1

Pattern Recognition

2

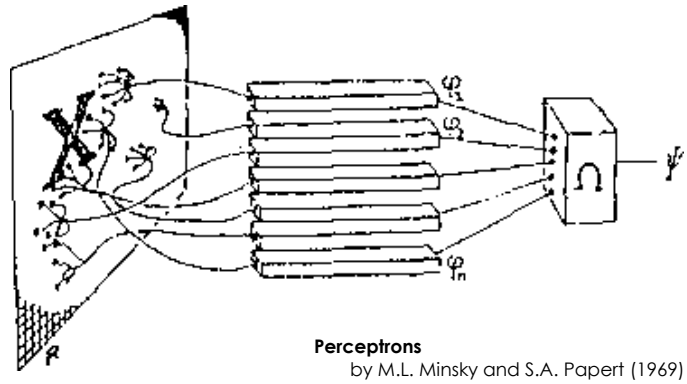

Perceptrons
by M.L. Minsky and S.A. Papert (1969)

Books:

Pattern Recognition, fourth Edition (Hardcover)
by Sergios Theodoridis, Konstantinos Koutroumbas
Publisher: Academic Press; 4th edition (2006, 2008)
Language: English
ISBN-10: 1597492728

4th Edition

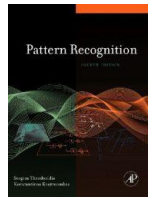

3 rd Edition

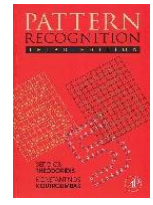

2nd Edition

Books:

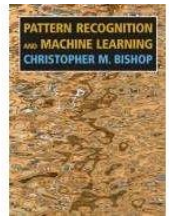

Pattern Recognition and Machine Learning
by Christopher Bishop
Publisher: Springer; 1 edition (August 17, 2006)
ISBN: 0387310738

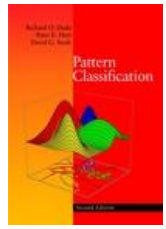

Pattern Classification, second Edition (Hardcover)
by Richard O. Duda, Peter E. Hart and David G. Stork
Publisher: Wiley Interscience 2 edition (2001)
Language: English
ISBN: 0-471-05669-3

Introduction to Pattern Recognition

Today:

- Machine Perception
- An Example
- Pattern Recognition Systems
- The Design Cycle
- Learning
- Conclusion

Pattern Recognition

Build a machine that can recognize patterns.

Machine Perception :

- Optical Character Recognition (OCR),
- Speech recognition,
- Email Spam Detection,
- Skin Detection based on pixel color,
- Texture classification,
- …..

Pattern Recognition

Base technology for:

- Image analysis,
- Speech understanding,
- Document analysis,
- Bioinformatics,
- Time series prediction.

An Example: Sea bass / Salmon

9

"Sorting incoming fish on a conveyor according to species using optical sensing."

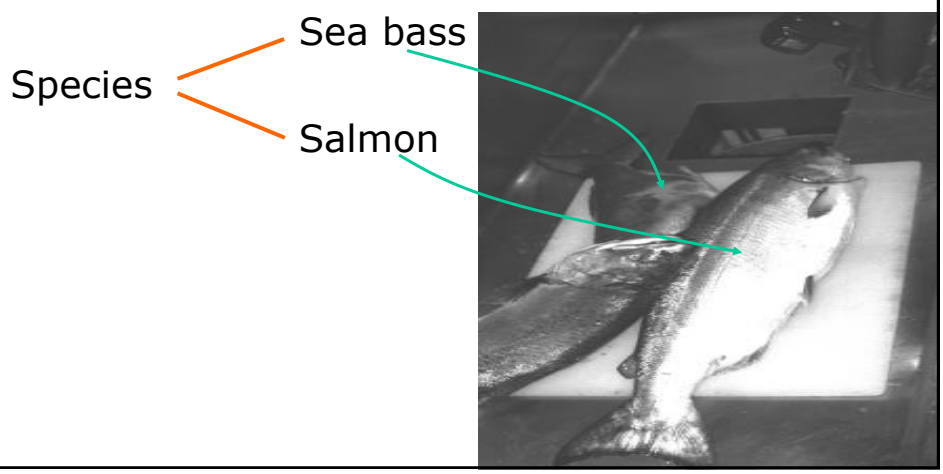

Sea bass / Salmon

10

Problem Analysis

Set up a camera and take some sample images to extract features:

- Length
- Lightness
- Width
- Number and shape of fins
- Position of the mouth, etc…

This is the set of all suggested features to explore for further use in our classification task!

Sea bass / Salmon

1. Preprocessing
   Use a segmentation operation to isolate fish from one another and from the background.

2. Feature extraction
   Information from a single fish is sent to a feature extractor whose purpose is to reduce the data by measuring certain features.
   (Mega Pixel -> few numbers)

3. The features are passed to a classifier.

Sea bass / Salmon

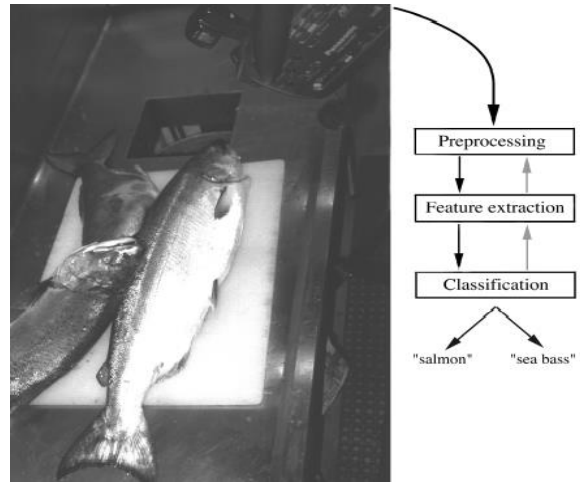

Sea bass / Salmon

Example of feature: length of the fish

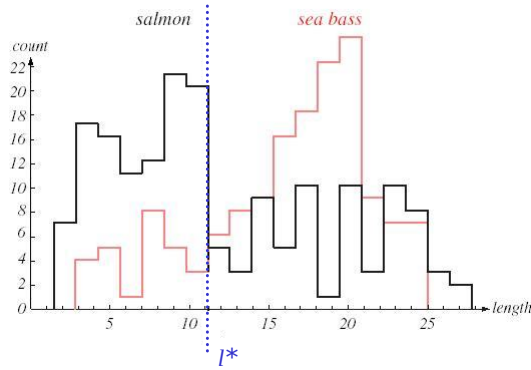

Decision: ***If*** length < l* ***then*** salmon ***else*** sea bass

Training error: 90 / 316 = 28%

Sea bass / Salmon

Training error: 90 / 316 = 28%

The length is a poor feature alone!

Select the lightness as a possible feature.

Sea bass / Salmon

Example of feature: lightness of the fish

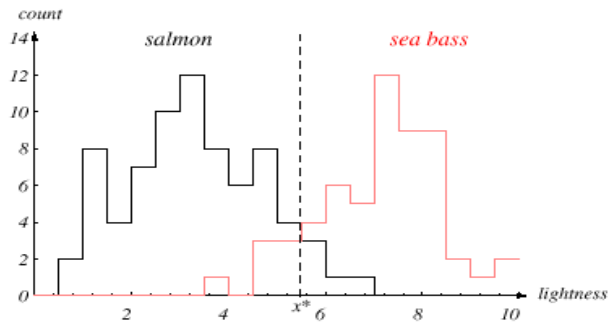

Decision: ***If*** lightn. < x*, ***then*** salmon ***else*** sea bass

Training error: 16 / 316 = 5%

Sea bass / Salmon

- Threshold decision boundary and cost relationship.
  - Move our decision boundary toward smaller values of lightness in order to minimize the cost (reduce the number of sea bass that are classified as salmon!).

Task of decision theory

Sea bass / Salmon

Now we use 2 features instead of 1:

Adopt the lightness and add the width of the fish.

Fish → $x^T = [x_1, x_2]$

x_1: Lightness

x_2: Width

Sea bass / Salmon

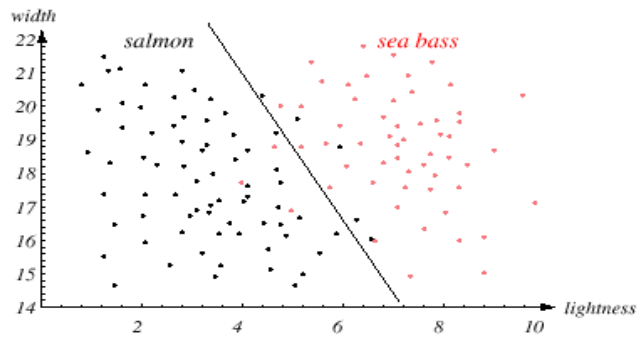

Linear decision function:

Training error: 8 / 316 = 2,5%

Sea bass / Salmon

- 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 "noisy features".

- Ideally, the best decision boundary should be the one which provides an optimal performance such as in the following figure:

Sea bass / Salmon

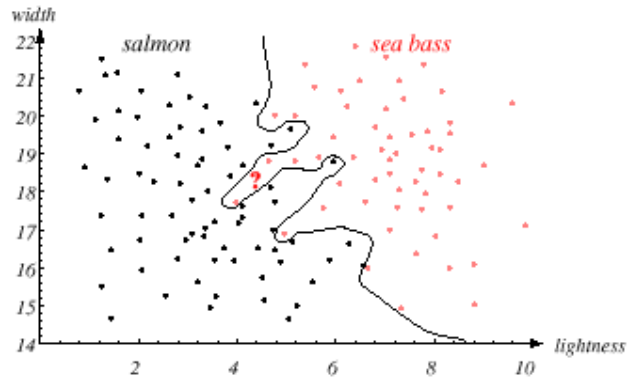

Complex decision function:

Training error: 0 / 316 = 0%

Is this good ?

Sea bass / Salmon

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

Issue of *generalization*!

Sea bass / Salmon

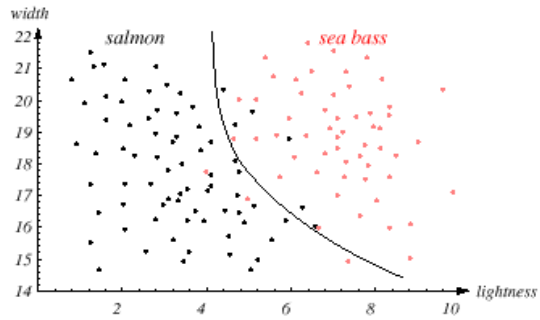

Quadratic decision function:

Training error: 9 / 316 = 2,5%

Pattern Recognition Systems:

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

Pattern Recognition Systems

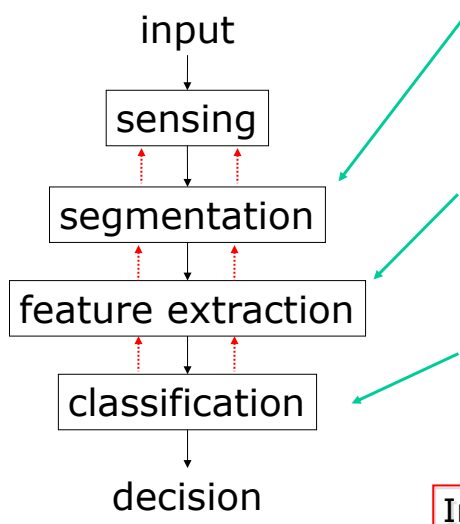

Segmentation:
- Isolate relevant data from the sensor output stream

Feature extraction:
- Discriminative
- Invariant to translation, rotation and scale….

Classification: Use a feature vector to assign the object to a category

Individual steps are in general not independent !!

The Design Cycle:

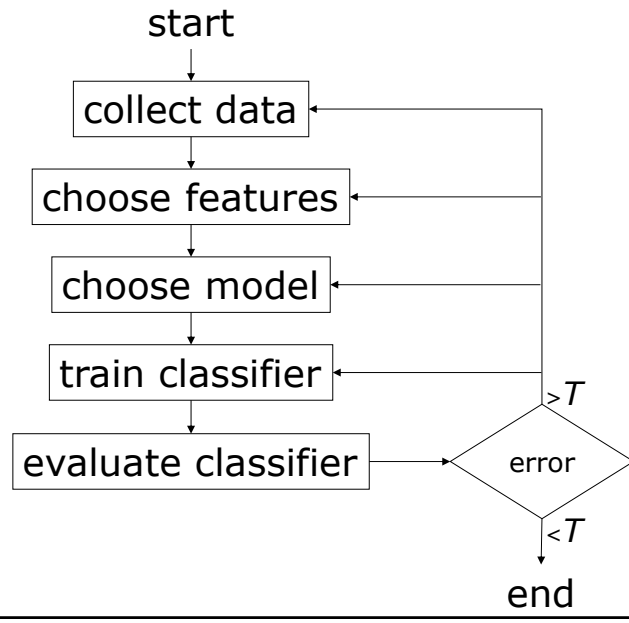

The Design Cycle

- Data Collection:
  - What type of sensor?
  - How do we know when we have collected an adequately large and representative set of examples for training and testing the system?
- Feature Choice: Depends on the characteristics of the problem domain.
  - simple to extract,
  - invariant to irrelevant transformation,
  - insensitive to noise and
  - best discrimination power.

The Design Cycle

- Model Choice:
  - e.g. should we use a linear or a quadratic decision function?
  - Can we estimate the probability distribution function that models the features?
- Training:
  - Depends on the model chosen.
  - Use data to determine the parameters of a classifier.
  - There are many different procedures for training classifiers and choosing models.

The Design Cycle

- Evaluation:
  - Measure the error rate on the validation set of examples that is different from the training set.
  - This tests the generalization performance.
  - If not good enough, go back to either of the design step.

The Design Cycle

Computational Complexity:

- More complex classifier are more computationally expensive.
- What is the optimal trade-off between computational ease and performance?
- (How does an algorithm scale as a function of the number of features, patterns or categories?)

Learning

- Supervised learning
  - A teacher provides a category label or cost for each pattern in the training set.
- Unsupervised learning
  - The system forms clusters or "natural groupings" of the input patterns.
  - Difficult: still the focus of intense research.
  - Will not be taught in this course.

Conclusion

- The number, complexity and magnitude of the sub-problems of Pattern Recognition appear often to be overwhelming.
- Many of these sub-problems can indeed be solved.
- Many fascinating unsolved problems still remain.