy(S) \equiv \sum_{i=1}^{n} -p_i \log_2 p_i$$

C

*p*_ the proportion of negative examples in *S*

Examples

Entropy $(S) \equiv -p_+ \log_2 p_+ - p_- \log_2 p_-$ [0 $\log_2 0 = 0$] Entropy $([14+, 0-]) = -14/14 \log_2 (14/14) - 0 \log_2 (0) = 0$ Entropy $([9+, 5-]) = -9/14 \log_2 (9/14) - 5/14 \log_2 (5/14) = 0,94$ Entropy $([7+, 7-]) = -7/14 \log_2 (7/14) - 7/14 \log_2 (7/14) = -1/2 + 1/2 = 1$

Example

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

Entropy



INFORMATION GAIN MEASURES THE EXPECTED REDUCTION IN ENTROPY

- Given entropy as a measure of the impurity in a collection of training examples, we can now define a measure of the effectiveness of an attribute in classifying the training data.
- Now, the *information gain*, is simply the expected reduction in entropy caused by partitioning the examples according to this attribute.
- More precisely, the information gain, *Gain(S, A)* of *an* attribute A, relative to a collection of examples S, is defined as,

$$Gain(S, A) \equiv Entropy(S) - \sum_{v \in Values(A)} \frac{|S_v|}{|S|} Entropy(S_v)$$

where Values(A) is the set of all possible values for attribute A, and S, is the subset of S for which attribute A has value v (i.e., S_v = {s ∈ S|A(s) = v})

• For example, suppose *S* is a collection of training-example days described by attributes including *Wind*, which can have the values *Weak* or *Strong*.

Values(Wind) = Weak, StrongS = [9+, 5-] $C_{SE}^{Sweak} \leftarrow [6+, 2-]$, Nidasoshi $Gain(S, Wind) = Entropy(S) - \sum_{v \in \{Weak, Strong\}} \frac{|S_v|}{|S|} Entropy(S_v)$ $= Entropy(S) - (8/14)Entropy(S_{Weak})$ $- (6/14) Entropy(S_{Strong})$ = 0.940 - (8/14)0.811 - (6/14)1.00= 0.048

• Information gain is precisely the measure used by ID3 to select the best attribute

at

- each step in growing the tree.
- The use of information gain to evaluate the relevance of attributes.
- Here the information gain of two different attributes, *Humidity* and *Wind*, is computed in order to determine which is the better attribute for classifying the training examples.



ID3(Examples, Target_attribute, Attributes)

Examples are the training examples. Target_attribute is the attribute whose value is to be predicted by the tree. Attributes is a list of other attributes that may be tested by the learned decision tree. Returns a decision tree that correctly classifies the given Examples.

- Create a *Root* node for the tree
- If all *Examples* are positive, Return the single-node tree *Root*, with label = +
- If all *Examples* are negative, Return the single-node tree *Root*, with label = -
- If Attributes is empty, Return the single-node tree Root, with label = most common value of Target_attribute in Examples
- Otherwise Begin
 - $A \leftarrow$ the attribute from Attributes that best* classifies Examples
 - The decision attribute for $Root \leftarrow A$
 - For each possible value, v_i , of A,
 - Add a new tree branch below *Root*, corresponding to the test $A = v_i$
 - Let $Examples_{v_i}$ be the subset of Examples that have value v_i for A
 - If $Examples_{v_i}$ is empty
 - Then below this new branch add a leaf node with label = most common value of *Target_attribute* in *Examples*
 - Else below this new branch add the subtree

 $ID3(Examples_{v_i}, Target_attribute, Attributes - \{A\}))$

- End
- Return Root
CSE, HIT, Nidasoshi

DECISION TREE – ID3 ALGORITHM NUMERICAL EXAMPLE

Day	Outlook	Тетр	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	SNSh
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

Day	Outlook	Тетр	Humidity	Wind	Play	Attribute: Outlook
					Iennis	Values (Outlook) = Sunny, Overcast, Rain
D1	Sunny	Hot	High	Weak	No	
D2	Sunny	Hot	High	Strong	No	$S = [9+, 5-] \qquad Entropy(S) = -\frac{9}{14} \log_2 \frac{9}{14} - \frac{5}{14} \log_2 \frac{5}{14} = 0.94$
D3	Overcast	Hot	High	Weak	Yes	$\begin{bmatrix} 2 & 2 & 3 \\ 2 & 2 & 3 \\ 2 & 3 & 2 \end{bmatrix} = \begin{bmatrix} 2 & 3 & 2 & 3 \\ 2 & 3 & 3 & 2 \\ 2 & 3 & 3 & 3 \\ 2 $
D4	Rain	Mild	High	Weak	Yes	$\begin{bmatrix} S_{Sunny} \leftarrow [2+, 3-] \\ \end{bmatrix} = -\frac{1}{5} \log_2 \frac{1}{5} - \frac{1}{5} \log_2 \frac{1}{5} = 0.9/1$
D5	Rain	Cool	Normal	Weak	Yes	$S_{Overcast} \leftarrow [4+, 0-] \qquad Entropy(S_{Overcast}) = -\frac{4}{2} \log_2 \frac{4}{2} - \frac{6}{2} \log_2 \frac{9}{2} = 0$
D6	Rain	Cool	Normal	Strong	No	
D7	Overcast	Cool	Normal	Strong	Yes	$S_{Rain} \leftarrow [3+,2-] \qquad Entropy(S_{Rain}) = -\frac{3}{5}log_2\frac{3}{5} - \frac{2}{5}log_2\frac{2}{5} = 0.971$
D8	Sunny	Mild	High	Weak	No	LIT Nideeebi
D9	Sunny	Cool	Normal	Weak	Yes	$Gain (S, Outlook) = Entropy(S) - \sum \frac{ S_v }{ S } Entropy(S_v)$
D10	Rain	Mild	Normal	Weak	Yes	$v \in \{Sunny, Overcast, Rain\}$
D11	Sunny	Mild	Normal	Strong	Yes	Gain(S, Outlook)
D12	Overcast	Mild	High	Strong	Yes	$-Entrony(S) = \frac{5}{2}Entrony(S_{2}) = \frac{4}{2}Entrony(S_{2})$
D13	Overcast	Hot	Normal	Weak	Yes	$\begin{bmatrix} -Lill(0)fy(3) & 14 \\ 14 \end{bmatrix} = \begin{bmatrix} 14 \\ 14 \end{bmatrix} = \begin{bmatrix} 14 \\ 14 \end{bmatrix}$
D14	Rain	Mild	High	Strong	No	$-\frac{5}{2}Entrony(S_{pair})$
D13	Rain	Hot Mild	High	Strong	res No	$-\frac{5}{14}Entropy(S_{Rain})$

$$Gain(S, Outlook) = 0.94 - \frac{5}{14}0.971 - \frac{4}{14}0 - \frac{5}{14}0.971 = 0.2464$$

Day	Outlook	Тетр	Humidity	Wind	Play Tennis	Attribute: Temp
D1	Sunny	Hot	High	Weak	No	Values (Temp) = Hot, Mild, Cool
D2	Sunny	Hot	High	Strong	No	$S = [9 + 5 -]$ Entrony(S) = $-\frac{9}{2} \log_2 \frac{9}{2} - \frac{5}{2} \log_2 \frac{5}{2} - 0.94$
D3	Overcast	Hot	High	Weak	Yes	$\begin{bmatrix} 3 & - [77, 3] \end{bmatrix} = \begin{bmatrix} 21110py(3) - & 1410y_{2} & 1410y_{2} & 14 \end{bmatrix} = \begin{bmatrix} 3 & - & 0 \end{bmatrix} = \begin{bmatrix} 2110py(3) - & 1410y_{2} & $
D4	Rain	Mild	High	Weak	Yes	$S_{Hot} \leftarrow [2+, 2-] \qquad Entropy(S_{Hot}) = -\frac{2}{2} \log_2 \frac{2}{2} - \frac{2}{2} \log_2 \frac{2}{2} = 1.0$
D5	Rain	Cool	Normal	Weak	Yes	
D6	Rain	Cool	Normal	Strong	No	$S_{Mild} \leftarrow [4+, 2-] \qquad Entropy(S_{Mild}) = -\frac{4}{6} \log_2 \frac{4}{6} - \frac{2}{6} \log_2 \frac{2}{6} = 0.9183$
D7	Overcast	Cool	Normal	Strong	Yes	
D8	Sunny	Mild	High	Weak	No	$S_{Cool} \leftarrow [3+, 1-] \qquad Entropy(S_{Cool}) = -\frac{3}{4}log_2\frac{3}{4} - \frac{1}{4}log_2\frac{1}{4} = 0.8113$
D9	Sunny	Cool	Normal	Weak	Yes	HII. NIdasoshi
D10	Rain	Mild	Normal	Weak	Yes	$Gain (S Temn) = Entrony(S) - \sum \frac{ S_v }{ S_v } Entrony(S_v)$
D11	Sunny	Mild	Normal	Strong	Yes	$\int u(s, r(t), p) = u(s, r(t),$
D12	Overcast	Mild	High	Strong	Yes	Gain(S,Temp)
D13	Overcast	Hot	Normal	Weak	Yes	4 6
D14	Rain	Mild	High	Strong	No	$= Entropy(S) - \frac{1}{14}Entropy(S_{Hot}) - \frac{3}{14}Entropy(S_{Mild})$
						$-\frac{4}{14}Entropy(S_{Cool})$

$$Gain(S, Temp) = 0.94 - \frac{4}{14}1.0 - \frac{6}{14}0.9183 - \frac{4}{14}0.8113 = 0.0289$$

Day	Outlook	Tomn	Lumidity	Wind	Play	Attrib
Day	Outlook	Temp	пиппиту	vvina	Tennis	
D1	Sunny	Hot	High	Weak	No	Values
D2	Sunny	Hot	High	Strong	No	s – [9
D3	Overcast	Hot	High	Weak	Yes	J J L J
D4	Rain	Mild	High	Weak	Yes	Suich +
D5	Rain	Cool	Normal	Weak	Yes	- nıyn
D6	Rain	Cool	Normal	Strong	No	S _{Normal}
D7	Overcast	Cool	Normal	Strong	Yes	
D8	Sunny	Mild	High	Weak	No	
D9	Sunny	Cool	Normal	Weak	Yes	Ga Ga
D10	Rain	Mild	Normal	Weak	Yes	Gain
D11	Sunny	Mild	Normal	Strong	Yes	
D12	Overcast	Mild	High	Strong	Yes	
D13	Overcast	Hot	Normal	Weak	Yes	
D14	Rain	Mild	High	Strong	No	

Attribute: Humidity

Values (Humidity) = High, Normal

$$S = [9+, 5-] \qquad Entropy(S) = -\frac{9}{14}log_{2}\frac{9}{14} - \frac{5}{14}log_{2}\frac{5}{14} = 0.94$$

$$S_{High} \leftarrow [3+, 4-] \qquad Entropy(S_{High}) = -\frac{3}{7}log_{2}\frac{3}{7} - \frac{4}{7}log_{2}\frac{4}{7} = 0.9852$$

$$S_{Normal} \leftarrow [6+, 1-] \qquad Entropy(S_{Normal}) = -\frac{6}{7}log_{2}\frac{6}{7} - \frac{1}{7}log_{2}\frac{1}{7} = 0.5916$$

$$Gain(S, Humidity) = Entropy(S) - \sum_{v \in \{High, Normal\}} \frac{|S_{v}|}{|S|} Entropy(S_{v})$$

$$Gain(S, Humidity) = Content + Con$$

Gain(S, Humidity)

$$= Entropy(S) - \frac{7}{14}Entropy(S_{High}) - \frac{7}{14}Entropy(S_{Normal})$$

$$Gain(S, Humidity) = 0.94 - \frac{7}{14}0.9852 - \frac{7}{14}0.5916 = 0.1516$$

Dav	Qutlook	Tomp	Lumidity	Wind	Play	Attri
Day	Outlook	Temp	питацу	wind	Tennis	
D1	Sunny	Hot	High	Weak	No	Value
D2	Sunny	Hot	High	Strong	No	s – T
D3	Overcast	Hot	High	Weak	Yes	
D4	Rain	Mild	High	Weak	Yes	Sstron
D5	Rain	Cool	Normal	Weak	Yes	
D6	Rain	Cool	Normal	Strong	No	S _{Weak}
D7	Overcast	Cool	Normal	Strong	Yes	
D8	Sunny	Mild	High	Weak	No	
D9	Sunny	Cool	Normal	Weak	Yes	IHI
D10	Rain	Mild	Normal	Weak	Yes	
D11	Sunny	Mild	Normal	Strong	Yes	Gain
D12	Overcast	Mild	High	Strong	Yes	
D13	Overcast	Hot	Normal	Weak	Yes	
D14	Rain	Mild	High	Strong	No	

Attribute: Wind

Values (Wind) = Strong, Weak

$$S = [9+, 5-] \qquad Entropy(S) = -\frac{9}{14}log_{2}\frac{9}{14} - \frac{5}{14}log_{2}\frac{5}{14} = 0.94$$

$$S_{Strong} \leftarrow [3+, 3-] \qquad Entropy(S_{Strong}) = 1.0$$

$$S_{Weak} \leftarrow [6+, 2-] \qquad Entropy(S_{Weak}) = -\frac{6}{8}log_{2}\frac{6}{8} - \frac{2}{8}log_{2}\frac{2}{8} = 0.8113$$

$$Gain(S, Wind) = Entropy(S) - \sum_{v \in \{Strong, Weak\}} \frac{|S_{v}|}{|S|} Entropy(S_{v})$$

$$Gain(S, Wind) = Entropy(S) - \frac{6}{14} Entropy(S_{Strong}) - \frac{8}{14} Entropy(S_{Weak})$$

$$Gain(S, Wind) = 0.94 - \frac{6}{14} \ 1.0 - \frac{8}{14} \ 0.8113 = 0.0478$$

Dav	Quitlook	Tama		\ A/in d	Play
Day	OULIOOK	Temp	Humidity	wind	Tennis
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

Gain(S, Outlook) = 0.2464

Gain(S, Temp) = 0.0289

Gain(S, Humidity) = 0.1516



Dav	Outlook	Tomn	Luna iditu	\A/ind	Play
Day	Outlook	Temp	numiaity	wind	Tennis
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

Dav	Tem	Lumidity	Wind	Play	Attribute: Temp	
Day	р	пиппиту	wind	Tennis	Values (Temp) = Hot. Mild. Cool	
D1	Hot	High	Weak	No		
D2	Hot	High	Strong	No	$S_{Sunny} = [2+, 3-]$	$Entropy(S_{Sunny}) = -\frac{2}{5}log_{2}\frac{2}{5} - \frac{3}{5}log_{2}\frac{3}{5} =$
D8	Mild	High	Weak	No	0.97	
D9	Cool	Normal	Weak	Yes	S (0 2]	$Entropy(\mathbf{S}) = 0.0$
D11	Mild	Normal	Strong	Yes	$\mathbf{S}_{Hot} \leftarrow [0+, 2-]$	$Entropy(S_{Hot}) = 0.0$
	_		·		$S_{Mild} \leftarrow [1+, 1-]$	$Entropy(S_{Mild}) = 1.0$
					$S_{Cool} \leftarrow [1+, 0-]$	$Entropy(S_{Cool}) = 0.0$
			(CS	E, $HIT, Nidas$ $Gain(S_{Sunny}, Temp) = Entropy(State)$	$Oshi_{Sunny} - \sum_{v \in \{Hot, Mild, Cool\}} \frac{ S_v }{ S } Entropy(S_v)$
					Gain(S _{Sunny} , Temp)	
					$= Entropy(S_{Sunny})$	$-\frac{2}{5}Entropy(S_{Hot})-\frac{2}{5}Entropy(S_{Mild})$
					$-\frac{1}{5}Entropy(S_{Cool})$	
					$Gain(S_{sunny}, Temp) = 0.9$	$97 - \frac{2}{5}0.0 - \frac{2}{5}1 - \frac{1}{5}0.0 = 0.570$

Dav	Tem	Uumiditu	\A/ind	Play
Day	р	пиппатту	wina	Tennis
DI	Hot	High	Weak	No
D2	Hot	High	Strong	No
D8	Mild	High	Weak	No
D9	Cool	Normal	Weak	Yes
DI1	Mild	Normal	Strong	Yes

Attribute: Humidity

Values (Humidity) = High, Normal

$$S_{Sunny} = [2+, 3-]$$

$$Entropy(S) = -\frac{2}{5}log_{2}\frac{2}{5} - \frac{3}{5}log_{2}\frac{3}{5} = 0.97$$

$$S_{high} \leftarrow [0+, 3-]$$

$$Entropy(S_{High}) = 0.0$$

$$S_{Normal} \leftarrow [2+, 0-]$$

$$Entropy(S_{Normal}) = 0.0$$

$$Gain(S_{Sunny}, Humidity) = Entropy(S_{Sunny}) - \sum_{v \in \{High, Normal\}} \frac{|S_v|}{|S|} Entropy(S_v)$$

Gain(S_{Sunny}, Humidity)

$$= Entropy(S_{Sunny}) - \frac{3}{5}Entropy(S_{High}) - \frac{2}{5}Entropy(S_{Normal})$$

$$Gain(S_{sunny}, Humidity) = 0.97 - \frac{3}{5} \ 0.0 - \frac{2}{5} \ 0.0 = 0.97$$

Dav	Tem	Uumidity	\A/ind	Play
Day	р	пиппатту	wina	Tennis
DI	Hot	High	Weak	No
D2	Hot	High	Strong	No
D8	Mild	High	Weak	No
D9	Cool	Normal	Weak	Yes
DI1	Mild	Normal	Strong	Yes

Attribute: Wind

Values (Wind) = Strong, Weak

$$S_{Sunny} = [2+, 3-] \qquad Entropy(S) = -\frac{2}{5}log_{2}\frac{2}{5} - \frac{3}{5}log_{2}\frac{3}{5} = 0.97$$

$$S_{Strong} \leftarrow [1+, 1-] \qquad Entropy(S_{Strong}) = 1.0$$

$$S_{Weak} \leftarrow [1+, 2-] \qquad Entropy(S_{Weak}) = -\frac{1}{3}log_{2}\frac{1}{3} - \frac{2}{3}log_{2}\frac{2}{3} = 0.97$$

0.9183

CSE, HIT, Nidasoshi $Gain(S_{Sunny}, Wind) = Entropy(S_{Sunny}) - \sum_{v \in \{Strong, Weak\}} \frac{|S_v|}{|S|} Entropy(S_v)$ $Gain(S_{Sunny}, Wind)$ $= Entropy(S_{Sunny}) - \frac{2}{5} Entropy(S_{Strong}) - \frac{3}{5} Entropy(S_{Weak})$ $Gain(S_{Sunny}, Wind) = 0.97 - \frac{2}{5}1.0 - \frac{3}{5}0.918 = 0.0192$

Dav	Tem	Uumidity	\A/ind	Play
Day	р	питнацу	wind	Tennis
D1	Hot	High	Weak	No
D2	Hot	High	Strong	No
D8	Mild	High	Weak	No
D9	Cool	Normal	Weak	Yes
D11	Mild	Normal	Strong	Yes

$$Gain(S_{sunny}, Temp) = 0.570$$

 $Gain(S_{sunny}, Humidity) = 0.97$

$$CSE HIT, Nidasoshi Gain(Ssunny, Wind) = 0.0192$$



Dav	Outlook	Tomn	Luna iditu	\A/ind	Play
Day	Outlook	Temp	numiaity	wind	Tennis
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

Dav	Tem	Humidity	Wind	Play
Day	р	пиппанту	wind	Tennis
D4	Mild	High	Weak	Yes
D5	Cool	Normal	Weak	Yes
D6	Cool	Normal	Strong	No
D10	Mild	Normal	Weak	Yes
D14	Mild	High	Strong	No

Attribute: Temp

Values (Temp) = Hot, Mild, Cool

Gain(S_{Rain}, Temp)

$$= Entropy(S_{Rain}) - \frac{0}{5}Entropy(S_{Hot}) - \frac{3}{5}Entropy(S_{Mild})$$

$$-\frac{2}{5}Entropy(S_{Cool})$$

$$Gain(S_{Rain}, Temp) = 0.97 - \frac{0}{5}0.0 - \frac{3}{5}0.918 - \frac{2}{5}1.0 = 0.0192$$

Dav	Tem	Uumidity	\A/ind	Play
Day	р	пиппанту	wina	Tennis
D4	Mild	High	Weak	Yes
D5	Cool	Normal	Weak	Yes
D6	Cool	Normal	Strong	No
DI0	Mild	Normal	Weak	Yes
DI4	Mild	High	Strong	No

Attribute: Humidity

Values (Humidity) = High, Normal

x
 Yes

$$S_{Rain} = [3+, 2-]$$
 $Entropy(S_{Sunny}) = -\frac{3}{5}log_2\frac{3}{5} - \frac{2}{5}log_2\frac{2}{5} =$

 0.97
 0.97
 $S_{High} \leftarrow [1+, 1-]$
 $Entropy(S_{High}) = 1.0$

 S_Normal \leftarrow [2+, 1-]
 $Entropy(S_{Normal}) = -\frac{2}{3}log_2\frac{2}{3} - \frac{1}{3}log_2\frac{1}{3} =$

 CSE, HIT, Nidasoshi

$$Gain(S_{Rain}, Humidity) = Entropy(S_{Rain}) - \sum_{v \in \{High, Normal\}} \frac{|S_v|}{|S|} Entropy(S_v)$$

 $Gain(S_{Rain}, Humidity)$

$$= Entropy(S_{Rain}) - \frac{2}{5}Entropy(S_{High}) - \frac{3}{5}Entropy(S_{Normal})$$

$$Gain(S_{Rain}, Humidity) = 0.97 - \frac{2}{5} 1.0 - \frac{3}{5} 0.918 = 0.0192$$

Day	Tem	Humidity	Wind	Play
Day	р	пиппату	vvina	Tennis
D4	Mild	High	Weak	Yes
D5	Cool	Normal	Weak	Yes
D6	Cool	Normal	Strong	No
DI0	Mild	Normal	Weak	Yes
DI4	Mild	High	Strong	No

Attribute: Wind

Values (wind) = Strong, Weak

$$S_{Rain} = [3+, 2-] \qquad Entropy(S_{Sunny}) = -\frac{3}{5}log_{2}\frac{3}{5} - \frac{2}{5}log_{2}\frac{2}{5} = 0.97$$

$$S_{Strong} \leftarrow [0+, 2-] \qquad Entropy(S_{Strong}) = 0.0$$

$$S_{Weak} \leftarrow [3+, 0-] \qquad Entropy(S_{weak}) = 0.0$$

$$SE, HIT, Nidasoshi_{Gain}(S_{Rain}, Wind) = Entropy(S_{Rain}) - \sum_{v \in \{Strong, Weak\}} \frac{|S_v|}{|S|} Entropy(S_v)$$
$$Gain(S_{Rain}, Wind) = Entropy(S_{Rain}) - \frac{2}{5} Entropy(S_{Strong}) - \frac{3}{5} Entropy(S_{Weak})$$

$$Gain(S_{Rain}, Wind) = 0.97 - \frac{2}{5} \ 0.0 - \frac{3}{5} \ 0.0 = 0.97$$

Dav	Tem	Uumidity	\A/ind	Play
Day	р	пиппанту	wina	Tennis
D4	Mild	High	Weak	Yes
D5	Cool	Normal	Weak	Yes
D6	Cool	Normal	Strong	No
DI0	Mild	Normal	Weak	Yes
DI4	Mild	High	Strong	No

 $Gain(S_{Rain}, Temp) = 0.0192$

 $Gain(S_{Rain}, Humidity) = 0.0192$

 $CSEGain(S_{Rain}, Wind) = 0.97$ Shi



CSE, HIT, Nidasoshi

DECISION TREE EXAMPLE



- What is the entropy of this collection of training examples with respect to the target function classification?
- 2. What is the information gain of *a2* relative to these training examples?
- 3. Draw decision tree for the given dataset.

Decision Tree Algorithm – ID3 Solved Example

- What is the entropy of this collection of training examples with respect to the target function classification?
- 2. What is the information gain of *a1* and *a2* relative to these training examples?
- 3. Draw decision tree for the given dataset.

Instance	Classification	al	a2	shi
1	+	Т	Т	
2	+	Т	Т	
3	-	Т	F	
4	+	F	F	
5	-	F	Т	
6	-	F	Т	

Instance	Classification	a1	a2	Attribute: a1	
1	+	Т	Т	Values(a1) = T, F	
2	+	Т	Т	$S = [3 \pm 3 =]$ Entrony(S) = 1.0	
3	-	Т	F	$\begin{bmatrix} 3 - [3+, 3-] \\ \end{bmatrix} = \begin{bmatrix} 2\pi i & 0 \\ 2\pi i & 0 \end{bmatrix}$	
4	+	F	F	$S_T = [2+, 1-] \qquad Entropy(S_T) = -\frac{2}{3}log_2\frac{2}{3} - \frac{1}{3}log_2\frac{1}{3} = 0.9183$	
5	-	F	Т		
6	-	F	Т	$S_F \leftarrow [1+, 2-] \qquad Entropy(S_F) = -\frac{1}{3}log_2\frac{1}{3} - \frac{2}{3}log_2\frac{2}{3} = 0.9183$	
CSE,				HIT, Nidasoshi	
	Example - 2			$Gain(S, a1) = Entropy(S) - \sum_{v \in \{T,F\}} \frac{ S_v }{ S } Entropy(S_v)$	
Decision Tree Algorithm – ID3 Solved Example				$Gain(S, a1) = Entropy(S) - \frac{3}{6}Entropy(S_T) - \frac{3}{6}Entropy(S_F)$	
				$Gain(S, a1) = 1.0 - \frac{3}{6} * 0.9183 - \frac{3}{6} * 0.9183 = 0.0817$	





CSE, HIT, Nidasoshi

DECISION TREE EXAMPLE

Instance	a1	a2	a3	Classification
1	True	Hot	High	No
2	True	Hot	High	No
3	False	Hot	High	Yes
4	False	Cool	Normal	Yes
SE	False	Cool	Normal	Yes
6,	True	Cool	High	No
7	True	Hot	High	No
8	True	Hot	Normal	Yes
9	False	Cool	Normal	Yes
10	False	Cool	High	Yes

1. Construct the decision tree for the following tree using ID3 Algorithm

Decision Tree Algorithm – ID3 Solved Example

Instance	a1	a2	a3	Classification
1	True	Hot	High	No
2	True	Hot	High	No
3	False	Hot	High	Yes
4	False	Cool	Normal	Yes
503	False	Cool	Normal	asoyes
6	True	Cool	High	No
7	True	Hot	High	No
8	True	Hot	Normal	Yes
9	False	Cool	Normal	Yes
10	False	Cool	High	Yes

Instance	a1	a2	a3	Classification	Attribute: a1
1	True	Hot	High	No	
2	True	Hot	High	No	Values (a1) = True, False
3	False	Hot	High	Yes	
4	False	Cool	Normal	Yes	$S = [6+, 4-] \qquad Entropy(S) = -\frac{6}{10} \log_2 \frac{6}{10} - \frac{4}{10} \log_2 \frac{4}{10} = 0.9709$
5	False	Cool	Normal	Yes	
6	True	Cool	High	No	$S_{True} = [1+, 4-] \qquad Entropy(S_{True}) = -\frac{1}{2}\log_2\frac{1}{2} - \frac{4}{2}\log_2\frac{4}{2} = 0.7219$
7	True	Hot	High	No	
8	True	Hot	Normal	Yes	$S_{Flase} \leftarrow [5+, 0-]$ $Entropy(S_{False}) = 0.0$
9	False	Cool	Normal	Yes	
10	False	Cool	High	Yes	$, \Pi \Pi, \Pi Uasos \Pi$
Example - 3 Decision Tree Algorithm – ID3 Solved Example			ple - 3		$Gain(S,a1) = Entropy(S) - \sum_{v \in \{True,False\}} \frac{ S_v }{ S } Entropy(S_v)$
			Algorith Examp	nm – ID3 le	$Gain(S, a1) = Entropy(S) - \frac{5}{10}Entropy(S_{True}) - \frac{5}{10}Entropy(S_{False})$
					$Gain(S, a1) = 0.9709 - \frac{5}{10} * 0.7219 - \frac{5}{10} * 1 = 0.6099$

Subscribe to Mahesh Huddar

Visit: vtupulse.com

Instance	a1	a2	a3	Classification	Attribute: a2
1	True	Hot	High	No	
2	True	Hot	High	No	Values (a2) = Hot, Cool
3	False	Hot	High	Yes	
4	False	Cool	Normal	Yes	$S = [6+, 4-] \qquad Entropy(S) = -\frac{6}{10} \log_2 \frac{6}{10} - \frac{4}{10} \log_2 \frac{4}{10} = 0.9709$
5	False	Cool	Normal	Yes	
6	True	Cool	High	No	$S_{Hot} = [2+,3-]$ $Entropy(S_{Hot}) = -\frac{2}{2}log_2\frac{2}{2} - \frac{3}{2}log_2\frac{3}{2} = 0.9709$
7	True	Hot	High	No	
8	True	Hot	Normal	Yes	$S = \{4, 1, 1, 1, 2, \dots, 1, 2, 2, 2, 2, \dots, 1, 2, 2, 2, 2, \dots, 1, 2, 2, 2, 2, \dots, 2, 2, 2, 2, \dots, 2, 2, \dots, 2, 2, \dots, 1, \dots, 1,$
9	False	Cool	Normal	Yes	$S_{Cool} \leftarrow [4+, 1-]$ $Entropy(S_{Cool}) = -\frac{1}{5}log_2\frac{1}{5} - \frac{1}{5}log_2\frac{1}{5} = 0.7219$
10	False	Cool	High	Yes	$, \Pi \Pi, \Pi Uasos \Pi$

Example - 3 Decision Tree Algorithm – ID3 Solved Example

$$Gain(S, a2) = Entropy(S) - \sum_{v \in \{Hot, Cool\}} \frac{|S_v|}{|S|} Entropy(S_v)$$

$$Gain(S, a2) = Entropy(S) - \frac{5}{10}Entropy(S_{Hot}) - \frac{5}{10}Entropy(S_{Cool})$$

$$Gain(S, a2) = 0.9709 - \frac{5}{10} * 0.9709 - \frac{5}{10} * 0.7219 = 0.1245$$

Visit: vtupulse.com

Instance	a1	a2	a3	Classification	Attribute: a3
1	True	Hot	High	No	
2	True	Hot	High	No	Values (a3) = High, Normal
3	False	Hot	High	Yes	
4	False	Cool	Normal	Yes	$S = [6+, 4-] \qquad Entropy(S) = -\frac{6}{10} \log_2 \frac{6}{10} - \frac{4}{10} \log_2 \frac{4}{10} = 0.9709$
5	False	Cool	Normal	Yes	
6	True	Cool	High	No	$\begin{bmatrix} S_{High} = [2+, 4-] \\ Entropy(S_{High}) = -\frac{2}{7} \log_2 \frac{2}{7} - \frac{4}{7} \log_2 \frac{4}{7} = 0.9183 \end{bmatrix}$
7	True	Hot	High	No	
8	True	Hot	Normal	Yes	$S_{Normal} \leftarrow [4+, 0-] \qquad Entropy(S_{Normal}) = 0.0$
9	False	Cool	Normal	Yes	
10	False	Cool	High	Yes	$I, \Pi II, INICASOS II$
					$Gain(S, a3) = Entropy(S) - \sum_{v \in \{High, Normal\}} \frac{ v }{ S } Entropy(S_v)$
Example - 3 Decision Tree Algorithm – ID3 Solved Example			ple - 3 Algorith Examp	ım – ID3 le	$Gain(S, a3) = Entropy(S) - \frac{6}{10}Entropy(S_{High}) - \frac{4}{10}Entropy(S_{Normal})$ $Gain(S, a3) = 0.9709 - \frac{6}{10} \times 0.9183 - \frac{4}{10} \times 0.0 = 0.4199$
Solved Lample					$uun(5, u5) = 0.9709 - \frac{10}{10} \times 0.9105 - \frac{10}{10} \times 0.0 = 0.4199$





Example - 3 Decision Tree Algorithm – ID3 Solved Example

Attribute: a2

Instance	a2	a3	Classification
1	Hot	High	No
2	Hot	High	No
6	Cool	High	No
7	Hot	High	No
8	Hot	Normal	Yes

Values(a2) = Hot, Cool

$$S_{a1} = [1+, 4-]$$
 Entropy $(S_{a1}) = -\frac{1}{5}log_2\frac{1}{5} - \frac{4}{5}log_2\frac{4}{5} = 0.7219$

$$S_{Hot} = [1+,3-]$$
 $Entropy(S_{Hot}) = -\frac{1}{4}log_2\frac{1}{4} - \frac{3}{4}log_2\frac{3}{4} = 0.8112$

$$CSE \underset{Gain(S,a2)}{S_{Cool} \leftarrow [0+, 1-]} \qquad Entropy(S_{Cool}) = 0.0$$

Example - 3 Decision Tree Algorithm – ID3 Solved Example

$$Gain(S, a2) = Entropy(S) - \frac{4}{5}Entropy(S_{Hot}) - \frac{1}{5}Entropy(S_{Cool})$$

$$Gain(S, a2) = 0.9709 - \frac{4}{5} * 0.8112 - \frac{1}{5} * 0.0 = 0.3219$$

Attribute: a3

Instance	a2	a3	Classification
1	Hot	High	No
2	Hot	High	No
6	Cool	High	No
7	Hot	High	No
8	Hot	Normal	Yes

Values (a3) = High, Normal

$$S_{a1} = [1+, 4-]$$
 $Entropy(S_{a1}) = -\frac{1}{5}log_2\frac{1}{5} - \frac{4}{5}log_2\frac{4}{5} = 0.7219$

 $S_{High} = [0+, 4-]$ Entropy $(S_{High}) = 0.0$

 $S_{Normal} \leftarrow [1+, 0-] \qquad Entropy(S_{Normal}) = 0.0$ $CSE_{Gain(S, a3)} = Entropy(S) - \underbrace{asos}_{v \in \{High, Normal\}} \underbrace{|S|}_{v \in \{High, Normal\}} Entropy(S_v)$

Example - 3 Decision Tree Algorithm – ID3 Solved Example

$$Gain(S, a3) = Entropy(S) - \frac{4}{5}Entropy(S_{High}) - \frac{1}{5}Entropy(S_{Normal})$$

$$Gain(S, a3) = 0.9709 - \frac{4}{5} * 0.0 - \frac{1}{5} * 0.0 = 0.7219$$

 $Gain(S_{a1}, a2) = 0.3219$

 $Gain(S_{a1}, a3) = 0.7219 - Maximum Gain$ Classification Instance a2 a3 High 1 Hot No 2 Hot High No a1 6 Cool High No 7 Hot High No False True 8 Hot Normal Yes CSE, HIT, 1,2,6,7,8,5,05,3,4,5,9,10 Yes a3 Example - 3 High Normal **Decision Tree Algorithm – ID3 Solved Example** 1, 2, 6, 7 8 No Yes
CSE, HIT, Nidasoshi

When to use Decision Trees

- Problem characteristics:
 - Instances can be described by attribute value pairs
 - Target function is discrete valued
 - Disjunctive hypothesis may be required
 - Possibly noisy training data samples
 Robust to errors in training data

 - Missing attribute values
- Different classification problems:
 - Equipment classification
 - Medical diagnosis
 - Credit risk analysis
 - Several tasks in natural language processing

Issues in decision trees learning

- Overfitting
 - Reduced error pruning
 - Rule post-pruning
- Extensions
 - Continuous valued attributes
 Continuous valued attributes
 - Alternative measures for selecting attributes
 - Handling training examples with missing attribute values
 - Handling attributes with different costs
 - Improving computational efficiency
 - Most of these improvements in C4.5 (Quinlan, 1993)

Overfitting: definition

- Building trees that "adapt too much" to the training examples may lead to "overfitting".
- Consider error of hypothesis h over
 - training data: $error_D(h)$ empirical error entire distribution X of data: $error_X(h)$ expected error SOSNI
- Hypothesis h overfits training data if there is an alternative hypothesis $h' \in$ H such that

 $error_{D}(h) < error_{D}(h')$ and $error_{x}(h') < error_{x}(h)$

i.e. h' behaves better over unseen data

Overfitting in decision tree learning



Avoid overfitting in Decision Trees

Two strategies:

- 1. Stop growing the tree earlier, before perfect classification
- 2. Allow the tree to *overfit* the data, and then *post-prune* the tree **CSE**, **HIT**, **Nidasoshi**
 - -Training and validation set: split the training in two parts (training and validation) and use validation to assess the utility of *post-pruning*
 - Reduced error pruning

Reduced-error pruning (Quinlan 1987)

- Each node is a candidate for pruning
- Pruning consists in removing a subtree rooted in a node: the node becomes a leaf and is assigned the most common classification
- Nodes are removed only if the resulting tree performs no worse on the validation set.
- Nodes are pruned iteratively: at each iteration the node whose removal most increases accuracy on the validation set is pruned.
- Pruning stops when no pruning increases accuracy

Effect of reduced error pruning



Rule post-pruning

- 1. Create the decision tree from the training set
- 2. Convert the tree into an equivalent set of rules
 - Each path corresponds to a rule
 - Each node along a path corresponds to a pre-condition
 - Each leaf classification to the post-condition
- 3. Prune (generalize) each rule by removing those preconditions whose removal improves accuracy over validation set
- 4. Sort the rules in estimated order of accuracy, and consider them in sequence when classifying new instances

Converting to rules



Rule Post-Pruning

- Convert tree to rules (one for each path from root to a leaf)
- For each antecedent in a rule, remove it if error rate on validation set does not decrease CSE, HIT, Nidasoshi
- Sort final rule set by accuracy

Outlook=sunny ^ humidity=high -> No Outlook=sunny ^ humidity=normal -> Yes Outlook=overcast -> Yes Outlook=rain ^ wind=strong -> No Outlook=rain ^ wind=weak -> Yes Compare first rule to: Outlook=sunny->No Humidity=high->No

Why converting to rules?

Each distinct path produces a different rule: a condition removal may be based on a local (contextual) criterion. Node pruning is global and affects all the rules
 In rule form, tests are not ordered and there is no book-keeping

involved when conditions (nodes) are removed

Converting to rules improves readability for humans

Dealing with continuous-valued attributes

- So far discrete values for attributes and for outcome.
- Given a continuous-valued attribute A, dynamically create a new attribute A_c

 A_c = True *if* A < c, False *otherwise*

- How to determine threshold value c?
- Example. Temperature in the PlayTennis example Oasoshi
 - Sort the examples according to *Temperature*

Temperature	40	48	60	72	80	90
PlayTennis	No	No	Yes	Yes	Yes	No

- Determine candidate thresholds by averaging consecutive values where there is a change in classification: (48+60)/2=54 and (80+90)/2=85
- Evaluate candidate thresholds (attributes) according to information gain. The best is *Temperature* $_{54}$. The new attribute competes with the other ones

Tid	Refund	Marital Status	Taxable Income	Cheat
1	Yes	Single	125K	No
2	No	Married	100K	No
3	No	Single	70K	No
4	Yes	Married	120K	No
5	No.S	Divorced	95K ICASO	Yes
6	No	Married	60K	No
7	Yes	Divorced	220К	No
8	No	Single	85K	Yes
9	No	Married	75K	No
10	No	Single	90К	Yes

Handling incomplete training data

- How to cope with the problem that the value of some attribute may be missing?
 - *Example*: Blood-Test-Result in a medical diagnosis problem
- The strategy: use other examples to guess attribute
 - 1. Assign the value that is most common among the training examples at the node
 - 2. Assign a probability to each value, based on frequencies, and assign values to missing attribute, according to this probability distribution
- Missing values in new instances to be classified are treated accordingly, and the most probable classification is chosen (C4.5)

Handling attributes with different costs

- Instance attributes may have an associated cost: we would prefer decision trees that use low-cost attributes
- ID3 can be modified to take into account costs:
 - 1. Tan and Schlimmer (1990) HIT, Nidasoshi $Gain^2(S, A)$

Cost(S, A)

2. Nunez (1988)

 $2^{Gain(S,A)} - 1$

 $(Cost(A) + 1)^w w \in [0,1]$

Search space in Decision Tree learning



- The search space is made by partial decision trees
- The algorithm is *hill-climbing*
- The evaluation function is information gain
- The hypotheses space is complete (represents all discrete-valued functions)
- No backtracking; no guarantee of optimality
- It uses all the available examples (not

incremental)

Inductive bias in decision tree learning

What is the inductive bias of DT learning?

1. Shorter trees are preferred over longer trees

Not enough. This is the bias exhibited by a simple breadth first algorithm generating all DT's e selecting the shorter one

2. Prefer trees that place high information gain attributes close to the root

S J P N Trust's



HIRASUGAR INSTITUTE OF TECHNOLOGY, NIDASOSHI.

Inculcating Values, Promoting Prosperity Approved by AICTE, Recognized by Govt. of Karnataka and Affiliated to VTU Belagavi. Accredited at 'A' Grade by NAAC Programmes Accredited by NBA: CSE, ECE, EEE & ME

Artificial Neural Networks

Prof. Mahesh G Huddar

Dept. of Computer Science and Engineering

Artificial Neural Networks

- Introduction
- Neural Network Representation
- Appropriate Problems for Neural Network Learning
- Perceptrons
- Multilayer Networks and BACKPROPAGATION Algorithms
- Remarks on the BACKPROPAGATION Algorithms

Artificial Neural Networks

- ANN learning well-suited to problems which the training data corresponds to noisy, complex data (inputs from cameras or microphones)
- Can also be used for problems with symbolic representations
- Most appropriate for problems where
 - Instances have many attribute-value pairs
 - Target function output may be discrete-valued, real-valued, or a vector of several real- or discrete-valued attributes

Nidasoshi

- Training examples may contain errors
- Long training times are acceptable
- Fast evaluation of the learned target function may be required
- The ability for humans to understand the learned target function is not important

Appropriate Problems – for ANN

- *Instances are represented by many attribute-value pairs.* The target function to be learned is defined over instances that can be described by a vector of predefined features, such as the pixel values in the ALVINN example. These input attributes may be highly correlated or independent of one another. Input values can be any real values.
- *The target function output may be discrete-valued, real-valued, or a vector of several real- or discrete-valued attributes.* For example, in the ALVINN system the output is a vector of 30 attributes, each corresponding to a recommendation regarding the steering direction. The value of each output is some real number between 0 and 1, which in this case corresponds to the confidence in predicting the corresponding steering direction. We can also train a single network to output both the steering command and suggested acceleration, simply by concatenating the vectors that encode these two output predictions.

Appropriate Problems – for ANN

The training examples may contain errors. ANN learning methods are quite robust to noise in the training data.

Long training times are acceptable. Network training algorithms typically require longer training times than, say, decision tree learning algorithms. Training times can range from a few seconds to many hours, depending on factors such as the number of weights in the network, the number of training examples considered, and the settings of various learning algorithm parameters.

Appropriate Problems – for ANN

Fast evaluation of the learned target function may be required. Although ANN

- learning times are relatively long, evaluating the learned network, in order to
- apply it to a subsequent instance, is typically very fast. For example, ALVINN
- applies its neural network several times per second to continually update its steering command as the vehicle drives forward.

The ability of humans to understand the learned target function is not important.

The weights learned by neural networks are often difficult for humans to

interpret. Learned neural networks are less easily communicated to humans than learned rules

Neural Network History

- History traces back to the 50's but became popular in the 80's with work by Rumelhart, Hinton, and Mclelland
 - A General Framework for Parallel Distributed Processing in Parallel Distributed Processing: Explorations in the Microstructure of Cognition
- Peaked in the 90's.:
 - Hundreds of variants
 - Less a model of the actual brain than a useful tool, but still some debate
- Numerous applications
 - Handwriting, face, speech recognition
 - Vehicles that drive themselves
 - Models of reading, sentence production, dreaming
- Debate for philosophers and cognitive scientists
 - Can human consciousness or cognitive abilities be explained by a connectionist model or does it require the manipulation of symbols?

Biological Motivation

- The study of artificial neural networks (ANNs) has been inspired by the observation that biological learning systems are built of very complex webs of interconnected *Neurons*
- Human information processing system consists of brain neuron: basic building block cell that communicates information to and from various parts of body
- Simplest model of a neuron: considered as a threshold unit –a processing element (PE)
- Collects inputs & produces output if the sum of the input exceeds an internal threshold value

Biological Motivation

- The human brain is made up of billions of simple processing units neurons.
- Inputs are received on dendrites, and if the input levels are over a threshold, the neuron fires, passing a signal through the axon to the synapse which then connects to another neuron.



Biological Motivation



Simplest Neural Network Artificial Neural Network



Simplest Neural Network



FIND-S: Step-2

Example	Sky	AirTemp	Humidity	Wind	Water	Forecast	EnjoySport
1	Sunny	Warm	Normal	Strong	Warm	Same	Yes
2	Sunny	Warm	High	Strong	Warm	Same	Yes
3	Rainy	Cold	High	Strong	Warm	Change	No
4	Sunny	Warm	High	Strong	Cool	Change	Yes



		Example	Sky	AirTemp	Humidity	Wind	Water	Forecast	EnjoySport	
		1 2 3	Sunny Sunny Rainy	Warm Warm Cold	Normal High High	Strong Strong Strong	Warm Warm Warm	Same Same Change	Yes Yes No	
		4	Sunny	Warm	High	Strong	Cool	Change	Yes	
S ₀ :				< 2	ð,Ø,Ø,Ø,Ø,Ø	$\emptyset. \varnothing angle$				
S ₁ :			⟨Sur	nny,Warm,	Normal, Str	rong, Wai	rm, Same	€>		
S ₂ :	S ₃ :		(Suni	ny,Warm, ?	, Strong, W	/arm, Sar	me>	sos	hi	
S_4			[⟨Sunny, W	/arm, ?, Str	ong, ?, ?)	,			
G ₄ :		[⟨Sunny,	?, ?, ?, ?, ?,	?>	P, Warm,	?, ?, ?, ?	\rangle		
G ₃ :	⟨ <i>Sunn</i> y	⟨, ?,?,?,?,? ,?⟩	,V</td <td>Varm,?,?,?,</td> <td>?> (?,</td> <td>?,Normal</td> <td>,?,?,?⟩</td> <td><pre>(?, ?,?,'</pre></td> <td>?,Cool,?></td> <td>⟨?,?,?,?,?,Same⟩</td>	Varm,?,?,?,	?> (?,	?,Normal	, ?,?,? ⟩	<pre>(?, ?,?,'</pre>	?,Cool,?>	⟨ ?,?,?,?,?,Same ⟩
G _{0:}	G _{1:}	G _{2:}			, ?, ?, ?, ?,</td <td>?, ?〉</td> <td></td> <td></td> <td></td> <td></td>	?, ?〉				

Decision Trees

• Decision trees represent a disjunction of conjunctions on constraints on the value of attributes:



Artificial Neurons

- Artificial neurons are based on biological neurons.
- Each neuron in the network receives one or more inputs.
- An activation function is applied to the inputs, which determines the output of the neuron – the activation level.

Artificial Neurons



Artificial Neurons


Artificial Neurons



Artificial Neurons

• A typical activation function works as follows:

$$X = \sum_{i=1}^{n} w_i x_i \qquad Y = \begin{cases} +1 & \text{for } X > t \\ 0 & \text{for } X \le t \end{cases}$$

- Each node *i* has a weight, *w_i* associated with it. CSE, HIT, Nidasoshi
- The input to node i is x_i .
- *t* is the threshold.
- So if the weighted sum of the inputs to the neuron is above the threshold, then the neuron fires.

Artificial Neurons

• The charts on the right show three typical activation functions.



Step function (Linear Threshold Unit)

step(x) = 1, if x >= threshold
0, if x < threshold</pre>

Sign function sign(x) = +1, if x >= 0 -1, if x < 0 Sigmoid function

 $sigmoid(x) = 1/(1+e^{-x})$

PERCEPTRON

- One type of ANN system is based on a unit called a perceptron.
- A perceptron takes a vector of real-valued inputs, calculates a linear combination of these inputs, then outputs a 1 if the result is greater than some threshold and -1 otherwise.
- More precisely, given inputs *x1* through *xn*, the output o(x1, ..., xn) computed by the perceptron is

$$o(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } w_0 + w_1 x_1 + w_2 x_2 + \dots + w_n x_n > 0 \\ -1 & \text{otherwise} \end{cases}$$

• where each *wi* is a real-valued constant, or weight, that determines the contribution of input *xi* to the perceptron output

PERCEPTRON



$$o(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } w_0 + w_1 x_1 + w_2 x_2 + \dots + w_n x_n > 0 \\ -1 & \text{otherwise} \end{cases}$$

PERCEPTRON

• For brevity, we will sometimes write the perceptron function as,

$$o(\vec{x}) = sgn(\vec{w} \cdot \vec{x})$$

 $sgn(y) = \begin{cases} 1 & \text{if } y > 0 \\ -1 & \text{otherwise} \end{cases}$

where

- Learning a perceptron involves choosing values for the weights *wo*, ..., *wn*.
- Therefore, the space H of candidate hypotheses considered in perceptron learning is the set of all possible real-valued weight vectors.

$$H = \{\vec{w} \mid \vec{w} \in \mathfrak{R}^{(n+1)}\}$$

- One way to learn an acceptable weight vector is to begin with random weights, then iteratively apply the perceptron to each training example, modifying the perceptron weights whenever it misclassifies an example.
- This process is repeated, iterating through the training examples as many times as needed until the perceptron classifies all training examples correctly.
- Weights are modified at each step according to the *perceptron training rule*, which revises the weight *wi* associated with input *xi* according to the rule

$$w_i \leftarrow w_i + \Delta w_i$$

where

$$\Delta w_i = \eta(t-o)x_i$$

 $w_i \leftarrow w_i + \Delta w_i$

where

• Here t is the 1

$$\Delta w_i = \eta(t - o) x_i$$
 perceptron, and *n* is a positive

constant called the *learning rate*.

- The role of the learning rate is to moderate the degree to which weights are changed at each step.
- It is usually set to some small value (e.g., 0.1) and is sometimes made to decay as the number of weight-tuning iterations increases.





- Why should this update rule converge toward successful weight values?
- To get an intuitive feel, consider some specific cases.
- Suppose the training example is correctly classified already by the perceptron. In this case, (t o) is zero, making Δwi zero, so that no weights are updated.
- Suppose the perceptron outputs a -1, when the target output is + 1.
- To make the perceptron output a + 1 instead of 1 in this case, the weights must be altered to increase the value of $\vec{w} \cdot \vec{x}$.
- For example, if xi > 0, then increasing wi will bring the perceptron closer to correctly classifying this example. Notice the training rule will increase wi in this case, because (t o), n, and xi are all positive.
- For example, if xi = .8, q = 0.1, t = 1, and o = -1, then the weight update will be $\Delta wi = q(t o)xi = 0.1(1 (-1))0.8 = 0.16$.
- On the other hand, if *t* = -1 and *o* = 1, then weights associated with positive *xi* will be decreased rather than increased.

A single perceptron can be used to represent many Boolean functions weights 0.6 and 0.6

AND function

If A=0 & B=0 \rightarrow 0*0.6 + 0*0.6 = 0

This is not greater than the threshold of 1, so the output = 0

If A=0 & B=1 \rightarrow 0*0.6 + 1*0.6 = 0.6

This is not greater than the threshold, so the output = 0

If A=1 & B=0 \rightarrow 1*0.6 + 0*0.6 = 0.6

This is not greater than the threshold, so the output = 0

If A=1 & B=1 \rightarrow 1*0.6 + 1*0.6 = 1.2

This exceeds the threshold, so the output = 1





A single perceptron can be used to represent many Boolean functions weights 1.2 and 0.6

AND function

If A=0 & B=0 \rightarrow 0*1.2 + 0*0.6 = 0

This is not greater than the threshold of 1, so the output = 0

If A=0 & B=1 \rightarrow 0*1.2 + 1*0.6 = 0.6

This is not greater than the threshold, so the output = 0

If A=1 & B=0 \rightarrow 1*1.2 + 0*0.6 = 1.2 E, HIT, Nidasos

This is greater than the threshold, so the output = 1

But the expected output is 0

А	В	A ^ B
0	0	0
0	1	0
1	0	0
1	1	1

A single perceptron can be used to represent many Boolean functions weights 1.2 and 0.6

AND function

Threshhold $\theta = 1$ Learning Rate n = 0.5

wi = wi + n(t - o)xi w1 = 1.2 + 0.5(0 - 1)1 = 0.7 w2 = 0.6 + 0.5(0 - 1)0 = 0.6, HIT, Nidasos $\frac{A \quad B \quad A^{A}B}{0 \quad 0 \quad 0}$ $\frac{1 \quad 0 \quad 0}{1 \quad 0}$

A single perceptron can be used to represent many Boolean functions weights 1.2 and 0.6

AND function

If A=0 & B=0 \rightarrow 0*0.7 + 0*0.6 = 0

This is not greater than the threshold of 1, so the output = 0

If A=0 & B=1 \rightarrow 0*0.7 + 1*0.6 = 0.6

This is not greater than the threshold, so the output = 0

If A=1 & B=0 \rightarrow 1*0.7 + 0*0.6 = 0.7

This is greater than the threshold, so the output = 0

If A=1 & B=0 \rightarrow 1*0.7 + 1*0.6 = 1.3

This is greater than the threshold, so the output = 0





A single perceptron can be used to represent many Boolean functions with initial weights 0.6 and

0.6

OR function

If A=0 & B=0 \rightarrow 0*0.6 + 0*0.6 = 0

This is not greater than the threshold of 1, so the output = 0 If A=0 & B=1 \rightarrow 0*0.6 + 1*0.6 = 0.6

This is not greater than the threshold, so the output = 0

But the expected output is 1

$$wi = wi + n(t - o)xi$$

$$w1 = 0.6 + 0.5(1 - 0)0 = 0.6$$

 $w^2 = 0.6 + 0.5(1 - 0)^2 = 1.1$

Α	B	Y=A+B
0	0	0
0	1	1
1	0	1
1	1	1

A single perceptron can be used to represent many Boolean functions

OR function

If A=0 & B=0 \rightarrow 0*0.6 + 0*1.1 = 0

This is not greater than the threshold of 1, so the output = 0

If A=0 & B=1 \rightarrow 0*0.6 + 1*1.1 = 1.1 This is greater than the threshold, so the output = 1 NiCaSOS

If A=1 & B= \rightarrow 1*1.1 + 0*1.1 = 1.1

This is greater than the threshold, so the output = 1

But the expected output is 1

wi = wi + n(t - o)xi w1 = 0.6 + 0.5(1 - 0)1 = 1.1w2 = 1.1 + 0.5(1 - 0)0 = 1.1

Α	B	Y=A+B
0	0	0
0	1	1
1	0	1
1	1	1

A single perceptron can be used to represent many Boolean functions

OR function

If A=0 & B=0 \rightarrow 0*1.1 + 0*1.1 = 0

This is not greater than the threshold of 1, so the output = 0

If A=0 & B=1 \rightarrow 0*1.1 + 1*1.1 = 1.1 HIT NICLES INCLESS IN This is greater than the threshold, so the output = 1.

If A=1 & B= $\rightarrow 1*1.1 + 0*1.1 = 1.1$

This is greater than the threshold, so the output = 1.

If A=1 & B= \rightarrow 1*1.1 + 1*1.1 = 2.2

This is greater than the threshold, so the output = 1.





A single perceptron can be used to represent many Boolean functions

NOT function





- We can view the perceptron as representing a hyperplane decision surface in the n-dimensional space of instances (i.e., points).
- The perceptron outputs a 1 for instances lying on one side of the hyperplane and outputs a -1 for instances lying on the other side.
- The equation for this decision hyperplane $\vec{w} \cdot \vec{x} = 0.05$
- Of course, some sets of positive and negative examples cannot be separated by any hyperplane.
- Those that can be separated are called linearly separable sets of examples.



- A single perceptron can be used to represent many boolean functions.
- For example, if we assume boolean values of 1 (true) and -1 (false), then one way to use a two-input perceptron to implement the AND function is to set the weights wo = -0.8, and w1 = w2 = 0.5.
- This perceptron can be made to represent the OR function instead by altering the threshold to wo = -0.3, ICASOSO
- In fact, AND and OR can be viewed as special cases of m-of-n functions: that is, functions where at least m of the n inputs to the perceptron must be true.
- The OR function corresponds to m = 1 and the AND function to m = n.
- Any m-of-n function is easily represented using a perceptron by setting all input weights to the same value (e.g., 0.5) and then setting the threshold *wo* accordingly.

- Perceptrons can represent all of the primitive boolean functions AND, OR, NAND, and NOR.
- Unfortunately, however, some boolean functions cannot be represented by a single perceptron, such as the XOR function whose value is 1 if and only if \mathbf{x} ! = $\mathbf{x}2$.
- Note the set of linearly nonseparable training examples shown in Figure (b) corresponds to this XOR function.

- The ability of perceptrons to represent AND, OR, NAND, and NOR is important because *every* boolean function can be represented by some network of interconnected units based on these primitives.
- In fact, every boolean function can be represented by some network of perceptrons only two levels deep, in which the inputs are fed to multiple units, and the outputs of these units are then input to a second, final stage.
- One way is to represent the boolean function in disjunctive normal form (i.e., as the disjunction (OR) of a set of conjunctions (ANDs) of the inputs and their negations).
- Note that the input to an AND perceptron can be negated simply by changing the sign of the corresponding input weight.
- Because networks of threshold units can represent a rich variety of functions and because single units alone cannot, we will generally be interested in learning multilayer networks of threshold units.



A perceptron can learn only examples that are called "linearly separable". These are examples that can be perfectly separated by a hyperplane.

CSE, HIT, Nidasoshi

- Although the perceptron rule finds a successful weight vector when the training examples are linearly separable, it can fail to converge if the examples are not linearly separable.
- A second training rule, called the *delta rule*, is designed to overcome this difficulty.
- If the training examples are not linearly separable, the delta rule converges toward a best-fit approximation to the target concept.
- The key idea behind the delta rule is to use *gradient descent* to search the hypothesis space of possible weight vectors to find the weights that best fit the training examples.
- This rule is important because gradient descent provides the basis for the BACKPROPAGATON algorithm, which can learn networks with many interconnected units.
- It is also important because gradient descent can serve as the basis for learning algorithms that must search through hypothesis spaces containing many different types of continuously parameterized hypotheses.

inputs









• The delta training rule is best understood by considering the task of training an *unthresholded* perceptron; that is, a *linear unit* for which the output *o* is given by

$$o = w_0 + w_1 x_1 + \dots + w_n x_n$$

- Thus, a linear unit corresponds to the mist stage ($o(\vec{x}) = \vec{w} \cdot \vec{x}$) the uneshold.
- In order to derive a weight learning rule for lin by specifying a measure for the *training error* of a hypothesis (weight vector), relative to the training examples.
- Although there are many ways to define this error, one common measure is
- where D is the set of training examples, *td* is the target output for training example *d*, and *od* is the output of the linear unit for training example d.

$$E(\vec{w}) \equiv \frac{1}{2} \sum_{d \in D} (t_d - o_d)^2$$

Derivation of Gradient Descent Rule

- How can we calculate the direction of steepest descent along the error surface?
- This direction can be found by computing the derivative of E with respect to each component of the vector \vec{w} .
- This vector derivative is called the *gradient* of *E* with respect to \vec{w} , written as $\nabla E(\vec{w})$.

$$\mathsf{CSE}_{\nabla E(\vec{w})} \equiv \begin{bmatrix} \frac{\partial E}{\partial w_0}, \frac{\partial E}{\partial w_1}, \frac{\partial E}{\partial w_n} \end{bmatrix} \mathsf{OShi}$$

• Since the gradient specifies the direction of steepest increase of *E*, the training rule for gradient descent is $\vec{w} \leftarrow \vec{w} + \Delta \vec{w}$

$$\Delta w_i = \eta(t-o)x_i$$

where

$$\Delta \vec{w} = -\eta \nabla E(\vec{w})$$

Derivation of Gradient Descent Rule

$\vec{w} \leftarrow \vec{w} + \Delta \vec{w}$

where

• Here *n* is a positive (

 $\Delta \vec{w} = -\eta \nabla E(\vec{w})$ in the gradient descent search. The

negative sign is present because we want to move the weight vector in the direction that *decreases* E.

• This training rule can also be written in its component form CSE, HIT, Nidasoshi

$$w_i \leftarrow w_i + \Delta w_i$$

where

$$\Delta w_i = -\eta \frac{\partial E}{\partial w_i}$$

$$\nabla E(\vec{w}) \equiv \left[\frac{\partial E}{\partial w_0}, \frac{\partial E}{\partial w_1}, \cdots, \frac{\partial E}{\partial w_n}\right]$$

Derivation of Gradient Descent Rule

$$\frac{\partial E}{\partial w_i} = \frac{\partial}{\partial w_i} \frac{1}{2} \sum_{d \in D} (t_d - o_d)^2 \qquad E(\vec{w}) \equiv \frac{1}{2} \sum_{d \in D} (t_d - o_d)^2 = \frac{1}{2} \sum_{d \in D} \frac{\partial}{\partial w_i} (t_d - o_d)^2 \qquad o(\vec{x}) = \vec{w} \cdot \vec{x} = \frac{1}{2} \sum_{d \in D} 2(t_d - o_d) \frac{\partial}{\partial w_i} (t_d - o_d) |\mathsf{T}| \qquad \mathsf{Nidasoshi} = \sum_{d \in D} (t_d - o_d) \frac{\partial}{\partial w_i} (t_d - \vec{w} \cdot \vec{x_d}) \qquad \Delta w_i = -\eta \frac{\partial E}{\partial w_i} \frac{\partial E}{\partial w_i} = \sum_{d \in D} (t_d - o_d) (-x_{id}) \qquad \Delta w_i = \eta \sum_{d \in D} (t_d - o_d) x_{id}$$

 $w_i \leftarrow w_i + \Delta w_i$

Gradient-Descent(training_examples, η)

Each training example is a pair of the form $\langle \vec{x}, t \rangle$, where \vec{x} is the vector of input values, and t is the target output value. η is the learning rate (e.g., .05).

- Initialize each w_i to some small random value
- Until the termination condition is met, Do
 - Initialize each Δw_i to zero. NICASOSNI For each $\langle \vec{x}, t \rangle$ in training_examples, Do
 - - * Input the instance \vec{x} to the unit and compute the output o
 - * For each linear unit weight w_i , Do

$$\Delta w_i := \Delta w_i + \eta (t - o) x_i$$

- For each linear unit weight w_i , Do

$$w_i := w_i + \Delta w_i$$
STOCHASTIC APPROXIMATION TO GRADIENT DESCENT

- Gradient descent is an important general paradigm for learning.
- It is a strategy for searching through a large or infinite hypothesis space that can be applied whenever
 - 1. the hypothesis space contains continuously parameterized hypotheses (e.g., the weights in a linear unit), and
 - 2. the error can be differentiated with respect to these hypothesis parameters.
- The key practical difficulties in applying gradient descent are
 - converging to a local minimum can sometimes be quite slow (i.e., it can require many thousands of gradient descent steps), and
 if there are multiple local minima in the error surface, then there is no guarantee that the procedure will find the global
 - 2. if there are multiple local minima in the error surface, then there is no guarantee that the procedure will find the global minimum.

STOCHASTIC APPROXIMATION TO GRADIENT DESCENT

Stochastic Gradient-Descent($training_examples, \eta$)

Each training example is a pair of the form $\langle \vec{x}, t \rangle$, where \vec{x} is the vector of input values, and t is the target output value. η is the learning rate (e.g., .05).

- Initialize each w_i to some small random value
- Until the termination condition is met, 6aSOShi
 - Initialize each Δw_i to zero.
 - For each $\langle \vec{x}, t \rangle$ in training_examples, Do
 - * Input the instance \vec{x} to the unit and compute the output o
 - * For each linear unit weight w_i , Do

 $w_i := w_i + \eta(t - o)x_i$

STOCHASTIC APPROXIMATION TO GRADIENT DESCENT

The key differences between standard gradient descent and stochastic gradient descent are:

- In standard gradient descent, the error is summed over all examples before updating weights, whereas in stochastic gradient descent weights are updated upon examining each training example.
- Summing over multiple examples in standard gradient descent requires more computation per weight update step. On the other hand, because it uses the true gradient, standard gradient descent is often used with a larger step size per weight update than stochastic gradient descent.
- In cases where there are multiple local minima with respect to $E(\vec{w})$, stochastic gradient descent can sometimes avoid falling into these local minima because it uses the various $\nabla E_d(\vec{w})$ rather than $\nabla E(\vec{w})$ to guide its search.

• Multilayer neural networks can classify a range of

functions, including non linearly separable ones.

- Each input layer neuron connects to all neurons in the hidden layer.
- The neurons in the hidden layer connect to all neurons in

the output layer.



A Differentiable Threshold Unit

- What type of unit shall we use as the basis for constructing multilayer networks?
- At first we might be tempted to choose the linear units discussed in the previous section, for which we have already derived a gradient descent learning rule.
- However, multiple layers of cascaded linear units still produce only linear functions, and we prefer networks capable of representing highly nonlinear functions.
- The perceptron unit is another possible choice, but its discontinuous threshold makes it undifferentiable and hence unsuitable for gradient descent.
- What we need is a unit whose output is a nonlinear function of its inputs, but whose output is also a differentiable function of its inputs.
- One solution is the sigmoid unit-a unit very much like a perceptron, but based on a smoothed, differentiable threshold function.

- The sigmoid unit is illustrated in below Figure. Like the perceptron, the sigmoid unit first computes a linear combination of its inputs, then applies a threshold to the result.
- In the case of the sigmoid unit, however, the threshold output is a continuous function of its input.SE, HIT, Nidasoshi x_{I} $x_0 = 1$ w0 $net = \sum_{i=0}^{n} w_i x_i$ $o = \sigma(net) =$



More precisely, the sigmoid unit computes its output *o* as

 $o=\sigma(\vec{w}\cdot\vec{x})$

CSE,
$$\prod_{y \in V} \frac{\sigma(y)}{dy} = \frac{1}{1 + \sigma^2} asoshi$$
$$\frac{d\sigma(\bar{y})}{dy} = \sigma(y) \cdot (1 - \sigma(y))$$

The BACKPROPAGATIAON Algorithm

- Multilayer neural networks learn in the same way as perceptrons.
- However, there are many more weights, and it is important to

assign credit (or blame) correctly when changing weights.

• *E* sums the errors over all of the network output units

$$E(\vec{w}) \equiv \frac{1}{2} \sum_{d \in D} \sum_{k \in outputs} (t_{kd} - o_{kd})^2$$

The BACKPROPAGATIAON Algorithm

- x_{ji} = the *i*th input to unit *j*
- w_{ji} = the weight associated with the *i*th input to unit *j*
- $net_j = \sum_i w_{ji} x_{ji}$ (the weighted sum of inputs for unit j)
- $o_j =$ the output computed by unit j, NidaSoShi $t_j =$ the target output for unit j
- σ = the sigmoid function
- outputs = the set of units in the final layer of the network
- Downstream(j) = the set of units whose immediate inputs include the output of unit j

















CSE, HIT, Nidasoshi

Back Propagation Algorithm

- Create a feed-forward network with n_i inputs, n_{hidden} hidden units, and n_{out} output units.
- Initialize all network weights to small random numbers
- Until the termination condition is met, Do
 - For each (x, t), in training examples, Do
 - Propagate the input forward through the network:
 1. Input the instance *x*, to the network and compute the output o_u of every unit u in the network.
 - Propagate the errors backward through the network

2. For each network unit k, calculate its error term $\boldsymbol{\delta}_k$

$$\delta_k \leftarrow o_k (1 - o_k) (t_k - o_k)$$

3. For each network unit h, calculate its error term δ_h

$$\delta_h \leftarrow o_h(1-o_h) \sum_{k \in outputs} w_{h,k} \delta_k$$

4. Update each network weight w_{ji} $w_{ji} \leftarrow w_{ji} + \Delta w_{ji}$

Where

$$\Delta \, \mathrm{w}_{\mathrm{ji}} = \eta \delta_j x_{\mathrm{ji}}$$

- To derive the equation for updating weights in back propagation algorithm, we use Stochastic gradient descent rule.
- Stochastic gradient descent involves iterating through the training examples one at a time, for each training example **d** descending the gradient of the error *Ed* with respect to this single example. **CSE, HIT, Nidasoshi**
- In other words, for each training example **d** every weight wji is updated by adding to it Δw_{ij} .

Where

• That is,

$$w_{ji} \leftarrow w_{ji} + \Delta w_{ji}$$

Where

$$W_{ji} \leftarrow W_{ji} + \Delta W_{ji}$$

 $\Delta w_{ji} = -\eta \frac{\partial E_d}{\partial w_{ji}}$

• where *Ed* is the error on training example **d**, that is half the squared difference between the target output and the actual output over all output units in the network,

$$E_d(\vec{w}) \equiv \frac{1}{2} \sum_{k \in outputs} (t_k - o_k)^2$$

• Here outputs is the set of output units in the network, \mathbf{t}_k is the target value of unit k for training example \mathbf{d} , and \mathbf{o}_k is the output of unit k given training example \mathbf{d} .

Notation Used:

 x_{ji} = the ith input to unit j

 w_{ji} = the weight associated with the ith input to unit j

- $net_j = \sum_i w_{ji} X_{ji}$ (the weighted sum of inputs for unit j)
- o_j = the output computed by unit j IT, Nidasoshi
- t_j = the target output for unit j
- σ = the sigmoid function

outputs = the set of units in the final layer of the network

Downstream(j) = the set of units whose immediate inputs include the output of unit j



Where

$$W_{ji} \leftarrow W_{ji} + \Delta W_{ji}$$

 $W_{ji} = -\eta \frac{\partial E_d}{\partial W_{ji}}$

• To begin, notice that weight *wji* can influence the rest of the network only through *netj*. Therefore, we can use the chain rule to write, $\frac{\partial E_d}{\partial w_{ji}} = \frac{\partial E_d}{\partial net_j} \frac{\partial net_j}{\partial w_{ji}}$ $= \frac{\partial E_d}{\partial net_i} x_{ji}$ $\frac{\partial net_j}{\partial w_{ji}} = x_{ji}$

$$\Delta w_{ji} = -\eta \, \frac{\partial E_d}{\partial net_j} \, x_{ji}$$

• Our remaining task is to derive a convenient expression for $\frac{\partial E_d}{\partial net_j}$

To derive a convenient expression for $\frac{\partial E_d}{\partial net_i}$

We consider two cases in turn:

- Case 1, where unit j is an output unit for the network, and OSNI
- Case 2, where unit j is an internal unit of the network.

Case 1: Training Rule for Output Unit Weights

Just as *wji* can influence the rest of the network only through *net_j*, *net_j* can influence the network only through *oj*. Therefore, we can invoke the chain rule again to write,

$$\frac{\partial E_d}{\partial net_j} = \frac{\partial E_d}{\partial o_j} \frac{\partial o_j}{\partial net_j} \qquad \frac{\partial E_d}{\partial o_j} = \frac{\partial}{\partial o_j} \frac{1}{2} \sum_{k \in outputs} (t_k - o_k)^2 , \qquad \frac{\partial o_j}{\partial (net_j)} = \frac{\partial \sigma(net_j)}{\partial (net_j)} \text{ for } \frac{\partial \sigma(x)}{\partial (x)} = \sigma(x) (1 - \sigma(x))$$

$$= \frac{\partial}{\partial o_j} \frac{1}{2} (t_j - o_j)^2 \qquad \qquad = \sigma(net_j) (1 - \sigma(net_j))$$

$$= \frac{1}{2} 2(t_j - o_j) \frac{\partial(t_j - o_j)}{\partial o_j} \qquad \qquad = o_j (1 - o_j)$$

$$= -(t_j - o_j) \qquad \qquad \frac{\partial E_d}{\partial net_j} = -(t_j - o_j) o_j (1 - o_j)$$

Case 1: Training Rule for Output Unit Weights

$$\Delta w_{ji} = -\eta \frac{\partial E_d}{\partial w_{ji}} \qquad \qquad \frac{\partial E_d}{\partial net_j} = -(t_j - o_j) o_j (1 - o_j)$$

$$\Delta w_{ji} = -\eta \frac{\partial E_d}{\partial net_j} x_{ji} CSE, HIT, Nidasoshi$$

$$\Delta w_{ji} = \eta (t_j - o_j) o_j (1 - o_j) x_{ji}$$

$$\delta_j = (t_j - o_j) o_j (1 - o_j).$$

 $\Delta w_{ji} = \eta \, \delta_j \, x_{ji}$

 $\mathbf{n} = \mathbf{n}$



Case 2: Training Rule for Hidden Unit Weights

Case 2: Training Rule for Hidden Unit Weights

$$\Delta w_{ji} = -\eta \frac{\partial E_d}{\partial net_j} x_{ji} \qquad \qquad \frac{\partial E_d}{\partial net_j} = \sum_{k \in Downstream(j)} -\delta_k w_{kj} o_j(1-o_j)$$
$$\Delta w_{ji} = \eta o_j(1-o_j) \sum_{k \in Downstream(j)} \delta_k w_{kj} x_{ji} , \text{Nidasoshi}$$

$$\Delta w_{ji} = \eta \ \delta_j \ x_{ji} \qquad \qquad \delta_j = o_j (1 - o_j) \sum_{k \in Downstream(j)} \delta_k \ w_{kj}$$