

0

These slides are edited by Roz from originals at

- <http://rii.ricoh.com/~stork/DHS.html>

Powerpoint lecture slides - DHSch3part2.ppt

0

Pattern Classification, Chapter 1

1

Chapter 3: Maximum-Likelihood and Bayesian Parameter Estimation (part 2)

- Bayesian Estimation (BE)
- Bayesian Parameter Estimation: Gaussian Case
- Bayesian Parameter Estimation: General Estimation
- Problems of Dimensionality

2

- Bayesian Estimation (Bayesian learning to pattern classification problems)
 - In MLE θ was supposed fix
 - In BE θ is a random variable
 - The computation of posterior probabilities $P(\omega_i | x)$ lies at the heart of Bayesian classification
 - Goal: compute $P(\omega_i | x, D)$
Given the sample D , Bayes formula can be written

$$P(\omega_i | x, D) = \frac{P(x | \omega_i, D) \cdot P(\omega_i | D)}{\sum_{j=1}^c P(x | \omega_j, D) \cdot P(\omega_j | D)}$$

2

Pattern Classification, Chapter 1

3

- To demonstrate the preceding equation, use:

$$P(x, D | \omega_i) = P(x | D, \omega_i) \cdot P(D | \omega_i)$$

$$P(x | D) = \sum_j P(x, \omega_j | D)$$

$$P(\omega_i) = P(\omega_i | D) \quad (\text{Training sample provides this!})$$

Thus :

$$P(\omega_i | x, D) = \frac{P(x | \omega_i, D) \cdot P(\omega_i)}{\sum_{j=1}^c P(x | \omega_j, D) \cdot P(\omega_j)}$$

3

Pattern Classification, Chapter 1

4

- Bayesian Parameter Estimation: Gaussian Case

Goal: Estimate θ using the a-posteriori density $P(\theta | \mathcal{D})$

- The univariate case: $P(\mu | \mathcal{D})$
 μ is the only unknown parameter

$$P(\mathbf{x} | \mu) \sim N(\mu, \sigma^2)$$

$$P(\mu) \sim N(\mu_0, \sigma_0^2)$$

(μ_0 and σ_0 are known!)

4

5

$$P(\mu | \mathcal{D}) = \frac{P(\mathcal{D} | \mu) \cdot P(\mu)}{\int P(\mathcal{D} | \mu) \cdot P(\mu) d\mu} \quad (1)$$

$$= \alpha \prod_{k=1}^{k=n} P(\mathbf{x}_k | \mu) \cdot P(\mu)$$

- Reproducing density

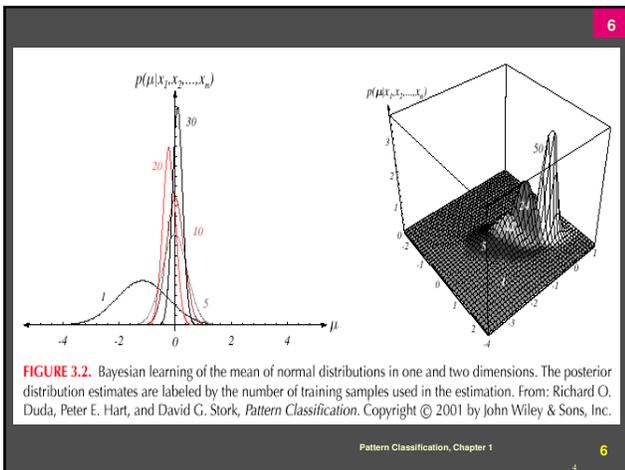
$$P(\mu | \mathcal{D}) \sim N(\mu_n, \sigma_n^2) \quad (2)$$

Identifying (1) and (2) yields:

$$\mu_n = \left(\frac{n\sigma_0^2}{n\sigma_0^2 + \sigma^2} \right) \hat{\mu}_n + \frac{\sigma^2}{n\sigma_0^2 + \sigma^2} \cdot \mu_0$$

and $\sigma_n^2 = \frac{\sigma_0^2 \sigma^2}{n\sigma_0^2 + \sigma^2}$

5



7

- The univariate case $P(\mathbf{x} | \mathcal{D})$
 - $P(\mu | \mathcal{D})$ computed
 - $P(\mathbf{x} | \mathcal{D})$ remains to be computed

$$P(\mathbf{x} | \mathcal{D}) = \int P(\mathbf{x} | \mu) \cdot P(\mu | \mathcal{D}) d\mu \text{ is Gaussian}$$

It provides:

$$P(\mathbf{x} | \mathcal{D}) \sim N(\mu_n, \sigma^2 + \sigma_n^2)$$

(Desired class-conditional density $P(\mathbf{x} | \omega_j, \omega_j)$)

Therefore: $P(\mathbf{x} | \omega_j, \omega_j)$ together with $P(\omega_j)$ and using Bayes formula, we obtain the Bayesian classification rule:

$$\text{Max}_{\omega_j} [P(\omega_j | \mathbf{x}, \mathcal{D})] = \text{Max}_{\omega_j} [P(\mathbf{x} | \omega_j, \mathcal{D}_j) \cdot P(\omega_j)]$$

7

Bayesian Parameter Estimation: General Theory

- $P(x | \Omega)$ computation can be applied to any situation in which the unknown density can be parametrized: the basic assumptions are:
 - The form of $P(x | \theta)$ is assumed known, but the value of θ is not known exactly
 - Our knowledge about θ is assumed to be contained in a known prior density $P(\theta)$
 - The rest of our knowledge θ is contained in a set Ω of n random variables x_1, x_2, \dots, x_n that follows $P(x)$

Pattern Classification, Chapter 1

The basic problem is:

“Compute the posterior density $P(\theta | \Omega)$ ”
then “Derive $P(x | \Omega)$ ”

Using Bayes formula, we have:

$$P(\theta | D) = \frac{P(D | \theta) \cdot P(\theta)}{\int P(D | \theta) \cdot P(\theta) d\theta}$$

And by independence assumption:

$$P(D | \theta) = \prod_{k=1}^{k=n} P(x_k | \theta)$$

Pattern Classification, Chapter 1

- Problems of Dimensionality
 - Problems involving 50 or 100 features (binary valued)
 - Classification accuracy depends upon the dimensionality and the amount of training data
 - Case of two classes multivariate normal with the same covariance

$$P(\text{error}) = \frac{1}{\sqrt{2\pi}} \int_{r/2}^{\infty} e^{-\frac{u^2}{2}} du$$

$$\text{where : } r^2 = (\mu_1 - \mu_2)^t \Sigma^{-1} (\mu_1 - \mu_2)$$

$$\lim_{r \rightarrow \infty} P(\text{error}) = 0$$

Pattern Classification, Chapter 1