# Lectures 6 and 7: Logistic Regression \& Decision-based Learning (Perceptrons) \& Decision Trees 

Prof. Krishna R. Pattipati<br>Dept. of Electrical and Computer Engineering<br>University of Connecticut<br>Contact: krishna@engr.uconn.edu (860) 486-2890

Fall 2018
October 22 \& 29, 2018

## Reading List

- Murphy Chapters 8, 21.6-21.8
- Ripley (LP formulation of Perceptron)
- Bishop, Chapter 4, Chapter 10.6
- S.Y. Kung, Fuzzy Perceptrons
- Sergios Theodiridis, Chapter 7
- Duda, Hart and Stork, Chapter 8.1-8.4


## Lecture Outline

- Classification as a Function Approximation Problem
- Discuss Generative versus Discriminative Learning
- Approximating Posterior Probabilities
- Gaussian Case: Two Class and Multiple Class Cases
- Decision Based Learning
- Perceptron Convergence Theorem
- Optimal Learning Rate \& Normalized Perceptron
- Extension to Multiple Classes
- Fuzzy Updates, LVQ as a Perceptron Update
- Decision Trees and Regularization
- Bagging, Random Forests and Gradient-based Boosting
- Summary


## Classification as a Function Approximation Problem-1

## $\square$ Classification as a Function Approximation Problem

Define a target vector $\underline{z}_{k}=\underline{e}_{k}=\left[\begin{array}{c}0 \\ \vdots \\ 1 \\ \vdots \\ 0\end{array}\right] \Rightarrow z=k ; z_{k}=\left[\begin{array}{c}\text { Alternate } \\ -1 /(C-1) \\ \vdots \\ 1 \\ \vdots \\ -1 /(C-1)\end{array}\right]$
want to find a mapping $\underline{\hat{z}}=\underline{y}(\underline{x}) \ni$ Mean Square Error is a min.
$J=\frac{1}{2} E\left\{(\underline{y}(\underline{x})-\underline{z})^{T}(\underline{y}(\underline{x})-\underline{z})\right\}=\frac{1}{2} \int_{\underline{x}} \sum_{\underline{z}}(\underline{y}(\underline{x})-\underline{z})^{T}(\underline{y}(\underline{x})-\underline{z}) p(\underline{x}, \underline{z}) d \underline{x}$ $=\frac{1}{2} \int_{\underline{x}} \sum_{\underline{\underline{1}}}(\underline{y}(\underline{x})-\underline{z})^{T}(\underline{y}(\underline{x})-\underline{z}) P(\underline{z} / \underline{x}) p(\underline{x}) d \underline{x}$

$$
\nabla_{\underline{\underline{y}}} J=0 \Rightarrow \quad \sum_{\underline{z}}(\underline{y}(\underline{x})-\underline{z}) P(\underline{z} / \underline{x})=0 \Rightarrow \underline{y}(\underline{x})=\sum_{\underline{z}} \underline{z} P(\underline{z} / \underline{x})=E[\underline{z} / \underline{x}]=\underline{\hat{z}}(\underline{x})
$$

conditional mean

## Classification as a Function Approximation Problem-2

- MMSE estimate $=$ Conditional Mean

$$
\underline{y}(\underline{x})=\sum_{k=1}^{c} e_{k} P(z=k \mid \underline{x})=\underline{P}=\left[\begin{array}{l}
P(z=1 \mid \underline{x}) \\
P(z=2 \mid \underline{x}) \\
P(z=C \mid \underline{x})
\end{array}\right]
$$



## Minimum Distance versus Maximum a Posteriori Decisions

$\square$ Minimum distance versus Maximum a Posteriori decisions
Select class $i$ if $\left\|\underline{P}-\underline{z}_{k}\right\|^{2}$ is a min

$$
d_{k}=\left(\underline{P}-\underline{z}_{k}\right)^{T}\left(\underline{P}-\underline{z}_{k}\right)=\|\underline{P}\|^{2}-2 \underline{P}^{T} \underline{z}_{k}+\underline{z}_{k}^{T} \underline{z}_{k}=1-2 P(z=k \mid \underline{x})+\|\underline{P}\|^{2}
$$

Alternate ：$d_{k}=\left(\underline{y}_{k}-\underline{z}_{k}\right)^{T}\left(\underline{y}_{k}-\underline{z}_{k}\right)=\frac{C}{C-1}\left[1-P(z=k \mid \underline{x})+C\|\underline{\underline{P}}\|^{2}\right]$

$$
\min _{k} d_{k} \Rightarrow \max _{k} P(z=k \mid \underline{x})
$$

－Learning posterior probabilities directly ．．．Discriminative Learning
$P(z=k \mid \underline{x}) \sim$ posterior probability of class $k$ given $\underline{x}$

$$
P(z=k \mid \underline{x})=\frac{P(\underline{x} \mid z=k) P(z=k)}{\sum_{i=1}^{C} P(\underline{x} \mid z=i) P(z=i)}
$$

## MAP Decisions: Gaussian Case

> Gaussian case

$$
\begin{aligned}
& P(z=k \mid \underline{x})=\frac{\frac{1}{\left|\Sigma_{k}\right|^{1 / 2}} e^{-\frac{1}{2}\left(\underline{x}-\underline{\mu}_{k}\right)^{T} \Sigma_{k}^{-1}\left(\underline{x}-\underline{\mu}_{k}\right)} P(z=k)}{\sum_{i=1}^{C} \frac{1}{\left|\Sigma_{i}\right|^{1 / 2}} e^{-\frac{1}{2}\left(\underline{x}-\underline{\mu}_{i}\right)^{T} \Sigma_{i}^{-1}\left(\underline{x}-\underline{\mu}_{i}\right)} P(z=i)} \\
& =\frac{e^{\left[-\frac{1}{2}\left(\underline{x}^{-\mu_{k}}\right)^{T} \Sigma_{k}^{-1}\left(\underline{x}^{-}-\underline{\mu}_{k}\right)-\frac{1}{2} \ln \left|\Sigma_{k}\right|+\ln P(z=k)\right]}}{\sum_{i=1}^{C} e^{\left[-\frac{1}{2}\left(\underline{x}-\underline{\mu}_{i}\right)^{T} \Sigma_{i}^{-1}\left(\underline{x}-\underline{\mu}_{i}\right)-\frac{1}{2} \ln \left|\Sigma_{i}\right|+\ln P(z=i)\right]}} \text { Why learn } \underline{\theta} ? ~ \\
& y(\underline{x})=\underline{x}^{T} A_{i} \underline{x}+\underline{b}_{i}^{T} \underline{x}+c_{i} \\
& \text { Why not learn weights directly? } \\
& =\frac{e^{y\left(\underline{x}, \underline{\theta}_{k}\right)}}{\sum_{i=1}^{C} e^{y\left(x, \theta_{i}\right)}} \quad \underline{\theta}=[\underbrace{\{\underbrace{}_{k}\},\left\{\Sigma_{k}\right\}, P\{z=k\}}_{\underline{\theta}_{k}}] \\
& \begin{array}{|c|}
\hline \text { "soft max" } \\
\text { "multiple logistic model" } \\
\text { "logistic discrimination" }
\end{array} \\
& y\left(\underline{x}, \underline{\theta}_{k}\right)=g_{k}(\underline{x}) \\
& \text { discriminant function }
\end{aligned}
$$

$>$ Equal covariance case

$$
\begin{aligned}
P(z=k \mid \underline{x}) & =\frac{e^{\underline{\mu}_{k}^{T} \Sigma^{-1} x-\frac{1}{2} \underline{\mu}_{k}^{T} \Sigma^{-1} \underline{\mu}_{k}+\ln P(z=k)}}{\sum_{i=1}^{C} e^{\underline{\mu}_{i}^{T} \Sigma^{-1} x-\frac{1}{2} \underline{\mu}_{i}^{T} \Sigma^{-1} \underline{\mu}_{i}+\ln P(z=i)}} \\
= & \frac{e^{y\left(\underline{x}, \underline{w}_{k}\right)}}{\sum_{i=1}^{C} e^{y\left(\underline{x}, \underline{w}_{i}\right)}}
\end{aligned}
$$

where $\mathrm{y}\left(\underline{x}, \underline{w}_{k}\right)=\underline{w}_{k}^{T} \underline{x}-w_{k 0}$

$$
\underline{w}_{k}=\Sigma^{-1} \underline{\mu}_{k} \quad w_{k 0}=\frac{1}{2} \underline{\mu}_{k}^{T} \Sigma^{-1} \underline{\mu}_{k}-\ln P(z=k)
$$

## MAP Decisions: Gaussian, Equal $\Sigma$, Two Class Case

$>$ Two class case
Let the classes be $z \in\{0,1\}$

Note: If $z \in\{-1,1\}$,
$P(z \mid \underline{x}, \underline{w})=\frac{1}{1+e^{-z y}}=g(y, z)$
$\mu=P\left(z=1 \mid \underline{x}, \underline{w}_{1}, \underline{w}_{0}\right)=\frac{e^{y\left(\underline{x}, \underline{w}_{1}\right)}}{e^{y\left(\underline{x}, \underline{w}_{1}\right)}+e^{y\left(x, w_{0}\right)}}=\frac{1}{1+e^{-\left(y\left(x, w_{1}\right)-y\left(\underline{x}, w_{0}\right)\right)}}$

| $=\frac{1}{1+e^{-y}}=g(y)=\mu$ | $\frac{d g}{d y}=g(1-g)$ |
| :--- | :--- |

where $y=y\left(\underline{x}, \underline{w}_{1}\right)-y\left(\underline{x}, \underline{w}_{0}\right)=\underline{w}_{1}^{T} \underline{x}-w_{10}-\underline{w}_{0}{ }^{T} \underline{x}+w_{00}$
$\frac{\mu}{1-\mu}=e^{v}$

$$
P(z \mid \underline{x}, \underline{w})=\mu^{z}(1-\mu)^{1-z}=\left(\frac{1}{1+e^{-y}}\right)^{z}\left(\frac{1}{1+e^{y}}\right)^{1-z}=\frac{1}{1+e^{-(2 z-1) y}}=e^{y z} g(-y)
$$

## One of the classes can be taken as a reference and set its weights to 0 !

The parameters $\underline{\tilde{w}}, \tilde{w}_{0}$ are directly estimated from the data. Assume $\underline{x}$ is augmented.

## Maximum Likelihood Estimation of Weights

- Suppose have data

$$
D=\left\{\left\{\underline{x}^{1}, z^{1}\right\}\left\{\underline{x}^{2}, z^{2}\right\}, \cdots \cdots,\left\{\underline{x}^{N}, z^{N}\right\}: z^{n} \in\{0,1\}\right\}
$$

- Likelihood

$$
\begin{aligned}
& L(\underline{w})=P\left(\left\{z^{n}\right\}_{n=1}^{N} \mid\left\{\underline{x}^{n}\right\}_{n=1}^{N}, \underline{w}\right)=\prod_{n=1}^{N} P\left(z^{n} \mid \underline{x}^{n}, \underline{w}\right)=\prod_{n=1}^{N}\left(\mu_{n}\right)^{z^{n}}\left(1-\mu_{n}\right)^{1-z^{n}} \\
& J(\underline{w})=-\ln L(\underline{w})=N L L(\underline{w})=-\sum_{n=1}^{N} J_{n} ; J_{n}=\left[z^{n} \ln \mu_{n}+\left(1-z^{n}\right) \ln \left(1-\mu_{n}\right)\right] \ldots-\text { ve Cross Entropy } \\
& X=\left[\begin{array}{c}
\left(\underline{x}^{1}\right)^{T} \\
\left(\underline{x}^{2}\right)^{T} \\
\cdot \\
\left(\underline{N}^{N}\right)^{T}
\end{array}\right] ; N \mathbf{x}(p+1) \\
& \nabla_{\underline{\underline{w}}}^{2} J=X^{T} \Sigma X>0 \text { where } \Sigma=\operatorname{Diag}\left[\mu_{n}\left(1-\mu_{n}\right)\right] \\
& \text { Re gularize by adding } \lambda \underline{w}^{T} \underline{w} \text { or } \lambda\|\underline{w}\|_{1} \text { to } J(\underline{w})
\end{aligned}
$$

## Iteratively Reweighted Least Squares (IRLS)

- Newton's Method

$$
\begin{array}{rlr}
\underline{w}_{i+1} & =\underline{w}_{i}-\left(\nabla_{\underline{w}_{i}}^{2} J\right)^{-1} \nabla_{\underline{w}_{i}} J \\
& =\underline{w}_{i}+\left(X^{T} \Sigma_{i} X\right)^{-1} X^{T}\left(\underline{z}-\underline{\mu}_{i}\right) \\
& =\left(X^{T} \Sigma_{i} X\right)^{-1}\left[X^{T} \Sigma_{i} X \underline{w}_{i}+X^{T}\left(\underline{z}-\underline{\mu}_{i}\right)\right] \\
& =\left(X^{T} \Sigma_{i} X\right)^{-1} X^{T} \Sigma_{i}\left[X \underline{w}_{i}+\Sigma_{i}^{-1}\left(\underline{z}-\underline{\mu}_{i}\right)\right] & y_{i n}=\underline{w}_{i}^{T} \underline{x}^{n}+\frac{z^{n}-\mu_{i n}}{\mu_{i n}\left(1-\mu_{i n}\right)}
\end{array}
$$

- This corresponds to the following weighted least squares problem at iteration $i$

$$
\min J_{i}=\frac{1}{2}\left(\underline{y}_{i}-X \underline{w}\right)^{T} \Sigma_{i}\left(\underline{y}_{i}-X \underline{w}\right)
$$

- Weighted least squares can be implemented in recursive way.... One sample at a time.


## Multi-class Case

- Convert the problem to one versus rest problem (one-hot coding)
- Multiclass problems
- $\quad M$ class problem: $M$ two-class tasks or $M(M-1) / 2$ two-class tasks
- Obtain a separate classifier for each pair
- What if more than one decision rule classifies? Majority rule
- Error Correcting Output Codes
- Each class: an $n$ bit pattern chosen so that Hamming distance is as large as possible among classes
- Learn each of the $n$ bits via binary classifiers
- Select nearest class (i.e., one with the least Hamming distance)
- Work Directly with Multi-class formulation
- Direct extension of the binary case
- Upper and lower bounds on sigmoid function
- Laplace approximation of posterior to quantify uncertainty in predictions


## Multiple Class Case-1

- What to do in multiple class case ?

$$
\begin{aligned}
& \text { Let } z_{k}^{n}=\left\{\begin{array}{c}
1 \text { if } z^{n}=k \\
0 ; \text { otherwise }
\end{array}\right. \text {.... one of C coding } \\
& P\left(z_{k}^{n}=1 \mid \underline{x}^{n},\left\{\underline{w}_{i}\right\}_{i=1}^{C}\right)=\frac{e^{y\left(\underline{x}^{n}, w_{k}\right)}}{\sum_{i=1}^{C} e^{y\left(\underline{\underline{x}}^{n}, \underline{w}_{i}\right)}}=\frac{e^{\underline{w}_{k}^{T} \underline{x}^{n}}}{\sum_{i=1}^{C} e^{w_{i}^{T} \underline{x}^{n}}}=\mu_{n k} \ldots \text { soft max function }
\end{aligned}
$$

$$
\text { Use class } C \text { as a reference. So, set } \underline{w}_{C}=\underline{0} \Rightarrow W^{T}=\left[\begin{array}{c}
\underline{w}_{1}^{T} \\
\underline{w}_{2}^{T} \\
\cdot \\
\underline{w}_{C-1}^{T}
\end{array}\right] ;(C-1) \mathrm{x}(p+1)
$$

$$
L=p\left(\left\{z^{n}\right\}_{n=1}^{N} \mid\left\{\underline{x}^{n}\right\}_{n=1}^{N}, W\right)=\prod_{n=1}^{N}\left(\prod_{k=1}^{C-1}\left(\mu_{n k}\right)^{z_{k}^{n}}\right)\left(1-\sum_{i=1}^{C-1} \mu_{n i}\right)^{\left(1-\sum_{i=1}^{C-1} z_{i}^{n}\right)}
$$

## Multiple Class Case－2

$$
J=-\ln L=\sum_{n=1}^{N}\left[\left\{-\sum_{k=1}^{C-1} z_{k}^{n} \ln \left(\frac{\mu_{n k}}{1-\sum_{i=1}^{C-1} \mu_{n i}}\right)\right\}-\ln \left(1-\sum_{i=1}^{C-1} \mu_{n i}\right)\right]
$$

Note $: \frac{\mu_{n k}}{1-\sum_{i=1}^{C-1} \mu_{n i}}=e^{w_{\underline{w}}^{T} \underline{x}^{n}}$ and $\left(1-\sum_{i=1}^{C-1} \mu_{n i}\right)=\left(1+\sum_{i=1}^{C-1} e^{\underline{w}_{i}^{T}} \underline{x}^{n}\right)^{-1} ; \operatorname{Recall} \underline{w}_{C}=\underline{0}$
$\Rightarrow J=\sum_{n=1}^{N}\left[\left(-\sum_{k=1}^{C-1} z_{k}^{n} \underline{w}_{k}^{T} \underline{x}^{n}\right)+\ln \left(1+\sum_{i=1}^{C-1} e^{w_{i}^{T}} \underline{x}^{n}\right)\right]$
$\nabla_{\underline{w}_{k}} J=\sum_{n=1}^{N}\left(\mu_{n k}-z_{k}^{n}\right) \underline{x}^{n} \because \nabla_{\underline{w}_{k}} \ln \left(1+\sum_{i=1}^{C-1} e^{w_{i}^{T}} \underline{x}^{n}\right)=\left(\frac{e^{e_{k}^{T} \underline{x}^{n}}}{1+\sum_{i=1}^{C-1} e^{w_{i}^{T}} \underline{x}^{n}}\right) \underline{x}^{n}=\mu_{n k} \underline{x}^{n}$
$\nabla_{\underline{w}_{k} \underline{w}_{k}}^{2} J=\sum_{n=1}^{N} \mu_{n k}\left(1-\mu_{n k}\right) \underline{x}^{n}\left(\underline{x}^{n}\right)^{T} ; \nabla_{\underline{w}_{k} \underline{w}_{j}}^{2} J=-\sum_{n=1}^{N} \mu_{n k} \mu_{n j} \underline{x}^{n}\left(\underline{x}^{n}\right)^{T} ; k \neq j \quad$ Big Matrix

## Variational Upper Bound on Sigmoid

- Upper bounds on sigmoid function (helps when C is large)

Consider the problem of max imizing $\ln g(x)=-\ln \left(1+e^{-x}\right)$... concave in $x$
Let $y=e^{-x} \Rightarrow \ln y=-x$
Consider $\max \{-\ln (1+y)\}$ subject to $\ln y+x=0$
$L(x, y, \lambda)=-\ln (1+y)+\lambda(\ln y+x) \Rightarrow y^{*}=\frac{\lambda}{1-\lambda}$
$L\left(x, y^{*}, \lambda\right)=\underbrace{(1-\lambda) \ln (1-\lambda)+\lambda \ln \lambda}_{-H(\lambda)}+\lambda x$
$\Rightarrow \ln g(x) \leq \min _{\lambda}[\underbrace{(1-\lambda) \ln (1-\lambda)+\lambda \ln \lambda}_{-H(\lambda)}+\lambda x] \forall x$


So, $g(x) \leq \exp (\lambda x-H(\lambda)) \forall x, \lambda$
This can be used to obtain an upper bound on LL and use primal-dual methods for optimization. We will illustrate its application in the multi-class case.

## Log-Sum-Exponent bound \& Multi-class Case

- Recall
$l=\ln p\left(\left\{z^{n}\right\}_{n=1}^{N} \mid\left\{\underline{x}^{n}\right\}_{n=1}^{N}, W\right)=\sum_{n=1}^{N}\left[\left(\sum_{k=1}^{c-1} z_{k}^{n} \underline{w}_{k}^{T} \underline{x}^{n}\right)-\ln \left(1+\sum_{i=1}^{C-1} e^{w_{i}^{T} \underline{x}^{n}}\right)\right]$
$\&-\ln \left(1+e^{x}\right)$ is concave in $x$
Consider $\max -\left\{\ln \left(1+\sum_{i=1}^{C-1} e^{y_{i}}\right)\right\} \Rightarrow$ Lagrangian $L\left(\left\{y_{i}\right\}, \lambda, v\right\}=-\ln (1+v)+\lambda\left(v-\sum_{i=1}^{C-1} e^{y_{i}}\right)$
$\Rightarrow v=(1-\lambda) / \lambda \Rightarrow-\ln \left(1+\sum_{i=1}^{c-1} e^{y_{i}}\right) \leq \ln \lambda+1-\lambda\left(1+\sum_{i=1}^{c-1} e^{y_{i}}\right)$
so, $-\ln \left(1+\sum_{i=1}^{C-1} e^{w_{i}^{T} x^{n}}\right) \leq \ln \lambda_{n}+1-\lambda_{n}\left(1+\sum_{i=1}^{C-1} e^{w_{i}^{T} \underline{x}^{n}}\right)$
So, $l \leq \sum_{n=1}^{N}\left[\left(\sum_{k=1}^{c-1} z_{k}^{n} \underline{w}_{k}^{T} \underline{x}^{n}\right)+\ln \lambda_{n}+1-\lambda_{n}\left(1+\sum_{i=1}^{c-1} e^{w_{i}^{T} \underline{x}^{n}}\right)\right] \cdots$ maximize UB wrt $\left\{\underline{w}_{k}\right\}$ and minimize wrt $\lambda$
The UB is a separable problem for each class given $\left\{\lambda_{n}\right\}$ !
$\left.\tilde{l}_{k}=\sum_{n=1}^{N}\left[z_{k}^{n} \underline{w}_{k}^{T} \underline{x}^{n}-\lambda_{n} e^{w_{k}^{T} \underline{x}^{n}}\right)\right] . .$. .Needs NLP optimizer
Next $\lambda_{n}: \lambda_{n}=1 /\left(1+\sum_{i=1}^{C-1} e^{w_{i}^{T} \underline{x}^{n}}\right)$
Iterate


## Jensen + Log-normal + Variational in Multi-class Case

- Recall MAP formulation (Note: NLL)
$J=-\ln p\left(\left\{z^{n}\right\}_{n=1}^{N} \mid\left\{\underline{x}^{n}\right\}_{n=1}^{N}, W\right)-\ln \underbrace{p(W)}_{\text {prior }}$
$=\sum_{n=1}^{N}\left[\left(-\sum_{k=1}^{C-1} z_{k}^{n} \underline{w}_{k}^{T} \underline{x}^{n}\right)+\ln \left(1+\sum_{i=1}^{C-1} e^{\underline{w}_{i}^{T} \underline{x}^{n}}\right)\right]-\frac{1}{2}\left(\sum_{k=1}^{C-1} \ln \left|\Sigma_{k 0}\right|+\left(\underline{w}_{k}-\underline{w}_{k 0}\right)^{T} \Sigma_{k 0}^{-1}\left(\underline{w}_{k}-\underline{w}_{k 0}\right)+(p+1) \ln 2 \pi\right)$
Let $q(W)=\prod_{k=1}^{C-1} q_{k}\left(\underline{w}_{k}\right)$ where $q_{k}\left(\underline{w}_{k}\right) \sim N\left(\underline{w}_{k} ; \underline{w}_{k N}, \Sigma_{k N}\right) ; p\left(\underline{w}_{k}\right) \sim N\left(\underline{w}_{k} ; \underline{w}_{k 0}, \Sigma_{k 0}\right)$
$\& \ln \left(1+e^{x}\right)$ is convex in $x$
$\Rightarrow E_{q}\left[\ln \left(1+\sum_{k=1}^{C-1} e^{\underline{w}_{k}^{T} \underline{x}^{n}}\right)\right] \leq \ln \left(1+\sum_{k=1}^{C-1} E_{q_{k}}\left(e^{\underline{w}_{k}^{T} \underline{x}^{n}}\right)\right)=\ln \left(1+\sum_{k=1}^{C-1} e^{\left(\underline{w}_{k N}^{T} \underline{x}^{n}+\frac{1}{2} \underline{x}^{\underline{x}^{T}} \Sigma_{k N} \underline{x}^{n}\right)}\right) \begin{aligned} & \text { Recall Moment } \\ & \text { Generating Function }\end{aligned}$
$\operatorname{so}, \ln \left(1+\sum_{k=1}^{C-1} e^{\left(w_{k N}^{T} \underline{\underline{n}}^{n}+\frac{1}{2} \underline{\underline{n}}^{n^{T}} \Sigma_{k N} \underline{\underline{x}}^{n}\right)}\right) \leq \ln \lambda_{n}+1-\lambda_{n}\left(1+\sum_{k=1}^{C-1} e^{\left(w_{k N N}^{T} \underline{\Sigma}^{n}+\frac{1}{2^{n}}{ }^{T} \sum_{k N} \underline{\Sigma}^{n}\right)}\right)$
So, $J \leq \sum_{n=1}^{N}\left[\left(-\sum_{k=1}^{C-1} z_{k}^{n} \underline{w}_{k N}^{T} \underline{x}^{n}\right)+\ln \lambda_{n}+1-\lambda_{n}\left(1+\sum_{k=1}^{C-1} e^{\left(w_{k N N}^{T} \underline{\underline{n}}^{n}+\frac{1}{2^{2}} T^{T} \sum_{k N} \underline{x}^{n}\right)}\right)\right]+$
$\frac{1}{2} \sum_{k=1}^{C-1}\left(\operatorname{tr}\left(\Sigma_{k N} \Sigma_{k 0}^{-1}\right)+\ln \left|\Sigma_{k 0}\right|+\left(\underline{w}_{k N}-\underline{w}_{k 0}\right)^{T} \Sigma_{k 0}^{-1}\left(\underline{w}_{k N}-\underline{w}_{k 0}\right)+(p+1) \ln 2 \pi\right)$


## Jensen＋Log－normal＋Variational in Multi－class Case

$$
\begin{aligned}
& \frac{1}{2} \sum_{k=1}^{C-1}\left(\operatorname{tr}\left(\Sigma_{k N} \Sigma_{k 0}^{-1}\right)+\ln \left|\Sigma_{k 0}\right|+\left(\underline{w}_{k N}-\underline{w}_{k 0}\right)^{T} \Sigma_{k 0}^{-1}\left(\underline{w}_{k N}-\underline{w}_{k 0}\right)+(p+1) \ln 2 \pi\right)
\end{aligned}
$$

The LB is a separable problem for each class！

$$
\left.\tilde{J}_{k}=\sum_{n=1}^{N}\left[-z_{k}^{n} \underline{w}_{k N}^{T} \underline{x}^{n}-\lambda_{n} e^{\left(w_{k N N}^{T} \Sigma^{n}+\frac{1}{2} \underline{z}^{T} x^{T} k N \underline{N}^{n}\right)}\right)\right]+\left(\operatorname{tr}\left(\Sigma_{k N} \Sigma_{k 0}^{-1}\right)+\ln \left|\Sigma_{k 0}\right|+\left(\underline{w}_{k N}-\underline{w}_{k 0}\right)^{T} \Sigma_{k 0}^{-1}\left(\underline{w}_{k N}-\underline{w}_{k 0}\right)\right)
$$

．．．．．needs NLP
Alternately，can approximate exponential upto first order at current iteration and obtain closed form solutions．


## Variational Logistic Regression -1

- Lower bound on sigmoid function

Consider $\ln g(x)=-\ln \left(1+e^{-x}\right)=-\ln \left[e^{-x / 2}\left(e^{x / 2}+e^{-x / 2}\right)\right]=\frac{x}{2}-\ln \left(e^{x / 2}+e^{-x / 2}\right)$
$f(x)=-\ln \left(e^{x / 2}+e^{-x / 2}\right)$ is convex in $x^{2}$. Why?
Let $x=\sqrt{y} \Rightarrow f(y)=-\ln \left(e^{\sqrt{y} / 2}+e^{-\sqrt{y} / 2}\right) ; y \geq 0$
$\frac{d f}{d y}=\frac{-1}{4 \sqrt{y}} \tanh \left(\frac{\sqrt{y}}{2}\right)<0 ; \frac{d^{2} f}{d y^{2}}=\frac{\frac{\sinh (\sqrt{y})}{y^{3 / 2}}-\frac{1}{y}}{8(\cosh (\sqrt{y})+1)}>0 \forall y \geq 0$
Re call for convex functions

$f(y) \geq f\left(y_{0}\right)+\left.\frac{d f}{d y}\right|_{y=y_{0}}\left(y-y_{0}\right) \forall y_{0}=x_{0}^{2}$
So, $f(y) \geq-\ln \left(e^{{\sqrt{y_{0}}}^{\prime} / 2}+e^{-\sqrt{y_{0}} / 2}\right)-\frac{1}{4 \sqrt{y_{0}}} \tanh \left(\frac{\sqrt{y_{0}}}{2}\right)\left(y-y_{0}\right)$
$\Rightarrow f(x) \geq-\ln \left(e^{x_{0} / 2}+e^{-x_{0} / 2}\right)-\underbrace{\frac{1}{4 x_{0}}}_{\lambda\left(x_{0}\right)} \tanh \left(\frac{x_{0}}{2}\right)\left(x^{2}-x_{0}^{2}\right)$
$\ln g(x) \geq \frac{x-x_{0}}{2}+\underbrace{\frac{x_{0}}{2}-\ln \left(e^{x_{0} / 2}+e^{-x_{0} / 2}\right)}_{\ln g\left(x_{0}\right)}-\lambda\left(x_{0}\right)\left(x^{2}-x_{0}^{2}\right)$
So, $g(x) \geq g\left(x_{0}\right) \exp \left\{\frac{x-x_{0}}{2}-\lambda\left(x_{0}\right)\left(x^{2}-x_{0}^{2}\right)\right\}$


## Variational Logistic Regression - 2

- Binary (Two class) Case using local variational lower bound Posterior distribution of $z$ for a given $x$

$$
\begin{aligned}
& y=\underline{w}^{T} \underline{x} \text { or } y=\underline{w}^{T} \phi(\underline{x}) \\
& \begin{aligned}
p(z \mid \underline{w}) & =g(y)^{z}(1-g(y))^{1-z}=\left(\frac{1}{1+e^{-y}}\right)^{z}\left(\frac{e^{-y}}{1+e^{-y}}\right)^{1-z} \\
& =e^{y z}\left(\frac{e^{-y}}{1+e^{-y}}\right)=e^{y z}\left(\frac{1}{1+e^{y}}\right)=e^{y z} g(-y)
\end{aligned}
\end{aligned}
$$

$\operatorname{Re}$ call $g(y) \geq g\left(y_{0}\right) \exp \left\{\left(y-y_{0}\right) / 2-\lambda\left(y_{0}\right)\left(y^{2}-y_{0}^{2}\right)\right\}$
$\lambda\left(y_{0}\right)=\frac{1}{4 y_{0}} \tanh \left(\frac{y_{0}}{2}\right)=\frac{1}{4 y_{0}}\left(\frac{1-e^{-y_{0}}}{1+e^{-y_{0}}}\right)=\frac{1}{2 y_{0}}\left[\frac{1}{2}-g\left(-y_{0}\right)\right]=\frac{1}{2 y_{0}}\left[g\left(y_{0}\right)-\frac{1}{2}\right]$
$g(-y) \geq g\left(y_{0}\right) \exp \left\{-\left(y+y_{0}\right) / 2-\lambda\left(y_{0}\right)\left(y^{2}-y_{0}^{2}\right)\right\} \forall y_{0}$
$p(z \mid \underline{w})=e^{y z} g(-y) \geq e^{y z} g\left(y_{0}\right) \exp \left\{-\left(y+y_{0}\right) / 2-\lambda\left(y_{0}\right)\left(y^{2}-y_{0}^{2}\right)\right\}$
$-\ln p(z \mid \underline{w}) \leq-z \underline{w}^{T} \underline{x}-\ln g\left(y_{0}\right)+\left(\underline{w}^{T} \underline{x}+y_{0}\right) / 2+\lambda\left(y_{0}\right)\left[\left(\underline{w}^{T} \underline{x}\right)^{2}-y_{0}^{2}\right]$
Quadratic function in $\underline{w} \Rightarrow$ Gaussian
posterior
Given Data, $D=\left\{\underline{x}^{n}, z^{n}\right\}_{n=1}^{N}$ and prior $p(\underline{w})=N\left(\underline{w} ; \underline{w}_{0}, \Sigma_{0}\right)$
$-\ln p(\underline{w} \mid D) \leq-\ln p(\underline{w})-\sum_{n=1}^{N}\left\{\ln \left(g\left(y_{0 n}\right)\right)+z^{n} \underline{w}^{T} \underline{x}^{n}-\left(\underline{w}^{T} \underline{x}^{n}+y_{0 n}\right) / 2-\lambda\left(y_{0 n}\right)\left[\left(\underline{w}^{T} \underline{x}^{n}\right)^{2}-y_{0 n}^{2}\right]\right\}$

## Variational EM for Logistic Regression

- Variational EM for minimizing the upper bound on NLL

Variational E-step :

$$
\begin{aligned}
& q(\underline{w})=p\left(\underline{w} \mid\left\{y_{0 n}\right\}^{o l d}\right)=N\left(\underline{w} ; \underline{w}_{N}, \Sigma_{N}\right) \\
& \left(\Sigma_{N}\right)^{-1}=\left(\Sigma_{0}\right)^{-1}+2 \sum_{n=1}^{N} \lambda\left(y_{0 n}^{\text {old }}\right) \underline{x}^{n} \underline{x}^{n^{T}}=\left(\Sigma_{0}\right)^{-1}+\sum_{n=1}^{N} \frac{1}{y_{0 n}}\left[g\left(y_{0 n}\right)-\frac{1}{2}\right] \underline{x}^{n} \underline{x}^{n^{T}} \\
& \underline{w}_{N}=\Sigma_{N}\left(\left(\Sigma_{0}\right)^{-1} \underline{w}_{0}+\sum_{n=1}^{N}\left(z^{n}-1 / 2\right) \underline{x}^{n}\right)
\end{aligned}
$$

Variational $M$ - step : decouples for each $y_{0 n}$

$$
\begin{aligned}
& Q_{i}\left(y_{o n}, y_{o n}^{\text {old }}\right)=E\left\{\ln \left(g\left(y_{0 n}\right)\right)-y_{0 n} / 2-\lambda\left(y_{0 n}\right)\left[\left(\underline{w}^{T} \underline{x}^{n}\right)^{2}-y_{0 n}^{2}\right]\right\} \\
& \frac{d Q_{i}\left(y_{o n}, y_{o n}^{o d}\right)}{d y_{o n}}=0 \Rightarrow \frac{1}{g\left(y_{0 n}\right)} g\left(y_{0 n}\right)\left[1-g\left(y_{0 n}\right)\right]-\frac{1}{2}+2 y_{0 n} \lambda\left(y_{0 n}\right)-\frac{d \lambda\left(y_{0 n}\right)}{d y_{o n}}\left[E\left(\underline{w}^{T} \underline{x}^{n}\right)^{2}-y_{0 n}^{2}\right] \\
& \quad=\left[\frac{1}{2}-g\left(y_{0 n}\right)\right]+\left[g\left(y_{0 n}\right)-\frac{1}{2}\right]-\frac{d \lambda\left(y_{0 n}\right)}{d y_{o n}}\left[E\left(\underline{w}^{T} \underline{x}^{n}\right)^{2}-y_{0 n}^{2}\right]=0 \\
& \Rightarrow E\left(\underline{w}^{T} \underline{x}^{n}\right)^{2}-y_{0 n}^{2}=0 \Rightarrow\left(y_{0 n}^{n e w}\right)^{2}=\left(\underline{x}^{n}\right)^{T}\left(\Sigma_{N}+\underline{w}_{N} \underline{w}_{N}^{T}\right) \underline{x}^{n}
\end{aligned}
$$

This is still too much work! Are there simpler algorithms? Perceptrons

## Laplace (Gaussian) Approximation of Posterior

Represent $W$ by a vector $\underline{w}=\left[\underline{w}_{1}^{T} \underline{w}_{2}^{T} \ldots ., \underline{w}_{C-1}^{T}\right]^{T}$, a $(p+1) \mathrm{x}(C-1)$ vector
Let $\underline{w}^{*}$ be the optimal solution minimizing the NLL function, $J$
Let $\nabla_{\underline{w} w}^{2} J=H$ be the Hessian evaluated at $\underline{w}^{*}$

- Laplace approximation of posterior weight distribution $p(\underline{w} \mid D) \sim N\left(\underline{w} ; \underline{w}^{*}, H^{-1}\right) ; H=$ Hessian is the Information matrix
- Posterior predictive distribution of $z$ for a given $\underline{x}$ ( $\Rightarrow$ testing)
$P(z \mid \underline{x}, D)=\int_{\underline{w}} P(z \mid \underline{x}, \underline{w}) p(\underline{w} \mid D) d \underline{w}$
(i) Plug-in MAP estimate $\Rightarrow p(\underline{w} \mid D)=\delta\left(\underline{w}-\underline{w}^{*}\right): P(z \mid \underline{x}, D) \approx P\left(z \mid \underline{x}, \underline{w}^{*}\right)$
(ii) Monte Carlo approximation: $P(z \mid \underline{x}, D) \approx \frac{1}{M} \sum_{m=1}^{M} P\left(z \mid \underline{x}, \underline{w}^{m}\right)$
$\underline{w}^{m}$ is drawn from the posterior distribution, $p(\underline{w} \mid D)$
(iii) Probit approximation: $P(z \mid \underline{x}, D) \approx \int_{\underline{w}} P(z \mid \underline{x}, \underline{w}) N\left(\underline{w} ; \underline{w}^{*}, H^{-1}\right) d \underline{w}$

Linearly Separable Data


MAP for Prior $N(\underline{w}, \underline{0}, 100 I)$


## Contours of Posterior Predictive Distribution

Plug-in Approximation


MC Approximation
MC approx of $p(Z=1 \mid x)$


Laplace + Probit Approximation of Posterior


- Plug-in underestimates uncertainty
- MC: Posterior mean decision boundary is equally far from both classes (similar to large margin classifiers, e.g., SVM)
- Laplace + probit is similar to MC


## Probit Approximation of Sigmoid - 1

- What is probit? CDF of normal distribution

$$
\Phi(a) \triangleq \int_{-\infty}^{a} N(x ; 0,1) d x \quad \text { cumulative distribution function }(C D F) ; \Phi(-\infty)=0 ; \Phi(\infty)=1
$$

- Sigmoid function $g(a)$ approximates probit $\Phi(\lambda a)$ well if slopes are matched at the origin. Note $g(-\infty)=0 ; g(\infty)=1$

$$
\begin{aligned}
& g^{\prime}(0)=g(0)[1-g(0)]=\frac{1}{4} ; \Phi^{\prime}(0)=\frac{\lambda}{\sqrt{2 \pi}} \Rightarrow \lambda=\sqrt{\frac{\pi}{8}} \begin{array}{l}
\text { Sofine } x=z-\lambda a ; z \sim N(z ; 0,1) \\
\Rightarrow p(x)=N\left(x ;-\lambda \mu_{a}, 1+\lambda^{2} \sigma_{a}^{2}\right) \\
P(z=1 \mid \underline{x}, D) \\
p(x \mid a)=N(x ;-\lambda a, 1) \\
P(x \leq 0)=\int_{a} P(x \leq 0 \mid a) N\left(a ; \mu_{a}, \sigma_{a}^{2}\right) d a \\
\\
=\int_{\underline{w}}\left(\frac{1}{1+e^{-\underline{w}^{T} \underline{x}}}\right) N\left(\underline{w} ; \underline{w^{*}}, H^{-1}\right) d \underline{w} \\
=\int_{a} \Phi(\lambda a) N\left(a ; \mu_{a}, \sigma_{a}^{2}\right) d a=\Phi\left(\frac{\mu_{a}}{\left(\sigma_{a}^{2}+8 / \pi\right)^{1 / 2}}\right)
\end{array}
\end{aligned}
$$

- So,

$$
\text { Let } a=\underline{w}^{T} \underline{x}=\underline{x}^{T} \underline{w} \Rightarrow a=N(a ; \underline{x}_{\mu_{a}}^{\underline{w}_{a}^{*}}, \underbrace{x^{T}}_{\sigma_{a}^{2}} H^{-1} \underline{x}) . S o, P(z=1 \mid \underline{x}, D) \approx \int_{a}^{\left(\frac{1}{\left(1+e^{-a}\right.}\right)} N\left(a ; \mu_{a(a)}, \sigma_{a}^{2}\right) d a
$$

$$
\approx \int_{a} \Phi(\sqrt{\pi / 8} a) N\left(a ; \mu_{a}, \sigma_{a}^{2}\right) d a=\Phi\left(\frac{\mu_{a}}{\left(\sigma_{a}^{2}+8 / \pi\right)^{1 / 2}}\right) \approx g\left(\frac{\mu_{a}}{\left(1+\pi \sigma_{a}^{2} / 8\right)^{1 / 2}}\right)
$$



## Non-parametric Learning of Discriminant Weights

Consider linear discriminants

$$
y\left(\underline{x}, \underline{w}_{k}\right)=\underline{w}_{k}{ }^{T} \underline{x}-w_{k 0}(\text { or }) y\left(\underline{x}, \underline{w}_{k}\right)=\sum_{j=1}^{M} w_{k j} \varphi_{k j}(\underline{x})-w_{k 0}
$$

Idea: Iteratively change wh


## Two Class Case

Let us consider two class case first
Select class 1 if $\underline{w}_{1}^{T} \underline{x}>\underline{w}_{2}^{T} \underline{x} ; \underline{x}=\left(-1 x_{1} x_{2} \cdots \cdots x_{p}\right)$

$$
\Rightarrow\left(\underline{w}_{1}-\underline{w}_{2}\right)^{T} \underline{x}>0 \Rightarrow \underline{w}^{T} \underline{x}>0
$$

For simplicity write $\quad y(\underline{x}, \underline{w})=\underline{w}^{T} \underline{x} \quad \begin{aligned} & \underline{w}^{T}=\left(w_{0} w_{1} \cdots \cdots w_{p}\right) \\ & \underline{x}=\left(-1 x_{1} x_{2} \cdots \cdots x_{p}\right)\end{aligned}$
The output $g=\left\{\begin{array}{l}1 \text { if } y>0 \\ 0 \text { if } y<0\end{array} \Rightarrow \begin{array}{r}g=I(y>0) \\ =I\left(\underline{w}^{T} \underline{x}>0\right)\end{array}\right.$

- Problem: Want to find a $\underline{w}$ э

$$
\begin{array}{|lr}
\underline{w}^{T} \underline{x}>0 \text { if } \underline{x} \text { belongs to class } 1 & \Rightarrow \theta<90^{\circ} \text { (acute angle) } \\
\underline{w}^{T} \underline{x}<0 \text { if } \underline{x} \text { belongs to class } 2 & \Rightarrow \theta>90^{\circ} \text { (obtuse angle) } \\
& \text { since } \underline{w}^{T} \underline{x}=\|w\| \underline{x} \| \cos \theta
\end{array}
$$

## Linear Programming (LP)Formulation-1

- Alternate problem

For class 2, let $\quad \underline{x}^{i}=-\underline{x}^{i} \quad i=n_{1}+1, \cdots \cdots, N \Rightarrow \underline{w}^{T} \underline{x}^{i}>0 \quad \forall i$
In fact, one can find $\underline{w} \ni \underline{w}^{T} \underline{x}^{i}>b_{i} \quad \forall i$
We have error if $\underline{w}^{T} \underline{x}^{i}<b_{i}$ or $b_{i}-\underline{w}^{T} \underline{x}^{i}>0$

- Why not minimize a measure of errors?

$$
\left.\begin{array}{rl}
\min _{\underline{w}} & \sum_{i=1}^{N} \max \left(0, b_{i}-\underline{w}^{T} \underline{x}^{i}\right) \\
\Rightarrow \min _{\underline{w}, \underline{u}} & \sum_{i=1}^{N} u_{i} \\
& u_{i} \geq 0 \\
\quad u_{i}-b_{i}+\underline{w}^{T} \underline{x}^{i} \geq 0 \quad i=1,2, \cdots \cdots, N
\end{array}\right\} \quad \mathrm{LP}
$$

Since $\underline{w}$ is unrestricted in sign, define $\underline{w}=\underline{w}^{+}-\underline{w}^{-}, \quad \underline{w}^{+} \geq \underline{0}, \underline{w}^{-} \geq \underline{0}$

## LP Formulation-2

Let

$$
\underline{u}_{a}=\left[\begin{array}{c}
\underline{u}^{+} \\
\underline{w}^{+} \\
\underline{w}^{-}
\end{array}\right]
$$

Then the LP becomes:

$$
\left.\begin{array}{ll} 
& \min \underline{c}^{T} \underline{u}_{a} \\
\text { s.t. } & A \underline{u}_{a} \geq \underline{b} \\
& \underline{u}_{a} \geq \underline{0}
\end{array}\right\} \quad \text { if } \sum_{i=1}^{N} u_{i}=0 \Rightarrow \text { Linearly separable }
$$

Will get an answer even if it is not

Initial Basic Feasible Solution: $u_{i}=b_{i}, \underline{w}^{+}=\underline{w}^{-}=\underline{0}$

## Dual Problem

－Dual of the LP problem

$$
\begin{gathered}
\max \underline{\lambda}^{T} \underline{b} \\
\text { s.t. } A^{T} \underline{\lambda} \leq \underline{c} \underline{i=1} \lambda_{i} b_{i} \\
\underline{\lambda} \geq \underline{0}
\end{gathered} \quad \Rightarrow \quad 0 \leq \lambda_{i} \leq 1 .
$$

This is also too much work！In 1950＇s，researchers were interested in finding iterative solutions．

## Decision-Based Learning - 1

$\square$ Rosenblatt's Perceptron learning rule $\sim$ Decision-based learning Note: Class 2 is scaled so that $\underline{x}^{i}=-\underline{x}^{i}$ Idea: If $\quad \underline{w}^{(n)^{T}} \underline{x}^{n}>0$ at iteration $n$ then do nothing

$$
\text { else } \underline{w}^{(n+1)}=\underline{w}^{(n)}+\eta \underline{x}^{\mathrm{n}} \quad \eta>0
$$

Does it make sense?

$$
\underline{w}^{(n+1) T} \underline{x}^{n}=\underline{w}^{(n) T} \underline{x}^{n}+\eta \underline{x}^{n T} \underline{x}^{n}>\underline{w}^{(n) T} \underline{x}^{n} \quad \Rightarrow \text { make it less negative }
$$

- Perceptron Algorithm ~ Supervisory learning (Reinforcement learning)

$$
\underline{w}=\underline{0}
$$

Do until no misclassification (or misclassification error becomes constant)
Do $n=1,2, \cdots \cdots, N$
If $\quad \underline{w}^{T} \underline{x}^{n}<0$

$$
\underline{w} \leftarrow \underline{w}+\eta \underline{x}^{n}
$$

End If
End Do
End Do

## Decision Based Learning - 2



* Note: When you do not scale class 2 features, $\underline{w} \leftarrow \underline{w}+\eta\left(z^{n}-g^{n}\right) \underline{x}^{n}$

$$
z^{n}= \begin{cases}1 & \text { if class } 1 \\ 0 & \text { otherwise }\end{cases}
$$

$\Rightarrow \underline{x}^{n}$ belongs to class $1\left(z^{n}=1\right)$ but assigned to class $2\left(g^{n}=0\right)$, then

$$
\underline{w} \leftarrow \underline{w}+\eta \underline{x}^{n} \quad \text { (reinforcement learning) }
$$

$\underline{x}^{n}$ belongs to class $2\left(z^{n}=0\right)$ but assigned to class $1\left(g^{n}=1\right)$, then

$$
\left.\underline{w} \leftarrow \underline{w}-\eta \underline{x}^{n} \quad \text { (antireinforcement learning }\right)
$$

- You see the reason why we scaled $\underline{x}^{n}=-\underline{x}^{n}$ for class 2


## Linear Perceptron Convergence Theorem-1

## $\square$ Linear Perceptron Convergence Theorem

If classes are linearly separable, i.e., if there exists a vector $\underline{w}^{k}$ and a number $b>0 \quad \ni \underline{w}^{T} \underline{x}^{n} \geq b \quad \forall n$ (when class 2 features are scaled), then the perceptron learning rule converges in finite number of iterations.
Proof 1:

$$
\begin{align*}
& \cos (\underline{u}, \underline{v})=\frac{\underline{u}^{T} \underline{v}}{\|\underline{u}\|_{2}\|\underline{v}\|_{2}} \leq 1 \\
& \underline{w}^{* T} \underline{w}^{(n+1)}=\underline{w}^{* T}\left(\underline{w}^{(n)}+\eta \underline{x}^{n}\right) \\
& \underline{w}^{* T} \underline{w}^{(n+1)} \geq \underline{w}^{* T} \underline{w}^{(n)}+\eta b \quad \text { Assuming } \underline{w}^{(0)}=\underline{0} \\
& \Rightarrow \underline{w}^{* T} \underline{w}^{(n+1)} \geq \eta n b \quad \ldots \ldots \ldots . \tag{1}
\end{align*}
$$

## Linear Perceptron Convergence Theorem－2

Proof 1 continued：

$$
\begin{aligned}
& \left\|\underline{w}^{(n+1)}\right\|^{2}=\left\|\underline{w}^{(n)}+\eta \underline{x}^{n}\right\|^{2}=\left\|\underline{w}^{(n)}\right\|^{2}+2 \eta \underline{w}^{(n) T} \underline{x}^{n}+\eta^{2}\left\|\underline{x}^{n}\right\|^{2}<\left\|\underline{w}^{(n)}\right\|^{2}+\eta^{2}\left\|\underline{x}^{n}\right\|^{2} \\
& \Rightarrow\left\|\underline{w}^{(n+1)}\right\|^{2}<n \eta^{2} \beta^{2} \text { where } \beta^{2}=\max _{i}\left\|\underline{x}^{i}\right\|^{2} \\
& \quad \Rightarrow \cos \left(\underline{w}^{(n+1)}, \underline{w}^{*}\right)=\frac{\underline{w}^{* T} \underline{w}^{(n+1)}}{\left\|\underline{w}^{*}\right\|_{2}\left\|\underline{w}^{(n+1)}\right\|_{2}} \geq \frac{\eta n b}{\sqrt{n \eta^{2} \beta^{2}}\left\|\underline{w}^{*}\right\|_{2}} \leq 1
\end{aligned}
$$

$$
\sqrt{n} b \leq\left\|w^{*}\right\|_{2} \beta
$$

$$
n \leq \frac{\left\|w^{*}\right\|^{2} \beta^{2}}{b^{2}}
$$

## Linear Perceptron Convergence Theorem-3

Proof 2:

$$
\begin{aligned}
\left\|\underline{w}^{(n+1)}-\underline{w}^{*}\right\|^{2} & =\left\|\underline{w}^{(n)}-\underline{w}^{*}+\eta \underline{\eta}^{n}\right\|^{2} \quad \text { know } \underline{w}^{(n) T} \underline{x}^{n}<0 \\
& =\left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}+2 \eta\left(\underline{w}^{(n)}-\underline{w}^{*}\right)^{T} \underline{x}^{n}+\eta^{2}\left\|\underline{x}^{n}\right\|^{2} \\
& <\left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}-2 \eta \mathrm{~b}+\eta^{2} \beta^{2} \\
& <\left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2} \text { if } \eta<\frac{2 b}{\beta^{2}}
\end{aligned}
$$

o For $\eta$ sufficiently small, linear term dominates the quadratic term.
o Every monotonic sequence has a limit. So, the sequence converges.

## Optimal Learning Rate

$\square$ Optimal Learning Rate $\eta_{\text {opt }}$

$$
\begin{aligned}
\left\|\underline{w}^{(n+1)}-\underline{w}^{*}\right\|^{2} & =\left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}+\eta^{2}\left\|\underline{x}^{n}\right\|^{2}+2 \eta\left(\underline{w}^{(n)}-\underline{w}^{*}\right)^{T} \underline{x}^{n} \\
& =\left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}+\eta^{2}\left\|\underline{x}^{n}\right\|^{2}-2 \eta\left[\left|\underline{w}^{* T} \underline{x}^{n}\right|+\left|\underline{w}^{(n) T} \underline{x}^{n}\right|\right]
\end{aligned}
$$

$$
\sqrt{ } \text { valid when don't scale also }
$$

$$
\frac{\partial}{\partial \eta}=0 \Rightarrow \quad \eta_{\text {opt }}=\frac{\left|\underline{w}^{* T} \underline{x}^{n}\right|+\left|\underline{w}^{(n) T} \underline{x}^{n}\right|}{\left\|\underline{x}^{n}\right\|^{2}}
$$

$$
=\frac{\left(\underline{w}^{*}-\underline{w}^{(n)}\right)^{T} \underline{x}^{n}}{\left\|\underline{x}^{n}\right\|^{2}}
$$

The optimal step size is chosen to move $\underline{w}^{(n+1)}$ as close to $\underline{w}^{*}$ as possible.


But do not know $\underline{w} \underline{w}^{*}$. In Relaxation method, $\underline{w}^{*} \underline{x}^{n}$ is replaced by a small positive number $b$.

$$
\eta_{n}=\frac{b+\left|\underline{w}^{(n) T} \underline{x}^{n}\right|}{\left\|\underline{x}^{n}\right\|^{2}}
$$

$\underline{w}^{(1)}$ is normalized（best method）

$$
\eta_{N}=\frac{2\left|\underline{w}^{(n) T} \underline{x}^{n}\right|}{\left\|\underline{x}^{n}\right\|^{2}}
$$



## Normalized Perceptron-2

Leads to finite convergence because

$$
\begin{aligned}
\left\|\underline{w}^{(n+1)}-\underline{w}^{*}\right\|^{2}= & \left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}-4\left|\underline{w}^{(n) T} \underline{x}^{n}\right|^{2}-4\left|\underline{w}^{(n) T} \underline{x}^{n}\right| \underline{w}^{* T} \underline{x}^{n} \\
& +4\left|\underline{w}^{(n) T} \underline{x}^{n}\right|^{2} \\
= & \left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}-4\left|\underline{w}^{(n) T} \underline{x}^{n}\right|\left|\underline{w}^{* T} \underline{x}^{n}\right| \\
< & \left\|\underline{w}^{(n)}-\underline{w}^{*}\right\|^{2}
\end{aligned}
$$

Also note that

$$
\left\|\underline{w}^{(n+1)}\right\|^{2}=\left\|\underline{w}^{(n)}+\eta_{N} \underline{x}^{n}\right\|^{2}=\left\|\underline{w}^{(n)}\right\|^{2}+2 \eta_{N} \underbrace{\underbrace{(n) T} \underline{x}^{n}}_{<0}+\eta_{N}{ }^{2}\left\|\underline{x}^{n}\right\|^{2}=\left\|\underline{w}^{(n)}\right\|^{2}
$$

## Extension to Multiple Classes

$\square$ Extension to Multiple Classes $\{1,2, \cdots \cdots \cdots, C\}$

$$
k=\arg \max _{i}\left(\underline{w}_{i}^{T} \underline{x}\right)
$$

$\square$ Learning rule
Suppose at present a pattern $\underline{x}^{n}$ belongs to class $j$ If $k=j$ do nothing

Else

$$
\begin{array}{ll}
\underline{w}_{j}^{(n+1)}=\underline{w}_{j}^{(n)}+\eta \underline{x}^{n} & \text { so that } \underline{w}_{j}^{(n+1) T} \underline{x}^{n} \uparrow \\
\underline{w}_{k}^{(n+1)}=\underline{w}_{k}^{(n)}-\eta \underline{x}^{n} & \text { so that } \underline{w}_{k}^{(n+1) T} \underline{x}^{n} \downarrow
\end{array}
$$

End If
Note: Leave other classes alone.
Note: Sometimes check $\underline{w}_{k}^{T} \underline{x}^{n}>\underline{w}_{i}^{T} \underline{x}^{n}+\varepsilon \quad \forall i \neq k$ for patterns belonging to class $k$, where $\varepsilon$ is called a vigilance parameter.

## Finite Convergence

$\underline{x}^{n}$ belongs to class $j$

$$
\underline{w}_{j}{ }^{* T} \underline{x}^{n}>\underline{w}_{k}{ }^{* T} \underline{x}^{n} \forall k \neq j
$$

This set of inequalities can be expressed as
$\underline{q}^{T}=\left[\underline{w}_{1}{ }^{* T} \underline{w}_{2}{ }^{*}{ }^{*} \cdots \cdots \cdots \underline{w}_{C}{ }^{* T}\right] \quad C(p+1)$ dimensional vector
$\underline{p}_{j k}{ }^{T}=\left[0^{T} \cdots\left(\underline{x}^{n}\right)^{T} \cdots-\left(\underline{x}^{n}\right)^{T} \cdots 0^{T}\right](p+1)$ dimensional vector
$\Rightarrow$ need $\quad \underline{q}^{T} \underline{p}_{j k}>0 \forall j \neq k$
If $\quad \underline{q}^{T} \underline{p}_{j k}<0$ for any $j \neq k \Rightarrow$ teacher says it's wrong
$\left.\begin{array}{l}\underline{w}_{j}^{(n+1)}=\underline{w}_{j}^{(n)}+\eta \underline{x}^{n} \\ \underline{w}_{k}^{(n+1)}=\underline{w}_{k}^{(n)}-\eta \underline{x}^{n}\end{array}\right\} \Rightarrow \underline{q}^{n e w}=\underline{q}^{\text {old }}+\eta(z-g) \underline{p}_{j k}$
Proof similar to binary case. See S. Y. Kung's Book
"Digital Neural Networks" for proof of convergence
o Only weights $j$ and $k$ are updated

## Key Properties of Perceptron

- Key properties of Perceptron:
- It employs distributed decision-based credit assignment
- Update weights only when misclassification occurs
- Distributed and localized: reinforce correct class, penalize wrong decision class
- Update depends on $y(\underline{x}, \underline{w})$

$$
\Delta \underline{w}= \pm \eta \nabla_{\underline{w}} y(\underline{x}, \underline{w})
$$

when $y(\underline{x}, \underline{w})=\underline{w}^{T} \underline{x} \Rightarrow \Delta \underline{w}= \pm \eta \underline{x}$
$\Rightarrow$ Perceptron learning rule (Incremental gradient, Stochastic gradient)

- Fuzzy updates
- Suppose $\underline{x}$ belongs to class $j$

In fuzzy update, one finds the challenger to correct class $j$

$$
k=\arg \max _{i \neq j} y\left(\underline{x}, \underline{w}_{i}\right)
$$

* Find a measure of misclassification

$$
d=-y\left(\underline{x}, \underline{w}_{j}\right)+y\left(\underline{x}, \underline{w}_{k}\right) \quad\left[\begin{array}{l}
r=\infty \Rightarrow \\
d=-y\left(\underline{x}, \underline{w}_{j}\right)+y\left(\underline{x}, \underline{w}_{k}\right)
\end{array}\right.
$$

more generally

$$
d=-y\left(\underline{x}, \underline{w}_{j}\right)+\left[\frac{1}{C-1} \sum_{i \neq j}\left[y\left(\underline{x}, \underline{w}_{i}\right)\right]^{r}\right]^{1 / r}, y>0
$$

$$
d>0 \quad \Rightarrow \text { misclassification }
$$

$$
d<0 \quad \Rightarrow \text { correct classification }
$$

## Fuzzy Updates - 2

Define a penalty function

- Logistic

$$
g(d)=\frac{1}{1+e^{-d / \xi}} \quad \begin{gathered}
\text { needed } \\
\xi \downarrow 0 \Rightarrow \text { step }
\end{gathered}
$$

- Linear

$$
g(d)=\left\{\begin{array}{lll}
0 & d \leq-\xi / 2 & g(d) \\
(d+\xi / 2) / \xi & -\xi / 2<d<\xi / 2 & \\
\hdashline & d \geq \xi / 2 & -\frac{\xi}{2} \\
1 & & \frac{\xi}{2}
\end{array}\right.
$$

$$
\begin{aligned}
& \underline{w}_{j}^{(n+1)}=\underline{w}_{j}^{(n)}+\eta g^{\prime}(d) \nabla y\left(\underline{x}^{n}, \underline{w}_{j}\right) \\
& \underline{w}_{k}^{(n+1)}=\underline{w}_{k}^{(n)}-\eta g^{\prime}(d) \nabla y\left(\underline{x}^{n}, \underline{w}_{k}\right)
\end{aligned}
$$

$\checkmark$ Logistic:

$$
g^{\prime}(d)=\frac{1}{\xi} g(1-g)
$$

$\checkmark$ Linear:
$g^{\prime}(d)=\left\{\begin{array}{lc|}0 & d \leq-\xi / 2 \\ 1 & -\xi / 2<d<\xi / 2 \\ 0 & \mathrm{~d} \geq \xi / 2\end{array}\right.$

Start with large $\xi$ and progressively reduce it.

## LVQ as Perceptron Update

- Another prominent example is the LVQ algorithm of Kohonen for supervised clustering

$$
y\left(\underline{x}, \underline{w}_{i}\right)=-\frac{\left\|\underline{x}-\underline{w}_{i}\right\|^{2}}{2} \quad i=1,2, \cdots \cdots, C
$$

Assign pattern $\underline{x}$ to ${ }^{2}$ class $k$ if

$$
\begin{gathered}
k=\arg \max _{i}-\frac{\left\|\underline{x}-\underline{w}_{i}\right\|^{2}}{2}=\arg \min _{i} \frac{\left\|\underline{x}-\underline{w}_{i}\right\|^{2}}{2} \\
\nabla_{w_{i}} y\left(\underline{x}, \underline{w}_{i}\right)=-\left(\underline{x}-\underline{w}_{i}\right) \text { for the min function }
\end{gathered}
$$

Suppose $\underline{x}^{n}$ belongs to class $j$, but the decision is class $k$

$$
\begin{array}{ll}
\underline{w}_{j}^{(n+1)}=\underline{w}_{j}^{(n)}+\eta\left(\underline{x}^{n}-\underline{w}_{j}^{(n)}\right) & \text { (reinforcement learning) } \\
\underline{w}_{k}^{(n+1)}=\underline{w}_{k}^{(n)}-\eta\left(\underline{x}^{n}-\underline{w}_{k}^{(n)}\right) & \text { (antireinforcement learning) }
\end{array}
$$

## Capacity Questions

- How many random patterns a Perceptron with $p$ inputs can learn reliably in a 2 class case? $\approx 2 p$

Suppose have $N$ patterns and randomly assign them to one of two classes.
They are linearly separable with probability $P(N, p)$

$$
P(N, p)= \begin{cases}1 & N \leq p+1 \\ \frac{2}{2^{N}} \sum_{i=0}^{p}\binom{N-1}{i} & N>p+1 \\ \approx \Phi\left(\frac{2 p-N}{\sqrt{N}}\right) & \text { for large } N \\ \hline\end{cases}
$$

$$
\begin{aligned}
& \begin{array}{l}
N=4 ; p=2 \\
\Rightarrow P(N, p)= \\
= \\
=\frac{1}{8}(1+3+3)
\end{array} \\
& \text { XOR Pattern cannot be correctly } \\
& \text { classified by a Perceptron }
\end{aligned}
$$

## Decision Trees and Instance-based Classifiers

- Decision Trees and Instance-based Classifiers (e.g., K nearest neighbor classifier)
- Common learning paradigms in AI and statistics
- Idea: Exploit regularities or features in observations and use them to predict on the basis of previously encountered situations
- Decision trees ...... Make shared properties explicit
- Instance-based ...... Find distance from the new and stored cases $\Rightarrow$ Nearest neighbor $\left\{I^{1}, I^{2}, \cdots \cdots, I^{N}\right\}$
o Set of instances, examples, cases
o Each instance $I^{k}$ is represented by $\left(x_{1}, x_{2}, \cdots \cdots, x_{p}\right)_{k}=\underline{x}^{k}$ attributes or features (continuous /discrete) and class $z^{k}$ Case-based reasoning


## Illustrative Example

| Case \# | Outlook | Temp ( ${ }^{\circ} \mathrm{F}$ ) | Humidity <br> $(\%)$ | Windy <br> (Taks <br> umbrella or <br> Not) |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 1. | Rain | 70 | 96 | False | Yes |
| 2. | Sunny | 80 | 90 | True | No |
| 3. | Overcast | 64 | 65 | True | Yes |
| 4. | Sunny | 75 | 70 | True | Yes |
| 5. | Sunny | 85 | 85 | False | No |
| 6. | Sunny | 72 | 95 | False | No |
| 7. | Rain | 75 | 80 | False | Yes |
| 8. | Ounny | 69 | 70 | False | Yes |
| 9. | Rain | 65 | 78 | False | Yes |
| 10. | Overcast | 72 | 70 | True | No |
| 11. | Overcast | 81 | 90 | True | Yes |
| 12. | Rain | 68 | 75 | False | Yes |
| 13. | Rain | 71 | 80 | False | Yes |
| 14. |  | 80 | True | No |  |

## Partitioning Data based on Attribute Outcomes



Outlook Temp Humidity Windy (Discrete)

$$
\begin{aligned}
& x_{1} \in\{\text { Sunny, Overcast, Rain }\} \\
& x_{2} \in R \\
& x_{3} \in R \\
& x_{4} \in\{\text { true, false }\}
\end{aligned}
$$

$$
\text { class } \in\{\text { yes, no }\} \quad \sim \quad \text { binary } Z=\left\{z_{1} z_{2}\right\}
$$

Take umbrella or not

Decision tree divides the description (feature) space into regions, each associated with one of the classes.


## History of Decision Trees

## -1. Decision Trees:

- Developed since $1960^{5}$ and popularized by Brieman and Quinlan CART C 4.5 Gradient Boosting Random Forests
(Classification and Regression trees)
- Have been used in industrial applications, particularly in diagnosis and control, quality control
- Invariant to scaling; can handle large datasets; easily ignore redundant variables; handle missing variables easily through surrogate splits; Easy to interpret
- Nearly half of the data mining competitions are won by using some variants of tree ensemble methods: Boosting, Random Forests, Bagging
- Used in statistics, pattern recognition, decision theory, signal processing, machine learning and artificial neural networks
- Three uses: Data description (compression, rules), classification, and generalization
- Survey: Srirama K. Murthy: "Automatic Construction of Decision Trees From Data: A Multi-disciplinary Survey," pp 1-49. Kulwer Academic Publishers. Data Mining and Knowledge Disocery 2 (4): 345-389 (1998)
- Also, look at my papers on sequential fault diagnosis since 1990. Primarily in IEEE T-SMC.


## Decision Tree for the Weather Problem



- Case: Outlook=sunny, Temp=82, Humidity $=85$, Windy $=$ True $\Rightarrow$ no
- This is a case of exact partitioning where we grow the tree until all the cases are classified correctly.
- In noisy situations, distributions of classes overlap. Here, we would like to stop growing the tree or prune it after constructing.



## Defining "Attribute Tests" - 1

Method for constructing trees ---- want to generate compact trees

- All employ recursive divide-and-conquer algorithm
- How to select tests? ....... Splitting rules
- Most tree construction techniques use single attribute (feature) tests
- Discrete attribute $x_{i}: x_{i} \in\left\{v_{1} v_{2} \cdots \cdots v_{m}\right\}$
$x_{i}=$ ? with $m$ outcomes
$x_{i}=v_{j}, m$ binary tests with outcomes true or false
- Continuous attribute $x_{i}$
$x_{i} \leq t$ (with outcomes true or false) for some constant $t$
Suppose have observations, $n_{1}<n_{2}<\cdots \cdots<n_{N}$
Need to select thresholds $t \in\left[n_{i}, n_{i+1}\right) ; e . g ., t=\left(n_{i}+n_{i+1}\right) / 2$


## Defining "Attribute Tests" - 2

- Select $t \ni$ mutual information $\operatorname{IG}(z \mid x, t)$ is maximum

$$
\begin{aligned}
& I G(z \mid x, t)=H(z)-H(z \mid x, t) \\
& H(z)=-\sum_{j=1}^{c} P(z=j) \log _{2} P(z=j) \\
& H(z \mid x, t)=P(x<t) H(z \mid x<t)+P(x \geq t) H(z \mid x \geq t)
\end{aligned}
$$

- Complex tests consider subsets or linear combination

$$
\underline{w}^{T} \underline{x}-w_{0} \geq 0 \quad \text { true or false }
$$

- Each test could be a classifier (see Error Correcting Codes later)


## Method for Constructing Trees



If the class of node decisions do not match the form of the training data，a very complicated decision tree will result，as shown at the top．However，if proper decision forms （＂Tests＂）are used，the tree can be quite simple．

## Criteria for Splitting Tests - 1

- Idea
- Suppose apply a test $T$ that partitions $D$ into $D_{1} D_{2}, \cdots \cdots, D_{m}$
- Want to select a test $T$ such that the child nodes $D_{1} \cdots \cdots D_{m}$ are "purer" than their parent $D$.
- Measure of impurity is such that it is zero if $P_{j}$ is concentrated on one class and is maximal if $P_{j}=1 / C$ is uniform. Recall the definition of entropy.
- So, stop building the tree if all cases belong to a single class
- Entropy: A measure of impurity (uncertainty)

$$
H(D)=-\sum_{j=1}^{C} P_{j} \log _{2} P_{j} \quad H(D)=\left\{\begin{array}{cl}
0 & \text { if } P_{j}=1 \\
\log _{2} C & \text { if } P_{j}=\frac{1}{C}
\end{array}\right.
$$

## Criteria for Splitting Tests - 2

- Gini index

$$
\begin{aligned}
& G(D)=1-\sum_{j=1}^{c} P_{j}^{2}=1-\underline{P}^{T} \underline{P}=\sum_{i} P_{i} \sum_{j \neq i} P_{j}=\sum_{i} P_{i}\left(1-P_{i}\right) \\
& G(D) \text { and } H(D) \text { are concave. } \\
& G(D)=\left\{\begin{array}{cc}
0 & \text { if } P_{j}=1 \\
1-\frac{1}{C} & \text { if } P_{j}=\frac{1}{C}
\end{array}\right.
\end{aligned}
$$

- Gini index $=$ Expected error rate if the class label is chosen randomly from the class distribution at the node

$$
G(D)=\sum_{i=1}^{C} P_{i} \sum_{j \neq i} P_{j}=1-\underline{P}^{T} \underline{P}
$$

- Variance Impurity (in two class case): $\mathrm{V}(D)=P_{1}\left(1-P_{1}\right)$. Maximum when $P_{1}=1 / 2$.


## Criteria for Splitting Tests－ 3

－Gini index is a generalization of variance impurity to more than two classes
－Misclassification Impurity（MAP Classification Error）

$$
\begin{aligned}
& M(D)=1-\max _{j} P_{j} \\
& M(D)=\left\{\begin{array}{cc}
0 & \text { if } P_{j}=1 \\
1-\frac{1}{C} & \text { if } P_{j}=\frac{1}{C}
\end{array}\right.
\end{aligned}
$$

So，decrease in average impurity $=$ Gain in purity

$$
G P(D \mid T)=G(D)-\sum_{i=1}^{m} \frac{\left|D_{i}\right|}{D \mid} G\left(D_{i}\right)
$$



FIGURE For the two-category case, the impurity functions peak at equal class frequencies and the variance and the Gini impurity functions are identical. The entropy, variance, Gini, and misclassification impurities (given by Eqs. 1-4, respectively) have been adjusted in scale and offset to facilitate comparison here; such scale and offset do not directly affect learning or classification.

## Criteria for Splitting Tests - 4

- Information gain:

$$
I G(D, T)=H(D)-\sum_{i=1}^{m} \frac{\left|D_{i}\right|}{|D|} H\left(D_{i}\right)
$$

Stop building the tree if $\operatorname{IG}(D, T)=0$ for all $T$

Mutual Information

$$
\underbrace{}_{H(D \mid T)}
$$

- Another way:

$$
\begin{gathered}
H(D, T)=H(D)+H(T \mid D) \\
=H(T)+H(D \mid T) \\
\Rightarrow H(D)-H(D \mid T)=H(T)-H(T \mid D) \\
\Rightarrow I G(D, T)=H(D)-H(D \mid T) \\
=H(T)-H(T \mid D) \\
=H(T) \text { if } T \text { is perfect. } \\
\text { Recall } H(X \mid Y)=-\sum_{j} P\left(y_{j}\right) H\left(X \mid Y=y_{j}\right)=-\sum_{j} P\left(y_{j}\right) \sum_{i} P\left(x_{i} \mid y_{j}\right) \log _{2} P\left(x_{i} \mid y_{j}\right)
\end{gathered}
$$

## Example

## Attributes Labels

| $\boldsymbol{t}_{1}$ | $\boldsymbol{t}_{2}$ | $\boldsymbol{z}$ | Count |
| :--- | :--- | :--- | :--- |
| $T$ | $T$ | + | 2 |
| $T$ | $F$ | + | 2 |
| $F$ | $T$ | - | 5 |
| $F$ | $F$ | + | 1 |

Which one do we choose $t_{1}$ or $t_{2}$ ?
Prior: $z=+$ or - equally likely
$P\left(t_{1}=T\right)=0.4 ; P\left(t_{2}=T\right)=0.7$
$P\left(t_{1}=T \mid z=+\right)=0.8 ; P\left(t_{1}=T \mid z=-\right)=0$

$$
\begin{aligned}
I G\left(t_{1}, z\right) & =H(z)-H\left(z \mid t_{1}\right)=H\left(t_{1}\right)-H\left(t_{1} \mid z\right) \\
H(z) & =-(1 / 2) \log _{2}(1 / 2)-(1 / 2) \log _{2}(1 / 2)=1 \\
H\left(z \mid t_{1}\right) & =P\left(t_{1}=T\right) H\left(z \mid t_{1}=T\right)+P\left(t_{1}=F\right) H\left(z \mid t_{1}=F\right) \\
& =-4 / 10\left(1 \log _{2} 1+0 \log _{2} 0\right)-6 / 10\left(5 / 6 \log _{2} 5 / 6+1 / 6 \log _{2} 1 / 6\right) \\
& =0.39
\end{aligned}
$$

Information gain $\operatorname{IG}\left(t_{1}, z\right)=1-0.39=0.61$; Similarly $\operatorname{IG}\left(t_{2}, z\right)=0.12$
$H\left(t_{1}\right)=-(0.4) \log _{2}(0.4)-(0.6) \log _{2}(0.6)=0.971$
$H\left(t_{1} \mid z\right)=P(\mathrm{z}=+) H\left(t_{1} \mid z=+\right)+P(z=-) H\left(t_{1} \mid z=-\right)=-0.5\left(0.8 \log _{2} 0.8+0.2 \log _{2} 0.2\right)$ $=0.361$
$I G\left(t_{1}, z\right)=0.971-0.361=0.61$; Similarly, $I G\left(t_{2}, z\right)=0.12 \Rightarrow t_{1}$ is a better test

## Criteria for Splitting Tests - 5

- Properties:

$$
\left.\begin{array}{c}
G P(\mathbf{D} \mid \mathrm{T}) \geq 0 \\
I G(\mathbf{D}, T) \geq 0
\end{array}\right\} \quad \text { equal to zero if and only if } T \text { provides }
$$

Some authors consider the gain ratio

$$
0 \leq \frac{I G(\mathbf{D} \mid T)}{-\sum_{i=1}^{n} \frac{\left|D_{i}\right|}{\mathbf{D}} \log _{2} \frac{\left|D_{i}\right|}{\mathbf{D}}}=\frac{I G(\mathbf{D} \mid T)}{H(\mathbf{D})}=1-\frac{H(D \mid T)}{H(D)} \leq 1
$$

- Other Measures: G-statistic (related to mutual information), half-split partitioning ("two-ing" rule), linear discriminant splits
- "Deeper" Information- theoretic splits
- Joint Mutual Information (JMI)
- Conditional Mutual Information Maximization (CMIM)


## JMI \& CMIM Criteria

$T_{k}=$ Current Test set used on the path leading to node $k$ Joint Mutual Information (JMI) Criterion: (works well)

$$
\begin{aligned}
j_{k} & =\arg \max _{j} \sum_{l \in T_{k}} I\left(\left\{t_{l}, t_{j}\right\} ; z\right)=\arg \max _{j} \sum_{l \in T_{k}}\left[I\left(t_{l} ; z\right)+I\left(t_{j} ; z \mid t_{l}\right)\right] \\
& =\arg \max _{j} \sum_{l \in T_{k}} I\left(t_{j} ; z \mid t_{l}\right)
\end{aligned}
$$



Conditional Mutual Information Maximization (CMIM) Criterion:

$$
j_{k}=\arg \max _{j} \min _{l \in T_{k}}\left[I\left(t_{j} ; z \mid t_{l}\right)\right]
$$

Binary Test Case:
$P\left(t_{j} \mid z=i\right)=\left\{\begin{array}{c}1-\mu_{i j} \text { for } t_{j}=0 \\ \mu_{i j} \text { for } t_{j}=1\end{array} ; j=1,2, . ., p ; i=0,1,2, . ., C\right.$
$P(z=i)=P_{i} ; i=0,1,2, . ., C ; 0 \Rightarrow$ none of the classes
$P\left(t_{j}=1\right)=q_{j}=\sum_{i=0}^{C} P_{i} \mu_{i j}$

## Binary Test Case

$J M I: j_{k}=\arg \max _{j} \sum_{l \in T_{k}} I\left(t_{j} ; z \mid t_{l}\right)$
Recall $I\left(t_{j} ; z \mid t_{l}\right)=H\left(z \mid t_{l}\right)-H\left(z \mid t_{l}, t_{j}\right)=H\left(z, t_{l}\right)-H\left(t_{l}\right)-H\left(z, t_{l}, t_{j}\right)+H\left(t_{l}, t_{j}\right)$
where $H\left(z, t_{l}\right)=-\sum_{i=0}^{c}\left\{P_{i} \mu_{i l} \log _{2}\left[P_{i} \mu_{i l}\right]+P_{i}\left(1-\mu_{i l}\right) \log _{2}\left[P_{i}\left(1-\mu_{i l}\right)\right]\right\}$

$$
\begin{aligned}
& H\left(t_{l}\right)=-\left[q_{l} \log _{2} q_{l}+\left(1-q_{l}\right) \log _{2}\left(1-q_{l}\right)\right] ; q_{l}=\sum_{i=0}^{c} P_{i} \mu_{i l} \\
& \begin{aligned}
H\left(z, t_{l}, t_{j}\right)=- & \sum_{i=0}^{c}
\end{aligned} \quad\left\{P_{i} \mu_{i l} \mu_{i j} \log _{2}\left[P_{i} \mu_{i l} \mu_{i j}\right]+P_{i}\left(1-\mu_{i l}\right) \mu_{i j} \log _{2}\left[P_{i}\left(1-\mu_{i l}\right) \mu_{i j}\right]\right. \\
& \\
& \quad+\mathrm{P}_{i} \mu_{i l}\left(1-\mu_{i j}\right) \log _{2}\left[P_{i} \mu_{i l}\left(1-\mu_{i j}\right)\right]+P_{i}\left(1-\mu_{i l}\right)\left(1-\mu_{i j}\right) \log _{2}\left[P_{i}\left(1-\mu_{i l}\right)\left(1-\mu_{i j}\right)\right]
\end{aligned}
$$

$$
H\left(t_{l}, t_{j}\right)=-q_{l j}^{(00)} \log _{2} q_{l j}^{(00)}-q_{l j}^{(10)} \log _{2} q_{l j}^{(10)}-q_{l j}^{(01)} \log _{2} q_{l j}^{(01)}-q_{l j}^{(11)} \log _{2} q_{l j}^{(11)}
$$

where $_{l i j}^{(00)}=\sum_{i=0}^{C} P_{i}\left(1-\mu_{i l}\right)\left(1-\mu_{i j}\right) ; q_{l j}^{(10)}=\sum_{i=0}^{C} \mathrm{P}_{i} \mu_{i l}\left(1-\mu_{i j}\right) ; q_{l j}^{(01)}=\sum_{i=0}^{C} P_{i}\left(1-\mu_{i l}\right) \mu_{i j} ; q_{l j}^{(11)}=\sum_{i=0}^{C} P_{i} \mu_{i l} \mu_{i j}$
Simplifying

$$
\begin{aligned}
I\left(t_{j} ; z \mid t_{l}\right)=-\sum_{i=0}^{C}\{ & \left.\frac{p_{i} \mu_{i l}}{q_{l}} \log _{2}\left[\frac{p_{i} \mu_{i l}}{q_{l}}\right]+\frac{p_{i}\left(1-\mu_{i l}\right)}{\left(1-q_{l}\right)} \log _{2}\left[\frac{p_{i}\left(1-\mu_{i l}\right)}{\left(1-q_{l}\right)}\right]\right\} \\
& -\sum_{i=0}^{m}\left\{\frac{p_{i} \mu_{i l} \mu_{i j}}{q_{l j}^{(11)}} \log _{2}\left[\frac{p_{i} \mu_{i l} \mu_{i j}}{q_{l j}^{(11)}}\right]+\frac{p_{i}\left(1-\mu_{i l}\right) \mu_{i j}}{q_{l j}^{(01)}} \log _{2}\left[\frac{p_{i}\left(1-\mu_{i l}\right) \mu_{i j}}{q_{l j}^{(01)}}\right]\right. \\
& +\frac{p_{i} \mu_{i l}\left(1-\mu_{i j}\right)}{q_{l j}^{(10)}} \log _{2}\left[\frac{p_{i} \mu_{i l}\left(1-\mu_{i j}\right)}{q_{l j}^{(10)}}\right]+\frac{p_{i}\left(1-\mu_{i l}\right)\left(1-\mu_{i j}\right)}{q_{l j}^{(00)}} \log _{2}\left[\frac{p_{i}\left(1-\mu_{i l}\right)\left(1-\mu_{i j}\right)}{q_{l j}^{(00)}}\right]
\end{aligned}
$$

－Overfitting：Decision Trees Overfit（low bias and high variance）
－Real－life data is noisy，e．g．，continuous attributes have measurement noise，discrete attributes such as color depend on subjective interpretation，instances are misclassified，and mistakes are made in recording．
－Noisy data leads to larger trees ．．fit the noise in addition to the structure in the data ．．tend to be＂brittle＂or＂sensitive＂．＂High Variability and Low Bias＂．
－Can prevent overfitting by stopping（e．g．，fixed depth）or by pruning back the full tree to an appropriate size．Pruning is preferred．Also， ensemble trees reduce variance．
－Cost－Complexity Pruning：
－Suppose have a tree Z．
－Let $E(\mathbf{Z})$ be the number of classes misclassified out of $N$
－Let $L(\mathbf{Z})$ be the number of leaves of $Z$ Cost complexity of $Z$ is：$\frac{E(\mathbf{Z})}{N}+\alpha L(\mathbf{Z})$


## Pruning the Tree - 2

- Suppose we prune $Z$ by replacing a sub-tree $S$ by a leaf and identifying the most frequent class among the instances from which $S$ was constructed $\Rightarrow$ we now have a tree $\hat{Z}$ with $L(S)-1$ fewer leaves. That is,

$$
L(\hat{Z})=L(Z)-(L(S)-1)
$$

$>$ Suppose the new error rate is

$$
\frac{E(\hat{Z})}{N}
$$

$\Rightarrow$ Prune $S$ if

$$
\frac{E(Z)}{N}+\alpha L(Z)>\frac{E(\hat{Z})}{N}+\alpha L(\hat{Z})
$$

or

$$
\alpha(L(Z)-L(\hat{Z}))>\frac{E(\hat{Z})-E(Z)}{N}
$$

## Pruning the Tree - 3

or, $\quad \alpha>\frac{E(\hat{Z})-E(Z)}{\underbrace{N}_{g(Z, \hat{Z})} \frac{1}{(L(S)-1)}}$
So, if $g(Z, \hat{Z})=\frac{\Delta E}{N} \frac{1}{(L(S)-1)}> \begin{cases}\alpha & \text { prune } \\ \alpha & \text { do not prune }\end{cases}$

- Algorithm
- Set $k=0, \quad Z_{0}=Z$
-     - Set $\alpha=\infty$
- Visit non terminal nodes $t$ in bottom-up order and calculate $E\left(Z_{t}\right)$ and $L\left(Z_{t}\right)$ by summing over the descendant ${ }^{-}$

$$
g(t)=\frac{E\left(Z_{t}\right)-E\left(Z_{k}\right)}{N\left[L\left(Z_{k}\right)-L\left(Z_{t}\right)\right]} \text { and } \alpha=\min (\alpha, g(t))
$$

## Pruning the Tree - 4

- Visit the nodes in the top-down order and prune whenever $g(t)=\alpha$
$-\quad$ Set $k=k+1, \quad \alpha_{k}=\alpha, \quad Z_{k}=Z$
$-\quad$ If $Z$ is non-trivial ( $>1$ node), return to *
- Reduced error pruning :
- Record the number of errors at each leaf node and, for each internal node, the number of errors that would be made if that node were a leaf node $\Rightarrow$ use MAP rule to make decision.
- When all error counts are determined, each internal node is investigated starting from the bottom levels of the tree.


## Pruning the Tree - 5

- Compare the number of errors made by the subtree rooted at that node with the number of errors that would result from changing the node to a leaf. If the latter is not greater than the former, prune the tree.

- Since the total number of errors made by a tree is the sum of errors at leaves, it is clear that the final sub-tree minimizes the number of errors on the pruning set.


## Handling Missing Attribute Values－ 1

－Missing Attribute Values：
－Cases do not have all features
－This can impact decision tree methods at three stages：
－When comparing tests on attributes with different numbers of missing values
－When partitioning set $D$ on the outcomes of a chosen test，since the outcomes for some cases may not be known
－When classifying an unseen instance whose outcome for a test is undetermined．

## Handling Missing Attribute Values - 2

- How to handle missing values?
- Filling in missing values
- Discrete: pick the most frequent value
- Continuous: use mean, median,...
- Estimating test outcomes by some other means
- Backup (surrogate) tests

Test on a different attribute that produces a similar partition of $D$. So, when the value of a tested attribute is not known, the best surrogate split whose outcome is known is used to predict the outcome of the original test.

## Handling Missing Attribute Values - 3

- Treating test outcome probabilistically
- Case: Outlook = ?, Humidity = 85, Windy

| Sunny | Overcast | Rain |
| :---: | :---: | :---: |
| $5 / 14$ | $4 / 14$ | $5 / 14$ |



## Handling Missing Attribute Values - 4

$$
\begin{aligned}
& P(\text { yes } \mid \text { Humidity }=85, \text { Windy })=\sum_{\text {outlook }} P(\text { yes }, \text { outlook } \mid \text { Humidity }=85, \text { windy }) \\
& =P(\text { yes }, \text { outlook }=\text { sunny } \mid \text { Humidity }=85)+P(\text { yes }, \text { outlook }=\text { overcast })+P(\text { yes }, \text { outlook }=\text { rainy } \mid \text { windy }) \\
& =0+\frac{4}{14}+0=\frac{4}{14} \\
& P(\text { no } \mid \text { Humidity }=85, \text { Windy })=\sum_{\text {outlook }} P(\text { no,outlook } \mid \text { Humidity }=85, \text { windy }) \\
& =P(\text { no,outlook }=\text { sunny } \mid \text { Humidity }=85)+P(\text { no }, \text { outlook }=\text { overcast })+P(\text { no }, \text { outlook }=\text { rainy }, \text { windy }) \\
& =\frac{3}{3} \cdot \frac{5}{14}+\frac{2}{2} \cdot \frac{5}{14}=\frac{10}{14} \\
& \Rightarrow \text { prediction: no }
\end{aligned}
$$

## Ensemble Trees: Bagging

- Bagging: Bootstrap Aggregating
- Construct a decision tree from each bootstrap sample ("sample with replacement")
- Average out-of-bag sample error to evaluate performance
- Recall averaging reduces variance


Hastie et al.,"The Elements of Statistical Learning: Data Mining, Inference, and Prediction", Springer (2009)

## Ensemble Trees: Random Forests

Construct a decision tree from each bootstrap sample At each node in choosing the split feature, choose only among $m<p$ features
Create bootstrap samples from the training data


## Random Forests Algorithm

For $b=1$ to $B$
(a) Draw a bootstrap sample $D_{b}$ of size $N$ from the training data.
(b) Grow a random-forest tree to the bootstrapped data, by recursively repeating the following steps at each terminal node of the tree, until the minimum node size $n_{\text {min }}$ is reached.
i. Select $m$ variables at random from the $p$ variables.
ii. Pick the best variable/split-point among the $m$.
iii. Split the node into two child nodes.
end
Output the ensemble of trees.
To make a prediction at a new point $\underline{x}$ we do:

- Classification: $m=p^{1 / 2}$ or $\log _{2} p ; n_{\text {min }}=1$
- Regression: $m=p / 3$;
$n_{\text {min }}=5$
For regression: average the results
For classification: majority vote


## Boosting

Greedy algorithm that successively applies a weak learning algorithm on weighted versions of the data using bounds on the loss function
－Recall for binary classification $(z \in\{-1,1\})$ ，loss function is

$$
L(z, \hat{z})=1-\operatorname{sgn}(z \hat{z})=1-\operatorname{sgn}(\eta)=1-\delta_{z z} \dot{z} ; \eta=z \hat{z}
$$

－Hinge loss used in Support Vector Machines（SVM）

$$
L(z, \hat{z}) \leq L_{h}(z, \hat{z})=[1-z \hat{z}]_{+}
$$

－Exponential loss


Log Loss（Logistic loss）
$L(z, \hat{z}) \leq L_{l}(z, y)=\ln \left(1+e^{-z y}\right) / \ln 2$



## AdaBoost

- At step $m$, minimize
$L_{m}\left(f_{m}\right)=\sum_{n=1}^{N} \exp \left\{-z^{n}\left\{\hat{\bar{m}}_{m-1}\left(\underline{x}^{n}\right)+\alpha_{m} f_{m}\left(\underline{x}^{n}\right)\right]\right\}=\sum_{n=1}^{N} \sum_{n}^{w_{n}^{(m)}} e^{-\alpha_{n} m^{2} f_{m}\left(\underline{\underline{x}}^{n}\right)}$
$w_{n}^{(m)}=\exp \left\{-z^{n} \hat{z}_{m-1}\left(\underline{x}^{n}\right)\right\}=$ weight on data item $n$

$\Rightarrow$ Fit weak classifier $f_{m}$ by minimizing error function $J_{m}=\sum_{n=1}^{N} w_{n}^{(m)}\left(1-\delta_{z^{n} f\left(\underline{x}^{n}\right)}\right)$

$$
\begin{aligned}
& \text { Optimal } \alpha_{m} \text { from } e^{2 \alpha_{m}}=\frac{\sum_{n=1}^{N} w_{n}^{(m)} \delta_{z^{n} f\left(\underline{x}^{n}\right)}}{\sum_{n=1}^{N} w_{n}^{(m)}\left(1-\delta_{z^{n} f\left(\underline{x}^{n}\right)}\right)} \\
& \qquad=\frac{\sum_{n=1}^{N} w_{n}^{(m)} \delta_{z^{n} f\left(\underline{x}^{n}\right)} / \sum_{n=1}^{N} w_{n}^{(m)}}{\sum_{n=1}^{N} w_{n}^{(m)}\left(1-\delta_{z^{n} f\left(\underline{x}^{n}\right)}\right) / \sum_{n=1}^{N} w_{n}^{(m)}}=\frac{1-e_{m}}{e_{m}} \Rightarrow \alpha_{m}=\frac{1}{2} \ln \left(\frac{1-e_{m}}{e_{m}}\right) \\
& \Rightarrow w_{n}^{(m+1)}=w_{n}^{(m)} e^{-\alpha_{m} z^{n} f_{m}\left(\underline{x}^{n}\right)}=\left\{\begin{array}{l}
w_{n}^{(m)} e^{-\alpha_{m}} \text { if } f_{m}\left(\underline{x}^{n}\right)=z^{n} \\
w_{n}^{(m)} e^{\alpha_{m}} \text { if } f_{m}\left(\underline{x}^{n}\right) \neq z^{n}
\end{array}\right. \\
& \hat{z}=\operatorname{sgn}\left[\sum_{m=1}^{M} \alpha_{m} f_{m}\left(\underline{x}^{n}\right)\right] ; M \approx 10-20
\end{aligned}
$$

## AdaBoost in Action



## Gradient Boosting Algorithm

- $L_{2}$ Boosting for regression (recall residual = negative gradient wrt $\hat{z}_{m} \ldots$ Gradient boosting ... similar to iterative improvement of LS)

$$
L\left(z^{n}, \hat{z}_{m}^{n}\left(\underline{x}^{n}\right)\right)=\left(r_{m}^{n}-f_{m}\left(\underline{x}^{n}\right)\right)^{2} ; r_{m}^{n}=z^{n}-\hat{z}_{m-1}^{n}\left(\underline{x}^{n}\right) ; \hat{z}_{m}^{n}\left(\underline{x}^{n}\right)=\hat{z}_{m-1}^{n}\left(\underline{x}^{n}\right)+f_{m}\left(\underline{x}^{n}\right)
$$

- Exponential loss is sensitive to outliers. Instead use logloss ... Logitboost for binary .... Gradient boosting

Preselected network

$$
\begin{aligned}
& L_{m}\left(f_{m}\right)=\sum_{n=1}^{N} \ln \left\{1+\exp \left\{-z^{n}\left[\hat{z}_{m-1}\left(\underline{x}^{n}\right)+f_{m}\left(\underline{x}^{n}\right)\right]\right\}\right\}+\Omega\left(f_{m}\right) \ldots \Omega(.) \text { regularization term } \\
& \approx \sum_{n=1}^{N}\left[\ln \left\{1+\exp \left[-z^{n} \hat{z}_{m-1}\left(\underline{x}^{n}\right)\right]\right\}+f_{m}\left(\underline{x}^{n}\right)\left(\delta_{z^{n}, 1}\left[g\left(\hat{z}_{m-1}\left(\underline{x}^{n}\right)\right)-1\right]+\delta_{z^{n},-1} g\left(\hat{z}_{m-1}\left(\underline{x}^{n}\right)\right)\right)\right. \\
& \left.+\frac{1}{2}\left[f_{m}\left(\underline{x}^{n}\right)\right]^{2} g\left(\hat{z}_{m-1}\left(\underline{x}^{n}\right)\right)\left(1-g\left(\hat{z}_{m-1}\left(\underline{x}^{n}\right)\right)\right)\right]+\Omega\left(f_{m}\right)
\end{aligned}
$$

- Gradient boosting is valid for general loss functions (see Friedman, '99) $L_{m}\left(f_{m}\right)=\sum_{n=1}^{N} L\left(z^{n}, \hat{z}_{m-1}\left(\underline{x}^{n}\right)+f_{m}\left(\underline{x}^{n}\right)\right)+\Omega\left(f_{m}\right) \ldots \Omega($.$) regularization term$

$$
\approx \sum_{n=1}^{N}\left\{L\left(z^{n}, \hat{z}_{m-1}\left(\underline{x}^{n}\right)\right]+g_{n} f_{m}\left(\underline{x}^{n}\right)+\frac{1}{2} h_{n}\left[f_{m}\left(\underline{x}^{n}\right)\right]^{2}\right\}+\Omega\left(f_{m}\right) ; g_{n}=\frac{\partial L\left(z^{n}, \hat{z}_{m-1}\left(\underline{x}^{n}\right)\right)}{\partial \hat{z}_{m-1}\left(\underline{x}^{n}\right)} ; h_{n}=\frac{\partial^{2} L\left(z^{n}, \hat{z}_{m-1}\left(\underline{x}^{n}\right)\right)}{\partial\left[\hat{z}_{m-1}\left(\underline{x}^{n}\right)\right]^{2}}
$$

## Gradient Boosting and Regularization

－Complexity of a tree and regularization
$\Omega\left(f_{m}\right)=\gamma T+\frac{1}{2} \sum_{j=1}^{T} w_{j}^{2} ; T=\#$ of leaves
$w_{j}=$ weight of leaf $j$（e．g．，expected cost of path in the tree leading to leaf $j$ ）
$f_{m}\left(\underline{x}^{n}\right)=w_{j}$ if $\underline{x}^{n}$ is mapped to leaf $j$ ，i．e．，$q\left(\underline{x}^{n}\right)=j ; q$ is the mapping fuunction


$$
\begin{aligned}
& T=3 \\
& \begin{aligned}
\Omega\left(f_{m}\right) & =3 \gamma+\frac{1}{2} \lambda(4+0.01+1) \\
& =3 \gamma+2.505 \lambda
\end{aligned}
\end{aligned}
$$

Figure taken from Tianqi Chen＂Introduction to Boosted Trees，＂October 22， 2014.

## Gradient Boosting Implementation

- Define the data instance set in leaf $j$ as $I_{j}=\left\{n: q\left(\underline{x}^{n}\right)=j\right\}$

$$
\begin{aligned}
L_{m}\left(f_{m}\right) & \approx \sum_{n=1}^{N}\left\{L\left(z^{n}, \hat{z}_{m-1}\left(\underline{x}^{n}\right)\right]+g_{n} f_{m}\left(\underline{x}^{n}\right)+\frac{1}{2} h_{n}\left[f_{m}\left(\underline{x}^{n}\right)\right]^{2}\right\}+\gamma T+\frac{1}{2} \lambda \sum_{j=1}^{T} w_{j}^{2} \\
& =\sum_{j=1}^{T}\{\underbrace{\left(\sum_{n \in I_{j}} g_{n}\right)}_{G_{j}} w_{j}+\frac{1}{2}\left(\sum_{n \in l_{j}}^{H_{j}} h_{n}+\lambda\right) w_{j}^{2}+\gamma\}+\text { constant..... separable in } w_{j} \ldots . \text { but huge possibilities for }\left\{I_{j}\right\}!
\end{aligned}
$$

For known $\left\{I_{j}\right\}$,optimal $w_{j}=-\frac{G_{j}}{H_{j}+\lambda} ; L_{m}\left(f_{m}\right)=\frac{1}{2} \sum_{j=1}^{T}\left(\frac{-G_{j}^{2}}{H_{j}+\lambda}+2 \gamma\right)$

- Greedy construction of tree : Binary split at node $j$ with $G_{L j}, H_{L j}, G_{R j}, H_{R j}$ Gain $=\frac{1}{2}\left[\frac{G_{L j}^{2}}{H_{L j}+\lambda}+\frac{G_{R j}^{2}}{H_{R j}+\lambda}-\frac{G_{j}^{2}}{H_{j}+\lambda}\right]-\gamma ; G_{j}=G_{L j}+G_{R j} ; H_{j}=H_{L j}+H_{R j}$
- Continuous: Left to right linear scan over sorted instance is enough to decide the best split along the feature
- Categorical: one-hot coding and ask the question: Is it category $k$ ?
- Stop if the best gain is negative or recursively prune all the leaf splits with negative gain after the tree is built
- To prevent overfitting, use a shrinkage factor: $\hat{z}_{m}=\hat{z}_{m-1}+\varepsilon f_{m} ; \varepsilon \approx 0.1$



## Error Correcting Output Codes

- Multiclass problems
- $\quad M$ class problem: $M$ two-class tasks or $M(M-1) / 2$ two-class tasks
- Obtain a separate decision tree for each class
- What if more than one tree classifies? Majority rule
- Error Correcting Output Codes:
- Each class: a $p$ bit pattern chosen so that Hamming distance is as large as possible among classes
- Learn each of the $p$ bits via trees $\Rightarrow$ p trees
- Select nearest class in terms of least Hamming distance
- $L_{K}$-Boosted Tree (Friedman, 1999)
- Gradient and diagonal Hessian are easy to compute (see slide 13)
- Gradient boosting algorithm
- Implementation: https://github.com/tqchen/xgboost


## Summary

MMSE Estimate $=$ Conditional Mean
Approximating Posterior Probabilities（Discriminative Learning）
－Linear Discriminants in Equal Covariance Case
－Leads to Logistic Sigmoidal Function
－Learn Weights via NLP
－Single Layer Perceptrons
－Perceptron Learning Theorem
－Relaxation and Normalized Updates
－Extension to Multiple Classes
－Fuzzy Updates
－LVQ as a Perceptron
－Capacity of a Perceptron
－Decision Trees，Bagging，Random Forests，Gradient Boosting

