Statistical methods in recognition

- Basic steps in classifier design
  - collect training images
  - choose a classification model
  - estimate parameters of classification model from training images
  - evaluate model on training data and refine
  - collect test image data set
  - apply classifier to test data

Zoran Duric

Why is classification a problem?

- Because classes overlap in our (impoverished) representations
- Example: Classify a person as a male or female based on weight
  - Male training set :{ 155, 122, 135, 160, 240, 220, 180, 145}
  - Female training set: {95, 132, 115, 124, 145, 110, 150}
  - Unknown sample has weight 125. Male or female?

Zoran Duric

Factors that should influence our decision

- How likely is it that a person weighs 125 pounds given that the person is a male? Is a female?
  - Class-conditional probabilities
- How likely is it that an arbitrary person is a male? A female?
  - Prior class probabilities
- What are the costs of calling a male a female? A female a male?
  - Risks

Basic approaches to classification

1. Build probabilistic models of our training data, and compute the probability that an unknown sample belongs to each of our possible classes using these models.
2. Compare an unknown sample directly to each member of the training set, looking for the training element “most similar” to the unknown.
   Nearest neighbor classification
3. Train a neural network to recognize unknown samples by “teaching it” how to correctly train the elements of the training set.

A primer on probability

- Probability spaces - models of random phenomena
- Example: a box contains s balls labeled 1, ..., s
  - Experiment: Pick a ball, note its label and then replace it in the box. Repeat this experiment n times.
  - Let $N_n(k)$ be the number of times that a ball labeled k was chosen in an experiment of length n
  - example: s = 3, n = 20

  1 1 3 2 1 2 2 3 2 3 3 2 1 2 3 3 1 3 2 2
  - $N_{20}(1) = 5$ $N_{20}(2) = 8$ $N_{20}(3) = 7$

Zoran Duric

Primer on probability

- The relative frequencies of the outcomes 1,2,3 are
  - $N_{20}(1)/20 = .25$ $N_{20}(2)/20 = .40$ $N_{20}(3)/20 = .35$
  - As n gets large, these numbers should settle down to fixed numbers p_1, p_2, p_3
  - We say p_i is the probability that the i'th ball will be chosen when the experiment is performed once
- Mathematical model: Let Ω be a set having s points which we place into a 1-1 correspondence with the possible **outcomes** of an experiment.
  - Call the points ω_k
  - to each ω_k we associate $p_k = 1/s$ and call it the probability of ω_k.

Zoran Duric

Primer on probability

- Suppose: we color balls 1, ..., r red and balls r+1, .., s green
  - What is the probability of choosing a red ball?
  - Intuitively it is $r/s = \Sigma p_k$ where the sum is over all ω_k such that the k'th ball is red
- Let A be the subset of Ω consisting of all ω_k such that k is red.
  - A has r points
  - A is called an event
  - When we say that A has occurred we mean that an experiment has been run and the outcome is represented by a point in A.
- If A and B are events, then so are $A \cap B$, $A \cup B$ and A^c

Zoran Duric

Primer on probability

- Assigning probabilities to events:

$$P(B) = \sum_{\omega_{k \in B}} p_k$$

- A probability measure on a set Ω is a real valued function having domain 2^{Ω} satisfying
  - $P(\Omega) = 1$
  - $0 <= P(A) <= 1$, for all $A \subset \Omega$
  - If A_n are mutually disjoint sets then

$$P(\bigcup_{n=1}^{k} A_n) = \sum_{n=1}^{k} P(A_n)$$

Zoran Duric

Primer on probability

- Simple properties of probabilities
  - $P(A^c) = 1 - P(A)$
    - $P(\varnothing)=1-P(\Omega) = 1-1 = 0$
    - if A is a subset of B, then P(A) <= P(B)
    - $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

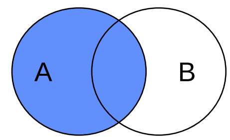

- Conditional probabilities
  - Our box has r red balls labeled 1, ..., r and b black balls labeled r+1, ..., r+b. If the ball drawn is known to be red, what is the probability that its label is 1?
    - A - event "red"
    - B - event "1"
    - interested in conditional probability of B knowing that A has occurred - P(B|A)

Primer on probability

- Let A and B be two events such that P(A) > 0. Then the conditional probability that B occurs given A, written P(B|A) is defined to be

$$P(B|A) = \frac{P(B \cap A)}{P(A)}$$

- Ball example: what is P("1"| "red")
  - Let r = 5 and b = 15
  - P(1 and red) = .05
  - P(red) = .25
  - So, P(1 | red) = .05/.25 = .20

Primer on probability

- General case
  - $A_1, ..., A_n$ are mutually disjoint events with union Ω.
    - think of the A_i as the possible identities of an object
  - B is an event with $P(B) > 0$
    - think of B as an observable event, like the area of a component in an image
  - $P(B|A_k)$ and $P(A_k)$ are known, $k = 1,..., n$
    - $P(B|A_k)$ is the probability that we would observe a component with area B if the identify of the object is A_i
  - Question: What is $P(A_i|B)$
    - What we will really be after - the probability that the identity of the object is A_i given that we make measurement B

$$P(B|A) = \frac{P(B \cap A)}{P(A)}$$

Primer on probability

$$B = B \cap \left(\bigcup_{k=1}^{n} A_k \right) = \bigcup_{k=1}^{n} (B \cap A_k)$$

So intersections are disjoint since the A_k are and

$$P(B) = \sum_{k=1}^{n} P(B \cap A_k)$$

But

$$P(B \cap A_k) = P(A_k) P(B|A_k)$$

Combining all this we get Bayes Rule

$$P(A_i|B) = \frac{P(A_i \cap B)}{P(B)} = \frac{P(A_i)P(B|A_i)}{\sum_{k=1}^{n} P(A_k)P(B|A_k)}$$

Training - computing $P(B|A_i)$

- Our training data is used to compute the $P(B|A_i)$, where B is the vector of features we plan to use to classify unknown images in the classes A_i
  - B might be (area, perimeter, moments)
- How might we represent $P(B|A_i)$?
  - as a table
    - quantize area, perimeter and average gray level suitably, and then use the training samples to fill in the three dimensional histogram.
    - analytically, by a standard probability density function such as the normal, uniform, ...

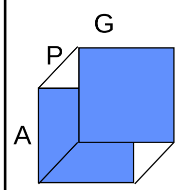

Zoran Duric

Primer on probability - training

- When we have many random variables it is usually impractical to create a table of the values of $P(B|A_i)$from our training set.
  - Example
    - 5 measurements
    - quantize each to 50 possible values
    - Then there are 50^5 possible 5-tuples we might observe in any element of the training set, and we would need to estimate this many probabilities to represent the conditional probability
      - too few training samples
      - too much storage required for the table

Zoran Duric

Primer on probability

- Instead, it is usually assumed that $P(B|A_i)$ has some simple mathematical form
  - uniform density function
    - each x_i takes on values only in the finite range $[a_i, b_i]$
    - $P(B|A_i)$ is constant for any realizable $(x_1, ..., x_n)$
    - for one random variable, $P(B|A_i)= 1/(b-a)$ for $a <= x <= b$ and 0 elsewhere
  - Normal distribution

$$f(x) = n(x;\mu,\sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(\frac{x-\mu}{\sigma})^2}$$

  - In any case, once the parameters of the assumed density function are estimated, its goodness of fit should also be evaluated.

Primer on probability

- Density function is called the Gaussian function and the error function
  - μ is called the location parameter
  - σ is called the scale parameter
- Generalization to multivariate density functions
  - mean vector
  - covariance matrix

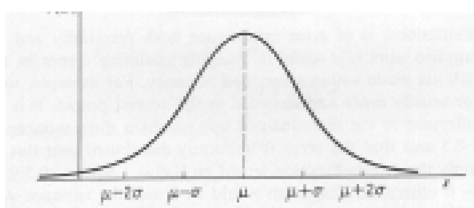

Fig. 3.7 Typical Normal Density Curve

Prior probabilities and their role in classification

- Prior probabilities of each object class
  - probabilities of the events: object is from class i ($P(A_i)$)
  - Example
    - two classes - A and B; two measurement outcomes: 0 and 1
    - prob(0|A) = .5, prob(1|A) = .5; prob(0|B) = .2 prob(1|B)=.8
  - Might guess that if we measure 0 we should decide that the class is A, but if we measure 1 we should decide B
  - But suppose that P(A) = .10 and P(B) = .90
    - Out of 100 samples, 90 will be B's and 18 of these (20% of those 90) will have measurement 0
      - We will classify these incorrectly as A's
      - Total error is nP(B)P(0|B)
    - 10 of these samples will be A's and 5 of them will have measurement 0 - these we'll get right
      - Total correct is nP(A)P(0|A)

Prior probabilities

- So, how do we balance the effects of the prior probabilities and the class conditional probabilities?
- We want a rule that will make the fewest errors
  - Errors in A proportional to P(A)P(x|A)
  - Errors in B proportional to P(B)P(x|B)
  - To minimize the number of errors choose A if P(A)P(x|A) > P(B)P(x|B); choose B otherwise
- The rule generalizes to many classes. Choose the C_i such that $P(C_i)P(x|C_i)$ is greatest.
- Of course, this is just Bayes' rule again

Bayes error

- The real formula for $P(C_i|x)$ is

$$P(C_i|x) = \frac{P(C_i)P(x|C_i)}{P(x)}$$

- where

$$P(x) = \sum_i P(C_i)P(x|C_i)$$

is a normalization factor that is the same for all classes.

- To evaluate the performance of our decision rule we can calculate the probability of error - probability that the sample is assigned to the wrong class.

Bayes error

- The **total error** which is called the **Bayes error** is defined as E[r(x)] =

$$\varepsilon = \int \min[P(C_1)P(x|C_1), P(C_2)P(x|C_2)]\, p(x)dx$$

$$= P(C_1)\int_{L_1} P(C_1|x)dx + P(C_2)\int_{L_2} P(C_2|x)dx$$

$$= P(C_1)\varepsilon_1 + P(C_2)\varepsilon_2$$

- The regions L_1 and L_2 are the regions where x is classified as C_1 and C_2 respectively.

Example

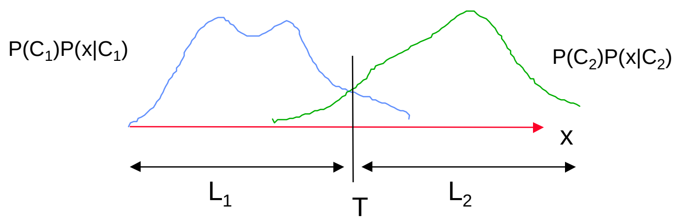

Moving T either left or right would increase the overall probability of error

Example - normal distributions

- In the case of normal distributions, the decision boundaries that provide the Bayes error can be shown to be quadratic functions - quadratic curves for two dimensional probability density functions
- In the special case where the classes have the same covariance matrix, decision boundary is a linear function - classes can be separated by a hyperplane

Bayes error for normal distributions

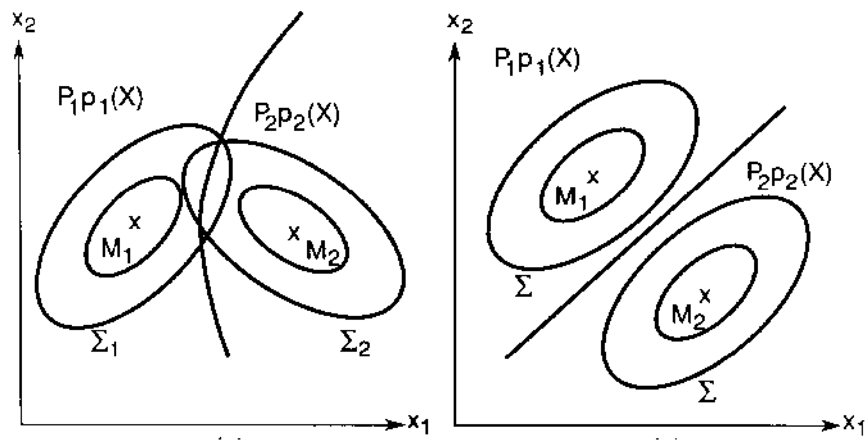

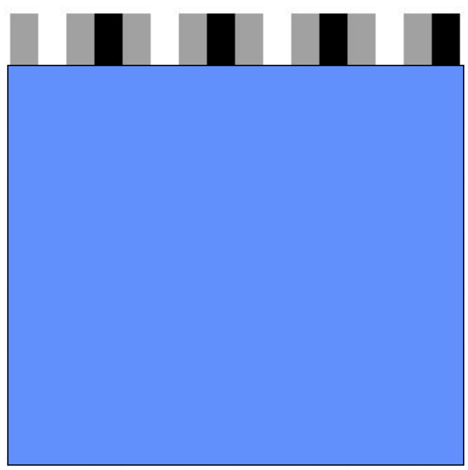

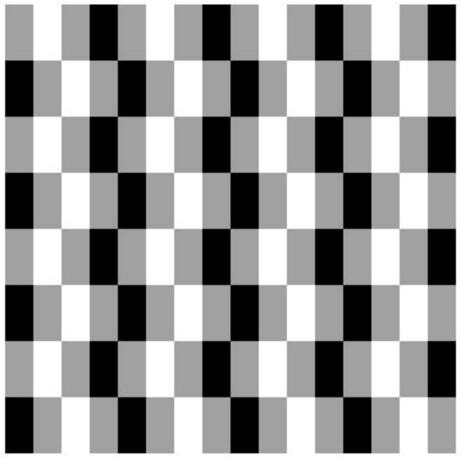

Adding risks

- Minimizing total number of errors does not take into account the cost of different types of errors
- Example: Screening X-rays for diagnosis
  - two classes - healthy and diseased
  - two types of errors
    - classifying a healthy patient as diseased - might lead to a human reviewing X-rays to verify computer classification
    - classifying diseased patient as healthy - might allow disease to progress to more threatening level
- Technically, including costs in the decision rule is accomplished by modifying the a priori probabilities

Nearest neighbor classifiers

- Can use the training set directly to classify objects from the test set.
  - Compare the new object to every element of the training set
    - need a measure of closeness between an object from the training set and a test object

$$D(x,y) = \sum_i \frac{(x_i - y_i)^2}{\sigma_i^2}$$

  - Choose the class corresponding to the closest element from the training set
  - Generalization - k nearest neighbors: find k nearest neighbors and perform a majority vote

Nearest neighbor classification

- Computational problems
  - Choosing a suitable similarity measurement
  - Efficient algorithms for computing nearest neighbors with large measurement sets (high dimensional spaces)
    - k-d trees
    - quadtrees
    - but must use a suitable similarity measure
  - Algorithms for “editing” the training set to produce a smaller set for comparisons
    - clustering: replace similar elements with a single element
    - removal: remove elements that are not chosen as nearest neighbors

Other classification models

- Neural networks
- Structural models
  - grammatical models
  - graph models
  - logical models
- Mixed models

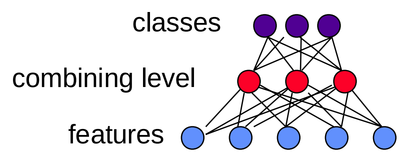

Zoran Duric

Primer on probability - random variables

- Example
  - Toss a coin three times with P[head] = p
  - If heads comes up we win $1
  - If tails come up we lose $1
  - Let X denote our winnings - it will be either 3,1,-1,-3 and it is what is actually observed in an experiment
- Can regard X as a function on the probability space
  - for ω in Ω, $X(\omega)$ is 1,3,-1 or -3
- Can compute P[X=c]
  - $P[X=3] = p^3$
  - $P[X=1] = 3p^2(1-p)$
- X is called a discrete random variable

ω	$X(\omega)$	$P[\omega]$
HHH	3	p^3
HHT	1	$p^2(1-p)$
HTH	1	$p^2(1-p)$
THH	1	$p^2(1-p)$
HTT	-1	$p(1-p)^2$
THT	-1	$p(1-p)^2$
TTH	-1	$p(1-p)^2$
TTT	-3	$(1-p)^3$

Zoran Duric

Primer on probability - random variables

- Definition: A discrete random variable, X, on a probability space (Ω,P) is a function with
  - domain Ω
  - range a finite or countably infinite set $\{x_1, x_2, ..., \}$
  - $P[X=x_i]$ means $P[\{\omega:X(\omega) = x_i\}]$
- The function, f, defined by
  - $f(x) = P[X=x]$
  - is called the discrete probability density function of X.
    - Example: If p = .5 then f(-3) = f(3) = 1/8 and f(-1) = f(1) = 3/8

Primer on probability - continuous random variables

- When we model “ideal” images we utilize random variables that are continuous (intensity, area, perimeter)
- A continuous random variable X on a probability space (Ω,P) is a function $X(\omega)$, ω in Ω, such that for
  - $-\infty < x < \infty$, $\{\omega|X(\omega) < x\}$ is an event, and
  - $P(X=x) = 0$
- The distribution function, F, of a continuous random variable X is $F(x) = P[X <= x]$
  - $0 <= F(x) <= 1$
  - F is nondecreasing in x
  - $F(-\infty) = 0$ and $F(\infty) = 1$

Primer on probability

- The value of f(x) at x is NOT the probability that X=x
  - if this were the case then the total probability of all events would be infinite!
  - just add up enough finite probabilties at an infinite number of points
  - we only talk about the probabilities of intervals - $P(x_1 < x < x_2)$ which we compute by integration.
- Most real world random variables are discrete (weights of individuals - there are only a few billion possibilities)
  - but we treat them as continuous random variables because there are simple mathematical formulas that we can then manipulate to compute probabilities

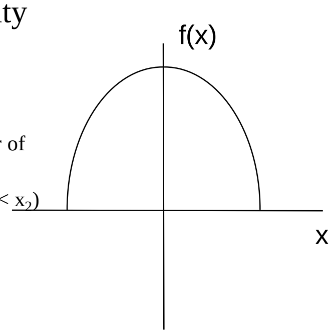

Primer on probability

- A density function is a nonnegative function f such that

$$\int_{-\infty}^{\infty} f(x)dx = 1$$

- If f is a density function, then F(x) defined below is a distribution function

$$F(x) = \int_{-\infty}^{x} f(y)dy$$

Primer on probability

- The value of f(x) at x is NOT the probability that X=x
  – if this were the case then the total probability of all events would be infinite!
  – just add up enough finite probabilties at an infinite number of points
  – we only talk about the probabilities of intervals - $P(x_1 < x < x_2)$ which we compute by integration.
- Most real world random variables are discrete (weights of individuals - there are only a few billion possibilities)
  – but we treat them as continuous random variables because there are simple mathematical formulas that we can then manipulate to compute probabilities

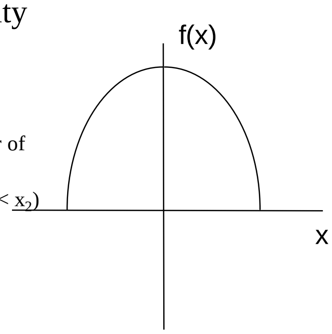

Primer on Probability

- Means, variance and covariances
- The mean of a continuous random variable, x, with pdf f is

$$\mu = E[x] = \int_{-\infty}^{\infty} x f(x) dx$$

  – For a discrete random variable the integral is replaced by a summation over the possible values of the variable.
- The variance of a continuous random variable is

$$\sigma^2 = E[(x-\mu)^2] = \int (x-\mu)^2 f(x) dx$$

Primer on probability

- Often, we have problems in which there are many random variables
  - area, perimeter, average gray level of a connected component
- Each has its own mean and variance, but each pair also has a covariance

$$\sigma_{x_i, x_j} = \int \cdots \int (x_i - \mu_i)(x_j - \mu_j) f(x_1, \ldots, x_n) dx_1 \cdots dx_n$$

  - Here, $f(x_1, \ldots, x_n)$ is the joint pdf of all of the random variables

Primer on Probability

- Example:
  - Population of a city is 40% male, 60% female
  - 50% of the males smoke
  - 30% of the females smoke
  - What is the probability that a smoker is male?

Primer on probability

- Define the events:
  - M - male
  - F - female
  - S - smoker
  - NS - nonsmoker
- Probabilities we are given are:
  - $P(M) = .4$
  - $P(F) = .6$
  - $P(S|M) = .5$
  - $P(S|F) = .3$
- Probability we want is $P(M|S) = P(M \cap S)/P(S)$

Primer on probability

- $P(M \cap S) = P(M)P(S|M) = (.4)\ (.5) = .2$
- Now, $S = (S \cap M) \cup (S \cap F)$ and $(S \cap M) \cap (S \cap F) = \varnothing$
  - so $P(S) = P(S \cap M) + P(S \cap F)$
  - $P(S \cap F) = P(F)P(S|F) = (.6)(.3) = .18$
  - $P(S \cap M) = .20$ from before
- So, $P(M|S) = .20/.38 = .53$