Current subjects in computer science

Seminar #2

PATTERN RECOGNITION

1

Seminar Description

- Introduction to pattern recognition systems
- Bayesian decision theory
- Feature Extraction and Selection Methods
- Example: Face Recognition
- Example: Robot Navigation

2

INTRODUCTION TO PATTERN RECOGNITION

3

Introduction

Our ability to recognize a face, to understand spoken words, to read handwritten characters… all these abilities belong to the complex processes of pattern recognition.

Definition of pattern recognition:

- The act of taking in raw data and making an action based on the "category" of the pattern.

4

Introduction

Machine perception:

- Build a machine that can recognize patterns:
  - Automatic speech recognition
  - Fingerprint identification
  - OCR (Optical Character Recognition)
  - DNA sequence identification
  - Face recognition
- Reliable pattern recognition machines would be extremely useful.

5

An example

- A fish packing plant wants to automate the process of <u>sorting incoming Fish on a conveyor according to species</u> using optical sensing

Species → Sea bass

Species → Salmon

6

Problem analysis

- We set up a camera and take some sample images to extract features that could help us to distinguish between the two species of fish
  - Length
  - Lightness
  - Width
  - Number and shape of fins
  - Position of the mouth, etc…

  - This is the set of all suggested features to explore for use in our classifier!

7

Problem analysis

1. The camera captures an image of the fish
2. The image is preprocessed to simplify subsequent operations without loosing relevant information.
   - Use a segmentation operation to isolate fishes from one another and from the background
3. The information from a single fish is sent to a feature extractor whose purpose is to reduce the data by measuring certain features
4. At the end, the features are passed to a classifier that evaluates the evidence presented and makes a final decision.

8

Problem analysis

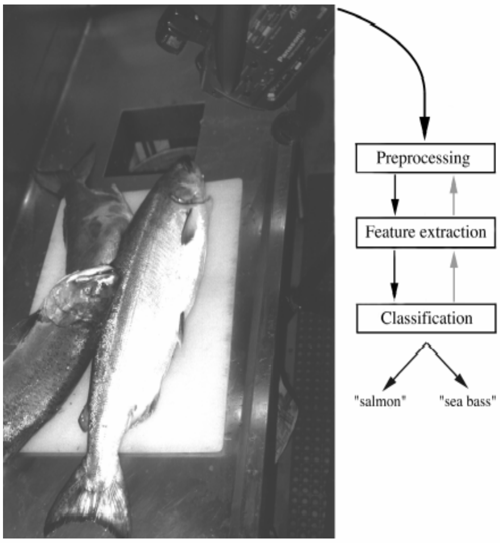

9

Problem analysis

Classification

- A sea bass is generally longer than a salmon: Sea bass has some typical length and this is greater than the salmon one. It gives us a preliminary model for the fish.
- Select the length of the fish as a possible feature for discrimination.
- We can classify the fish seeing whether or not the length l of the fish exceeds some critical value l*.

10

Problem analysis

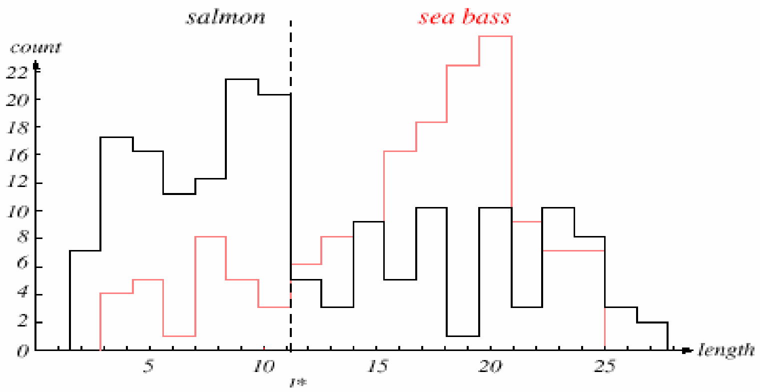

- Sea bass is longer than salmon, on average, but this simple criterion is quite poor.
- No matter how we choose l*, we cannot reliably separate sea bass from salmon by length only.

Problem analysis

The length is a poor feature alone!

Select the lightness as a possible feature (being careful to eliminate variations in illumination).

Problem analysis

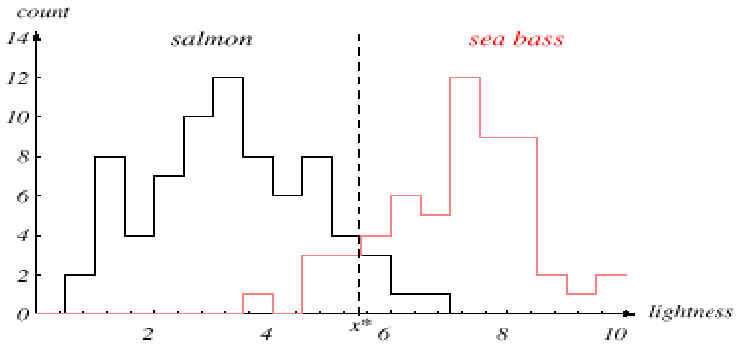

- The classes are better separated.
- No simple threshold value x* will unambiguously discriminate between the two categories. We can choose x* so that we have the smallest number of errors.

Problem analysis

- Threshold decision boundary and cost relationship
  - There is an overall cost associated with our decision and our true task is to make a decision rule that minimizes this cost. This is the central task of decision theory.
  - As a fish packing company, our customers will easily accept occasional pieces of salmon in the boxes of sea bass, but they will object if a piece of sea bass appears in a box of salmon.
  - Move our decision boundary toward smaller values of lightness in order to minimize the cost (reduce the number of sea bass that are classified salmon!)

Task of decision theory

Problem analysis

- To improve the recognition task, we can work with more than one feature.
- Adopt the lightness and add the width of the fish.
  - Sea bass is typically wider than salmon.

Fish $\Rightarrow$ $x^T = [x_1, x_2]$

Lightness   Width

15

Problem analysis

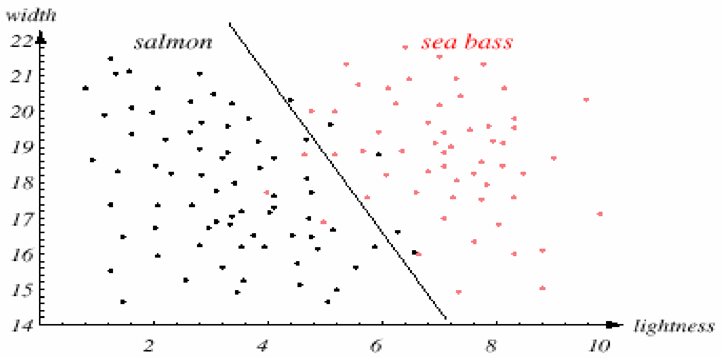

- Our problem is now to partition the feature space into two regions, where for all patterns in one region we will call the fish a sea bass, and all points in the other, salmon.
- The dark line might serve as a decision boundary of our classifier.

16

Problem analysis

- We might add other features that are not correlated with the ones we already have. A precaution should be taken not to reduce the performance by adding such "noisy features"
- Ideally, the best decision boundary should be the one which provides an optimal performance such as in the following figure:

17

Problem analysis

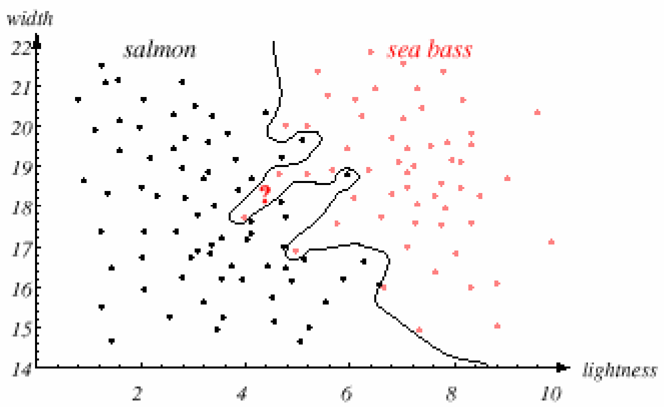

- In this case, our classifier would have a decision boundary more complex than the simple straight line.
- All the training patterns would be separated perfectly.

18

Problem analysis

- However, our satisfaction is premature because the central aim of designing a classifier is to correctly classify novel input

Issue of generalization!

- The boundary has been tuned to the particular training samples but it is unlikely that provides good generalization for future patterns.

19

Problem analysis

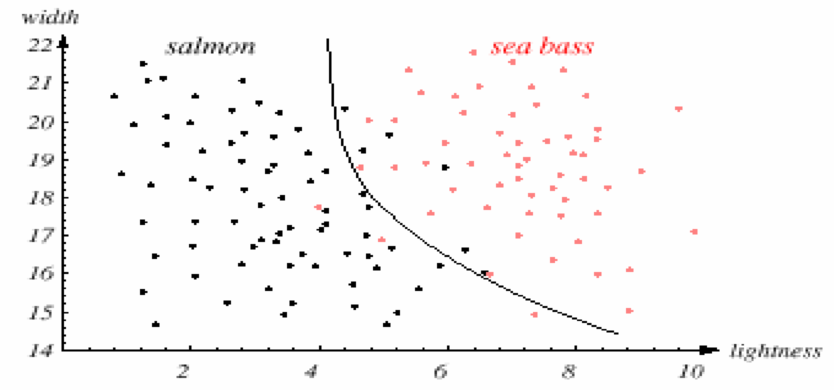

- We can simplify the recognizer, motivated by a belief that the classifier will have better performance on novel patterns.
- This is one of the central problems in statistical pattern recognition.

20

Pattern Recognition Systems

- Sensing
  - Use of a transducer (camera or microphone)
  - PR system depends of the bandwidth, the resolution, the sensitivity and the distortion of the transducer
- Segmentation and grouping
  - Patterns should be well separated and should not overlap

21

Pattern Recognition Systems

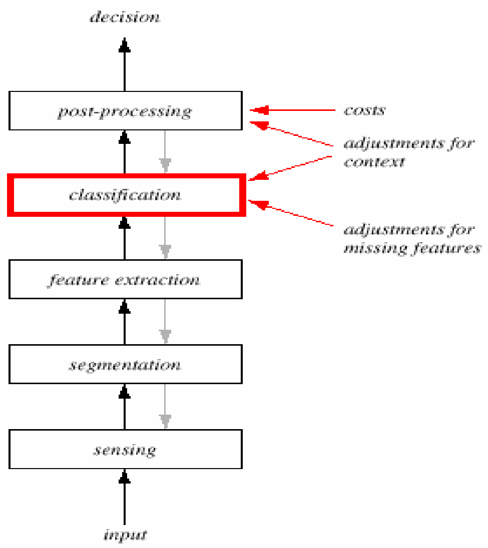

22

Pattern Recognition Systems

- Feature extraction
  - Discriminative features
  - Invariant features with respect to translation, rotation and scale.
- Classification
  - Use a feature vector provided by a feature extractor to assign the object to a category
- Post Processing
  - Exploit context: input dependent information other than from the target pattern itself to improve performance