

Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
 3RD EDITION  
 ETHEM ALPAYDIN  
 © The MIT Press, 2014  
 alpaydin@boun.edu.tr  
 http://www.cmpe.boun.edu.tr/~ethem/i2ml3e

CHAPTER 1:  
**INTRODUCTION**

- ### Big Data
- Widespread use of personal computers and wireless communication leads to "big data"
  - We are both producers and consumers of data
  - Data is not random, it has structure, e.g., customer behavior
  - We need "big theory" to extract that structure from data for
    - Understanding the process
    - Making predictions for the future

- ### Why "Learn" ?
- Machine learning is programming computers to optimize a performance criterion using example data or past experience.
  - There is no need to "learn" to calculate payroll
  - Learning is used when:
    - Human expertise does not exist (navigating on Mars),
    - Humans are unable to explain their expertise (speech recognition)
    - Solution changes in time (routing on a computer network)
    - Solution needs to be adapted to particular cases (user biometrics)


- ### What We Talk About When We Talk About "Learning"
- Learning general models from a data of particular examples
  - Data is cheap and abundant (data warehouses, data marts); knowledge is expensive and scarce.
  - Example in retail: Customer transactions to consumer behavior:
    - People who bought "Blink" also bought "Outliers" (www.amazon.com)
  - Build a model that is a *good and useful approximation* to the data.

- ### Data Mining
- Retail: Market basket analysis, Customer relationship management (CRM)
  - Finance: Credit scoring, fraud detection
  - Manufacturing: Control, robotics, troubleshooting
  - Medicine: Medical diagnosis
  - Telecommunications: Spam filters, intrusion detection
  - Bioinformatics: Motifs, alignment
  - Web mining: Search engines
  - ...

- ### What is Machine Learning?
- Optimize a performance criterion using example data or past experience.
  - Role of Statistics: Inference from a sample
  - Role of Computer science: Efficient algorithms to
    - Solve the optimization problem
    - Representing and evaluating the model for inference

- ### Applications
- Association
  - Supervised Learning
    - Classification
    - Regression
  - Unsupervised Learning
  - Reinforcement Learning

- ### Learning Associations
- Basket analysis:
    - $P(Y | X)$  probability that somebody who buys  $X$  also buys  $Y$  where  $X$  and  $Y$  are products/services.
  - Example:  $P(\text{chips} | \text{beer}) = 0.7$



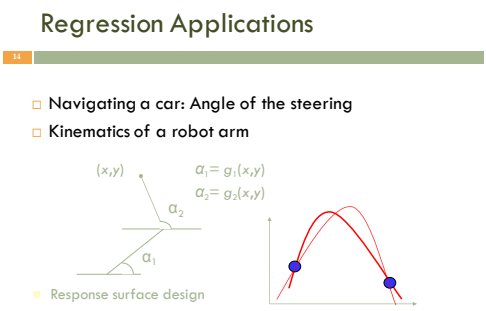

- ### Classification: Applications
- Aka Pattern recognition
  - Face recognition: Pose, lighting, occlusion (glasses, beard), make-up, hair style
  - Character recognition: Different handwriting styles.
  - Speech recognition: Temporal dependency.
  - Medical diagnosis: From symptoms to illnesses
  - Biometrics: Recognition/authentication using physical and/or behavioral characteristics: Face, iris, signature, etc
  - Outlier/novelty detection:

### Face Recognition

Training examples of a person

Test images

ORL dataset, AT&T Laboratories, Cambridge UK



- ### Supervised Learning: Uses
- Prediction of future cases: Use the rule to predict the output for future inputs
  - Knowledge extraction: The rule is easy to understand
  - Compression: The rule is simpler than the data it explains
  - Outlier detection: Exceptions that are not covered by the rule, e.g., fraud

## Unsupervised Learning

- Learning "what normally happens"
- No output
- Clustering: Grouping similar instances
- Example applications
  - Customer segmentation in CRM
  - Image compression: Color quantization
  - Bioinformatics: Learning motifs

## Reinforcement Learning

- Learning a policy: A sequence of outputs
- No supervised output but delayed reward
- Credit assignment problem
- Game playing
- Robot in a maze
- Multiple agents, partial observability, ...

## Resources: Datasets

- UCI Repository: <http://www.ics.uci.edu/~mllearn/MLRepository.html>
- Statlib: <http://lib.stat.cmu.edu/>


## Resources: Journals

- Journal of Machine Learning Research [www.jmlr.org](http://www.jmlr.org)
- Machine Learning
- Neural Computation
- Neural Networks
- IEEE Trans on Neural Networks and Learning Systems
- IEEE Trans on Pattern Analysis and Machine Intelligence
- Journals on Statistics/Data Mining/Signal Processing/Natural Language Processing/Bioinformatics/...

## Resources: Conferences

- International Conference on Machine Learning (ICML)
- European Conference on Machine Learning (ECML)
- Neural Information Processing Systems (NIPS)
- Uncertainty in Artificial Intelligence (UAI)
- Computational Learning Theory (COLT)
- International Conference on Artificial Neural Networks (ICANN)
- International Conference on AI & Statistics (AISTATS)
- International Conference on Pattern Recognition (ICPR)
- ...


## i2ml3e-chap02.pdf




## Learning a Class from Examples

- Class C of a "family car"
  - Prediction: Is car x a family car?
  - Knowledge extraction: What do people expect from a family car?
- Output:
  - Positive (+) and negative (-) examples
- Input representation:
  - $x_1$ : price,  $x_2$ : engine power


## Training set $\mathcal{X}$




## Class C



## Hypothesis class $\mathcal{H}$




## S, G, and the Version Space




## Margin

- Choose  $h$  with largest margin




## VC Dimension

- $N$  points can be labeled in  $2^N$  ways as +/-
- $\mathcal{H}$  shatters  $N$  if there exists  $h \in \mathcal{H}$  consistent for any of these:
  - $VC(\mathcal{H}) = N$



## Probably Approximately Correct (PAC) Learning


- How many training examples  $N$  should we have, such that with probability  $1 - \delta$ ,  $h$  has error at most  $\epsilon$ ? (Blumer et al., 1989)
- Each strip is at most  $\epsilon/4$
- Pr that we miss a strip  $1 - \epsilon/4$
- Pr that  $N$  instances miss a strip  $(1 - \epsilon/4)^N$
- Pr that  $N$  instances miss 4 strips  $4(1 - \epsilon/4)^N$
- $4(1 - \epsilon/4)^N \leq \delta$  and  $(1 - x) \leq \exp(-x)$
- $4\exp(-N\epsilon/4) \leq \delta$  and  $N \geq (4/\epsilon)\log(4/\delta)$



## Noise and Model Complexity

### Use the simpler one because

- Simpler to use (lower computational complexity)
- Easier to train (lower space complexity)
- Easier to explain (more interpretable)
- Generalizes better (lower variance - Occam's razor)



## Triple Trade-Off

- There is a trade-off between three factors (Dietterich, 2003):
  1. Complexity of  $\mathcal{H}$ ,  $c(\mathcal{H})$ ,
  2. Training set size,  $N$ ,
  3. Generalization error,  $E$ , on new data
- As  $N \uparrow$ ,  $E \downarrow$
- As  $c(\mathcal{H}) \uparrow$ , first  $E \downarrow$  and then  $E \uparrow$

## Lecture Slides for INTRODUCTION TO MACHINE LEARNING 3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014  
alpaydin@boun.edu.tr  
http://www.cmp.e.boun.edu.tr/~ethem/i2ml3e

## Bayes' Rule

$$P(C|\mathbf{x}) = \frac{P(C)P(\mathbf{x}|C)}{P(\mathbf{x})}$$


Labels: prior (P(C)), likelihood (P(x|C)), evidence (P(x)), posterior (P(C|x)).

$$P(C=0) + P(C=1) = 1$$

$$P(\mathbf{x}) = P(\mathbf{x}|C=1)P(C=1) + P(\mathbf{x}|C=0)P(C=0)$$

$$P(C=0|\mathbf{x}) + P(C=1|\mathbf{x}) = 1$$

## Multiple Classes, $C_i, i=1, \dots, K$



$$\mathcal{X} = \{\mathbf{x}^t, r^t\}_{t=1}^N$$

$$r^t = \begin{cases} 1 & \text{if } \mathbf{x}^t \in C_i \\ 0 & \text{if } \mathbf{x}^t \in C_j, j \neq i \end{cases}$$


Train hypotheses  $h_i(\mathbf{x}), i=1, \dots, K$ :

$$h_i(\mathbf{x}^t) = \begin{cases} 1 & \text{if } \mathbf{x}^t \in C_i \\ 0 & \text{if } \mathbf{x}^t \in C_j, j \neq i \end{cases}$$

## Cross-Validation

- To estimate generalization error, we need data unseen during training. We split the data as
  - Training set (50%)
  - Validation set (25%)
  - Test (publication) set (25%)
- Resampling when there is few data

## Regression



$$\mathcal{X} = \{\mathbf{x}^t, r^t\}_{t=1}^N$$

$$r^t \in \mathbb{R}$$

$$r^t = f(\mathbf{x}^t) + \varepsilon$$

$$E(g|\mathcal{X}) = \frac{1}{N} \sum_{t=1}^N [r^t - g(\mathbf{x}^t)]^2$$

$$E(w_1, w_0 | \mathcal{X}) = \frac{1}{N} \sum_{t=1}^N [r^t - (w_1 \mathbf{x}^t + w_0)]^2$$

## Dimensions of a Supervised Learner

1. Model:  $g(\mathbf{x}|\theta)$
2. Loss function:  $E(\theta|\mathcal{X}) = \sum_t L(r^t, g(\mathbf{x}^t|\theta))$
3. Optimization procedure:  $\theta^* = \text{argmin}_{\theta} E(\theta|\mathcal{X})$

## Probability and Inference

- Result of tossing a coin is  $\in \{\text{Heads}, \text{Tails}\}$
- Random var  $X \in \{1, 0\}$   
Bernoulli:  $P\{X=1\} = p_0^X (1-p_0)^{(1-X)}$
- Sample:  $\mathcal{X} = \{\mathbf{x}^t\}_{t=1}^N$   
Estimation:  $p_0 = \# \{\text{Heads}\} / \#\{\text{Tosses}\} = \sum_t x^t / N$
- Prediction of next toss:  
Heads if  $p_0 > 1/2$ , Tails otherwise

## Losses and Risks

- Actions:  $\alpha_i$
- Loss of  $\alpha_i$  when the state is  $C_k$ :  $\lambda_{ik}$
- Expected risk (Duda and Hart, 1973)  
 $R(\alpha_i|\mathbf{x}) = \sum_{k=1}^K \lambda_{ik} P(C_k|\mathbf{x})$   
choose  $\alpha_i$  if  $R(\alpha_i|\mathbf{x}) = \min_{\alpha} R(\alpha|\mathbf{x})$

## Model Selection & Generalization

- Learning is an ill-posed problem: data is not sufficient to find a unique solution
- The need for inductive bias: assumptions about  $\mathcal{H}$
- Generalization: How well a model performs on new data
- Overfitting:  $\mathcal{H}$  more complex than  $C$  or  $f$
- Underfitting:  $\mathcal{H}$  less complex than  $C$  or  $f$

1 Mon Sep 24 12:39:54 2018 3

## i2ml3e-chap03.pdf

## Classification

- Credit scoring: Inputs are income and savings. Output is low-risk vs high-risk
- Input:  $\mathbf{x} = [x_1, x_2]^T$ , Output:  $C \in \{0, 1\}$
- Prediction: choose  $\begin{cases} C=1 & \text{if } P(C=1|x_1, x_2) > 0.5 \\ C=0 & \text{otherwise} \end{cases}$   
or choose  $\begin{cases} C=1 & \text{if } P(C=1|x_1, x_2) > P(C=0|x_1, x_2) \\ C=0 & \text{otherwise} \end{cases}$

## Losses and Risks: 0/1 Loss

$$\lambda_{ik} = \begin{cases} 0 & \text{if } i=k \\ 1 & \text{if } i \neq k \end{cases}$$

$$R(\alpha_i|\mathbf{x}) = \sum_{k=1}^K \lambda_{ik} P(C_k|\mathbf{x})$$

$$= \sum_{k \neq i} P(C_k|\mathbf{x})$$

$$= 1 - P(C_i|\mathbf{x})$$

For minimum risk, choose the most probable class

## Losses and Risks: Reject

$$\lambda_{ik} = \begin{cases} 0 & \text{if } i = k \\ \lambda & \text{if } i = K+1, \quad 0 < \lambda < 1 \\ 1 & \text{otherwise} \end{cases}$$


$$R(\alpha_{K+1} | \mathbf{x}) = \sum_{k=1}^K \lambda P(C_k | \mathbf{x}) = \lambda$$

$$R(\alpha_i | \mathbf{x}) = \sum_{k \neq i} P(C_k | \mathbf{x}) = 1 - P(C_i | \mathbf{x})$$


choose  $C_i$  if  $P(C_i | \mathbf{x}) > P(C_k | \mathbf{x}) \forall k \neq i$  and  $P(C_i | \mathbf{x}) > 1 - \lambda$   
 reject otherwise

## Different Losses and Reject


Equal losses



Unequal losses



With reject



## Utility Theory

- Prob of state  $k$  given evidence  $\mathbf{x}$ :  $P(S_k | \mathbf{x})$
- Utility of  $\alpha_i$  when state is  $k$ :  $U_{ik}$
- Expected utility:  
 $EU(\alpha_i | \mathbf{x}) = \sum_k U_{ik} P(S_k | \mathbf{x})$   
 Choose  $\alpha_i$  if  $EU(\alpha_i | \mathbf{x}) = \max_j EU(\alpha_j | \mathbf{x})$

## Association Rules

- Association rule:  $X \rightarrow Y$
- People who buy/click/visit/enjoy  $X$  are also likely to buy/click/visit/enjoy  $Y$ .
- A rule implies association, not necessarily causation.


## Discriminant Functions

choose  $C_i$  if  $g_i(\mathbf{x}) = \max_k g_k(\mathbf{x})$

$$g_i(\mathbf{x}) = \begin{cases} -R(\alpha_i | \mathbf{x}) \\ P(C_i | \mathbf{x}) \\ \rho(\mathbf{x} | C_i) P(C_i) \end{cases}$$

$K$  decision regions  $\mathcal{R}_1, \dots, \mathcal{R}_K$

$$\mathcal{R}_i = \{\mathbf{x} | g_i(\mathbf{x}) = \max_k g_k(\mathbf{x})\}$$



## K=2 Classes

- Dichotomizer ( $K=2$ ) vs Polychotomizer ( $K>2$ )
- $g(\mathbf{x}) = g_1(\mathbf{x}) - g_2(\mathbf{x})$   
 choose  $C_1$  if  $g(\mathbf{x}) > 0$   
 $C_2$  otherwise
- Log odds:  
 $\log \frac{P(C_1 | \mathbf{x})}{P(C_2 | \mathbf{x})}$

## Apriori algorithm (Agrawal et al, 1996)

- For  $(X, Y, Z)$ , a 3-item set, to be frequent (have enough support),  $(X, Y)$ ,  $(X, Z)$ , and  $(Y, Z)$  should be frequent.
- If  $(X, Y)$  is not frequent, none of its supersets can be frequent.
- Once we find the frequent  $k$ -item sets, we convert them to rules:  $X, Y \rightarrow Z, \dots$  and  $X \rightarrow Y, Z, \dots$

## i2m13e-chap04.pdf



## Example

Transaction	Items in basket
1	milk, bananas, chocolate
2	milk, chocolate
3	milk, bananas
4	chocolate
5	chocolate
6	milk, chocolate

SOLUTION:

- milk - bananas : Support = 2/6, Confidence = 2/4
- bananas - milk : Support = 2/6, Confidence = 2/2
- milk - chocolate : Support = 3/6, Confidence = 3/4
- chocolate - milk : Support = 3/6, Confidence = 3/5

## Parametric Estimation

- $\mathcal{X} = \{x^i\}$ , where  $x^i \sim p(x)$
- Parametric estimation:  
 Assume a form for  $p(x | \theta)$  and estimate  $\theta$ , its sufficient statistics, using  $\mathcal{X}$   
 e.g.,  $N(\mu, \sigma^2)$  where  $\theta = \{\mu, \sigma^2\}$


## Maximum Likelihood Estimation

- Likelihood of  $\theta$  given the sample  $\mathcal{X}$   
 $l(\theta | \mathcal{X}) = p(\mathcal{X} | \theta) = \prod_i p(x^i | \theta)$
- Log likelihood  
 $\mathcal{L}(\theta | \mathcal{X}) = \log l(\theta | \mathcal{X}) = \sum_i \log p(x^i | \theta)$
- Maximum likelihood estimator (MLE)  
 $\theta^* = \text{argmax}_{\theta} \mathcal{L}(\theta | \mathcal{X})$

## Examples: Bernoulli/Multinomial

- Bernoulli: Two states, failure/success,  $x$  in  $\{0, 1\}$   
 $P(x) = p_0^x (1 - p_0)^{(1-x)}$   
 $\mathcal{L}(p_0 | \mathcal{X}) = \log \prod_i p_0^{x^i} (1 - p_0)^{(1-x^i)}$   
 MLE:  $p_0 = \sum_i x^i / N$
- Multinomial:  $K > 2$  states,  $x_i$  in  $\{0, 1\}$   
 $P(x_1, x_2, \dots, x_K) = \prod_i p_i^{x_i}$   
 $\mathcal{L}(p_1, p_2, \dots, p_K | \mathcal{X}) = \log \prod_i \prod_j p_j^{x_j^i}$   
 MLE:  $p_i = \sum_i x_i^i / N$

## Gaussian (Normal) Distribution



- $p(x) = \mathcal{N}(x | \mu, \sigma^2)$   


$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right]$$
- MLE for  $\mu$  and  $\sigma^2$ :

## Bias and Variance

Unknown parameter  $\theta$

Estimator  $d_i = d(X_i)$  on sample  $X_i$

Bias:  $b_i(d) = E[d] - \theta$   
Variance:  $E[(d - E[d])^2]$



Mean square error:

$$r(d, \theta) = E[(d - \theta)^2] = E[(d - E[d])^2] + E[(E[d] - \theta)^2] = \text{Bias}^2 + \text{Variance}$$

Given the sample  $\mathcal{X} = \{x^t, r^t\}_{t=1}^N$

$$x \in \mathfrak{R} \quad r_i^t = \begin{cases} 1 & \text{if } x^t \in C_i \\ 0 & \text{if } x^t \in C_j, j \neq i \end{cases}$$

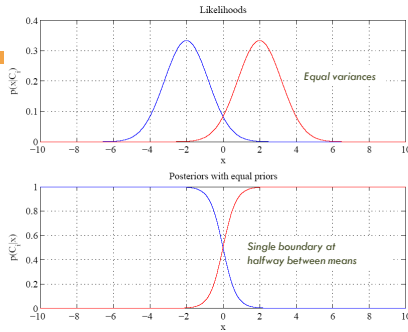
ML estimates are

$$\hat{P}(C_i) = \frac{\sum_t r_i^t}{N} \quad m_i = \frac{\sum_t x^t r_i^t}{\sum_t r_i^t} \quad s_i^2 = \frac{\sum_t (x^t - m_i)^2 r_i^t}{\sum_t r_i^t}$$

Discriminant  $g_i(x) = -\frac{1}{2} \log 2\pi - \log s_i - \frac{(x - m_i)^2}{2s_i^2} + \log \hat{P}(C_i)$

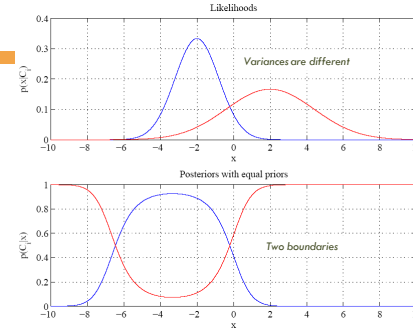
## Bayes' Estimator

- Treat  $\theta$  as a random var with prior  $p(\theta)$
- Bayes' rule:  $p(\theta | \mathcal{X}) = p(\mathcal{X} | \theta) p(\theta) / p(\mathcal{X})$
- Full:  $p(x | \mathcal{X}) = \int p(x | \theta) p(\theta | \mathcal{X}) d\theta$
- Maximum a Posteriori (MAP):  $\theta_{\text{MAP}} = \text{argmax}_{\theta} p(\theta | \mathcal{X})$
- Maximum Likelihood (ML):  $\theta_{\text{ML}} = \text{argmax}_{\theta} p(\mathcal{X} | \theta)$
- Bayes':  $\theta_{\text{Bayes}} = E[\theta | \mathcal{X}] = \int \theta p(\theta | \mathcal{X}) d\theta$



## Bayes' Estimator: Example

- $x^t \sim \mathcal{N}(\theta, \sigma_0^2)$  and  $\theta \sim \mathcal{N}(\mu, \sigma^2)$
- $\theta_{\text{ML}} = m$
- $\theta_{\text{MAP}} = \theta_{\text{Bayes}} = E[\theta | \mathcal{X}] = \frac{N/\sigma_0^2}{N/\sigma_0^2 + 1/\sigma^2} m + \frac{1/\sigma^2}{N/\sigma_0^2 + 1/\sigma^2} \mu$



## Parametric Classification

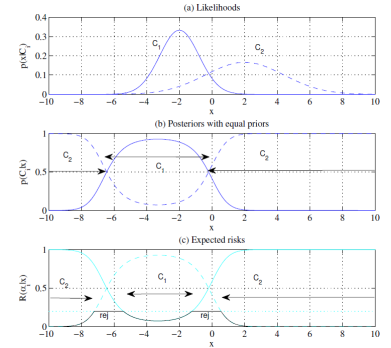
$$g_i(x) = p(x | C_i) P(C_i)$$

or

$$g_i(x) = \log p(x | C_i) + \log P(C_i)$$

$$p(x | C_i) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left[-\frac{(x - \mu_i)^2}{2\sigma_i^2}\right]$$

$$g_i(x) = -\frac{1}{2} \log 2\pi - \log \sigma_i - \frac{(x - \mu_i)^2}{2\sigma_i^2} + \log P(C_i)$$



## Regression

$$r = f(x) + \varepsilon$$


estimator  $g(x | \theta)$

$$\varepsilon \sim \mathcal{N}(0, \sigma^2)$$

$$p(r | x) \sim \mathcal{N}(g(x | \theta), \sigma^2)$$

$$L(\theta | \mathcal{X}) = \log \prod_{t=1}^N p(x^t, r^t)$$

$$= \log \prod_{t=1}^N p(r^t | x^t) + \log \prod_{t=1}^N p(x^t)$$



## Regression: From LogL to Error

$$L(\theta | \mathcal{X}) = \log \prod_{t=1}^N \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{[r^t - g(x^t | \theta)]^2}{2\sigma^2}\right]$$

$$= -N \log \sqrt{2\pi}\sigma - \frac{1}{2\sigma^2} \sum_{t=1}^N [r^t - g(x^t | \theta)]^2$$

$$E(\theta | \mathcal{X}) = \frac{1}{2} \sum_{t=1}^N [r^t - g(x^t | \theta)]^2$$

$$E[(r - g(x))^2 | x] = E[(r - E[r | x])^2 | x] + (E[r | x] - g(x))^2$$

noise                      squared error

$$E_x[E[r | x] - g(x)]^2 = (E[r | x] - E_x[g(x)])^2 + E_x[(g(x) - E_x[g(x)])^2]$$

bias                                      variance

## Other Error Measures

Square Error:  $E(\theta | \mathcal{X}) = \frac{1}{2} \sum_{t=1}^N [r^t - g(x^t | \theta)]^2$

Relative Square Error:  $E(\theta | \mathcal{X}) = \frac{\sum_{t=1}^N [r^t - g(x^t | \theta)]^2}{\sum_{t=1}^N [r^t - \bar{r}]^2}$

Absolute Error:  $E(\theta | \mathcal{X}) = \sum_i |r^i - g(x^i | \theta)|$

$\varepsilon$ -sensitive Error:

$$E(\theta | \mathcal{X}) = \sum_i 1(|r^i - g(x^i | \theta)| > \varepsilon) (|r^i - g(x^i | \theta)| - \varepsilon)$$

## Bias and Variance

## Linear Regression

$$g(x^t | w_1, w_0) = w_1 x^t + w_0$$

$$\sum_t r^t = N w_0 + w_1 \sum_t x^t$$

$$\sum_t r^t x^t = w_0 \sum_t x^t + w_1 \sum_t (x^t)^2$$

$$\mathbf{A} = \begin{bmatrix} N & \sum_t x^t \\ \sum_t x^t & \sum_t (x^t)^2 \end{bmatrix} \quad \mathbf{w} = \begin{bmatrix} w_0 \\ w_1 \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} \sum_t r^t \\ \sum_t r^t x^t \end{bmatrix}$$

$$\mathbf{w} = \mathbf{A}^{-1} \mathbf{y}$$

## Estimating Bias and Variance

$M$  samples  $X_i = \{x^i, r^i\}$ ,  $i = 1, \dots, M$  are used to fit  $g_i(x)$ ,  $i = 1, \dots, M$

$$\text{Bias}^2(g) = \frac{1}{N} \sum_t [\bar{g}(x^t) - f(x^t)]^2$$

$$\text{Variance}(g) = \frac{1}{NM} \sum_t \sum_i [g_i(x^t) - \bar{g}(x^t)]^2$$

$$\bar{g}(x) = \frac{1}{M} \sum_i g_i(x)$$

## Polynomial Regression

$$g(x^t | w_k, \dots, w_2, w_1, w_0) = w_k (x^t)^k + \dots + w_2 (x^t)^2 + w_1 x^t + w_0$$

$$\mathbf{D} = \begin{bmatrix} 1 & x^1 & (x^1)^2 & \dots & (x^1)^k \\ 1 & x^2 & (x^2)^2 & \dots & (x^2)^k \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x^N & (x^N)^2 & \dots & (x^N)^k \end{bmatrix} \quad \mathbf{r} = \begin{bmatrix} r^1 \\ r^2 \\ \vdots \\ r^N \end{bmatrix}$$

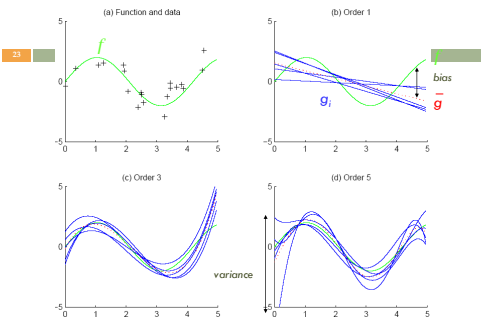
$$\mathbf{w} = (\mathbf{D}^T \mathbf{D})^{-1} \mathbf{D}^T \mathbf{r}$$

## Bias/Variance Dilemma

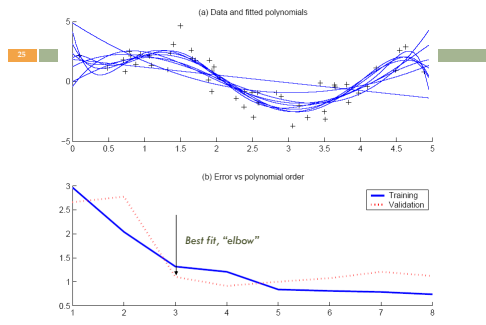

Example:  $g_i(x) = 2$  has no variance and high bias  
 $g_i(x) = \sum_t r^t / N$  has lower bias with variance

As we increase complexity, bias decreases (a better fit to data) and variance increases (fit varies more with data)

Bias/Variance dilemma: (Geman et al., 1992)



## Polynomial Regression




## Model Selection

- Cross-validation: Measure generalization accuracy by testing on data unused during training
- Regularization: Penalize complex models
  - $E'$  = error on data +  $\lambda$  model complexity
- Akaike's information criterion (AIC), Bayesian information criterion (BIC)
- Minimum description length (MDL): Kolmogorov complexity, shortest description of data
- Structural risk minimization (SRM)

## Bayesian Model Selection

- Prior on models,  $p(\text{model})$
- $$p(\text{model} | \text{data}) = \frac{p(\text{data} | \text{model}) p(\text{model})}{p(\text{data})}$$
- Regularization, when prior favors simpler models
- Bayes, MAP of the posterior,  $p(\text{model} | \text{data})$
- Average over a number of models with high posterior (voting, ensembles: Chapter 17)

## Regression example



$$\text{Regularization (L2): } E(\mathbf{w} | \mathbf{X}) = \frac{1}{2} \sum_{t=1}^N [r^t - g(x^t | \mathbf{w})]^2 + \lambda \sum w_i^2$$

## i2m13e-chap05.pdf

Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION

ETHEM ALPAYDM  
© The MIT Press, 2014

alpaydm@boun.edu.tr  
<http://www.cmp.e.boun.edu.tr/~ethem/i2m13e>

CHAPTER 5:  
**MULTIVARIATE METHODS**

## Multivariate Data

- Multiple measurements (sensors)
- $d$  inputs/features/attributes:  $d$ -variate
- $N$  instances/observations/examples

$$\mathbf{X} = \begin{bmatrix} X_1^1 & X_2^1 & \dots & X_d^1 \\ X_1^2 & X_2^2 & \dots & X_d^2 \\ \vdots & \vdots & \ddots & \vdots \\ X_1^N & X_2^N & \dots & X_d^N \end{bmatrix}$$

## Multivariate Parameters

Mean:  $E[\mathbf{x}] = \boldsymbol{\mu} = [\mu_1, \dots, \mu_d]^T$

Covariance:  $\sigma_{ij} \equiv \text{Cov}(X_i, X_j)$

Correlation:  $\text{Corr}(X_i, X_j) = \rho_{ij} = \frac{\sigma_{ij}}{\sigma_i \sigma_j}$

$$\Sigma \equiv \text{Cov}(\mathbf{X}) = E[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T] = \begin{bmatrix} \sigma_1^2 & \sigma_{12} & \dots & \sigma_{1d} \\ \sigma_{21} & \sigma_2^2 & \dots & \sigma_{2d} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{d1} & \sigma_{d2} & \dots & \sigma_d^2 \end{bmatrix}$$

## Parameter Estimation

Sample mean  $\mathbf{m}$ :  $m_i = \frac{\sum_{t=1}^N X_i^t}{N}, i = 1, \dots, d$


Covariance matrix  $\mathbf{S}$ :  $s_{ij} = \frac{\sum_{t=1}^N (X_i^t - m_i)(X_j^t - m_j)}{N}$

Correlation matrix  $\mathbf{R}$ :  $r_{ij} = \frac{s_{ij}}{s_i s_j}$

## Estimation of Missing Values

- What to do if certain instances have missing attributes?
- Ignore those instances: not a good idea if the sample is small
- Use 'missing' as an attribute: may give information
- Imputation: Fill in the missing value
  - Mean imputation: Use the most likely value (e.g., mean)
  - Imputation by regression: Predict based on other attributes

## Multivariate Normal Distribution



## Multivariate Normal Distribution


- Mahalanobis distance:  $(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})$  measures the distance from  $\mathbf{x}$  to  $\boldsymbol{\mu}$  in terms of  $\Sigma$  (normalized for difference in variances and correlations)
- Bivariate:  $d = 2$

$$\Sigma = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}$$

$$p(x_1, x_2) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} \exp\left[-\frac{1}{2(1-\rho^2)} (z_1^2 - 2\rho z_1 z_2 + z_2^2)\right]$$


$$z_i = (x_i - \mu_i) / \sigma_i$$

## Bivariate Normal



$$\text{Cov}(x_1, x_2) = 0, \text{Var}(x_1) = \text{Var}(x_2)$$

$$\text{Cov}(x_1, x_2) = 0, \text{Var}(x_1) \neq \text{Var}(x_2)$$




## Independent Inputs: Naive Bayes

- If  $x_i$  are independent, offdiagonals of  $\Sigma$  are 0, Mahalanobis distance reduces to weighted (by  $1/\sigma_i$ ) Euclidean distance:


$$p(\mathbf{x}) = \prod_{i=1}^d p_i(x_i) = \frac{1}{(2\pi)^{d/2} \prod_{i=1}^d \sigma_i} \exp\left[-\frac{1}{2} \sum_{i=1}^d \left(\frac{x_i - \mu_i}{\sigma_i}\right)^2\right]$$

- If variances are also equal, reduces to Euclidean distance



15

## Diagonal S



23

## Parametric Classification

- If  $p(\mathbf{x} | C_i) \sim N(\boldsymbol{\mu}_i, \Sigma_i)$

$$p(\mathbf{x} | C_i) = \frac{1}{(2\pi)^{d/2} |\Sigma_i|^{1/2}} \exp\left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_i)^T \Sigma_i^{-1} (\mathbf{x} - \boldsymbol{\mu}_i)\right]$$

- Discriminant functions

$$g_i(\mathbf{x}) = \log p(\mathbf{x} | C_i) + \log P(C_i) \\ = -\frac{d}{2} \log 2\pi - \frac{1}{2} \log |\Sigma_i| - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_i)^T \Sigma_i^{-1} (\mathbf{x} - \boldsymbol{\mu}_i) + \log P(C_i)$$

## Common Covariance Matrix S

- Shared common sample covariance S

$$S = \sum_i \hat{P}(C_i) S_i$$

- Discriminant reduces to

$$g_i(\mathbf{x}) = -\frac{1}{2} (\mathbf{x} - \mathbf{m}_i)^T S^{-1} (\mathbf{x} - \mathbf{m}_i) + \log \hat{P}(C_i)$$

which is a linear discriminant

$$g_i(\mathbf{x}) = \mathbf{w}_i^T \mathbf{x} + w_{i0}$$

where

$$\mathbf{w}_i = S^{-1} \mathbf{m}_i, w_{i0} = -\frac{1}{2} \mathbf{m}_i^T S^{-1} \mathbf{m}_i + \log \hat{P}(C_i)$$

## Diagonal S, equal variances

- Nearest mean classifier: Classify based on Euclidean distance to the nearest mean

$$g_i(\mathbf{x}) = -\frac{\|\mathbf{x} - \mathbf{m}_i\|^2}{2s^2} + \log \hat{P}(C_i) \\ = -\frac{1}{2s^2} \sum_{j=1}^d (x_j - m_{ij})^2 + \log \hat{P}(C_i)$$

- Each mean can be considered a prototype or template and this is template matching

## Discrete Features

- Binary features:  $p_{ij} \equiv p(x_j = 1 | C_i)$

if  $x_j$  are independent (Naive Bayes')

$$p(\mathbf{x} | C_i) = \prod_{j=1}^d p_{ij}^{x_j} (1 - p_{ij})^{1-x_j}$$

the discriminant is linear

$$g_i(\mathbf{x}) = \log p(\mathbf{x} | C_i) + \log P(C_i) \\ = \sum_j [x_j \log p_{ij} + (1 - x_j) \log (1 - p_{ij})] + \log P(C_i)$$

Estimated parameters

$$\hat{p}_{ij} = \frac{\sum_i x_j^i r_i^i}{\sum_i r_i^i}$$

## Estimation of Parameters


$$\hat{P}(C_i) = \frac{\sum_i r_i^i}{N}$$

$$\mathbf{m}_i = \frac{\sum_i r_i^i \mathbf{x}^i}{\sum_i r_i^i}$$


$$S_i = \frac{\sum_i r_i^i (\mathbf{x}^i - \mathbf{m}_i)(\mathbf{x}^i - \mathbf{m}_i)^T}{\sum_i r_i^i}$$

$$g_i(\mathbf{x}) = -\frac{1}{2} \log |S_i| - \frac{1}{2} (\mathbf{x} - \mathbf{m}_i)^T S_i^{-1} (\mathbf{x} - \mathbf{m}_i) + \log \hat{P}(C_i)$$

## Common Covariance Matrix S



## Diagonal S, equal variances



## Different S<sub>i</sub>

- Quadratic discriminant

$$g_i(\mathbf{x}) = -\frac{1}{2} \log |S_i| - \frac{1}{2} (\mathbf{x}^T S_i^{-1} \mathbf{x} - 2\mathbf{x}^T S_i^{-1} \mathbf{m}_i + \mathbf{m}_i^T S_i^{-1} \mathbf{m}_i) + \log \hat{P}(C_i) \\ = \mathbf{x}^T \mathbf{W}_i \mathbf{x} + \mathbf{w}_i^T \mathbf{x} + w_{i0}$$

where

$$\mathbf{W}_i = -\frac{1}{2} S_i^{-1}$$

$$\mathbf{w}_i = S_i^{-1} \mathbf{m}_i$$

$$w_{i0} = -\frac{1}{2} \mathbf{m}_i^T S_i^{-1} \mathbf{m}_i - \frac{1}{2} \log |S_i| + \log \hat{P}(C_i)$$

## Diagonal S

- When  $x_j, j = 1, \dots, d$ , are independent,  $\Sigma$  is diagonal  $p(\mathbf{x} | C_i) = \prod_j p(x_j | C_i)$  (Naive Bayes' assumption)

$$g_i(\mathbf{x}) = -\frac{1}{2} \sum_{j=1}^d \left( \frac{x_j - m_{ij}}{s_j} \right)^2 + \log \hat{P}(C_i)$$

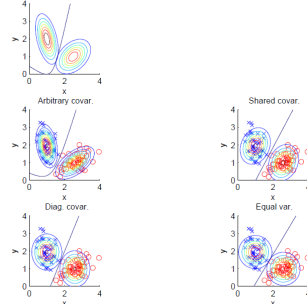
Classify based on weighted Euclidean distance (in  $s_j$  units) to the nearest mean

## Model Selection

Assumption	Covariance matrix	No of parameters
Shared, Hyperspheric	$S_i = S = s^2 I$	1
Shared, Axis-aligned	$S_i = S$ , with $s_i = 0$	$d$
Shared, Hyperellipsoidal	$S_i = S$	$d(d+1)/2$
Different, Hyperellipsoidal	$S_i$	$K d(d+1)/2$

- As we increase complexity (less restricted S), bias decreases and variance increases
- Assume simple models (allow some bias) to control variance (regularization)

Population likelihoods and posteriors



24

25

$$r^t = g(x^t | w_0, w_1, \dots, w_d) + \epsilon$$

Multivariate linear model

$$w_0 + w_1 x_1^t + w_2 x_2^t + \dots + w_d x_d^t$$

$$E(w_0, w_1, \dots, w_d | X) = \frac{1}{2} \sum_t [r^t - w_0 - w_1 x_1^t - \dots - w_d x_d^t]^2$$

Multivariate polynomial model:

Define new higher-order variables

$$z_1 = x_1, z_2 = x_2, z_3 = x_1^2, z_4 = x_2^2, z_5 = x_1 x_2$$

and use the linear model in this new  $z$  space  
(basis functions, kernel trick: Chapter 13)



Lecture Slides for  
**INTRODUCTION  
TO  
MACHINE  
LEARNING**  
3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014

alpaydin@boun.edu.tr  
http://www.cmp.e.boun.edu.tr/~ethem/i2ml3e

CHAPTER 6:  
**DIMENSIONALITY  
REDUCTION**

### Why Reduce Dimensionality?

- Reduces time complexity: Less computation
- Reduces space complexity: Fewer parameters
- Saves the cost of observing the feature
- Simpler models are more robust on small datasets
- More interpretable; simpler explanation
- Data visualization (structure, groups, outliers, etc) if plotted in 2 or 3 dimensions


### Feature Selection vs Extraction

- Feature selection:** Choosing  $k < d$  important features, ignoring the remaining  $d - k$ 
  - Subset selection algorithms
- Feature extraction:** Project the original  $x_i, i = 1, \dots, d$  dimensions to new  $k < d$  dimensions,  $z_j, j = 1, \dots, k$


### Subset Selection

- There are  $2^d$  subsets of  $d$  features
- Forward search: Add the best feature at each step
  - Set of features  $F$  initially  $\emptyset$ .
  - At each iteration, find the best new feature  $j = \text{argmin}_j E(F \cup x_j)$
  - Add  $x_j$  to  $F$  if  $E(F \cup x_j) < E(F)$
- Hill-climbing  $O(d^2)$  algorithm
- Backward search: Start with all features and remove one at a time, if possible.
- Floating search (Add  $k$ , remove  $l$ )

### Iris data: Single feature



### Iris data: Add one more feature to F4



### Principal Components Analysis

- Find a low-dimensional space such that when  $x$  is projected there, information loss is minimized.
- The projection of  $x$  on the direction of  $w$  is:  $z = w^T x$
- Find  $w$  such that  $\text{Var}(z)$  is maximized
 
$$\begin{aligned} \text{Var}(z) &= \text{Var}(w^T x) = E[(w^T x - w^T \mu)^2] \\ &= E[(w^T x - w^T \mu)(w^T x - w^T \mu)^T] \\ &= E[w^T (x - \mu)(x - \mu)^T w] \\ &= w^T E[(x - \mu)(x - \mu)^T] w = w^T \Sigma w \end{aligned}$$
 where  $\text{Var}(x) = E[(x - \mu)(x - \mu)^T] = \Sigma$

- Maximize  $\text{Var}(z)$  subject to  $\|w\| = 1$

$$\max_{w_1} w_1^T \Sigma w_1 - \alpha (w_1^T w_1 - 1)$$

$\Sigma w_1 = \alpha w_1$ , that is,  $w_1$  is an eigenvector of  $\Sigma$   
Choose the one with the largest eigenvalue for  $\text{Var}(z)$  to be max

- Second principal component: Max  $\text{Var}(z_2)$ , s.t.,  $\|w_2\| = 1$  and orthogonal to  $w_1$

$$\max_{w_2} w_2^T \Sigma w_2 - \alpha (w_2^T w_2 - 1) - \beta (w_2^T w_1 - 0)$$


$\Sigma w_2 = \alpha w_2$  that is,  $w_2$  is another eigenvector of  $\Sigma$  and so on.

### What PCA does

$$z = W^T(x - m)$$

where the columns of  $W$  are the eigenvectors of  $\Sigma$  and  $m$  is sample mean

Centers the data at the origin and rotates the axes





### How to choose $k$ ?

- Proportion of Variance (PoV) explained

$$\frac{\lambda_1 + \lambda_2 + \dots + \lambda_k}{\lambda_1 + \lambda_2 + \dots + \lambda_k + \dots + \lambda_d}$$

when  $\lambda_i$  are sorted in descending order

- Typically, stop at  $\text{PoV} > 0.9$
- Scree graph plots of  $\text{PoV}$  vs  $k$ , stop at "elbow"



### Feature Embedding

- When  $X$  is the  $N \times d$  data matrix,  $X^T X$  is the  $d \times d$  matrix (covariance of features, if mean-centered)
- $XX^T$  is the  $N \times N$  matrix (pairwise similarities of instances)
- PCA uses the eigenvectors of  $X^T X$  which are  $d$ -dim and can be used for projection
- Feature embedding uses the eigenvectors of  $XX^T$  which are  $N$ -dim and which give directly the coordinates after projection
- Sometimes, we can define pairwise similarities (or distances) between instances, then we can use feature embedding without needing to represent instances as vectors.



## Factor Analysis

- Find a small number of factors  $z$ , which when combined generate  $x$ :

$$x_i - \mu_i = v_{i1}z_1 + v_{i2}z_2 + \dots + v_{ik}z_k + \epsilon_i$$

where  $z_j, j=1, \dots, k$  are the latent factors with

$$E[z_j] = 0, \text{Var}(z_j) = 1, \text{Cov}(z_i, z_j) = 0, i \neq j,$$

$\epsilon_i$  are the noise sources


$$E[\epsilon_i] = \psi_i, \text{Cov}(\epsilon_i, \epsilon_j) = 0, i \neq j, \text{Cov}(\epsilon_i, z_j) = 0,$$

and  $v_{ij}$  are the factor loadings

## PCA vs FA


PCA From  $x$  to  $z \quad z = W^T(x - \mu)$

FA From  $z$  to  $x \quad x - \mu = Vz + \epsilon$



## Factor Analysis

- In FA, factors  $z_j$  are stretched, rotated and translated to generate  $x$




## Singular Value Decomposition and Matrix Factorization

- Singular value decomposition:  $X = VAW^T$   
 $V$  is  $N \times N$  and contains the eigenvectors of  $XX^T$   
 $W$  is  $d \times d$  and contains the eigenvectors of  $X^T X$   
 $A$  is  $N \times d$  and contains singular values on its first  $k$  diagonal  
 $X = u_1 v_1^T + \dots + u_k v_k^T$  where  $k$  is the rank of  $X$

## Matrix Factorization

- Matrix factorization:  $X = FG$

$F$  is  $N \times k$  and  $G$  is  $k \times d$



$$X_{ij} = F_i^T G_j = \sum_{l=1}^k F_{il} G_{lj} \quad \text{Latent semantic indexing}$$

## Multidimensional Scaling

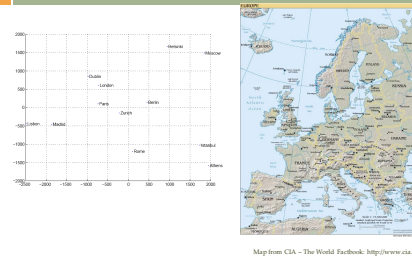
- Given pairwise distances between  $N$  points,  $d_{ij}, i, j = 1, \dots, N$   
 place on a low-dim map s.t. distances are preserved (by feature embedding)

- $z = g(x | \theta)$  Find  $\theta$  that min Sammon stress

$$E(\theta | X) = \sum_{r,s} \frac{(\|z^r - z^s\| - \|x^r - x^s\|)^2}{\|x^r - x^s\|^2}$$

$$= \sum_{r,s} \frac{(\|g(x^r | \theta) - g(x^s | \theta)\| - \|x^r - x^s\|)^2}{\|x^r - x^s\|^2}$$

## Map of Europe by MDS



## Linear Discriminant Analysis

- Find a low-dimensional space such that when  $x$  is projected, classes are well-separated.
- Find  $w$  that maximizes

$$J(w) = \frac{(m_1 - m_2)^2}{S_1^2 + S_2^2}$$

$$m_1 = \frac{\sum_t x^t r^t}{\sum_t r^t} \quad S_1^2 = \sum_t (w^T x^t - m_1)^2 r^t$$

- Between-class scatter:

$$(m_1 - m_2)^2 = (w^T m_1 - w^T m_2)^2$$

$$= w^T (m_1 - m_2)(m_1 - m_2)^T w$$

$$= w^T S_B w \quad \text{where } S_B = (m_1 - m_2)(m_1 - m_2)^T$$

- Within-class scatter:

$$S_1^2 = \sum_t (w^T x^t - m_1)^2 r^t$$

$$= \sum_t w^T (x^t - m_1)(x^t - m_1)^T w r^t = w^T S_1 w$$

where  $S_1 = \sum_t (x^t - m_1)(x^t - m_1)^T r^t$

$$S_1^2 + S_2^2 = w^T S_W w \quad \text{where } S_W = S_1 + S_2$$

## Fisher's Linear Discriminant

- Find  $w$  that max

$$J(w) = \frac{w^T S_B w}{w^T S_W w} = \frac{w^T (m_1 - m_2)(m_1 - m_2)^T w}{w^T S_W w}$$

- LDA soln:

$$w = c \cdot S_W^{-1} (m_1 - m_2)$$

- Parametric soln:

$$w = \Sigma^{-1} (\mu_1 - \mu_2)$$

when  $p(x | C_i) \sim \mathcal{N}(\mu_i, \Sigma)$

## K>2 Classes

- Within-class scatter:


$$S_W = \sum_{i=1}^K S_i \quad S_i = \sum_t r^t (x^t - m_i)(x^t - m_i)^T$$

- Between-class scatter:


$$S_B = \sum_{i=1}^K N_i (m_i - m)(m_i - m)^T \quad m = \frac{1}{K} \sum_{i=1}^K m_i$$

- Find  $W$  that max  $J(W) = \frac{W^T S_B W}{W^T S_W W}$

The largest eigenvectors of  $S_W^{-1} S_B$ , maximum rank of  $K-1$



## PCA vs LDA



## Canonical Correlation Analysis


- $X = \{x^t, y^t\}$ ; two sets of variables  $x$  and  $y$
- We want to find two projections  $w$  and  $v$  st when  $x$  is projected along  $w$  and  $y$  is projected along  $v$ , the correlation is maximized:

$$\rho = \text{Corr}(w^T x, v^T y) = \frac{\text{Cov}(w^T x, v^T y)}{\sqrt{\text{Var}(w^T x)} \sqrt{\text{Var}(v^T y)}}$$

$$= \frac{w^T \text{Cov}(x, y) v}{\sqrt{w^T \text{Var}(x) w} \sqrt{v^T \text{Var}(y) v}} = \frac{w^T S_{xy} v}{\sqrt{w^T S_{xx} w} \sqrt{v^T S_{yy} v}}$$


## CCA

- $x$  and  $y$  may be two different views or modalities; e.g., image and word tags, and CCA does a joint mapping




## Isomap

- Geodesic distance is the distance along the manifold that the data lies in, as opposed to the Euclidean distance in the input space




## Isomap

- Instances  $r$  and  $s$  are connected in the graph if  $||x^r - x^s|| < \epsilon$  or if  $x^s$  is one of the  $k$  neighbors of  $x^r$ . The edge length is  $||x^r - x^s||$ .
- For two nodes  $r$  and  $s$  not connected, the distance is equal to the shortest path between them.
- Once the  $N \times N$  distance matrix is thus formed, use MDS to find a lower-dimensional mapping.



32

## LLE on Opltdigits



Matlab source from <http://www.cs.toronto.edu/~rowley/le/code.html>

## Laplacian Eigenmaps

- Let  $r$  and  $s$  be two instances and  $B_{rs}$  is their similarity, we want to find  $z^r$  and  $z^s$  that

$$\min \sum_{r,s} ||z^r - z^s||^2 B_{rs}$$

- $B_{rs}$  can be defined in terms of similarity in an original space: 0 if  $x^r$  and  $x^s$  are too far, otherwise


$$B_{rs} = \exp \left[ -\frac{||x^r - x^s||^2}{2\sigma^2} \right]$$

- Defines a graph Laplacian, and feature embedding returns  $z^r$

## Locally Linear Embedding

- Given  $x^r$  find its neighbors  $x^{s(r)}$
  - Find  $W_{rs}$  that minimize
- $$E(W | X) = \sum_r ||x^r - \sum_s W_{rs} x^{s(r)}||^2$$
- Find the new coordinates  $z^r$  that minimize

$$E(z | W) = \sum_r ||z^r - \sum_s W_{rs} z^{s(r)}||^2$$




34

1 Mon Sep 24 12:33:54 2018 7

## i2ml3e-chap07.pdf

## Laplacian Eigenmaps on Iris



Spectral clustering (chapter 7)

## Semiparametric Density Estimation

- Parametric: Assume a single model for  $p(x | C_i)$  (Chapters 4 and 5)
- Semiparametric:  $p(x | C_i)$  is a mixture of densities. Multiple possible explanations/prototypes: Different handwriting styles, accents in speech
- Nonparametric: No model; data speaks for itself (Chapter 8)

## Mixture Densities

$$p(x) = \sum_{i=1}^k p(x | G_i) P(G_i)$$

- where  $G_i$  the components/groups/clusters,  $P(G_i)$  mixture proportions (priors),  $p(x | G_i)$  component densities
- Gaussian mixture where  $p(x | G_i) \sim N(\mu_i, \Sigma_i)$
- parameters  $\Phi = \{P(G_i), \mu_i, \Sigma_i\}_{i=1}^k$
- unlabeled sample  $X = \{x^t\}$ , (unsupervised learning)


## Classes vs. Clusters

- Supervised:  $X = \{x^t, r^t\}$
  - Classes  $C_i, i=1, \dots, K$
  - $p(x) = \sum_{i=1}^K p(x | C_i) P(C_i)$
  - where  $p(x | C_i) \sim N(\mu_i, \Sigma_i)$
  - $\Phi = \{P(C_i), \mu_i, \Sigma_i\}_{i=1}^K$
  - Unsupervised:  $X = \{x^t\}$
  - Clusters  $G_i, i=1, \dots, k$
  - $p(x) = \sum_{i=1}^k p(x | G_i) P(G_i)$
  - where  $p(x | G_i) \sim N(\mu_i, \Sigma_i)$
  - $\Phi = \{P(G_i), \mu_i, \Sigma_i\}_{i=1}^k$
  - Labels  $r^t$ ?
- $$\hat{P}(C_i) = \frac{\sum_{t=1}^N r^t_i}{N}, \quad \hat{m}_i = \frac{\sum_{t=1}^N r^t_i x^t}{\sum_{t=1}^N r^t_i}$$
- $$S_i = \frac{\sum_{t=1}^N (x^t - \hat{m}_i)(x^t - \hat{m}_i)^T}{\sum_{t=1}^N r^t_i}$$

## k-Means Clustering


- Find  $k$  reference vectors (prototypes/codebook vectors/codewords) which best represent data
- Reference vectors,  $m_j, j=1, \dots, k$
- Use nearest (most similar) reference:  $||x^t - m_j|| = \min_j ||x^t - m_j||$
- Reconstruction error  $E(\{m_j\}_{j=1}^k | X) = \sum_t \sum_j b_j^t ||x^t - m_j||$
- $b_j^t = \begin{cases} 1 & \text{if } ||x^t - m_j|| = \min_j ||x^t - m_j|| \\ 0 & \text{otherwise} \end{cases}$

## Encoding/Decoding



## k-means Clustering

- Initialize  $m_i, i=1, \dots, k$ , for example, to  $k$  random  $x^t$
- Repeat
- For all  $x^t \in X$
- $b_j^t = \begin{cases} 1 & \text{if } ||x^t - m_j|| = \min_j ||x^t - m_j|| \\ 0 & \text{otherwise} \end{cases}$
- For all  $m_i, i=1, \dots, k$
- $m_i = \frac{\sum_t b_i^t x^t}{\sum_t b_i^t}$
- Until  $m_i$  converge



## Expectation-Maximization (EM)

- Log likelihood with a mixture model

$$\mathcal{L}(\Phi | X) = \log \prod_i p(x^i | \Phi) = \sum_i \log \sum_j p(x^i | G_j) p(G_j)$$

- Assume hidden variables  $z$ , which when known, make optimization much simpler
- Complete likelihood,  $L_c(\Phi | X, Z)$ , in terms of  $x$  and  $z$
- Incomplete likelihood,  $L(\Phi | X)$ , in terms of  $x$


## Mixtures of Latent Variable Models

Regularize clusters

- Assume shared/diagonal covariance matrices
- Use PCA/FA to decrease dimensionality: Mixtures of PCA/FA

$$p(x_i | G_i) = \mathcal{N}(m_i, V_i V_i^T + \psi_i)$$

Can use EM to learn  $V_i$  (Ghahramani and Hinton, 1997; Tipping and Bishop, 1999)



## E- and M-steps

Iterate the two steps

- E-step: Estimate  $z$  given  $X$  and current  $\Phi$
- M-step: Find new  $\Phi'$  given  $z$ ,  $X$ , and old  $\Phi$ .

$$\text{E-step: } Q(\Phi | \Phi') = E[\mathcal{L}_c(\Phi | X, Z) | X, \Phi']$$

$$\text{M-step: } \Phi^{i+1} = \text{argmax}_{\Phi} Q(\Phi | \Phi')$$

An increase in  $Q$  increases incomplete likelihood

$$\mathcal{L}(\Phi^{i+1} | X) \geq \mathcal{L}(\Phi^i | X)$$

## After Clustering

- Dimensionality reduction methods find correlations between features and group features
- Clustering methods find similarities between instances and group instances
- Allows knowledge extraction through number of clusters, prior probabilities, cluster parameters, i.e., center, range of features. Example: CRM, customer segmentation

## EM in Gaussian Mixtures

- $z_i^j = 1$  if  $x^i$  belongs to  $G_j$ , 0 otherwise (labels  $r^i$  of supervised learning); assume  $p(x | G_j) \sim \mathcal{N}(\mu_j, \Sigma_j)$

$$\text{E-step: } E[z_i^j | X, \Phi^i] = \frac{p(x^i | G_j, \Phi^i) p(G_j)}{\sum_r p(x^i | G_r, \Phi^i) p(G_r)} = P(G_j | x^i, \Phi^i) = h_i^j$$

$$\text{M-step: } p(G_j) = \frac{\sum_i h_i^j}{N} \quad m_j^{i+1} = \frac{\sum_i h_i^j x^i}{\sum_i h_i^j} \quad s_j^{i+1} = \frac{\sum_i h_i^j (x^i - m_j^{i+1})(x^i - m_j^{i+1})^T}{\sum_i h_i^j}$$

*Use estimated labels in place of unknown labels*

## Clustering as Preprocessing

- Estimated group labels  $h_i$  (soft) or  $b_i$  (hard) may be seen as the dimensions of a new  $k$  dimensional space, where we can then learn our discriminant or regressor.
- Local representation (only one  $b_i$  is 1, all others are 0; only few  $h_i$  are nonzero) vs Distributed representation (After PCA; all  $z_i$  are nonzero)


## Mixture of Mixtures

- In classification, the input comes from a mixture of classes (supervised).
- If each class is also a mixture, e.g., of Gaussians, (unsupervised), we have a mixture of mixtures:

$$p(x | C_i) = \sum_{j=1}^K p(x | G_{ij}) p(G_{ij})$$

$$p(x) = \sum_{i=1}^K p(x | C_i) p(C_i)$$

## Example: Single-Link Clustering



## Spectral Clustering

- Cluster using predefined pairwise similarities  $B_{rs}$  instead of using Euclidean or Mahalanobis distance
- Can be used even if instances not vectorially represented
- Steps:
  - Use Laplacian Eigenmaps (chapter 6) to map to a new  $z$  space using  $B_{rs}$
  - Use  $k$ -means in this new  $z$  space for clustering

## Choosing $k$

- Defined by the application, e.g., image quantization
- Plot data (after PCA) and check for clusters
- Incremental (leader-cluster) algorithm: Add one at a time until "elbow" (reconstruction error/log likelihood/intergroup distances)
- Manually check for meaning

## Hierarchical Clustering

- Cluster based on similarities/distances
- Distance measure between instances  $x^r$  and  $x^s$

$$\text{Minkowski } (L_p) \text{ (Euclidean for } p = 2) \\ d_m(x^r, x^s) = \left[ \sum_{j=1}^d |x_j^r - x_j^s|^p \right]^{1/p}$$

City-block distance

$$d_{cb}(x^r, x^s) = \sum_{j=1}^d |x_j^r - x_j^s|$$

## Agglomerative Clustering

- Start with  $N$  groups each with one instance and merge two closest groups at each iteration
- Distance between two groups  $G_i$  and  $G_j$ :

$$\text{Single-link: } d(G_i, G_j) = \min_{x^r \in G_i, x^s \in G_j} d(x^r, x^s)$$

$$\text{Complete-link: } d(G_i, G_j) = \max_{x^r \in G_i, x^s \in G_j} d(x^r, x^s)$$

$$\text{Average-link, centroid} \\ d(G_i, G_j) = \text{ave}_{x^r \in G_i, x^s \in G_j} d(x^r, x^s)$$

L Mon Sep 24 11:39:54 2018

## i2ml3e-chap08.pdf

Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION


ETHEM ALPAYDM  
© The MIT Press, 2014

alpaydm@boun.edu.tr  
<http://www.cmp.e.boun.edu.tr/~ethem/i2ml3e>

CHAPTER 8:  
NONPARAMETRIC  
METHODS

### Nonparametric Estimation

- Parametric (single global model), semiparametric (small number of local models)
- Nonparametric: Similar inputs have similar outputs
- Functions (pdf, discriminant, regression) change smoothly
- Keep the training data; "let the data speak for itself"
- Given  $x$ , find a small number of closest training instances and interpolate from these
- Aka lazy/memory-based/case-based/instance-based learning



### Kernel Estimator

- Kernel function, e.g., Gaussian kernel:

$$K(u) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{u^2}{2}\right]$$

- Kernel estimator (Parzen windows)

$$\hat{p}(x) = \frac{1}{Nh} \sum_{t=1}^N K\left(\frac{x-x^t}{h}\right)$$



### Density Estimation

- Given the training set  $X=\{x^t\}$ , drawn iid from  $p(x)$
- Divide data into bins of size  $h$

Histogram: 
$$\hat{p}(x) = \frac{\#\{x^t \text{ in the same bin as } x\}}{Nh}$$

Naive estimator: 
$$\hat{p}(x) = \frac{\#\{x-h < x^t \leq x+h\}}{2Nh}$$

or 
$$\hat{p}(x) = \frac{1}{Nh} \sum_{t=1}^N w\left(\frac{x-x^t}{h}\right) \quad w(u) = \begin{cases} 1/2 & \text{if } |u| < 1 \\ 0 & \text{otherwise} \end{cases}$$




### k-Nearest Neighbor Estimator

- Instead of fixing bin width  $h$  and counting the number of instances, fix the instances (neighbors)  $k$  and check bin width

$$\hat{p}(x) = \frac{k}{2Nd_k(x)}$$

$d_k(x)$ , distance to  $k$ th closest instance to  $x$



### Multivariate Data

- Kernel density estimator

$$\hat{p}(x) = \frac{1}{Nh^d} \sum_{t=1}^N K\left(\frac{x-x^t}{h}\right)$$

Multivariate Gaussian kernel

spheric 
$$K(u) = \left(\frac{1}{\sqrt{2\pi}}\right)^d \exp\left[-\frac{\|u\|^2}{2}\right]$$

ellipsoid 
$$K(u) = \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp\left[-\frac{1}{2} u^T \Sigma^{-1} u\right]$$

### Nonparametric Classification

- Estimate  $p(x|C_i)$  and use Bayes' rule

- Kernel estimator

$$\hat{p}(x|C_i) = \frac{1}{Nh_i^d} \sum_{t=1}^{N_i} K\left(\frac{x-x^t}{h}\right) \quad \hat{p}(C_i) = \frac{N_i}{N}$$


$$g_i(x) = \hat{p}(x|C_i) \hat{p}(C_i) = \frac{1}{Nh_i^d} \sum_{t=1}^{N_i} K\left(\frac{x-x^t}{h}\right) \frac{N_i}{N}$$

- k-NN estimator

$$\hat{p}(x|C_i) = \frac{k_i}{N_i V^d(x)} \quad \hat{p}(C_i|x) = \frac{\hat{p}(x|C_i) \hat{p}(C_i)}{\hat{p}(x)} = \frac{k_i}{k}$$

### Condensed Nearest Neighbor

- Time/space complexity of  $k$ -NN is  $O(N)$
- Find a subset  $Z$  of  $X$  that is small and is accurate in classifying  $X$  (Hart, 1968)



### Condensed Nearest Neighbor

- Incremental algorithm: Add instance if needed

```
Z ← {}
Repeat
  For all x ∈ X (in random order)
    Find x' ∈ Z s.t. ||x - x'|| = min_{x' ∈ Z} ||x - x'||
    If class(x) ≠ class(x') add x to Z
Until Z does not change
```

### Distance-based Classification

- Find a distance function  $D(x^t, x')$  such that if  $x^t$  and  $x'$  belong to the same class, distance is small and if they belong to different classes, distance is large
- Assume a parametric model and learn its parameters using data, e.g.,


$$D(x, x^t|M) = (x - x^t)^T M (x - x^t)$$

### Learning a Distance Function

- The three-way relationship between distances, dimensionality reduction, and feature extraction.
- $M=L^T L$  is  $d \times d$  and  $L$  is  $k \times d$

$$\begin{aligned} D(x, x^t|M) &= (x - x^t)^T M (x - x^t) = (x - x^t)^T L^T L (x - x^t) \\ &= (L(x - x^t))^T (L(x - x^t)) = (Lx - Lx^t)^T (Lx - Lx^t) \\ &= (z - z^t)^T (z - z^t) = \|z - z^t\|^2 \end{aligned}$$

- Similarity-based representation using similarity scores
- Large-margin nearest neighbor (chapter 13)




Euclidean distance (circle) is not suitable, Mahalanobis distance using an  $M$  (ellipse) is suitable. After the data is projected along  $L$ , Euclidean distance can be used.

## Outlier Detection

- Find outlier/novelty points
- Not a two-class problem because outliers are very few, of many types, and seldom labeled
- Instead, one-class classification problem: Find instances that have low probability
- In nonparametric case: Find instances far away from other instances

## Local Outlier Factor

$$LOF(x) = \frac{d_k(x)}{\sum_{s \in N(x)} d_k(s) / |N(x)|}$$



## Running Mean/Kernel Smoother


- Running mean smoother
- Kernel smoother
- Running line smoother
- Additive models (Hastie and Tibshirani, 1990)

$$\hat{g}(x) = \frac{\sum_{i=1}^N w \left( \frac{x-x^i}{h} \right) r^i}{\sum_{i=1}^N w \left( \frac{x-x^i}{h} \right)}$$


$$\hat{g}(x) = \frac{\sum_{i=1}^N K \left( \frac{x-x^i}{h} \right) r^i}{\sum_{i=1}^N K \left( \frac{x-x^i}{h} \right)}$$

where  $w(u) = \begin{cases} 1 & \text{if } |u| < 1 \\ 0 & \text{otherwise} \end{cases}$

where  $K(\cdot)$  is Gaussian



22



26

Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014

alpaydm@boun.edu.tr  
http://www.cmlpe.boun.edu.tr/~ethem/i2ml3e

## How to Choose k or h ?

- When k or h is small, single instances matter; bias is small, variance is large (undersmoothing): High complexity
- As k or h increases, we average over more instances and variance decreases but bias increases (oversmoothing): Low complexity
- Cross-validation is used to finetune k or h.


27

## Nonparametric Regression


- Aka smoothing models
- Regression

$$\hat{g}(x) = \frac{\sum_{i=1}^N b(x, x^i) r^i}{\sum_{i=1}^N b(x, x^i)}$$


where  $b(x, x^i) = \begin{cases} 1 & \text{if } x^i \text{ is in the same bin with } x \\ 0 & \text{otherwise} \end{cases}$



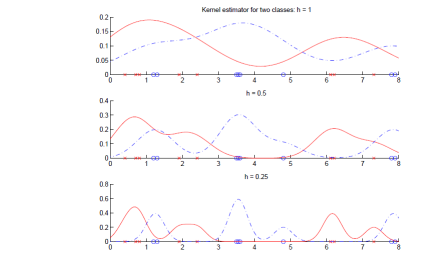
24



21




25



28

1 Min Sep 24 12:39:54 2018 9  
**i2ml3e-chap09.pdf**

## Tree Uses Nodes and Leaves



## Divide and Conquer

- Internal decision nodes
  - Univariate: Uses a single attribute,  $x_i$ 
    - Numeric  $x_i$ : Binary split:  $x_i > w_m$
    - Discrete  $x_i$ : n-way split for n possible values
  - Multivariate: Uses all attributes,  $x$
- Leaves
  - Classification: Class labels, or proportions
  - Regression: Numeric; r average, or local fit
- Learning is greedy; find the best split recursively (Breiman et al, 1984; Quinlan, 1986, 1993)

3

CHAPTER 9:  
**DECISION TREES**


# Classification Trees (ID3,CART,C4.5)

For node  $m$ ,  $N_m$  instances reach  $m$ ,  $N_m^i$  belong to  $C_i$

$$\hat{p}(C_i | \mathbf{x}, m) \equiv p_m^i = \frac{N_m^i}{N_m}$$

- Node  $m$  is pure if  $p_m^i$  is 0 or 1
- Measure of impurity is entropy

$$I_m = -\sum_{i=1}^K p_m^i \log_2 p_m^i$$



# Best Split

- If node  $m$  is pure, generate a leaf and stop, otherwise split and continue recursively
- Impurity after split:  $N_{mj}$  of  $N_m$  take branch  $j$ ,  $N_{mj}^i$  belong to  $C_i$

$$\hat{p}(C_i | \mathbf{x}, m, j) \equiv p_{mj}^i = \frac{N_{mj}^i}{N_{mj}} \quad I_m = -\sum_{j=1}^d \sum_{i=1}^K \frac{N_{mj}^i}{N_m} \log_2 \frac{N_{mj}^i}{N_{mj}}$$

- Find the variable and split that min impurity (among all variables -- and split positions for numeric variables)

```

GenerateTree(X)
  If NodeEntropy(X) < #e /* eq. 9.3
    Create leaf labeled by majority class in X
    Return
  i ← SplitAttribute(X)
  For each branch of #e
    Find X_i falling in branch
    GenerateTree(X_i)
  SplitAttribute(X)
  MinEnt ← MAX
  For all attributes i = 1, ..., d
    If #e_i is discrete with n values
      Split X into X_1, ..., X_n by #e_i
      e ← SplitEntropy(X_1, ..., X_n) /* eq. 9.8 */
      If e < MinEnt MinEnt ← e; bestf ← i
    Else /* #e_i is numeric */
      For all possible splits
        Split X into X_1, X_2 on #e_i
        e ← SplitEntropy(X_1, X_2)
        If e < MinEnt MinEnt ← e; bestf ← i
  Return bestf
    
```

# Regression Trees

Error at node  $m$ :

$$b_m(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in X_m : \mathbf{x} \text{ reaches node } m \\ 0 & \text{otherwise} \end{cases}$$



$$E_m = \frac{1}{N_m} \sum_{i=1}^n (r^i - g_m)^2 b_m(\mathbf{x}^i) \quad g_m = \frac{\sum_{i=1}^n b_m(\mathbf{x}^i) r^i}{\sum_{i=1}^n b_m(\mathbf{x}^i)}$$

After splitting:

$$b_m(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in X_m : \mathbf{x} \text{ reaches node } m \text{ and branch } j \\ 0 & \text{otherwise} \end{cases}$$

$$E'_m = \frac{1}{N_m} \sum_{j=1}^d \sum_{i=1}^n (r^i - g_{mj})^2 b_{mj}(\mathbf{x}^i) \quad g_{mj} = \frac{\sum_{i=1}^n b_{mj}(\mathbf{x}^i) r^i}{\sum_{i=1}^n b_{mj}(\mathbf{x}^i)}$$

# Model Selection in Trees



# Pruning Trees

- Remove subtrees for better generalization (decrease variance)
  - Prepruning: Early stopping
  - Postpruning: Grow the whole tree then prune subtrees that overfit on the pruning set
- Prepruning is faster, postpruning is more accurate (requires a separate pruning set)

# Rule Extraction from Trees

C4.5 Rules (Quinlan, 1993)

R1: IF (age > 38.5) AND (years-in-job > 2.5) THEN y = 0.8  
 R2: IF (age > 38.5) AND (years-in-job <= 2.5) THEN y = 0.6  
 R3: IF (age <= 38.5) AND (job-type='A') THEN y = 0.4  
 R4: IF (age <= 38.5) AND (job-type='B') THEN y = 0.3  
 R5: IF (age <= 38.5) AND (job-type='C') THEN y = 0.2

# Learning Rules

- Rule induction is similar to tree induction but
  - tree induction is breadth-first,
  - rule induction is depth-first; one rule at a time
- Rule set contains rules; rules are conjunctions of terms
- Rule covers an example if all terms of the rule evaluate to true for the example
- Sequential covering: Generate rules one at a time until all positive examples are covered
- IREP (Fürnkranz and Widmer, 1994), Ripper (Cohen, 1995)


```

Ripper(Pos, Neg, k)
  RuleSet ← LearnRuleSet(Pos, Neg)
  For k times
    RuleSet ← OptimizeRuleSet(RuleSet, Pos, Neg)
  LearnRuleSet(Pos, Neg)
  RuleSet ← {}
  DL ← DescLen(RuleSet, Pos, Neg)
  Repeat
    Rule ← LearnRule(Pos, Neg)
    Add Rule to RuleSet
    DL' ← DescLen(RuleSet, Pos, Neg)
    If DL' > DL + #e4
      PruneRuleSet(RuleSet, Pos, Neg)
      Return RuleSet
    If DL' < DL DL ← DL'
    Delete instances covered from Pos and Neg
  Until Pos = {}
  Return RuleSet
    
```

```

PruneRuleSet(RuleSet, Pos, Neg)
  For each Rule ∈ RuleSet in reverse order
    DL ← DescLen(RuleSet, Pos, Neg)
    DL' ← DescLen(RuleSet - Rule, Pos, Neg)
    If DL' < DL Delete Rule from RuleSet
  Return RuleSet
  OptimizeRuleSet(RuleSet, Pos, Neg)
  For each Rule ∈ RuleSet
    DL0 ← DescLen(RuleSet, Pos, Neg)
    DL1 ← DescLen(RuleSet - Rule + ReplaceRule(RuleSet, Pos, Neg), Pos, Neg)
    DL2 ← DescLen(RuleSet - Rule + ReviseRule(RuleSet, Rule, Pos, Neg), Pos, Neg)
    If DL1 = min(DL0, DL1, DL2)
      Delete Rule from RuleSet and
      add ReplaceRule(RuleSet, Pos, Neg)
    Else If DL2 = min(DL0, DL1, DL2)
      Delete Rule from RuleSet and
      add ReviseRule(RuleSet, Rule, Pos, Neg)
  Return RuleSet
    
```

# Multivariate Trees



i2ml3e-chap10.pdf

Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
 3RD EDITION

ETHEM ALPAYDIN  
 © The MIT Press, 2014

alpaydin@boun.edu.tr  
 http://www.cmlpe.boun.edu.tr/~ethem/i2ml3e

CHAPTER 10:  
**LINEAR DISCRIMINATION**

# Likelihood- vs. Discriminant-based Classification

- Likelihood-based: Assume a model for  $p(\mathbf{x} | C_i)$ , use Bayes' rule to calculate  $P(C_i | \mathbf{x})$ 

$$g_i(\mathbf{x}) = \log P(C_i | \mathbf{x})$$
- Discriminant-based: Assume a model for  $g_i(\mathbf{x} | \Phi_i)$ ; no density estimation
- Estimating the boundaries is enough; no need to accurately estimate the densities inside the boundaries

# Linear Discriminant

- Linear discriminant:
 
$$g(\mathbf{x} | \mathbf{w}_1, \mathbf{w}_0) = \mathbf{w}_1^T \mathbf{x} + w_{10} = \sum_{j=1}^d w_{1j} x_j + w_{10}$$
- Advantages:
  - Simple:  $O(d)$  space/computation
  - Knowledge extraction: Weighted sum of attributes; positive/negative weights, magnitudes (credit scoring)
  - Optimal when  $p(\mathbf{x} | C_i)$  are Gaussian with shared cov matrix; useful when classes are (almost) linearly separable

## Generalized Linear Model

- Quadratic discriminant:

$$g_i(\mathbf{x} | \mathbf{W}_i, \mathbf{w}_i, w_{i0}) = \mathbf{x}^T \mathbf{W}_i \mathbf{x} + \mathbf{w}_i^T \mathbf{x} + w_{i0}$$


- Higher-order (product) terms:

$$Z_1 = X_1, Z_2 = X_2, Z_3 = X_1^2, Z_4 = X_2^2, Z_5 = X_1 X_2$$


Map from  $\mathbf{x}$  to  $\mathbf{z}$  using nonlinear basis functions and use a linear discriminant in  $\mathbf{z}$ -space

$$g(\mathbf{x}) = \sum_{j=1}^k w_j \phi_j(\mathbf{x})$$


## Two Classes




## Geometry



## Multiple Classes



## Pairwise Separation



## From Discriminants to Posteriors

When  $p(\mathbf{x} | C_i) \sim N(\boldsymbol{\mu}_i, \Sigma)$

$$g_i(\mathbf{x} | \mathbf{w}_i, w_{i0}) = \mathbf{w}_i^T \mathbf{x} + w_{i0}$$

$$\mathbf{w}_i = \Sigma^{-1} \boldsymbol{\mu}_i, w_{i0} = -\frac{1}{2} \boldsymbol{\mu}_i^T \Sigma^{-1} \boldsymbol{\mu}_i + \log P(C_i)$$

$$y = P(C_1 | \mathbf{x}) \text{ and } P(C_2 | \mathbf{x}) = 1 - y$$

$$\text{choose } C_1 \text{ if } \begin{cases} y > 0.5 \\ \log[y/(1-y)] > 0 \end{cases} \text{ and } C_2 \text{ otherwise}$$

$$\begin{aligned} \log \frac{P(C_1 | \mathbf{x})}{1 - P(C_1 | \mathbf{x})} &= \log \frac{P(C_1 | \mathbf{x})}{P(C_2 | \mathbf{x})} \\ &= \log \frac{p(\mathbf{x} | C_1)}{p(\mathbf{x} | C_2)} + \log \frac{P(C_1)}{P(C_2)} \\ &= \log \frac{(2\pi)^{-d/2} |\Sigma|^{-1/2} \exp[-(1/2)(\mathbf{x} - \boldsymbol{\mu}_1)^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}_1)]}{(2\pi)^{-d/2} |\Sigma|^{-1/2} \exp[-(1/2)(\mathbf{x} - \boldsymbol{\mu}_2)^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}_2)]} + \log \frac{P(C_1)}{P(C_2)} \\ &= \mathbf{w}^T \mathbf{x} + w_0 \end{aligned}$$


where  $\mathbf{w} = \Sigma^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)$ ,  $w_0 = -\frac{1}{2}(\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2)^T \Sigma^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)$

The inverse of logit

$$\log \frac{P(C_1 | \mathbf{x})}{1 - P(C_1 | \mathbf{x})} = \mathbf{w}^T \mathbf{x} + w_0$$

$$P(C_1 | \mathbf{x}) = \text{sigmoid}(\mathbf{w}^T \mathbf{x} + w_0) = \frac{1}{1 + \exp[-(\mathbf{w}^T \mathbf{x} + w_0)]}$$

## Sigmoid (Logistic) Function



Calculate  $g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0$  and choose  $C_1$  if  $g(\mathbf{x}) > 0$ , or  
Calculate  $y = \text{sigmoid}(\mathbf{w}^T \mathbf{x} + w_0)$  and choose  $C_1$  if  $y > 0.5$

## Gradient-Descent

- $E(\mathbf{w} | X)$  is error with parameters  $\mathbf{w}$  on sample  $X$   
 $\mathbf{w}^* = \arg \min_{\mathbf{w}} E(\mathbf{w} | X)$


Gradient

$$\nabla_{\mathbf{w}} E = \begin{bmatrix} \frac{\partial E}{\partial w_1} & \frac{\partial E}{\partial w_2} & \dots & \frac{\partial E}{\partial w_d} \end{bmatrix}^T$$

- Gradient-descent:

Starts from random  $\mathbf{w}$  and updates  $\mathbf{w}$  iteratively in the negative direction of gradient

## Gradient-Descent



## Logistic Discrimination

Two classes: Assume log likelihood ratio is linear

$$\log \frac{p(\mathbf{x} | C_1)}{p(\mathbf{x} | C_2)} = \mathbf{w}^T \mathbf{x} + w_0$$

$$\log \frac{P(C_1 | \mathbf{x})}{1 - P(C_1 | \mathbf{x})} = \log \frac{p(\mathbf{x} | C_1)}{p(\mathbf{x} | C_2)} + \log \frac{P(C_1)}{P(C_2)}$$

$$= \mathbf{w}^T \mathbf{x} + w_0$$

where  $w_0 = w_0 + \log \frac{P(C_1)}{P(C_2)}$

$$y = \hat{P}(C_1 | \mathbf{x}) = \frac{1}{1 + \exp[-(\mathbf{w}^T \mathbf{x} + w_0)]}$$

## Training: Two Classes

$$\mathcal{X} = \{\mathbf{x}^t, \mathbf{r}^t\}_t, \mathbf{r}^t | \mathbf{x}^t \sim \text{Bernoulli}(y^t)$$

$$y = P(C_1 | \mathbf{x}) = \frac{1}{1 + \exp[-(\mathbf{w}^T \mathbf{x} + w_0)]}$$

$$l(\mathbf{w}, w_0 | \mathcal{X}) = \prod_t (y^t)^{r^t} (1 - y^t)^{1 - r^t}$$

$$E = -\log l$$

$$E(\mathbf{w}, w_0 | \mathcal{X}) = -\sum_t r^t \log y^t + (1 - r^t) \log(1 - y^t)$$

## Training: Gradient-Descent

$$E(\mathbf{w}, w_0 | \mathcal{X}) = -\sum_t r^t \log y^t + (1 - r^t) \log(1 - y^t)$$

If  $y = \text{sigmoid}(a)$ ,  $\frac{dy}{da} = y(1 - y)$


$$\Delta w_j = -\eta \frac{\partial E}{\partial w_j} = \eta \sum_t \left( \frac{r^t}{y^t} - \frac{1 - r^t}{1 - y^t} \right) y^t (1 - y^t) x_j^t$$

$$= \eta \sum_t (r^t - y^t) x_j^t, j = 1, \dots, d$$

$$\Delta w_0 = -\eta \frac{\partial E}{\partial w_0} = \eta \sum_t (r^t - y^t)$$

```

For j = 0, ..., d
  w_j ← rand(-0.01, 0.01)
Repeat
  For j = 0, ..., d
    Δw_j ← 0
  For t = 1, ..., N
    o ← 0
    For j = 0, ..., d
      o ← o + w_j x_j^t
    y ← sigmoid(o)
    Δw_j ← Δw_j + (r^t - y) x_j^t
  For j = 0, ..., d
    w_j ← w_j + η Δw_j
Until convergence
    
```



## K > 2 Classes

$$\mathcal{X} = \{\mathbf{x}^t, \mathbf{r}^t\}_t, \mathbf{r}^t | \mathbf{x}^t \sim \text{Mult}_K(\mathbf{1}, y^t)$$

$$\log \frac{p(\mathbf{x} | C_j)}{p(\mathbf{x} | C_k)} = \mathbf{w}^T \mathbf{x} + w_0$$

$$y = \hat{P}(C_j | \mathbf{x}) = \frac{\exp[\mathbf{w}_j^T \mathbf{x} + w_{j0}]}{\sum_{i=1}^K \exp[\mathbf{w}_i^T \mathbf{x} + w_{i0}]}, j = 1, \dots, K \quad \text{softmax}$$

$$l(\{\mathbf{w}_i, w_{i0}\} | \mathcal{X}) = \prod_t \prod_i (y_i^t)^{r_i^t}$$

$$E(\{\mathbf{w}_i, w_{i0}\} | \mathcal{X}) = -\sum_t \sum_i r_i^t \log y_i^t$$


$$\Delta w_j = \eta \sum_t (r_j^t - y_j^t) \mathbf{x}^t, \Delta w_{j0} = \eta \sum_t (r_j^t - y_j^t)$$

```

For i = 1, ..., K, For j = 0, ..., d, wij ← rand(-0.01, 0.01)
Repeat
  For i = 1, ..., K, For j = 0, ..., d, Δwij ← 0
  For t = 1, ..., N
    For i = 1, ..., K
      oi ← 0
      For j = 0, ..., d
        oi ← oi + wijxtj
    For i = 1, ..., K
      yt ← exp(o1) / (∑k exp(ok))
    For i = 1, ..., K
      For j = 0, ..., d
        Δwij ← Δwij + (rt - yt)xtj
  For i = 1, ..., K
    For j = 0, ..., d
      wij ← wij + ηΔwij
Until convergence

```

### Example



### Generalizing the Linear Model

- Quadratic:  $\log \frac{p(\mathbf{x}|C_i)}{p(\mathbf{x}|C_k)} = \mathbf{x}^T \mathbf{W}_i \mathbf{x} + \mathbf{w}_i^T \mathbf{x} + w_{i0}$
  - Sum of basis functions:  $\log \frac{p(\mathbf{x}|C_i)}{p(\mathbf{x}|C_k)} = \mathbf{w}_i^T \phi(\mathbf{x}) + w_{i0}$
- where  $\phi(\mathbf{x})$  are basis functions. Examples:
- Hidden units in neural networks (Chapters 11 and 12)
  - Kernels in SVM (Chapter 13)

### Discrimination by Regression


- Classes are NOT mutually exclusive and exhaustive
- $r^i = y^i + \varepsilon$  where  $\varepsilon \sim \mathcal{N}(0, \sigma^2)$
- $y^i = \text{sigmoid}(\mathbf{w}^T \mathbf{x}^i + w_0) = \frac{1}{1 + \exp[-(\mathbf{w}^T \mathbf{x}^i + w_0)]}$
- $J(\mathbf{w}, w_0 | \mathcal{X}) = \prod_i \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(r^i - y^i)^2}{2\sigma^2}\right]$
- $E(\mathbf{w}, w_0 | \mathcal{X}) = \frac{1}{2} \sum_i (r^i - y^i)^2$
- $\Delta \mathbf{w} = \eta \sum_i (r^i - y^i) \mathbf{x}^i (1 - y^i) \mathbf{x}^i$

### Learning to Rank

- Ranking: A different problem than classification or regression
- Let us say  $\mathbf{x}^u$  and  $\mathbf{x}^v$  are two instances, e.g., two movies
- We prefer  $u$  to  $v$  implies that  $g(\mathbf{x}^u) > g(\mathbf{x}^v)$  where  $g(\mathbf{x})$  is a score function, here linear:  $g(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$
- Find a direction  $\mathbf{w}$  such that we get the desired ranks when instances are projected along  $\mathbf{w}$

### Ranking Error

- We prefer  $u$  to  $v$  implies that  $g(\mathbf{x}^u) > g(\mathbf{x}^v)$ , so error is  $g(\mathbf{x}^v) - g(\mathbf{x}^u)$ , if  $g(\mathbf{x}^u) < g(\mathbf{x}^v)$
- $E(\mathbf{w} | \{r^u, r^v\}) = \sum_{r^u < r^v} [g(\mathbf{x}^v | \theta) - g(\mathbf{x}^u | \theta)]_+$
- where  $a_+$  is equal to  $a$  if  $a \geq 0$  and 0 otherwise.



Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
 3RD EDITION


ETHEM ALPAYDIN  
 © The MIT Press, 2014

alpaydin@boun.edu.tr  
<http://www.cmppe.boun.edu.tr/~ethem/i2ml3e>

CHAPTER 11:  
**MULTILAYER PERCEPTRONS**

### Neural Networks


- Networks of processing units (neurons) with connections (synapses) between them
- Large number of neurons:  $10^{10}$
- Large connectivity:  $10^5$
- Parallel processing
- Distributed computation/memory
- Robust to noise, failures




### Understanding the Brain

- Levels of analysis (Marr, 1982)
  - Computational theory
  - Representation and algorithm
  - Hardware implementation
- Reverse engineering: From hardware to theory
- Parallel processing: SIMD vs MIMD
- Neural net: SIMD with modifiable local memory
- Learning: Update by training/experience

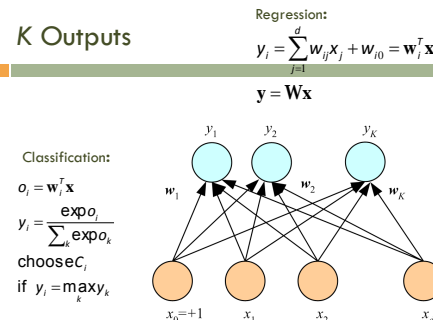
### Perceptron



### What a Perceptron Does



### K Outputs



### Training

- Online (instances seen one by one) vs batch (whole sample) learning:
  - No need to store the whole sample
  - Problem may change in time
  - Wear and degradation in system components
- Stochastic gradient-descent: Update after a single pattern
- Generic update rule (LMS rule):  $\Delta w_{ij}^t = \eta (r_i^t - y_i^t) \mathbf{x}_j^t$
- Update = Learning Factor \* (Desired Output - Actual Output) \* Input



## Training a Perceptron: Regression

- Regression (Linear output):

$$E^t(\mathbf{w} | \mathbf{x}^t, r^t) = \frac{1}{2} (r^t - y^t)^2 = \frac{1}{2} [r^t - (\mathbf{w}^T \mathbf{x}^t)]^2$$

$$\Delta w_j^t = \eta (r^t - y^t) x_j^t$$

## Classification

- Single sigmoid output

$$y^t = \text{sigmoid}(\mathbf{w}^T \mathbf{x}^t)$$

$$E^t(\mathbf{w} | \mathbf{x}^t, r^t) = -r^t \log y^t - (1 - r^t) \log(1 - y^t)$$


$$\Delta w_j^t = \eta (r^t - y^t) x_j^t$$

- K > 2 softmax outputs


$$y^t = \frac{\exp \mathbf{w}_i^T \mathbf{x}^t}{\sum_k \exp \mathbf{w}_k^T \mathbf{x}^t} \quad E^t(\mathbf{w}_i | \mathbf{x}^t, r^t) = -\sum_i r_i^t \log y_i^t$$

$$\Delta w_{ij}^t = \eta (r_i^t - y_i^t) x_j^t$$



## Learning Boolean AND




## XOR



## Multilayer Perceptrons



## Backpropagation



## Regression

$$E(\mathbf{W}, \mathbf{v} | \mathcal{X}) = \frac{1}{2} \sum_t (r^t - y^t)^2$$

$$y^t = \sum_{h=1}^H v_{jh} z_h^t + v_{j0}$$


$$\Delta v_{jh} = \sum_t (r^t - y^t) z_h^t$$

$$\Delta w_{hj} = -\eta \frac{\partial E}{\partial w_{hj}} = -\eta \sum_t \frac{\partial E}{\partial y^t} \frac{\partial y^t}{\partial z_h^t} \frac{\partial z_h^t}{\partial w_{hj}} = -\eta \sum_t (r^t - y^t) v_{jh} z_h^t (1 - z_h^t) x_j^t = \eta \sum_t (r^t - y^t) v_{jh} z_h^t (1 - z_h^t) x_j^t$$



Forward:  $z_h = \text{sigmoid}(\mathbf{w}_h^T \mathbf{x})$

Backward:  $\Delta w_{hj} = -\eta \frac{\partial E}{\partial w_{hj}}$

## Regression with Multiple Outputs



Initialize all  $v_{ih}$  and  $w_{hj}$  to  $\text{rand}(-0.01, 0.01)$   
 Repeat  
 For all  $(\mathbf{x}^t, r^t) \in \mathcal{X}$  in random order  
 For  $h = 1, \dots, H$   
 $z_h = \text{sigmoid}(\mathbf{w}_h^T \mathbf{x}^t)$   
 For  $i = 1, \dots, K$   
 $y_i = \mathbf{v}_i^T \mathbf{z}$   
 For  $i = 1, \dots, K$   
 $\Delta v_i = \eta (r_i^t - y_i^t) \mathbf{z}$   
 For  $h = 1, \dots, H$   
 $\Delta w_{hj} = \eta \left( \sum_i (r_i^t - y_i^t) v_{jh} \right) z_h (1 - z_h) \mathbf{x}^t$   
 For  $i = 1, \dots, K$   
 $\mathbf{v}_i \leftarrow \mathbf{v}_i + \Delta \mathbf{v}_i$   
 For  $h = 1, \dots, H$   
 $\mathbf{w}_h \leftarrow \mathbf{w}_h + \Delta \mathbf{w}_h$   
 Until convergence



## Two-Class Discrimination

- One sigmoid output  $y^t$  for  $P(C_1 | \mathbf{x}^t)$  and  $P(C_2 | \mathbf{x}^t) \equiv 1 - y^t$

$$y^t = \text{sigmoid}\left(\sum_{h=1}^H v_{jh} z_h^t + v_{j0}\right)$$

$$E(\mathbf{W}, \mathbf{v} | \mathcal{X}) = -\sum_t r^t \log y^t + (1 - r^t) \log(1 - y^t)$$

$$\Delta v_{j0} = \eta \sum_t (r^t - y^t) z_{j0}^t$$

$$\Delta w_{hj} = \eta \sum_t (r^t - y^t) v_{jh} z_h^t (1 - z_h^t) x_j^t$$

## K > 2 Classes

$$o_k^t = \sum_{h=1}^H v_{kh} z_h^t + v_{k0} \quad y^t = \frac{\exp o_k^t}{\sum_k \exp o_k^t} \equiv P(C_k | \mathbf{x}^t)$$

$$E(\mathbf{W}, \mathbf{v} | \mathcal{X}) = -\sum_t \sum_i r_i^t \log y_i^t$$

$$\Delta v_{k0} = \eta \sum_t (r_k^t - y_k^t) z_{k0}^t$$

$$\Delta w_{hj} = \eta \sum_t \left[ \sum_i (r_i^t - y_i^t) v_{jh} \right] z_h^t (1 - z_h^t) x_j^t$$

## Multiple Hidden Layers

- MLP with one hidden layer is a universal approximator (Hornik et al., 1989), but using multiple layers may lead to simpler networks

$$z_{1h} = \text{sigmoid}(\mathbf{w}_{1h}^T \mathbf{x}) = \text{sigmoid}\left(\sum_{j=1}^J w_{1hj} x_j + w_{1h0}\right), h = 1, \dots, H_1$$

$$z_{2l} = \text{sigmoid}(\mathbf{w}_{2l}^T \mathbf{z}_1) = \text{sigmoid}\left(\sum_{h=1}^{H_1} w_{2lh} z_{1h} + w_{2l0}\right), l = 1, \dots, H_2$$

$$y = \mathbf{v}^T \mathbf{z}_2 = \sum_{l=1}^{H_2} v_{jl} z_{2l} + v_{j0}$$

## Improving Convergence

- Momentum



$$\Delta w_i^t = -\eta \frac{\partial E^t}{\partial w_i} + \alpha \Delta w_i^{t-1}$$

- Adaptive learning rate

$$\Delta \eta = \begin{cases} +a & \text{if } E^{t+\tau} < E^t \\ -b\eta & \text{otherwise} \end{cases}$$


## Overfitting/Overtraining

Number of weights:  $H(d+1)+(H+1)K$



## Hints

Invariance to translation, rotation, size



Virtual examples (Abu-Mostafa, 1995)

Augmented error:  $E' = E + \lambda_i E_i$

If  $x'$  and  $x$  are the "same":  $E_i = [g(x|\theta) - g(x'|\theta)]^2$

Approximation hint:

$$E_i = \begin{cases} 0 & \text{if } g(x|\theta) \in [a, b] \\ (g(x|\theta) - a)^2 & \text{if } g(x|\theta) < a \\ (g(x|\theta) - b)^2 & \text{if } g(x|\theta) > b \end{cases}$$

## Tuning the Network Size

Destructive


Weight decay:

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

$$E' = E + \frac{\lambda}{2} \sum_i w_i^2$$


Constructive

Growing networks




## Structured MLP

Convolutional networks (Deep learning)




(Le Cun et al, 1989)

## Weight Sharing



## Hints

Invariance to translation, rotation, size



Virtual examples (Abu-Mostafa, 1995)

Augmented error:  $E' = E + \lambda_i E_i$

If  $x'$  and  $x$  are the "same":  $E_i = [g(x|\theta) - g(x'|\theta)]^2$

Approximation hint:

$$E_i = \begin{cases} 0 & \text{if } g(x|\theta) \in [a, b] \\ (g(x|\theta) - a)^2 & \text{if } g(x|\theta) < a \\ (g(x|\theta) - b)^2 & \text{if } g(x|\theta) > b \end{cases}$$

## Bayesian Learning

Consider weights  $w_i$  as random vars, prior  $p(w_i)$

$$p(w|\mathcal{X}) = \frac{p(\mathcal{X}|w)p(w)}{p(\mathcal{X})} \quad \hat{w}_{MAP} = \arg \max_w \log p(w|\mathcal{X})$$

$$\log p(w|\mathcal{X}) = \log p(\mathcal{X}|w) + \log p(w) + C$$

$$p(w) = \prod_i p(w_i) \quad \text{where } p(w_i) = c \cdot \exp\left[-\frac{w_i^2}{2(1/2\lambda)}\right]$$


$$E' = E + \lambda \|w\|^2$$

Weight decay, ridge regression, regularization

cost = data-misfit +  $\lambda$  complexity

More about Bayesian methods in chapter 14

## Dimensionality Reduction



Autoencoder networks

## Learning Time


Applications:

- Sequence recognition: Speech recognition
- Sequence reproduction: Time-series prediction
- Sequence association


Network architectures

- Time-delay networks (Waibel et al., 1989)
- Recurrent networks (Rumelhart et al., 1986)

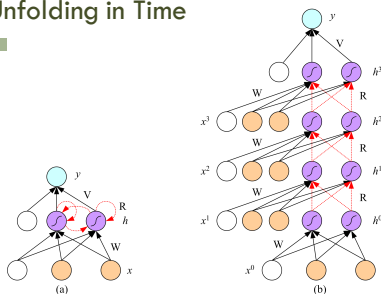
## Time-Delay Neural Networks



## Recurrent Networks



## Unfolding in Time



## Deep Networks

Layers of feature extraction units

Can have local receptive fields as in convolution networks, or can be fully connected

Can be trained layer by layer using an autoencoder in an unsupervised manner

No need to craft the right features or the right basis functions or the right dimensionality reduction method; learns multiple layers of abstraction all by itself given a lot of data and a lot of computation

Applications in vision, language processing, ...

1 Mon Sep 24 11:33:54 2018 12

i2ml3e-chap12.pdf


Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014

alpaydm@boun.edu.tr  
<http://www.cmpe.boun.edu.tr/~ethem/i2ml3e>

### Introduction

- Divide the input space into local regions and learn simple (constant/linear) models in each patch



- Unsupervised: Competitive, online clustering
- Supervised: Radial-basis functions, mixture of experts

### Competitive Learning


$$E(\{m_i\}_{i=1}^k | \mathcal{X}) = \sum_i \sum_j b_j^i \|x^t - m_i\|$$

$$b_j^i = \begin{cases} 1 & \text{if } \|x^t - m_i\| = \min_l \|x^t - m_l\| \\ 0 & \text{otherwise} \end{cases}$$

$$\text{Batch } k\text{-means: } m_i = \frac{\sum_j b_j^i x^t}{\sum_j b_j^i}$$


Online  $k$ -means:

$$\Delta m_{ij} = -\eta \frac{\partial E^t}{\partial m_{ij}} = \eta b_j^i (x_j^t - m_{ij})$$



Initialize  $m_i, i = 1, \dots, k$ , for example, to  $k$  random  $x^t$   
Repeat  
For all  $x^t \in \mathcal{X}$  in random order  
 $i \leftarrow \arg \min_j \|x^t - m_j\|$   
 $m_i \leftarrow m_i + \eta(x^t - m_i)$   
Until  $m_i$  converge

Winner-take-all network




### Adaptive Resonance Theory

- Incremental; add a new cluster if not covered; defined by vigilance,  $\rho$

$$b_j^i = \|x^t - m_i\| = -\min_{l=1}^k \|x^t - m_l\|$$

$$\begin{cases} m_{i+1} \leftarrow x^t & \text{if } b_j > \rho \\ \Delta m_i = \eta(x^t - m_i) & \text{otherwise} \end{cases}$$



(Carpenter and Grossberg, 1988)

### Self-Organizing Maps


- Units have a neighborhood defined;  $m_i$  is "between"  $m_{i-1}$  and  $m_{i+1}$ , and are all updated together

- One-dim map:

(Kohonen, 1990)

$$\Delta m_i = \eta e(l, i) (x^t - m_i)$$


$$e(l, i) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left[-\frac{(l-i)^2}{2\sigma^2}\right]$$




### Radial-Basis Functions

- Locally-tuned units:


$$p_h^i = \exp\left[-\frac{\|x^t - m_h\|^2}{2s_h^2}\right]$$



$$y^t = \sum_{h=1}^H w_h p_h^i + w_0$$



### Local vs Distributed Representation



Local representation in the space of  $(p_1, p_2, p_3)$

$x^a: (1.0, 0.0, 0.0)$   
 $x^b: (0.0, 0.0, 1.0)$   
 $x^c: (1.0, 1.0, 0.0)$

Distributed representation in the space of  $(h_1, h_2)$

$x^a: (1.0, 1.0)$   
 $x^b: (0.0, 1.0)$   
 $x^c: (1.0, 0.0)$

### Training RBF

- Hybrid learning:

- First layer centers and spreads: Unsupervised  $k$ -means
- Second layer weights: Supervised gradient-descent

- Fully supervised

(Broomhead and Lowe, 1988; Moody and Darken, 1989)

### Regression

$$E(\{m_h, s_h, w_{h,i}\}_{h=1}^H | \mathcal{X}) = \frac{1}{2} \sum_t \sum_i (r_i^t - y_i^t)^2$$

$$y_i^t = \sum_{h=1}^H w_{h,i} p_h^i + w_{i0}$$

$$\Delta w_{h,i} = \eta \sum_t (r_i^t - y_i^t) p_h^i$$

$$\Delta m_{hj} = \eta \sum_t \left[ \sum_i (r_i^t - y_i^t) w_{h,i} \right] p_h^i \frac{(x_j^t - m_{hj})}{s_h^2}$$

$$\Delta s_h = \eta \sum_t \left[ \sum_i (r_i^t - y_i^t) w_{h,i} \right] p_h^i \frac{\|x^t - m_h\|^2}{s_h^3}$$

### Classification

$$E(\{m_h, s_h, w_{h,i}\}_{h=1}^H | \mathcal{X}) = -\sum_t \sum_i r_i^t \log y_i^t$$

$$y_i^t = \frac{\exp\left[\sum_{h=1}^H w_{h,i} p_h^i + w_{i0}\right]}{\sum_{k=1}^K \exp\left[\sum_{h=1}^H w_{h,i} p_h^i + w_{i0}\right]}$$

### Rule-Based Knowledge

IF  $(x_1 \approx a)$  AND  $(x_2 \approx b)$  OR  $(x_3 \approx c)$  THEN  $y = 0.1$

$$p_1 = \exp\left[-\frac{(x_1 - a)^2}{2s_1^2}\right] \cdot \exp\left[-\frac{(x_2 - b)^2}{2s_2^2}\right] \text{ with } w_1 = 0.1$$

$$p_2 = \exp\left[-\frac{(x_3 - c)^2}{2s_3^2}\right] \text{ with } w_2 = 0.1$$

- Incorporation of prior knowledge (before training)
- Rule extraction (after training) (Tresp et al., 1997)
- Fuzzy membership functions and fuzzy rules

### Normalized Basis Functions


$$g_h^i = \frac{p_h^i}{\sum_{l=1}^H p_l^i}$$

$$= \frac{\exp\left[-\frac{\|x^t - m_h\|^2}{2s_h^2}\right]}{\sum_l \exp\left[-\frac{\|x^t - m_l\|^2}{2s_l^2}\right]}$$

$$y_i^t = \sum_{h=1}^H w_{h,i} g_h^i$$

$$\Delta w_{h,i} = \eta \sum_t (r_i^t - y_i^t) g_h^i$$

$$\Delta m_{hj} = \eta \sum_t \sum_i (r_i^t - y_i^t) (w_{h,i} - y_i^t) g_h^i \frac{(x_j^t - m_{hj})}{s_h^2}$$




### Competitive Basis Functions

- Mixture model:  $p(r^t | x^t) = \sum_{h=1}^H p(h | x^t) p(r^t | h, x^t)$

$$p(h | x^t) = \frac{p(x^t | h) p(h)}{\sum_l p(x^t | l) p(l)}$$

$$g_h^i = \frac{\alpha_h \exp\left[-\frac{\|x^t - m_h\|^2}{2s_h^2}\right]}{\sum_l \alpha_l \exp\left[-\frac{\|x^t - m_l\|^2}{2s_l^2}\right]}$$

### Rules and Exceptions



### Regression

$$p(r^t | x^t) = \prod_{i=1}^I \frac{1}{\sqrt{2\pi\sigma}} \exp\left[-\frac{(r_i^t - y_i^t)^2}{2\sigma^2}\right]$$

$$\mathcal{L}(\{m_h, s_h, w_{h,i}\}_{h=1}^H | \mathcal{X}) = \sum_t \log \sum_{h=1}^H g_h^i \exp\left[-\frac{1}{2} \sum_i (r_i^t - y_i^t)^2\right]$$

$y_{i0}^t = w_{i0}$  is the constant fit

$$\Delta w_{h,i} = \eta \sum_t (r_i^t - y_i^t) f_h^i \quad \Delta m_{hj} = \eta \sum_t (f_h^i - g_h^i) \frac{(x_j^t - m_{hj})}{s_h^2}$$

$$f_h^i = \frac{g_h^i \exp\left[-(1/2) \sum_i (r_i^t - y_i^t)^2\right]}{\sum_l g_l^i \exp\left[-(1/2) \sum_i (r_i^t - y_i^t)^2\right]}$$

$$p(h | r, x) = \frac{p(h | x) p(r | h, x)}{\sum_l p(l | x) p(r | l, x)}$$

## Classification

$$\begin{aligned} \mathcal{L}(\{\mathbf{m}_n, s_n, \mathbf{w}_{ih}\}_{i,h} | \mathcal{X}) &= \sum_t \log \sum_n g_n^t \prod_i (y_{ih}^t)^{r_i^t} \\ &= \sum_t \log \sum_n g_n^t \exp \left[ \sum_i r_i^t \log y_{ih}^t \right] \\ y_{ih}^t &= \frac{\exp w_{ih}}{\sum_k \exp w_{kh}} \\ f_h^t &= \frac{g_h^t \exp \left[ \sum_i r_i^t \log y_{ih}^t \right]}{\sum_i g_i^t \exp \left[ \sum_i r_i^t \log y_{ih}^t \right]} \end{aligned}$$

## EM for RBF (Supervised EM)

□ E-step:  $f_h^t \equiv \rho(\mathbf{r}^t | h, \mathbf{x}^t)$


□ M-step: 
$$\mathbf{m}_n = \frac{\sum_i f_h^t \mathbf{x}^t}{\sum_i f_h^t}$$

$$s_n = \frac{\sum_i f_h^t (\mathbf{x}^t - \mathbf{m}_n)(\mathbf{x}^t - \mathbf{m}_n)^T}{\sum_i f_h^t}$$

$$\mathbf{w}_{ih} = \frac{\sum_i f_h^t r_i^t}{\sum_i f_h^t}$$


## Learning Vector Quantization

- $H$  units per class prelabeled (Kohonen, 1990)
- Given  $\mathbf{x}$ ,  $\mathbf{m}_j$  is the closest:
 
$$\begin{cases} \Delta \mathbf{m}_j = \eta(\mathbf{x}^t - \mathbf{m}_j) & \text{if label}(\mathbf{x}^t) = \text{label}(\mathbf{m}_j) \\ \Delta \mathbf{m}_j = -\eta(\mathbf{x}^t - \mathbf{m}_j) & \text{otherwise} \end{cases}$$




## Mixture of Experts

- In RBF, each local fit is a constant,  $w_{ih}$ , second layer weight
- In MoE, each local fit is a linear function of  $\mathbf{x}$ , a local expert:  $w_{ih} \equiv \mathbf{v}_{ih}^T \mathbf{x}^t$



(Jacobs et al., 1991)

## MoE as Models Combined



## Cooperative MoE

□ Regression

$$E(\{\mathbf{m}_n, s_n, \mathbf{w}_{ih}\}_{i,h} | \mathcal{X}) = \frac{1}{2} \sum_t \sum_i (r_i^t - y_i^t)^2$$

$$\Delta \mathbf{v}_{ih} = \eta \sum_t (r_i^t - y_i^t) g_h^t \mathbf{x}^t$$

$$\Delta \mathbf{m}_{hj} = \eta \sum_t (r_i^t - y_i^t) (\mathbf{w}_{ih} - g_h^t) \mathbf{x}^t$$

## Competitive MoE: Regression

$$\mathcal{L}(\{\mathbf{m}_n, s_n, \mathbf{w}_{ih}\}_{i,h} | \mathcal{X}) = \sum_t \log \sum_n g_n^t \exp \left[ -\frac{1}{2} \sum_i (r_i^t - y_{ih}^t)^2 \right]$$

$$y_{ih}^t = w_{ih} = \mathbf{v}_{ih}^T \mathbf{x}^t$$

$$\Delta \mathbf{v}_{ih} = \eta \sum_t (r_i^t - y_{ih}^t) g_h^t \mathbf{x}^t$$

$$\Delta \mathbf{m}_n = \eta \sum_t (f_n^t - g_n^t) \mathbf{x}^t$$


## Competitive MoE: Classification

$$\begin{aligned} \mathcal{L}(\{\mathbf{m}_n, s_n, \mathbf{w}_{ih}\}_{i,h} | \mathcal{X}) &= \sum_t \log \sum_n g_n^t \prod_i (y_{ih}^t)^{r_i^t} \\ &= \sum_t \log \sum_n g_n^t \exp \left[ \sum_i r_i^t \log y_{ih}^t \right] \\ y_{ih}^t &= \frac{\exp w_{ih}}{\sum_k \exp w_{kh}} = \frac{\exp \mathbf{v}_{ih}^T \mathbf{x}^t}{\sum_k \exp \mathbf{v}_{kh}^T \mathbf{x}^t} \end{aligned}$$

## Hierarchical Mixture of Experts

- Tree of MoE where each MoE is an expert in a higher-level MoE
- Soft decision tree: Takes a weighted (gating) average of all leaves (experts), as opposed to using a single path and a single leaf
- Can be trained using EM (Jordan and Jacobs, 1994)

## i2ml3e-chap13.pdf



## Kernel Machines

- Discriminant-based: No need to estimate densities first
- Define the discriminant in terms of support vectors
- The use of kernel functions, application-specific measures of similarity
- No need to represent instances as vectors
- Convex optimization problems with a unique solution

## Optimal Separating Hyperplane

$$\mathcal{X} = \{\mathbf{x}^t, r^t\}, \text{ where } r^t = \begin{cases} +1 & \text{if } \mathbf{x}^t \in C_1 \\ -1 & \text{if } \mathbf{x}^t \in C_2 \end{cases}$$

find  $\mathbf{w}$  and  $w_0$  such that

$$\mathbf{w}^T \mathbf{x}^t + w_0 \geq +1 \text{ for } r^t = +1$$

$$\mathbf{w}^T \mathbf{x}^t + w_0 \leq -1 \text{ for } r^t = -1$$

which can be rewritten as


$$r^t (\mathbf{w}^T \mathbf{x}^t + w_0) \geq +1$$

(Cortes and Vapnik, 1995; Vapnik, 1995)

## Margin

- Distance from the discriminant to the closest instances on either side
  - Distance of  $\mathbf{x}$  to the hyperplane is  $\frac{|\mathbf{w}^T \mathbf{x}^t + w_0|}{\|\mathbf{w}\|}$
  - We require  $\frac{r^t (\mathbf{w}^T \mathbf{x}^t + w_0)}{\|\mathbf{w}\|} \geq \rho, \forall t$
  - For a unique soln, fix  $\rho | \|\mathbf{w}\| = 1$ , and to max margin
- $$\min \frac{1}{2} \|\mathbf{w}\|^2 \text{ subject to } r^t (\mathbf{w}^T \mathbf{x}^t + w_0) \geq +1, \forall t$$

## Margin



$$\min \frac{1}{2} \|\mathbf{w}\|^2 \text{ subject to } r^i(\mathbf{w}^T \mathbf{x}^i + w_0) \geq +1, \forall i$$

$$L_p = \frac{1}{2} \|\mathbf{w}\|^2 - \sum_{i=1}^N \alpha^i [r^i(\mathbf{w}^T \mathbf{x}^i + w_0) - 1]$$

$$= \frac{1}{2} \|\mathbf{w}\|^2 - \sum_{i=1}^N \alpha^i r^i(\mathbf{w}^T \mathbf{x}^i + w_0) + \sum_{i=1}^N \alpha^i$$

$$\frac{\partial L_p}{\partial \mathbf{w}} = 0 \Rightarrow \mathbf{w} = \sum_{i=1}^N \alpha^i r^i \mathbf{x}^i$$

$$\frac{\partial L_p}{\partial w_0} = 0 \Rightarrow \sum_{i=1}^N \alpha^i r^i = 0$$

$$L_d = \frac{1}{2} (\mathbf{w}^T \mathbf{w}) - \mathbf{w}^T \sum_{i=1}^N \alpha^i r^i \mathbf{x}^i - w_0 \sum_{i=1}^N \alpha^i r^i + \sum_{i=1}^N \alpha^i$$

$$= -\frac{1}{2} (\mathbf{w}^T \mathbf{w}) + \sum_{i=1}^N \alpha^i$$

$$= -\frac{1}{2} \sum_{i=1}^N \sum_{s=1}^N \alpha^i \alpha^s r^i r^s (\mathbf{x}^i)^T \mathbf{x}^s + \sum_{i=1}^N \alpha^i$$

subject to  $\sum_{i=1}^N \alpha^i r^i = 0$  and  $\alpha^i \geq 0, \forall i$

Most  $\alpha^i$  are 0 and only a small number have  $\alpha^i > 0$ ; they are the support vectors

## Soft Margin Hyperplane

- Not linearly separable


$$r^i(\mathbf{w}^T \mathbf{x}^i + w_0) \geq 1 - \xi^i$$

- Soft error


$$\sum_{i=1}^N \xi^i$$

- New primal is

$$L_p = \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^N \xi^i - \sum_{i=1}^N \alpha^i [r^i(\mathbf{w}^T \mathbf{x}^i + w_0) - 1 + \xi^i] - \sum_{i=1}^N \mu^i \xi^i$$



## Hinge Loss



## v-SVM

$$\min \frac{1}{2} \|\mathbf{w}\|^2 - \nu \rho + \frac{1}{N} \sum_{i=1}^N \xi^i$$

subject to

$$r^i(\mathbf{w}^T \mathbf{x}^i + w_0) \geq \rho - \xi^i, \xi^i \geq 0, \rho \geq 0$$

$$L_p = -\frac{1}{2} \sum_{i=1}^N \sum_{s=1}^N \alpha^i \alpha^s r^i r^s (\mathbf{x}^i)^T \mathbf{x}^s$$

subject to

$$\sum_{i=1}^N \alpha^i r^i = 0, 0 \leq \alpha^i \leq \frac{1}{N}, \sum_{i=1}^N \alpha^i \leq \nu$$

$\nu$  controls the fraction of support vectors

## Kernel Trick

- Preprocess input  $\mathbf{x}$  by basis functions

$$\mathbf{z} = \boldsymbol{\varphi}(\mathbf{x}) \quad g(\mathbf{z}) = \mathbf{w}^T \mathbf{z}$$

$$g(\mathbf{x}) = \mathbf{w}^T \boldsymbol{\varphi}(\mathbf{x})$$

- The SVM solution

$$\mathbf{w} = \sum_{i=1}^N \alpha^i r^i \mathbf{z}^i = \sum_{i=1}^N \alpha^i r^i \boldsymbol{\varphi}(\mathbf{x}^i)$$

$$g(\mathbf{x}) = \mathbf{w}^T \boldsymbol{\varphi}(\mathbf{x}) = \sum_{i=1}^N \alpha^i r^i \boldsymbol{\varphi}(\mathbf{x}^i)^T \boldsymbol{\varphi}(\mathbf{x})$$

$$g(\mathbf{x}) = \sum_{i=1}^N \alpha^i r^i K(\mathbf{x}^i, \mathbf{x})$$

## Vectorial Kernels

- Polynomials of degree  $q$ :


$$K(\mathbf{x}^i, \mathbf{x}) = (\mathbf{x}^i \mathbf{x}^T + 1)^q$$

$$K(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^T \mathbf{y} + 1)^2$$

$$= (x_1 y_1 + x_2 y_2 + 1)^2$$

$$= 1 + 2x_1 y_1 + 2x_2 y_2 + 2x_1 x_2 y_1 y_2 + x_1^2 y_1^2 + x_2^2 y_2^2$$


$$\boldsymbol{\phi}(\mathbf{x}) = [1, \sqrt{2}x_1, \sqrt{2}x_2, \sqrt{2}x_1 x_2, x_1^2, x_2^2]^T$$



## Vectorial Kernels

- Radial-basis functions:

$$K(\mathbf{x}^i, \mathbf{x}) = \exp\left[-\frac{\|\mathbf{x}^i - \mathbf{x}\|^2}{2s^2}\right]$$



## Defining kernels

- Kernel "engineering"
  - Defining good measures of similarity
  - String kernels, graph kernels, image kernels, ...
  - Empirical kernel map: Define a set of templates  $\mathbf{m}_i$  and score function  $s(\mathbf{x}, \mathbf{m}_i)$  and
- $$\boldsymbol{\phi}(\mathbf{x}^i) = [s(\mathbf{x}^i, \mathbf{m}_1), s(\mathbf{x}^i, \mathbf{m}_2), \dots, s(\mathbf{x}^i, \mathbf{m}_M)]$$
- and
- $$K(\mathbf{x}, \mathbf{x}') = \boldsymbol{\phi}(\mathbf{x})^T \boldsymbol{\phi}(\mathbf{x}')$$

## Multiple Kernel Learning

- Fixed kernel combination

$$K(\mathbf{x}, \mathbf{y}) = \begin{cases} cK(\mathbf{x}, \mathbf{y}) \\ K_1(\mathbf{x}, \mathbf{y}) + K_2(\mathbf{x}, \mathbf{y}) \\ K_1(\mathbf{x}, \mathbf{y}) K_2(\mathbf{x}, \mathbf{y}) \end{cases}$$

- Adaptive kernel combination

$$K(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^K \eta_i K_i(\mathbf{x}, \mathbf{y})$$

$$L_p = \sum_{i=1}^K \alpha^i - \frac{1}{2} \sum_{i=1}^K \sum_{s=1}^K \alpha^i \alpha^s r^i r^s \sum_{i=1}^K \eta_i K_i(\mathbf{x}^i, \mathbf{x}^s)$$

$$g(\mathbf{x}) = \sum_{i=1}^K \alpha^i r^i \sum_{i=1}^K \eta_i K_i(\mathbf{x}^i, \mathbf{x})$$

- Localized kernel combination  $g(\mathbf{x}) = \sum_{i=1}^K \alpha^i r^i \sum_{i=1}^K \eta_i(\mathbf{x}) \theta_i K_i(\mathbf{x}^i, \mathbf{x})$

## Multiclass Kernel Machines

- 1-vs-all
- Pairwise separation
- Error-Correcting Output Codes (section 17.5)
- Single multiclass optimization

$$\min \frac{1}{2} \sum_{i=1}^K \|\mathbf{w}_i\|^2 + C \sum_{i=1}^K \sum_{i=1}^K \xi^i$$

subject to

$$\mathbf{w}_i^T \mathbf{x}^i + w_{i0} \geq \mathbf{w}_j^T \mathbf{x}^i + w_{j0} + 2 - \xi^i, \forall i \neq j, \xi^i \geq 0$$

## SVM for Regression

- Use a linear model (possibly kernelized)

$$f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0$$

- Use the  $\epsilon$ -sensitive error function



$$e_\epsilon(r^i, f(\mathbf{x}^i)) = \begin{cases} 0 & \text{if } |r^i - f(\mathbf{x}^i)| < \epsilon \\ |r^i - f(\mathbf{x}^i)| - \epsilon & \text{otherwise} \end{cases}$$

$$\min \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^N (\xi_+^i + \xi_-^i)$$

$$r^i - (\mathbf{w}^T \mathbf{x} + w_0) \leq \epsilon + \xi_+^i$$


$$(\mathbf{w}^T \mathbf{x} + w_0) - r^i \leq \epsilon + \xi_-^i$$

$$\xi_+^i, \xi_-^i \geq 0$$



## Kernel Regression

- Polynomial kernel
- Gaussian kernel



## Kernel Machines for Ranking

- We require not only that scores be correct order but at least +1 unit margin.
- Linear case:

$$\min \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^K \sum_{i=1}^K \xi^i$$

subject to

$$\mathbf{w}^T \mathbf{x}^u \geq \mathbf{w}^T \mathbf{x}^v + 1 - \xi^t, \forall t: r^u < r^v, \xi^t \geq 0$$

## One-Class Kernel Machines

- Consider a sphere with center  $a$  and radius  $R$

$$\min R^2 + C \sum_i \xi_i^2$$



subject to

$$\|x^i - a\| \leq R^2 + \xi_i^2, \xi_i \geq 0$$

$$L_d = \sum_i \alpha^i \|x^i\|^2 - \sum_{i=1}^M \sum_{j=1}^M \alpha^i \alpha^j r^i r^j (x^i \cdot x^j)$$

subject to

$$0 \leq \alpha^i \leq C, \sum_i \alpha^i = 1$$




24

## Kernel Dimensionality Reduction

- Kernel PCA does PCA on the kernel matrix (equal to canonical PCA with a linear kernel)


- Kernel LDA, CCA



## Graphical Models

- Aka Bayesian networks, probabilistic networks
- Nodes are hypotheses (random vars) and the probabilities corresponds to our belief in the truth of the hypothesis
- Arcs are direct influences between hypotheses
- The structure is represented as a directed acyclic graph (DAG)
- The parameters are the conditional probabilities in the arcs (Pearl, 1988, 2000; Jensen, 1996; Lauritzen, 1996)

## Causes and Bayes' Rule



Diagnostic inference: Knowing that the grass is wet, what is the probability that rain is the cause?


$$P(R|W) = \frac{P(W|R)P(R)}{P(W)}$$

$$= \frac{P(W|R)P(R)}{P(W|R)P(R) + P(W|\sim R)P(\sim R)}$$

$$= \frac{0.9 \times 0.4}{0.9 \times 0.4 + 0.2 \times 0.6} = 0.75$$


## Case 2: Tail-to-Tail

- $P(X, Y, Z) = P(X)P(Y|X)P(Z|X, Y)$



## Case 3: Head-to-Head

- $P(X, Y, Z) = P(X)P(Y)P(Z|X, Y)$




## Large Margin Nearest Neighbor

- Learns the matrix  $M$  of Mahalanobis metric  $D(x^i, x^j) = (x^i - x^j)^T M (x^i - x^j)$
- For three instances  $i, j$ , and  $l$ , where  $i$  and  $j$  are of the same class and  $l$  different, we require  $D(x^i, x^l) > D(x^j, x^l) + 1$  and if this is not satisfied, we have a slack for the difference and we learn  $M$  to minimize the sum of such slacks over all  $i, j, l$  triples ( $j$  and  $l$  being one of  $k$  neighbors of  $i$ , over all  $i$ )

## Learning a Distance Measure

- LMNN algorithm (Weinberger and Saul 2009)  $(1 - \mu) \sum_{i,j} D(x^i, x^j) + \mu \sum_{i,j} (1 - \gamma_{ij}) \xi_{ij}$  subject to  $D(x^i, x^j) \geq D(x^i, x^l) + 1 - \xi_{ij}$ , if  $r^i = r^j$  and  $r^i \neq r^l$   $\xi_{ij} \geq 0$
- LMCA algorithm (Torresani and Lee 2007) uses a similar approach where  $M = L^T L$  and learns  $L$

i2ml3e-chap14.pdf





## Conditional Independence

- $X$  and  $Y$  are independent if  $P(X, Y) = P(X)P(Y)$
- $X$  and  $Y$  are conditionally independent given  $Z$  if  $P(X, Y | Z) = P(X | Z)P(Y | Z)$
- or  $P(X | Y, Z) = P(X | Z)$
- Three canonical cases: Head-to-tail, Tail-to-tail, head-to-head


## Case 1: Head-to-Head

- $P(X, Y, Z) = P(X)P(Y|X)P(Z|Y)$



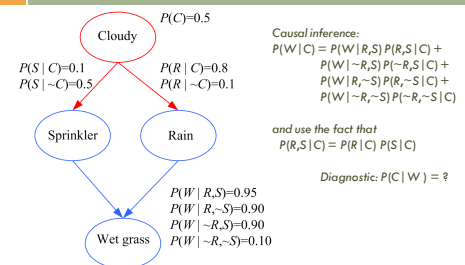
- $P(W|C) = P(W|R)P(R|C) + P(W|\sim R)P(\sim R|C)$

## Causal vs Diagnostic Inference



Diagnostic inference: If the grass is wet, what is the probability that the sprinkler is on?  $P(S|W) = 0.35 > 0.2 P(S)$   $P(S|R, W) = 0.21$  Explaining away: Knowing that it has rained decreases the probability that the sprinkler is on.

## Causes

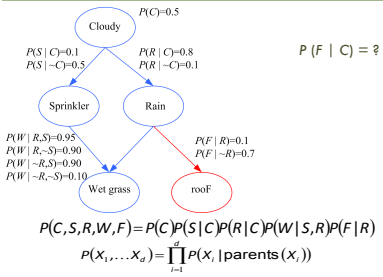


Causal inference:  $P(W|C) = P(W|R, S)P(R, S|C) + P(W|\sim R, S)P(\sim R, S|C) + P(W|R, \sim S)P(R, \sim S|C) + P(W|\sim R, \sim S)P(\sim R, \sim S|C)$


and use the fact that  $P(R, S|C) = P(R|C)P(S|C)$

Diagnostic:  $P(C|W) = ?$


## Exploiting the Local Structure



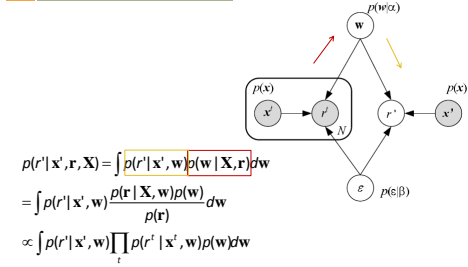
## Classification



## Naive Bayes' Classifier




## Linear Regression




## d-Separation


- A path from node A to node B is blocked if
  - The directions of edges on the path meet head-to-tail (case 1) or tail-to-tail (case 2) and the node is in C, or
  - The directions of edges meet head-to-head (case 3) and neither that node nor any of its descendants is in C.
- If all paths are blocked, A and B are d-separated (conditionally independent) given C.




## Belief Propagation (Pearl, 1988)



## Trees




## Polytrees



## Junction Trees

- If  $X$  does not separate  $E^+$  and  $E^-$ , we convert it into a junction tree and then apply the polytree algorithm




## Undirected Graphs: Markov Random Fields

- In a Markov random field, dependencies are symmetric, for example, pixels in an image
- In an undirected graph, A and B are independent if removing C makes them unconnected.
- Potential function  $\psi_c(X_c)$  shows how favorable is the particular configuration X over the clique C
- The joint is defined in terms of the clique potentials


$$p(X) = \frac{1}{Z} \prod_c \psi_c(X_c) \text{ where normalizer } Z = \sum_X \prod_c \psi_c(X_c)$$

## Factor Graphs

- Define new factor nodes and write the joint in terms of them



## Influence Diagrams



i2ml3e-chap15.pdf

Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION

ETHEM ALPAYDM  
© The MIT Press, 2014

alpaydm@boun.edu.tr  
<http://www.cmp.e.boun.edu.tr/~ethem/i2ml3e>

CHAPTER 15:  
**HIDDEN MARKOV MODELS**


## Introduction

- Modeling dependencies in input; no longer iid
- Sequences:
  - Temporal: In speech; phonemes in a word (dictionary), words in a sentence (syntax, semantics of the language). In handwriting, pen movements
  - Spatial: In a DNA sequence; base pairs

## Discrete Markov Process

- $N$  states:  $S_1, S_2, \dots, S_N$  State at "time"  $t$ ,  $q_t = S_i$
- First-order Markov
 
$$P(q_{t+1}=S_j | q_t=S_i, q_{t-1}=S_k, \dots) = P(q_{t+1}=S_j | q_t=S_i)$$
- Transition probabilities
 
$$\alpha_{ij} \equiv P(q_{t+1}=S_j | q_t=S_i) \quad \alpha_{ij} \geq 0 \text{ and } \sum_{j=1}^N \alpha_{ij} = 1$$
- Initial probabilities
 
$$\pi_i \equiv P(q_1=S_i) \quad \sum_{i=1}^N \pi_i = 1$$

## Stochastic Automaton



## Example: Balls and Urns

- Three urns each full of balls of one color
    - $S_1$ : red,  $S_2$ : blue,  $S_3$ : green
- $$\Pi = [0.5, 0.2, 0.3]^T \quad \mathbf{A} = \begin{bmatrix} 0.4 & 0.3 & 0.3 \\ 0.2 & 0.6 & 0.2 \\ 0.1 & 0.1 & 0.8 \end{bmatrix}$$
- $$\mathbf{O} = \{S_1, S_2, S_3\}$$
- $$P(\mathbf{O} | \mathbf{A}, \Pi) = P(S_1) \cdot P(S_1 | S_1) \cdot P(S_1 | S_2) \cdot P(S_1 | S_3)$$
- $$= \pi_1 \cdot \alpha_{11} \cdot \alpha_{13} \cdot \alpha_{33}$$
- $$= 0.5 \cdot 0.4 \cdot 0.3 \cdot 0.8 = 0.048$$

## Balls and Urns: Learning

- Given  $K$  example sequences of length  $T$


$$\hat{\pi}_i = \frac{\#\{\text{sequences starting with } S_i\}}{\#\{\text{sequences}\}} = \frac{\sum_k I(q_1^k = S_i)}{K}$$

$$\hat{\alpha}_{ij} = \frac{\#\{\text{transitions from } S_i \text{ to } S_j\}}{\#\{\text{transitions from } S_i\}} = \frac{\sum_k \sum_{t=1}^{T-1} I(q_t^k = S_i \text{ and } q_{t+1}^k = S_j)}{\sum_k \sum_{t=1}^{T-1} I(q_t^k = S_i)}$$

## Hidden Markov Models

- States are not observable
- Discrete observations  $\{v_1, v_2, \dots, v_M\}$  are recorded; a probabilistic function of the state
- Emission probabilities
 
$$b_j(m) \equiv P(O_t = v_m | q_t = S_j)$$
- Example: In each urn, there are balls of different colors, but with different probabilities.
- For each observation sequence, there are multiple state sequences

## HMM Unfolded in Time



## Elements of an HMM

- $N$ : Number of states
- $M$ : Number of observation symbols
- $\mathbf{A} = [\alpha_{ij}]$ :  $N$  by  $N$  state transition probability matrix
- $\mathbf{B} = [b_j(m)]$ :  $N$  by  $M$  observation probability matrix
- $\Pi = [\pi_i]$ :  $N$  by 1 initial state probability vector

$$\lambda = (\mathbf{A}, \mathbf{B}, \Pi), \text{ parameter set of HMM}$$

## Three Basic Problems of HMMs

1. Evaluation: Given  $\lambda$ , and  $\mathbf{O}$ , calculate  $P(\mathbf{O} | \lambda)$
2. State sequence: Given  $\lambda$ , and  $\mathbf{O}$ , find  $Q^*$  such that
 
$$P(Q^* | \mathbf{O}, \lambda) = \max_Q P(Q | \mathbf{O}, \lambda)$$
3. Learning: Given  $X = \{\mathbf{O}^k\}_k$ , find  $\lambda^*$  such that
 
$$P(X | \lambda^*) = \max_\lambda P(X | \lambda)$$

(Rabiner, 1989)

## Evaluation

- Forward variable:

$$\alpha_t(i) \equiv P(O_1, \dots, O_t, q_t = S_i | \lambda)$$


Initialization:

$$\alpha_1(i) = \pi_i b_i(O_1)$$

Recursion:

$$\alpha_{t+1}(j) = \sum_{i=1}^N \alpha_t(i) a_{ij} b_j(O_{t+1})$$

$$P(\mathbf{O} | \lambda) = \sum_{i=1}^N \alpha_T(i)$$



Backward variable:


$$\beta_t(i) \equiv P(O_{t+1}, \dots, O_T | q_t = S_i, \lambda)$$

Initialization:

$$\beta_T(i) = 1$$

Recursion:

$$\beta_t(i) = \sum_{j=1}^N a_{ij} b_j(O_{t+1}) \beta_{t+1}(j)$$




13

## Finding the State Sequence

$$\gamma_t(i) \equiv P(q_t = S_i | \mathbf{O}, \lambda)$$

$$= \frac{\alpha_t(i) \beta_t(i)}{\sum_{j=1}^N \alpha_t(j) \beta_t(j)}$$



Choose the state that has the highest probability, for each time step:

$$q_t^* = \arg \max_j \gamma_t(j)$$

No!

## Viterbi's Algorithm


$$\bar{\delta}_t(i) \equiv \max_{q_1, q_2, \dots, q_{t-1}} P(q_1, q_2, \dots, q_{t-1}, q_t = S_i, O_1, \dots, O_t | \lambda)$$

- Initialization:
 
$$\bar{\delta}_1(i) = \pi_i b_i(O_1), \psi_1(i) = 0$$
- Recursion:
 
$$\bar{\delta}_t(i) = \max_j \bar{\delta}_{t-1}(j) \alpha_{ji} b_i(O_t), \psi_t(i) = \arg \max_j \bar{\delta}_{t-1}(j)$$
- Termination:
 
$$p^* = \max_i \bar{\delta}_T(i), q_T^* = \arg \max_i \bar{\delta}_T(i)$$
- Path backtracking:
 
$$q_t^* = \psi_{t+1}(q_{t+1}^*), t = T-1, T-2, \dots, 1$$

## Learning

$$\xi_t(i, j) \equiv P(q_t = S_i, q_{t+1} = S_j | \mathbf{O}, \lambda)$$

$$\hat{\xi}_t(i, j) = \frac{\alpha_t(i) a_{ij} b_j(O_{t+1}) \beta_{t+1}(j)}{\sum_k \sum_l \alpha_t(k) a_{kl} b_l(O_{t+1}) \beta_{t+1}(l)}$$



Baum-Welch algorithm (EM):

$$z_t^i = \begin{cases} 1 & \text{if } q_t = S_i \\ 0 & \text{otherwise} \end{cases} \quad z_{t+1}^j = \begin{cases} 1 & \text{if } q_t = S_i \text{ and } q_{t+1} = S_j \\ 0 & \text{otherwise} \end{cases}$$

## Baum-Welch (EM)

$$E\text{-step: } E[z_t^i] = \gamma_t(i) \quad E[z_{t+1}^j] = \xi_t(i, j)$$

M-step:

$$\hat{\pi}_i = \frac{\sum_{k=1}^K z_1^k(i)}{K} \quad \hat{\alpha}_{ij} = \frac{\sum_{k=1}^K \sum_{t=1}^{T-1} z_t^k(i) z_{t+1}^k(j)}{\sum_{k=1}^K \sum_{t=1}^{T-1} z_t^k(i)}$$

$$\hat{b}_j(m) = \frac{\sum_{k=1}^K \sum_{t=1}^{T-1} z_t^k(j) I(O_{t+1}^k = v_m)}{\sum_{k=1}^K \sum_{t=1}^{T-1} z_t^k(j)}$$

## Continuous Observations

- Discrete:
 
$$P(O_t | q_t = S_j, \lambda) = \prod_{m=1}^M b_j(m)^{o_{tm}} \quad b_j(m) = \begin{cases} 1 & \text{if } O_t = v_m \\ 0 & \text{otherwise} \end{cases}$$
- Gaussian mixture (Discretize using  $k$ -means):
 
$$P(O_t | q_t = S_j, \lambda) = \sum_{g=1}^G P(g_j) p(O_t | q_t = S_j, \mathcal{G}_j, \lambda) \sim \mathcal{N}(\mu_j, \Sigma_j)$$
- Continuous:
 
$$P(O_t | q_t = S_j, \lambda) \sim \mathcal{N}(\mu_j, \sigma_j^2)$$

Use EM to learn parameters, e.g.,

$$\hat{\mu}_j = \frac{\sum_t \gamma_t(j) O_t}{\sum_t \gamma_t(j)}$$



## HMM with Input

- Input-dependent observations:

$$P(O_t | q_t = S, x^t, \lambda) \sim \mathcal{N}(g(x^t | \theta), \sigma^2)$$

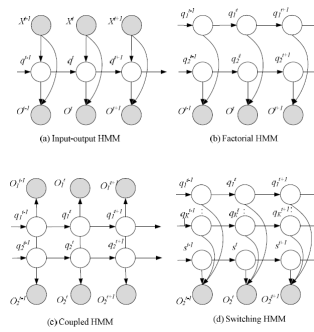

- Input-dependent transitions (Meila and Jordan, 1996; Bengio and Frasconi, 1996):

$$P(q_{t+1} = S, | q_t = S, x^t)$$

- Time-delay input:

$$x^t = \mathbf{f}(O_{t-r}, \dots, O_{t-1})$$

## HMM as a Graphical Model




21

## Model Selection in HMM

- Left-to-right HMMs:

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} & 0 \\ 0 & a_{22} & a_{23} & a_{24} \\ 0 & 0 & a_{33} & a_{34} \\ 0 & 0 & 0 & a_{44} \end{bmatrix}$$



- In classification, for each  $C_j$ , estimate  $P(O | \lambda_j)$  by a separate HMM and use Bayes' rule

$$P(\lambda_j | O) = \frac{P(O | \lambda_j) P(\lambda_j)}{\sum_l P(O | \lambda_l) P(\lambda_l)}$$

## Rationale

- Parameters  $\theta$  not constant, but random variables with a prior,  $p(\theta)$

- Bayes' Rule:  $p(\theta | X) = \frac{p(\theta) p(X | \theta)}{p(X)}$

## i2ml3e-chap16.pdf


Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014

alpaydm@buon.edu.fr  
http://www.cmlpe.buon.edu.fr/~ethem/i2ml3e

CHAPTER 16:  
**BAYESIAN ESTIMATION**

## Generative Model



$$p(x^1 | X) = \frac{p(x^1, X)}{p(X)} = \frac{\int p(x^1, X, \theta) d\theta}{p(X)} = \frac{\int p(\theta) p(x^1 | \theta) p(x^2 | \theta) \dots p(x^N | \theta) d\theta}{p(X)} = \int p(x^1 | \theta) p(\theta | X) d\theta$$

## Bayesian Approach

- Prior  $p(\theta)$  allows us to concentrate on region where  $\theta$  is likely to lie, ignoring regions where it's unlikely
- Instead of a single estimate with a single  $\theta$ , we generate several estimates using several  $\theta$  and average, weighted by how their probabilities

Even if prior  $p(\theta)$  is uninformative, (2) still helps.

MAP estimator does not make use of (2):

$$\theta_{MAP} = \arg \max_{\theta} p(\theta | X)$$


## Estimating the Parameters of a Distribution: Continuous case

- $p(x^t) \sim \mathcal{N}(\mu, \sigma^2)$
- Gaussian prior for  $\mu$ ,  $p(\mu) \sim \mathcal{N}(\mu_0, \sigma_0^2)$
- Posterior is also Gaussian  $p(\mu | X) \sim \mathcal{N}(\mu_N, \sigma_N^2)$

where

$$\mu_N = \frac{\sigma^2}{N\sigma_0^2 + \sigma^2} \mu_0 + \frac{N\sigma_0^2}{N\sigma_0^2 + \sigma^2} \bar{x}$$

$$\frac{1}{\sigma_N^2} = \frac{1}{\sigma_0^2} + \frac{N}{\sigma^2}$$



## Gaussian: Prior on Variance

- Let's define a prior (gamma) on precision  $\lambda = 1/\sigma^2$

$$p(\lambda) \sim \text{gamma}(a_0, b_0) = \frac{1}{\Gamma(a_0)} b_0^{a_0} \lambda^{a_0-1} \exp(-b_0 \lambda)$$

$$p(X | \lambda) = \prod_t \frac{\lambda^{1/2}}{\sqrt{2\pi}} \exp\left[-\frac{\lambda}{2}(x^t - \mu)^2\right]$$

$$= \lambda^{N/2} (2\pi)^{-N/2} \exp\left[-\frac{\lambda}{2} \sum_t (x^t - \mu)^2\right]$$

$$p(\lambda | X) \propto p(X | \lambda) p(\lambda) \sim \text{gamma}(a_N, b_N)$$

$$a_N = a_0 + N/2 = \frac{v_0 + N}{2}$$

$$b_N = b_0 + \frac{N}{2} s^2 = \frac{v_0}{2} s_0^2 + \frac{N}{2} s^2$$

## Bayesian Approach

$$p(x^1 | X) = \int p(x^1 | \theta) p(\theta | X) d\theta$$

- In certain cases, it is easy to integrate
- Conjugate prior: Posterior has the same density as prior
- Sampling (Markov Chain Monte Carlo): Sample from the posterior and average
- Approximation: Approximate the posterior with a model easier to integrate
  - Laplace approximation: Use a Gaussian
  - Variational approximation: Split the multivariate density into a set of simpler densities using independencies

## Joint Prior and Making a Prediction

$$p(\mu, \lambda) = p(\mu | \lambda) p(\lambda)$$

$$p(\mu, \lambda | X) \sim \text{normal-gamma}(\mu_N, \kappa_N, a_N, b_N)$$

where

$$\kappa_N = \kappa_0 + N$$

$$\mu_N = \frac{\kappa_0 \mu_0 + N \bar{x}}{\kappa_N}$$

$$a_N = a_0 + N/2$$

$$b_N = b_0 + \frac{N}{2} s^2 + \frac{\kappa_0 N}{2 \kappa_N} (m - \mu_0)^2$$

$$p(x | X) = \iint p(x | \mu, \lambda) p(\mu, \lambda | X) d\mu d\lambda \sim t_{2a_N} \left( \mu_N, \frac{b_N (\kappa_N + 1)}{a_N \kappa_N} \right)$$

## Estimating the Parameters of a Distribution: Discrete case

- $x_t^i = 1$  if in instance  $t$  is in state  $i$ , probability of state  $i$  is  $q_i$
- Dirichlet prior,  $\alpha$ , are hyperparameters

$$\text{Sample likelihood} \quad \text{Dirichlet}(\mathbf{q} | \mathbf{a}) = \frac{\Gamma(\sum \alpha_i)}{\prod \Gamma(\alpha_i)} \prod q_i^{\alpha_i - 1}$$

$$p(X | \mathbf{q}) = \prod_{t=1}^N \prod_{i=1}^K q_i^{x_t^i}$$

$$\text{Posterior} \quad p(\mathbf{q} | \mathbf{a}) = \frac{\Gamma(\sum \alpha_i + N)}{\Gamma(\sum \alpha_i) \Gamma(N)} \prod_{i=1}^K q_i^{\alpha_i + n_i - 1} = \text{Dirichlet}(\mathbf{q} | \mathbf{a} + \mathbf{n})$$

- Dirichlet is a conjugate prior
- With  $K=2$ , Dirichlet reduced to Beta

## Multivariate Gaussian

$$p(\mathbf{x}) \sim \mathcal{N}_d(\boldsymbol{\mu}, \Lambda) \quad p(\boldsymbol{\mu} | \Lambda) \sim \mathcal{N}_d(\boldsymbol{\mu}_0, (1/\kappa_0)\Lambda) \quad p(\Lambda) \sim \text{Wishart}(v_0, \mathbf{V}_0)$$

$$p(\boldsymbol{\mu}, \Lambda) = p(\boldsymbol{\mu} | \Lambda) p(\Lambda)$$

$$\sim \text{normal-Wishart}(\boldsymbol{\mu}_0, \kappa_0, v_0, \mathbf{V}_0)$$

$$p(\boldsymbol{\mu}, \Lambda | X) \sim \text{normal-Wishart}(\boldsymbol{\mu}_N, \kappa_N, v_N, \mathbf{V}_N)$$

$$\kappa_N = \kappa_0 + N$$

$$\boldsymbol{\mu}_N = \frac{\kappa_0 \boldsymbol{\mu}_0 + N \bar{\mathbf{x}}}{\kappa_N}$$

$$v_N = v_0 + N$$

$$\mathbf{V}_N = \left( \mathbf{V}_0^{-1} + \mathbf{C} + \frac{\kappa_0 N}{\kappa_N} (\mathbf{m} - \boldsymbol{\mu}_0)(\mathbf{m} - \boldsymbol{\mu}_0)^T \right)^{-1}$$

$$p(\mathbf{x} | X) = \iint p(\mathbf{x} | \boldsymbol{\mu}, \Lambda) p(\boldsymbol{\mu}, \Lambda | X) d\boldsymbol{\mu} d\Lambda$$

$$\sim t_{v_N - d + 1} \left( \boldsymbol{\mu}_N, \frac{\kappa_N + 1}{\kappa_N (v_N - d + 1)} (\mathbf{V}_N)^{-1} \right)$$

## Estimating the Parameters of a Function: Regression

- $r = w^T x + \epsilon$ ,  $p(\epsilon) \sim \mathcal{N}(0, 1/\beta)$ , and  $p(r^T | x^T, w, \beta) \sim \mathcal{N}(w^T x^T, 1/\beta)$
  - Log likelihood  $\ell(r | X, w, \beta) = \log \prod_t p(r^T | x^T, w, \beta)$
- $$= -N \log \sqrt{2\pi} - \frac{\beta}{2} \sum_t (r^T - w^T x^T)^2$$


ML solution  $w_{ML} = (X^T X)^{-1} X^T r$

- Gaussian conjugate prior:  $p(w) \sim \mathcal{N}(0, 1/\alpha)$
- Posterior:  $p(w | X) \sim \mathcal{N}(\mu_N, \Sigma_N)$  where

$$\mu_N = \beta \Sigma_N X^T r$$

$$\Sigma_N = (\alpha I + \beta X^T X)^{-1}$$

Aka ridge regression/parameter shrinkage/  
L2 regularization/weight decay



13

## Basis/Kernel Functions

- For new  $x^t$ , the estimate  $r^t$  is calculated as

$$r^t = (x^t)^T$$

$$= \beta (x^t)^T \Sigma_N X^T r \quad \text{Dual representation}$$

$$= \sum_t \beta (x^t)^T \Sigma_N x^t r^t$$

- Linear kernel
- For any other  $\phi(x)$ , we can write  $K(x^t, x^s) = \phi(x^t)^T \phi(x^s)$

$$r^t = \sum_t \beta (x^t)^T \Sigma_N x^t r^t \sum_t \beta K(x^t, x^s) r^t$$

16

## Mixture Model

$$p(x) = \sum_{i=1}^k p(\mathcal{G}_i) p(x | \mathcal{G}_i)$$

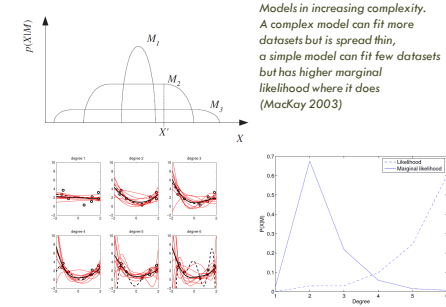

$$p(\Phi) = p(\pi) \prod_i p(\mu_i, \Lambda_i)$$

$$= \text{Dirichlet}(\pi | \alpha) \prod_i \text{normal-Wishart}(\mu_i, \nu_i, V_i)$$

$$Q(\Phi | \Phi^t) = \sum_t \sum_i h_i^t \log \pi_i + \sum_t \sum_i h_i^t \log p_i(x^t | \Phi^t) + \log p(\pi) + \sum_t \log p(\mu_i, \Lambda_i)$$

$$m_i^{t+1} = \frac{\alpha_i + N_i - 1}{\sum_j \alpha_j + N - k}$$

$$\Lambda_i^{t+1} = \left( \frac{K_0 \mu_0 + N_t m_t}{K_0 + N_t}, \frac{V_0^{-1} + C_t + S_t}{\nu_0 + N_t + d + 2} \right)^{-1}$$




21

## Dirichlet Processes

- Nonparametric Bayesian approach for clustering
- Chinese restaurant process
- Customers arrive and either join one of the existing tables or start a new one, based on the table occupancies:

$$\text{Join existing table } i \text{ with } P(z_i = 1) = \frac{n_i}{\alpha + n - 1}, i = 1, \dots, k$$

$$\text{Start new table with } P(z_{k+1} = 1) = \frac{\alpha}{\alpha + n - 1}$$



24

## Prior on Noise Variance

$$p(\beta) \sim \text{gamma}(a_0, b_0) \quad p(w | \beta) \sim \mathcal{N}(\mu_0, \beta \Sigma_0)$$

$$p(w, \beta) = p(\beta) p(w | \beta) \sim \text{normal-gamma}(\mu_0, \Sigma_0, a_0, b_0)$$

$$p(w, \beta | X, r) \sim \text{normal-gamma}(\mu_N, \Sigma_N, a_N, b_N)$$

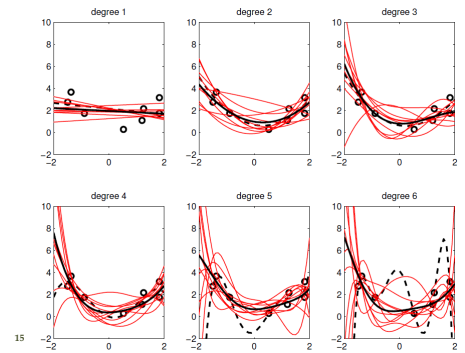
$$\Sigma_N = (X^T X + \Sigma_0)^{-1}$$

$$\mu_N = \Sigma_N (X^T r + \Sigma_0 \mu_0)$$

$$a_N = a_0 + N/2$$

$$b_N = b_0 + \frac{1}{2} (r^T r + \mu_0^T \Sigma_0 \mu_0 - \mu_N^T \Sigma_N \mu_N)$$

Markov Chain Monte Carlo (MCMC) sampling



15

## What's in a Prior?

- Defining a prior is subjective
- Uninformative prior if no prior preference
- How high to go?
  - Level I:  $p(x|X) = \int p(x|\theta) p(\theta|X) d\theta$
  - Level II:  $p(x|X) = \int p(x|\theta) p(\theta|X, \alpha) p(\alpha) d\theta d\alpha$
- Empirical Bayes: Use one good  $\alpha^*$ 
  - Level II ML:  $p(x|X) = \int p(x|\theta) p(\theta|X, \alpha^*) d\theta$

## Nonparametric Bayes

- Model complexity can increase with more data (in practice up to  $N$ , potentially to infinity)
- Similar to  $k$ -NN and Parzen windows we saw before where training set is the parameters

## Nonparametric Gaussian Mixture

- Tables are Gaussian components and decisions based both on prior and also on input  $x$ :

$$\text{Join component } i \text{ with } P(z_i^t = 1) \propto \frac{n_i}{\alpha + n - 1} p(x^t | X_i), i = 1, \dots, k$$

$$\text{Start new component with } P(z_{k+1}^t = 1) \propto \frac{\alpha}{\alpha + n - 1} p(x^t)$$

## Bayesian Model Comparison

- Marginal likelihood of a model:
 
$$p(X|M) = \int p(X|\theta, M) p(\theta|M) d\theta$$
- Posterior probability of model given data:
 
$$p(M|X) = \frac{p(X|M) p(M)}{p(X)}$$
- Bayes' factor:
 
$$\frac{P(M_1|X)}{P(M_0|X)} = \frac{P(X|M_1) P(M_1)}{P(X|M_0) P(M_0)}$$
- Approximations:
  - BIC:  $\log p(X|M) \approx \text{BIC} \approx \log p(X|\theta_{ML}, M) - \frac{|M|}{2} \log N$
  - AIC:  $\text{AIC} \approx \log p(X|\theta_{ML}, M) - |M|$

## Gaussian Processes

- Nonparametric model for supervised learning
- Assume Gaussian prior  $p(w) \sim \mathcal{N}(0, 1/\alpha)$
- $y = Xw$ , where  $E[y] = 0$  and  $\text{Cov}[y] = K$  with  $K_{ij} = (x^i)^T x^j$
- $K$  is the covariance function, here linear
- With basis function  $\phi(x)$ ,  $K_{ij} = (\phi(x^i))^T \phi(x^j)$
- $r \sim \mathcal{N}_N(0, C_N)$  where  $C_{ij} = (1/\beta) I + K$
- With new  $x^t$  added as  $x_{N+1}$ ,  $r_{N+1} \sim \mathcal{N}_{N+1}(0, C_{N+1})$


$$C_{N+1} = \begin{bmatrix} C_N & \mathbf{k} \\ \mathbf{k} & c \end{bmatrix}$$

where  $\mathbf{k} = [K(x^t, x^i)]^T$  and  $c = K(x^t, x^t) + 1/\beta$ .

$$p(r^t | x^t, X, r) \sim \mathcal{N}(k^T C_{N+1}^{-1} r, c - k^T C_{N+1}^{-1} k)$$

## Latent Dirichlet Allocation

- Bayesian feature extraction



23

Nonparametric Bayesian approach for feature extraction

Matrix factorization:

$$X = ZA \quad z_j^i = \begin{cases} 1 & \text{with probability } \mu_j \\ 0 & \text{with probability } 1 - \mu_j \end{cases}$$

$$\mu_j \sim \text{beta}(\alpha, 1)$$

- Nonparametric version: Allow  $j$  to increase with more data probabilistically
- Indian buffet process: Customer can take one of the existing dishes with prob  $\mu_j$  or add a new dish to the buffet

Rationale

- No Free Lunch Theorem: There is no algorithm that is always the most accurate
- Generate a group of base-learners which when combined has higher accuracy
- Different learners use different
  - Algorithms
  - Hyperparameters
  - Representations /Modalities/Views
  - Training sets
  - Subproblems
- Diversity vs accuracy

Error-Correcting Output Codes

- $K$  classes;  $L$  problems (Dietterich and Bakiri, 1995)
- Code matrix  $W$  codes classes in terms of learners

- One per class  $L=K$

$$W = \begin{bmatrix} +1 & -1 & -1 & -1 \\ -1 & +1 & -1 & -1 \\ -1 & -1 & +1 & -1 \\ -1 & -1 & -1 & +1 \end{bmatrix}$$

- Pairwise  $L=K(K-1)/2$


$$W = \begin{bmatrix} +1 & +1 & +1 & 0 & 0 & 0 \\ -1 & 0 & 0 & +1 & +1 & 0 \\ 0 & -1 & 0 & -1 & 0 & +1 \\ 0 & 0 & -1 & 0 & -1 & -1 \end{bmatrix}$$

Mixture of Experts

Voting where weights are input-dependent (gating)

$$y = \sum_{j=1}^L w_j d_j$$

(Jacobs et al., 1991) Experts or gating can be nonlinear



Voting


- Linear combination

$$y = \sum_{j=1}^L w_j d_j$$

$$w_j \geq 0 \text{ and } \sum_{j=1}^L w_j = 1$$

- Classification

$$y_i = \sum_{j=1}^L w_j d_{ji}$$



- Full code  $L=2^{(K-1)}$

$$W = \begin{bmatrix} -1 & -1 & -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & +1 & +1 & +1 & +1 \\ -1 & +1 & +1 & -1 & -1 & +1 & +1 \\ +1 & -1 & +1 & -1 & +1 & -1 & +1 \end{bmatrix}$$

- With reasonable  $L$ , find  $W$  such that the Hamming distance btw rows and columns are maximized.


- Voting scheme

$$y_i = \sum_{j=1}^L w_j d_{ji}$$

- Subproblems may be more difficult than one-per-K

Stacking

- Combiner  $f()$  is another learner (Wolpert, 1992)



Lecture Slides for  
INTRODUCTION  
TO  
MACHINE  
LEARNING  
3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014

alpaydin@boun.edu.tr  
http://www.cmp.e.boun.edu.tr/~ethem/i2ml3e

CHAPTER 17:

COMBINING MULTIPLE  
LEARNERS

- Bayesian perspective:

$$P(C_i | x) = \sum_{\text{all models } \mathcal{M}_j} P(C_i | x, \mathcal{M}_j) P(\mathcal{M}_j)$$

- If  $d_j$  are iid

$$E[y] = E\left[\sum_{j=1}^L d_j\right] = \frac{1}{L} \cdot E[d_j] = E[d_j]$$

$$\text{Var}(y) = \text{Var}\left(\sum_{j=1}^L d_j\right) = \frac{1}{L^2} \cdot \text{Var}\left(\sum_{j=1}^L d_j\right) = \frac{1}{L^2} \cdot L \cdot \text{Var}(d_j) = \frac{1}{L} \text{Var}(d_j)$$

Bias does not change, variance decreases by  $L$

- If dependent, error increase with positive correlation

$$\text{Var}(y) = \frac{1}{L^2} \text{Var}\left(\sum_{j=1}^L d_j\right) = \frac{1}{L^2} \left[ \sum_{j=1}^L \text{Var}(d_j) + 2 \sum_{i < j} \text{Cov}(d_i, d_j) \right]$$

Fixed Combination Rules

Rule	Fusion function $f(\cdot)$
Sum	$y_i = \sum_{j=1}^L d_{ji}$
Weighted sum	$y_i = \sum_{j=1}^L w_j d_{ji}, w_j \geq 0, \sum_{j=1}^L w_j = 1$
Median	$y_i = \text{median}_j d_{ji}$
Minimum	$y_i = \min_j d_{ji}$
Maximum	$y_i = \max_j d_{ji}$
Product	$y_i = \prod_j d_{ji}$

	$C_1$	$C_2$	$C_3$
$d_1$	0.2	0.5	0.3
$d_2$	0.0	0.6	0.4
$d_3$	0.4	0.4	0.2
Sum	0.2	0.5	0.3
Median	0.2	0.5	0.4
Minimum	0.0	0.4	0.2
Maximum	0.4	0.6	0.4
Product	0.0	0.12	0.032

Bagging

- Use bootstrapping to generate  $L$  training sets and train one base-learner with each (Breiman, 1996)
- Use voting (Average or median with regression)
- Unstable algorithms profit from bagging

AdaBoost

Generate a sequence of base-learners each focusing on previous one's errors (Freund and Schapire, 1996)

Training:

- For all  $\{x^i, y^i\}_{i=1}^N \in X$ , initialize  $p_1^i = 1/N$
- For all base-learners  $j = 1, \dots, L$
- Randomly draw  $x_j$  from  $X$  with probabilities  $p_j^i$
- Train  $d_j$  using  $x_j$
- For each  $\{x^i, y^i\}$ , calculate  $w_j^i = d_j(x^i)$
- Calculate error rate:  $\epsilon_j = \sum_{i=1}^N p_j^i \cdot 1(y_j^i \neq y^i)$
- If  $\epsilon_j > 1/2$ , then  $L = j - 1$ ; stop
- $\beta_j = \epsilon_j / (1 - \epsilon_j)$
- For each  $\{x^i, y^i\}$ , decrease probabilities if correct:
  - If  $y^i = d_j(x^i)$ ,  $p_{j+1}^i = \beta_j p_j^i$
  - Else  $p_{j+1}^i = p_j^i$
- Normalize probabilities:
  - $Z_j = \sum_{i=1}^N p_{j+1}^i$
  - $p_{j+1}^i = p_j^i / Z_j$

Testing:


- Given  $x$ , calculate  $d_j(x), j = 1, \dots, L$
- Calculate class outputs,  $i = 1, \dots, K$ :
  - $n_i = \sum_{j=1}^L \left( \text{arg max}_{y^i} \frac{1}{\beta_j} \right) d_j(x)$

Fine-Tuning an Ensemble

- Given an ensemble of dependent classifiers, do not use it as is, try to get independence
- Subset selection: Forward (growing)/Backward (pruning) approaches to improve accuracy/diversity/independence
- Train meta-classifiers: From the output of correlated classifiers, extract new combinations that are uncorrelated. Using PCA, we get "eigenlearners."
- Similar to feature selection vs feature extraction

Cascading

Use  $d_j$  only if preceding ones are not confident




Cascade learners in order of complexity

- Early integration: Concat all features and train a single learner
- Late integration: With each feature set, train one learner, then either use a fixed rule or stacking to combine decisions
- Intermediate integration: With each feature set, calculate a kernel, then use a single SVM with multiple kernels
- Combining features vs decisions vs kernels

Introduction

- Game-playing: Sequence of moves to win a game
- Robot in a maze: Sequence of actions to find a goal
- Agent has a state in an environment, takes an action and sometimes receives reward and the state changes
- Credit-assignment
- Learn a policy



$$V^*(s_t) = \max_{\pi} V^{\pi}(s_t), \forall s_t$$

$$= \max_{\pi} E \left[ \sum_{t=1}^{\infty} \gamma^{t-1} r_{t+i} \right]$$

$$= \max_{\pi} E \left[ r_{t+1} + \gamma \sum_{t=1}^{\infty} \gamma^{t-1} r_{t+i+1} \right]$$

$$= \max_{\pi} E [r_{t+1} + \gamma V^*(s_{t+1})] \quad \text{Bellman's equation}$$

$$V^*(s_t) = \max_{a_t} \left( E[r_{t+1}] + \gamma \sum_{s_{t+1}} P(s_{t+1} | s_t, a_t) V^*(s_{t+1}) \right)$$

$V^*(s_t) = \max_{a_t} Q^*(s_t, a_t)$  Value of  $a_t$  in  $s_t$

$$Q^*(s_t, a_t) = E[r_{t+1}] + \gamma \sum_{s_{t+1}} P(s_{t+1} | s_t, a_t) \max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1})$$

Temporal Difference Learning

- Environment,  $P(s_{t+1} | s_t, a_t)$ ,  $\rho(r_{t+1} | s_t, a_t)$ , is not known; model-free learning
- There is need for exploration to sample from  $P(s_{t+1} | s_t, a_t)$  and  $\rho(r_{t+1} | s_t, a_t)$
- Use the reward received in the next time step to update the value of current state (action)
- The temporal difference between the value of the current action and the value discounted from the next state

Single State: K-armed Bandit


- Among  $K$  levers, choose the one that pays best
- $Q(a)$ : value of action  $a$
- Reward is  $r_a$
- Set  $Q(a) = r_a$
- Choose  $a^*$  if  $Q(a^*) = \max_a Q(a)$
- Rewards stochastic (keep an expected reward):  $Q_{t+1}(a) \leftarrow Q_t(a) + \eta[r_{t+1}(a) - Q_t(a)]$

Model-Based Learning

- Environment,  $P(s_{t+1} | s_t, a_t)$ ,  $\rho(r_{t+1} | s_t, a_t)$  known
- There is no need for exploration
- Can be solved using dynamic programming
- Solve for  $V^*(s_t) = \max_{a_t} \left( E[r_{t+1}] + \gamma \sum_{s_{t+1}} P(s_{t+1} | s_t, a_t) V^*(s_{t+1}) \right)$
- Optimal policy  $\pi^*(s_t) = \arg \max_{a_t} \left( E[r_{t+1}] + \gamma \sum_{s_{t+1}} P(s_{t+1} | s_t, a_t) V^*(s_{t+1}) \right)$

Exploration Strategies

- $\epsilon$ -greedy: With pr  $\epsilon$ , choose one action at random uniformly; and choose the best action with pr  $1 - \epsilon$
- Probabilistic:  $P(a|s) = \frac{\exp Q(s,a)}{\sum_{b=1}^A \exp Q(s,b)}$
- Move smoothly from exploration/exploitation.
- Decrease  $\epsilon$
- Annealing  $P(a|s) = \frac{\exp(Q(s,a)/T)}{\sum_{b=1}^A \exp(Q(s,b)/T)}$



Elements of RL (Markov Decision Processes)

- $s_t$ : State of agent at time  $t$
- $a_t$ : Action taken at time  $t$
- In  $s_t$ , action  $a_t$  is taken, clock ticks and reward  $r_{t+1}$  is received and state changes to  $s_{t+1}$
- Next state prob:  $P(s_{t+1} | s_t, a_t)$
- Reward prob:  $\rho(r_{t+1} | s_t, a_t)$
- Initial state(s), goal state(s)
- Episode (trial) of actions from initial state to goal
- (Sutton and Barto, 1998; Kaelbling et al., 1996)

Value Iteration

Initialize  $V(s)$  to arbitrary values

Repeat

For all  $s \in S$

For all  $a \in A$

$$Q(s, a) \leftarrow E[r | s, a] + \gamma \sum_{s' \in S} P(s' | s, a) V(s')$$

$$V(s) \leftarrow \max_a Q(s, a)$$

Until  $V(s)$  converge

Deterministic Rewards and Actions

$$Q^*(s_t, a_t) = E[r_{t+1}] + \gamma \sum_{s_{t+1}} P(s_{t+1} | s_t, a_t) \max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1})$$

- Deterministic: single possible reward and next state  $Q(s_t, a_t) = r_{t+1} + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})$
- used as an update rule (backup)  $\hat{Q}(s_t, a_t) \leftarrow r_{t+1} + \gamma \max_{a_{t+1}} \hat{Q}(s_{t+1}, a_{t+1})$
- Starting at zero,  $Q$  values increase, never decrease

Policy and Cumulative Reward

- Policy,  $\pi: S \rightarrow \mathcal{A}$   $a_t = \pi(s_t)$
- Value of a policy,  $V^{\pi}(s_t)$
- Finite-horizon:  $V^{\pi}(s_t) = E[r_{t+1} + r_{t+2} + \dots + r_{t+T}] = E \left[ \sum_{i=1}^T r_{t+i} \right]$
- Infinite horizon:  $V^{\pi}(s_t) = E[r_{t+1} + \gamma r_{t+2} + \gamma^2 r_{t+3} + \dots] = E \left[ \sum_{i=1}^{\infty} \gamma^{i-1} r_{t+i} \right]$
- $0 \leq \gamma < 1$  is the discount rate

Policy Iteration

Initialize a policy  $\pi$  arbitrarily

Repeat

$\pi \leftarrow \pi'$


Compute the values using  $\pi$  by solving the linear equations

$$V^{\pi}(s) = E[r | s, \pi(s)] + \gamma \sum_{s' \in S} P(s' | s, \pi(s)) V^{\pi}(s')$$

Improve the policy at each state

$$\pi'(s) \leftarrow \arg \max_a (E[r | s, a] + \gamma \sum_{s' \in S} P(s' | s, a) V^{\pi}(s'))$$

Until  $\pi = \pi'$



Consider the value of action marked by \*:

If path A is seen first,  $Q^*(*) = 0.9 * \max(0, 81) = 73$

Then B is seen,  $Q^*(*) = 0.9 * \max(100, 81) = 90$

Or,

If path B is seen first,  $Q^*(*) = 0.9 * \max(100, 0) = 90$

Then A is seen,  $Q^*(*) = 0.9 * \max(100, 81) = 90$

$Q$  values increase but never decrease

# Nondeterministic Rewards and Actions

- When next states and rewards are nondeterministic (there is an opponent or randomness in the environment), we keep averages (expected values) instead as assignments

- Q-learning (Watkins and Dayan, 1992):

$$\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \eta \left( r_{t+1} + \gamma \max_{a'} \hat{Q}(s_{t+1}, a') - \hat{Q}(s, a) \right)$$

- Off-policy vs on-policy (Sarsa) backup
- Learning V (TD-learning: Sutton, 1988)

$$V(s_t) \leftarrow V(s_t) + \eta (r_{t+1} + \gamma V(s_{t+1}) - V(s_t))$$

# Q-learning

```

Initialize all Q(s, a) arbitrarily
For all episodes
  Initialize s
  Repeat
    Choose a using policy derived from Q, e.g., ε-greedy
    Take action a, observe r and s'
    Update Q(s, a):
      Q(s, a) ← Q(s, a) + η(r + γ max_{a'} Q(s', a') - Q(s, a))
    s ← s'
  Until s is terminal state
    
```

# Sarsa

```

Initialize all Q(s, a) arbitrarily
For all episodes
  Initialize s
  Choose a using policy derived from Q, e.g., ε-greedy
  Repeat
    Take action a, observe r and s'
    Choose a' using policy derived from Q, e.g., ε-greedy
    Update Q(s, a):
      Q(s, a) ← Q(s, a) + η(r + γ Q(s', a') - Q(s, a))
    s ← s', a ← a'
  Until s is terminal state
    
```


# Eligibility Traces

Keep a record of previously visited states (actions)

$$e_t(s, a) = \begin{cases} 1 & \text{if } s = s_t \text{ and } a = a_t \\ \gamma \lambda e_{t-1}(s, a) & \text{otherwise} \end{cases}$$

$$\delta_t = r_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)$$

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \eta \delta_t e_t(s, a), \forall s, a$$



# Sarsa (λ)

```

Initialize all Q(s, a) arbitrarily, e(s, a) ← 0, ∀ s, a
For all episodes
  Initialize s
  Choose a using policy derived from Q, e.g., ε-greedy
  Repeat
    Take action a, observe r and s'
    Choose a' using policy derived from Q, e.g., ε-greedy
    δ ← r + γ Q(s', a') - Q(s, a)
    e(s, a) ← 1
    For all s, a:
      Q(s, a) ← Q(s, a) + η δ e(s, a)
    e(s, a) ← γ λ e(s, a)
    s ← s', a ← a'
  Until s is terminal state
    
```

# Generalization

- Tabular: Q(s, a) or V(s) stored in a table
- Regressor: Use a learner to estimate Q(s, a) or V(s)

$$E^t(\theta) = [r_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)]^2$$

$$\Delta \theta = \eta [r_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)] \nabla_{\theta} Q(s_t, a_t)$$

Eligibility



$$\Delta \theta = \eta \delta_t e_t$$

$$\delta_t = r_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)$$

$$e_t = \gamma \lambda e_{t-1} + \nabla_{\theta} Q(s_t, a_t) \text{ with } e_0, \text{ all zeros}$$

# Partially Observable States

- The agent does not know its state but receives an observation  $p(o_{t+1} | s_t, a_t)$  which can be used to infer a belief about states
- Partially observable MDP



# The Tiger Problem

- Two doors, behind one of which there is a tiger
- p: prob that tiger is behind the left door


R(A, Z)	Tiger left	Tiger right
Open left	-100	+80
Open right	+90	-100

- $R(a_L) = -100p + 80(1-p)$ ,  $R(a_R) = 90p - 100(1-p)$
- We can sense with a reward of  $R(o_S) = -1$
- We have unreliable sensors
  - $P(o_L|Z_L) = 0.7$ ,  $P(o_L|Z_R) = 0.3$
  - $P(o_R|Z_L) = 0.3$ ,  $P(o_R|Z_R) = 0.7$

- Let us say the tiger can move from one room to the other with prob 0.8

$$p' = 0.2p + 0.8(1-p)$$

$$V = \max \begin{pmatrix} -100p' + 80(1-p') \\ 33p' + 26(1-p') \\ 90p' - 100(1-p') \end{pmatrix}$$



- If we sense  $o_L$ , our belief in tiger's position changes

$$p' = P(z_L | o_L) = \frac{P(o_L | z_L) P(z_L)}{P(o_L)} = \frac{0.7p}{0.7p + 0.3(1-p)}$$

$$R(a_L | o_L) = r(a_L, z_L) P(z_L | o_L) + r(a_L, z_R) P(z_R | o_L)$$

$$= -100p' + 80(1-p')$$

$$= -100 \frac{0.7p}{0.7p + 0.3(1-p)} + 80 \frac{0.3(1-p)}{0.7p + 0.3(1-p)}$$

$$R(a_R | o_L) = r(a_R, z_L) P(z_L | o_L) + r(a_R, z_R) P(z_R | o_L)$$

$$= 90p' - 100(1-p')$$

$$= 90 \frac{0.7p}{0.7p + 0.3(1-p)} - 100 \frac{0.3(1-p)}{0.7p + 0.3(1-p)}$$


$$R(o_S | o_L) = -1$$

$$V^i = \sum_j \max_a R(a, |o_i) P(o_i)$$

$$= \max \begin{pmatrix} -100p + 80(1-p) \\ -43p + 26(1-p) \\ 90p - 100(1-p) \end{pmatrix}$$

# i2ml3e-chap19.pdf

- When planning for episodes of two, we can take  $a_L$ ,  $a_R$ , or sense and wait:



Lecture Slides for  
**INTRODUCTION TO MACHINE LEARNING**  
3RD EDITION

ETHEM ALPAYDIN  
© The MIT Press, 2014

alpaydin@boun.edu.tr  
<http://www.cmlpe.boun.edu.tr/~ethem/i2ml3e>

CHAPTER 19:  
**DESIGN AND ANALYSIS OF MACHINE LEARNING EXPERIMENTS**

## Introduction


- Questions:
  - Assessment of the expected error of a learning algorithm: Is the error rate of 1-NN less than 2%?
  - Comparing the expected errors of two algorithms: Is k-NN more accurate than MLP?
- Training/validation/test sets
- Resampling methods: K-fold cross-validation

## Algorithm Preference

- Criteria (Application-dependent):
  - Misclassification error, or risk (loss functions)
  - Training time/space complexity
  - Testing time/space complexity
  - Interpretability
  - Easy programmability
- Cost-sensitive learning


## Factors and Response

- Response function based on output to be maximized
- Depends on controllable factors
- Uncontrollable factors introduce randomness
- Find the configuration of controllable factors that maximizes response and minimally affected by uncontrollable factors



## Strategies of Experimentation

### How to search the factor space?



Response surface design for approximating and maximizing the response function in terms of the controllable factors

## Guidelines for ML experiments

- Aim of the study
- Selection of the response variable
- Choice of factors and levels
- Choice of experimental design
- Performing the experiment
- Statistical Analysis of the Data
- Conclusions and Recommendations

## Resampling and K-Fold Cross-Validation

- The need for multiple training/validation sets  $\{X_i, V_i\}$ ; Training/validation sets of fold  $i$
- K-fold cross-validation: Divide X into  $k$ ,  $X_{i,j} = 1, \dots, K$ 

$$V_1 = X_1 \quad T_1 = X_2 \cup X_3 \cup \dots \cup X_K$$

$$V_2 = X_2 \quad T_2 = X_1 \cup X_3 \cup \dots \cup X_K$$

$$\vdots$$

$$V_K = X_K \quad T_K = X_1 \cup X_2 \cup \dots \cup X_{K-1}$$
- $T_i$  share K-2 parts

## 5x2 Cross-Validation

- 5 times 2 fold cross-validation (Dietterich, 1998)

$$\begin{aligned} T_1 &= X_1^{(1)} & V_1 &= X_1^{(2)} \\ T_2 &= X_1^{(2)} & V_2 &= X_1^{(1)} \\ T_3 &= X_2^{(1)} & V_3 &= X_2^{(2)} \\ T_4 &= X_2^{(2)} & V_4 &= X_2^{(1)} \\ &\vdots & & \\ T_9 &= X_5^{(1)} & V_9 &= X_5^{(2)} \\ T_{10} &= X_5^{(2)} & V_{10} &= X_5^{(1)} \end{aligned}$$

## Bootstrapping

- Draw instances from a dataset with replacement
- Prob that we do not pick an instance after N draws

$$\left(1 - \frac{1}{N}\right)^N \approx e^{-1} = 0.368$$



that is, only 36.8% is new!

## Performance Measures

True Class	Predicted class	
	Yes	No
Yes	TP: True Positive	FN: False Negative
No	FP: False Positive	TN: True Negative


- Error rate = # of errors / # of instances = (FN+FP) / N
- Recall = # of found positives / # of positives = TP / (TP+FN) = sensitivity = hit rate
- Precision = # of found positives / # of found = TP / (TP+FP)
- Specificity = TN / (TN+FP)
- False alarm rate = FP / (FP+TN) = 1 - Specificity

## ROC Curve




13

## Precision and Recall



## Interval Estimation

- $X = \{x^i\}$ , where  $x^i \sim N(\mu, \sigma^2)$
- $m \sim N(\mu, \sigma^2/N)$



$$P\left\{\sqrt{N} \frac{(m-\mu)}{\sigma} < 1.64\right\} = 0.95$$

$$P\left\{m - 1.64 \frac{\sigma}{\sqrt{N}} < \mu\right\} = 0.95$$

$$P\left\{m - z_{\alpha} \frac{\sigma}{\sqrt{N}} < \mu\right\} = 1 - \alpha$$

When  $\sigma^2$  is not known:

$$S^2 = \sum_i (x^i - m)^2 / (N-1) \quad \frac{\sqrt{N}(m-\mu)}{S} \sim t_{N-1}$$

$$P\left\{m - t_{\alpha/2, N-1} \frac{S}{\sqrt{N}} < \mu < m + t_{\alpha/2, N-1} \frac{S}{\sqrt{N}}\right\} = 1 - \alpha$$

16

## Hypothesis Testing

- Reject a null hypothesis if not supported by the sample with enough confidence
- $X = \{x^i\}$ , where  $x^i \sim N(\mu, \sigma^2)$
- $H_0: \mu = \mu_0$  vs.  $H_1: \mu \neq \mu_0$
- Accept  $H_0$  with level of significance  $\alpha$  if  $\mu_0$  is in the  $100(1-\alpha)$  confidence interval
- Two-sided test


Truth	Decision	
	Accept	Reject
True	Correct	Type I error
False	Type II error	Correct (Power)

- One-sided test:  $H_0: \mu \leq \mu_0$  vs.  $H_1: \mu > \mu_0$
- Accept if  $\frac{\sqrt{N}(m-\mu_0)}{\sigma} \in (-\infty, z_{\alpha})$
- Variance unknown: Use  $t$ , instead of  $z$
- Accept  $H_0: \mu = \mu_0$  if  $\frac{\sqrt{N}(m-\mu_0)}{S} \in (-t_{\alpha/2, N-1}, t_{\alpha/2, N-1})$

18


## Assessing Error: $H_0: p \leq p_0$ vs. $H_1: p > p_0$

- Single training/validation set: Binomial Test
- If error prob is  $p_0$ , prob that there are  $e$  errors or less in  $N$  validation trials is



## Normal Approximation to the Binomial

- Number of errors  $X$  is approx  $N$  with mean  $Np_0$  and var  $Np_0(1-p_0)$



## K-Fold CV Paired t Test

- Use  $K$ -fold cv to get  $K$  training/validation folds
- $p_i^1, p_i^2$ : Errors of classifiers 1 and 2 on fold  $i$
- $p_i = p_i^1 - p_i^2$ : Paired difference on fold  $i$
- The null hypothesis is whether  $p_i$  has mean 0

$H_0: \mu = 0$  vs.  $H_1: \mu \neq 0$

$$m = \frac{\sum_{i=1}^K p_i}{K} \quad s^2 = \frac{\sum_{i=1}^K (p_i - m)^2}{K-1}$$

$$\frac{\sqrt{K}(m-0)}{s} = \frac{\sqrt{K} \cdot m}{s} \sim t_{K-1} \text{ Accept if in } (-t_{\alpha/2, K-1}, t_{\alpha/2, K-1})$$

## 5x2 cv Paired t Test

- Use 5x2 cv to get 2 folds of 5 tra/val replications (Dietterich, 1998)
- $p_i^{(j)}$ : difference btw errors of 1 and 2 on fold  $j=1, 2$  of replication  $i=1, \dots, 5$

$$\bar{p}_i = (p_i^{(1)} + p_i^{(2)})/2 \quad s_i^2 = (p_i^{(1)} - \bar{p}_i)^2 + (p_i^{(2)} - \bar{p}_i)^2$$

$$\frac{\bar{p}_i^{(j)}}{\sqrt{\sum_{i=1}^5 s_i^2 / 5}} \sim t_5$$

Two-sided test: Accept  $H_0: \mu_0 = \mu_1$  if in  $(-t_{\alpha/2, 5}, t_{\alpha/2, 5})$   
 One-sided test: Accept  $H_0: \mu_0 \leq \mu_1$  if  $< t_{\alpha, 5}$

If  $H_0$  is true:

$$m_j = \frac{\sum_{i=1}^L X_{ij}}{L} \sim \mathcal{N}(\mu_j, \sigma^2 / K)$$

$$m = \frac{\sum_{j=1}^L m_j}{L} \quad s^2 = \frac{\sum_{j=1}^L (m_j - m)^2}{L-1}$$

Thus an estimator of  $\sigma^2$  is  $K \cdot S^2$ , namely,

$$\hat{\sigma}^2 = K \sum_{j=1}^L \frac{(m_j - m)^2}{L-1}$$

$$\sum_j \frac{(m_j - m)^2}{\sigma^2 / K} \sim \chi_{L-1}^2 \quad Ssb \equiv K \sum_j (m_j - m)^2$$

So when  $H_0$  is true, we have

$$\frac{Ssb}{\hat{\sigma}^2} \sim \chi_{L-1}^2$$

Regardless of  $H_0$ , our second estimator to  $\sigma^2$  is the average of group variances  $S_j^2$ :

$$S_j^2 = \frac{\sum_{i=1}^K (X_{ij} - m_j)^2}{K-1} \quad \hat{\sigma}^2 = \sum_{j=1}^L \frac{S_j^2}{L} = \sum_j \sum_i \frac{(X_{ij} - m_j)^2}{L(K-1)}$$

$$SSw \equiv \sum_j \sum_i (X_{ij} - m_j)^2$$

$$(K-1) \frac{S_j^2}{\sigma^2} \sim \chi_{K-1}^2 \quad \frac{SSw}{\sigma^2} \sim \chi_{L(K-1)}^2$$

$$\left( \frac{Ssb / \sigma^2}{L-1} \right) / \left( \frac{SSw / \sigma^2}{L(K-1)} \right) = \frac{Ssb / (L-1)}{SSw / (L(K-1))} \sim F_{L-1, L(K-1)}$$

$H_0: \mu_1 = \mu_2 = \dots = \mu_L$  if  $< F_{\alpha, L-1, L(K-1)}$

## Multivariate Tests

- Instead of testing using a single performance measure, e.g., error, use multiple measures for better discrimination, e.g., [fp-rate, fn-rate]
- Compare  $p$ -dimensional distributions
- Parametric case: Assume  $p$ -variate Gaussians

$$H_0: \mu_1 = \mu_2 \text{ vs. } H_1: \mu_1 \neq \mu_2$$

## Multivariate Pairwise Comparison

- Paired differences:  $d_i = x_{1i} - x_{2i}$

$$H_0: \mu_d = 0 \text{ vs. } H_1: \mu_d \neq 0$$

- Hotelling's multivariate  $T^2$  test

$$T^{*2} = K m^T S^{-1} m$$

- For  $p=1$ , reduces to paired t test

## Paired t Test

- Multiple training/validation sets
- $x_i^t = 1$  if instance  $t$  misclassified on fold  $i$
- Error rate of fold  $i$ :  $p_i = \frac{\sum_{t=1}^N x_i^t}{N}$
- With  $m$  and  $s^2$  average and var of  $p_i$ , we accept  $p_0$  or less error if

$$\frac{\sqrt{K}(m-p_0)}{s} \sim t_{K-1}$$

is less than  $t_{\alpha, K-1}$

## 5x2 cv Paired F Test

$$\frac{\sum_{i=1}^5 \sum_{j=1}^2 (p_i^{(j)})^2}{2 \sum_{i=1}^5 s_i^2} \sim F_{10, 5}$$

Two-sided test: Accept  $H_0: \mu_0 = \mu_1$  if  $< F_{\alpha, 10, 5}$

## ANOVA table

Source of variation	Sum of squares	Degrees of freedom	Mean square	$F_0$
Between groups	$SS_b \equiv K \sum_j (m_j - m)^2$	$L - 1$	$MS_b = \frac{SS_b}{L-1}$	$\frac{MS_b}{MS_w}$
Within groups	$SS_w \equiv \sum_j \sum_i (X_{ij} - m_j)^2$	$L(K-1)$	$MS_w = \frac{SS_w}{L(K-1)}$	
Total	$SS_T \equiv \sum_j \sum_i (X_{ij} - m)^2$	$L \cdot K - 1$		

If ANOVA rejects, we do pairwise posthoc tests  
 $H_0: \mu_i = \mu_j$  vs  $H_1: \mu_i \neq \mu_j$   
 $t = \frac{m_i - m_j}{\sqrt{2\sigma_w^2}} \sim t_{L(K-1)}$

## Multivariate ANOVA

- Comparison of  $L > 2$  algorithms

$$H_0: \mu_1 = \mu_2 = \dots = \mu_L \text{ vs. } H_1: \mu_r \neq \mu_s \text{ for at least one pair } r, s$$

$$H = K \sum_{j=1}^L (m_j - m)(m_j - m)^T$$

$$E = \sum_{j=1}^L \sum_{i=1}^K (x_{ij} - m_j)(x_{ij} - m_j)^T$$

$$\Lambda' = \frac{|E|}{|E + H|}$$

is Wilks's  $\Lambda$  distributed with  $p, L(K-1), L-1$  degrees of freedom

## Comparing Classifiers: $H_0: \mu_0 = \mu_1$ vs. $H_1: \mu_0 \neq \mu_1$

- Single training/validation set: McNemar's Test

$e_{00}$ : Number of examples misclassified by both	$e_{01}$ : Number of examples misclassified by 1 but not 2
$e_{10}$ : Number of examples misclassified by 2 but not 1	$e_{11}$ : Number of examples correctly classified by both

- Under  $H_0$ , we expect  $e_{01} = e_{10} = (e_{01} + e_{10})/2$

$$\frac{(e_{01} - e_{10} - 1)^2}{e_{01} + e_{10}} \sim \chi_1^2$$

Accept if  $< \chi_{\alpha, 1}^2$

## Comparing $L > 2$ Algorithms: Analysis of Variance (Anova)

$$H_0: \mu_1 = \mu_2 = \dots = \mu_L$$

- Errors of  $L$  algorithms on  $K$  folds

$$X_{ij} \sim \mathcal{N}(\mu_j, \sigma^2), j=1, \dots, L, i=1, \dots, K$$

- We construct two estimators to  $\sigma^2$ .

One is valid if  $H_0$  is true, the other is always valid.

We reject  $H_0$  if the two estimators disagree.

## Comparison over Multiple Datasets

- Comparing two algorithms:

Sign test: Count how many times A beats B over  $N$  datasets, and check if this could have been by chance if A and B did have the same error rate

- Comparing multiple algorithms

Kruskal-Wallis test: Calculate the average rank of all algorithms on  $N$  datasets, and check if these could have been by chance if they all had equal error

If KW rejects, we do pairwise posthoc tests to find which ones have significant rank difference