Jaime S. Cardoso
jaime.cardoso@inescporto
INESC Porto, Faculdade de Engenharia, Universidade do Porto
December 12, 2011

MAPi – Computer Vision 2011/12
Lecture 1a – Pattern Recognition Concepts

INESC Porto & Speaker

PRELUDE

INESC Porto > UTM > Multimedia

1. Information Processing and **Pattern Recognition**
   i. **Computer vision**
   ii. Sound and music computing
   iii. Network information processing
2. Digital Media Technologies
   i. Management and distribution of multimedia content
   ii. Context-aware multimedia services
   iii. Multimedia content recommendation systems
   iv. Adaptable mobile multimedia applications

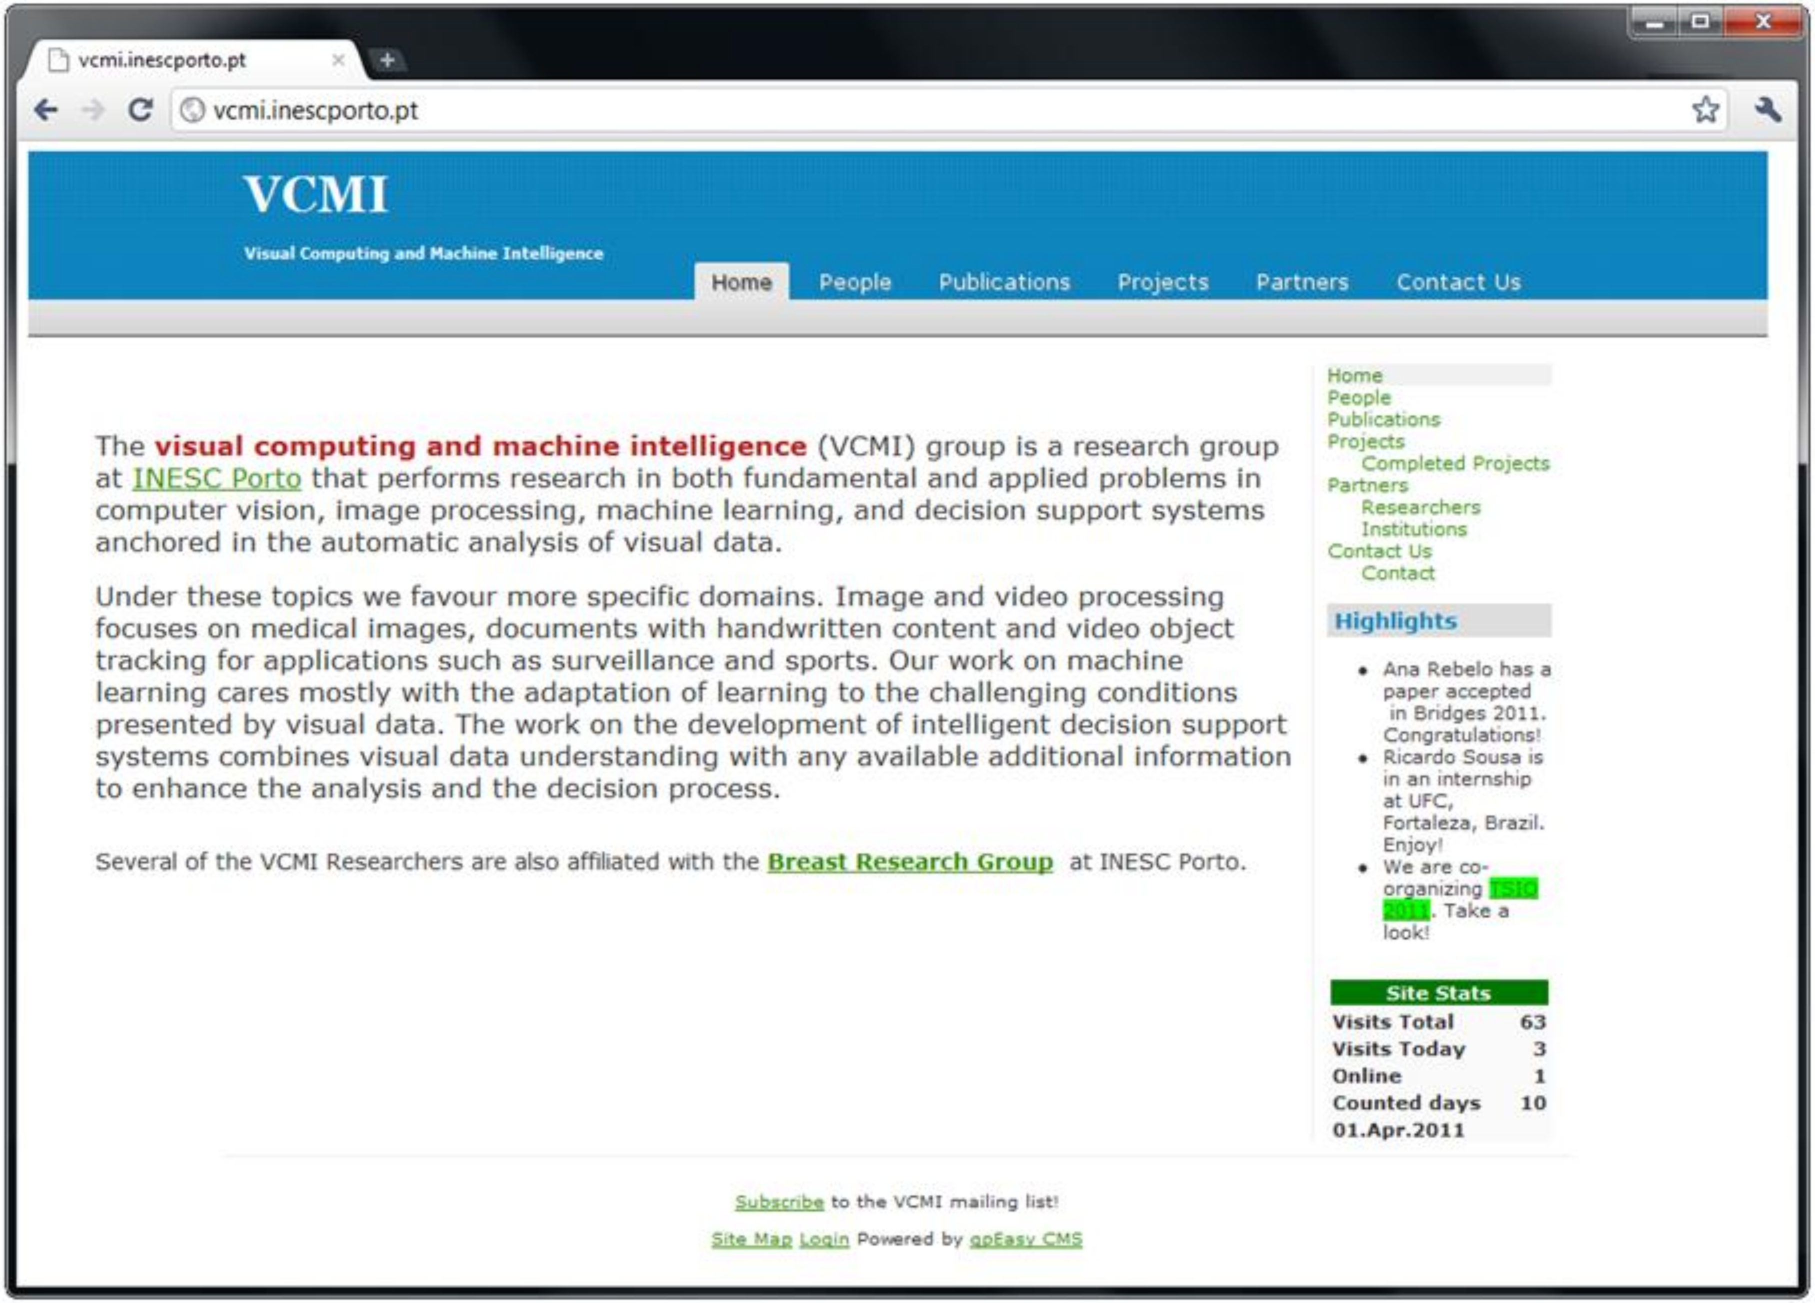

vcmi.inescporto.pt

vcmi.inescporto.pt

VCMI

Visual Computing and Machine Intelligence

Home | People | Publications | Projects | Partners | Contact Us

The **visual computing and machine intelligence** (VCMI) group is a research group at INESC Porto that performs research in both fundamental and applied problems in computer vision, image processing, machine learning, and decision support systems anchored in the automatic analysis of visual data.

Under these topics we favour more specific domains. Image and video processing focuses on medical images, documents with handwritten content and video object tracking for applications such as surveillance and sports. Our work on machine learning cares mostly with the adaptation of learning to the challenging conditions presented by visual data. The work on the development of intelligent decision support systems combines visual data understanding with any available additional information to enhance the analysis and the decision process.

Several of the VCMI Researchers are also affiliated with the **Breast Research Group** at INESC Porto.

- Home
- People
- Publications
- Projects
  - Completed Projects
- Partners
  - Researchers
  - Institutions
- Contact Us
  - Contact

Highlights

- Ana Rebelo has a paper accepted in Bridges 2011. Congratulations!
- Ricardo Sousa is in an internship at UFC, Fortaleza, Brazil. Enjoy!
- We are co-organizing TSIG 2011. Take a look!

Site Stats	
Visits Total	63
Visits Today	3
Online	1
Counted days	10
01.Apr.2011	

Subscribe to the VCMI mailing list!

Site Map Login Powered by gpEasy CMS

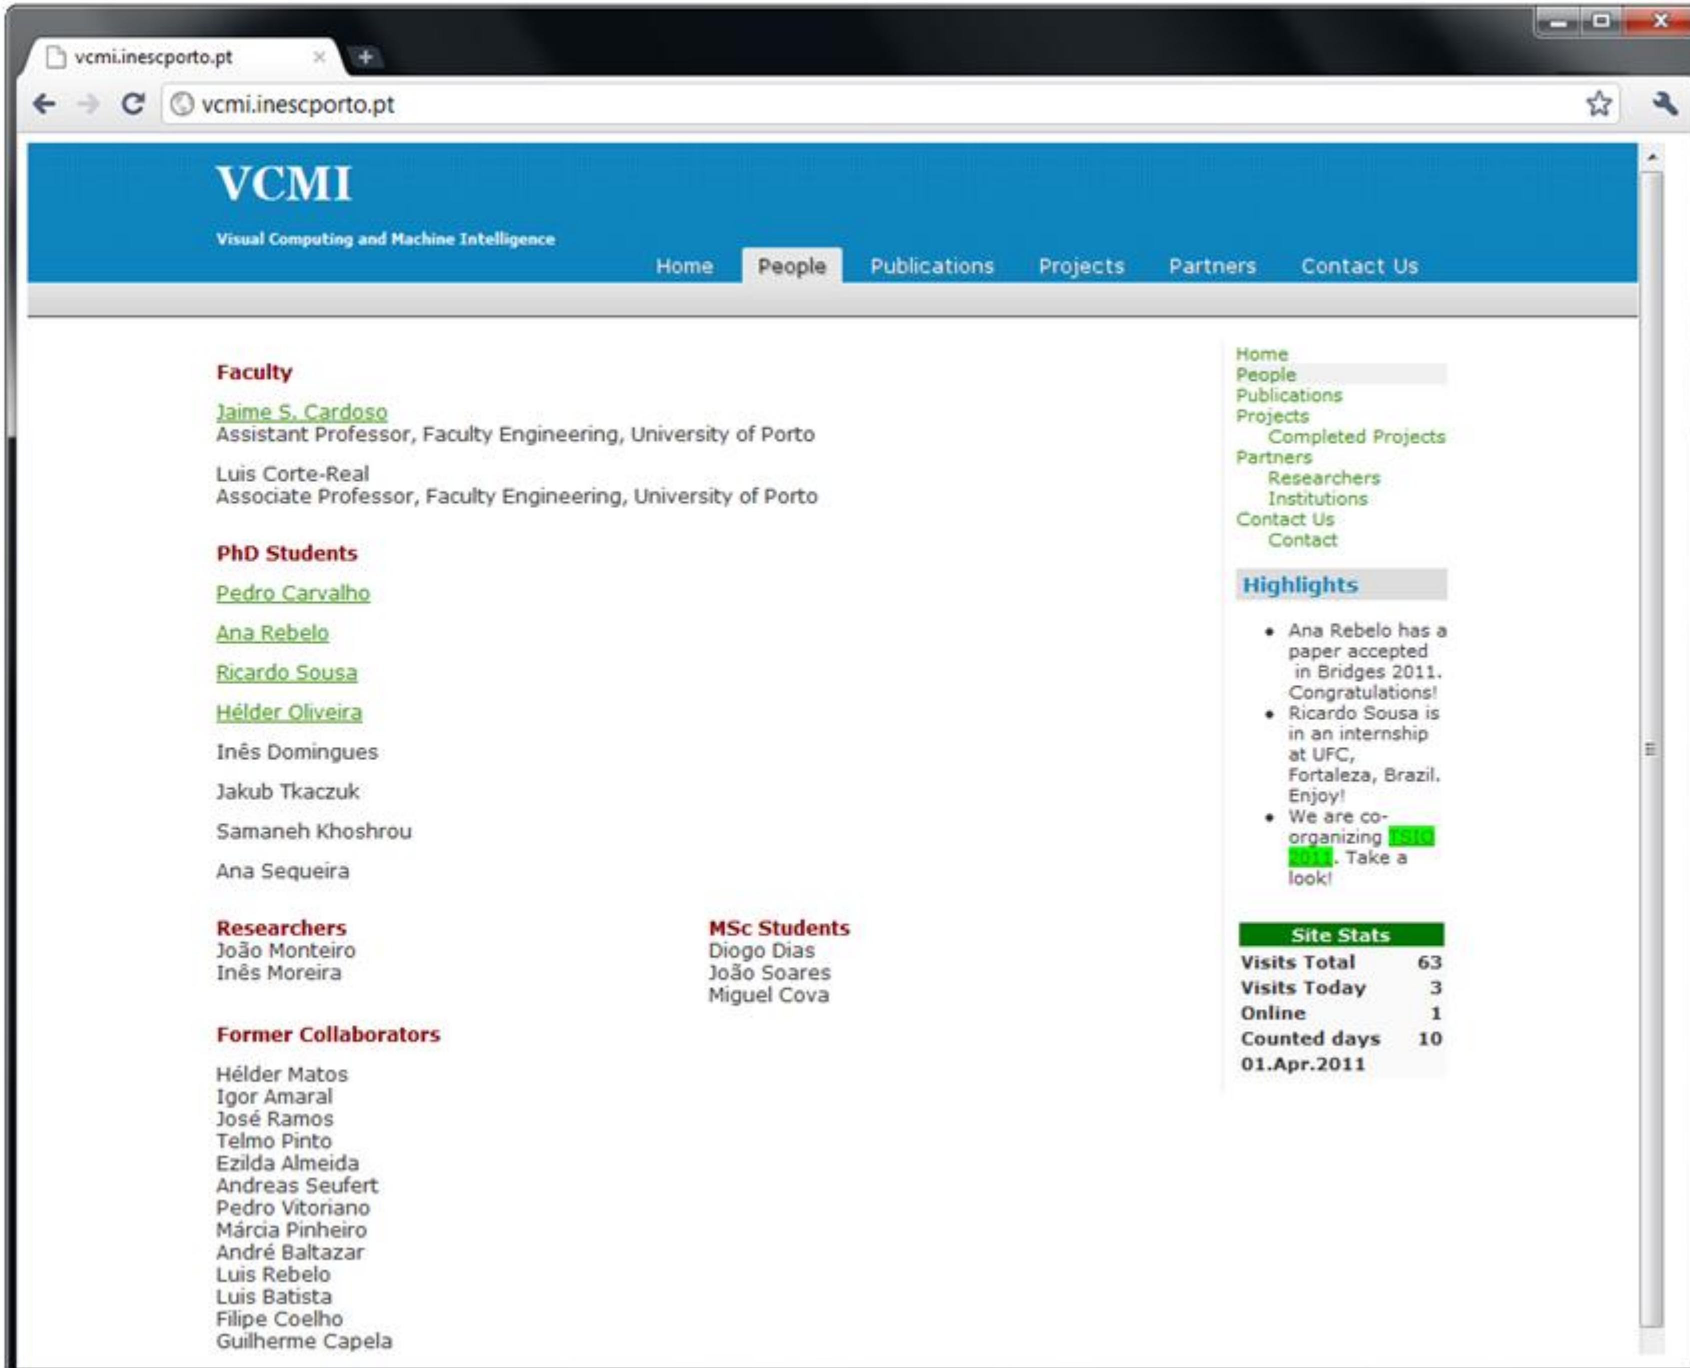
vcmi.inescporto.pt
vcmi.inescporto.pt
VCMI
Visual Computing and Machine Intelligence
Home
People
Publications
Projects
Partners
Contact Us
Faculty
Jaime S. Cardoso
Assistant Professor, Faculty Engineering, University of Porto
Luis Corte-Real
Associate Professor, Faculty Engineering, University of Porto
PhD Students
Pedro Carvalho
Ana Rebelo
Ricardo Sousa
Hélder Oliveira
Inês Domingues
Jakub Tkaczuk
Samaneh Khoshrou
Ana Sequeira
Researchers
João Monteiro
Inês Moreira
MSc Students
Diogo Dias
João Soares
Miguel Cova
Former Collaborators
Hélder Matos
Igor Amaral
José Ramos
Telmo Pinto
Ezilda Almeida
Andreas Seufert
Pedro Vitoriano
Márcia Pinheiro
André Baltazar
Luis Rebelo
Luis Batista
Filipe Coelho
Guilherme Capela
Home
People
Publications
Projects
Completed Projects
Partners
Researchers
Institutions
Contact Us
Contact
Highlights
Ana Rebelo has a paper accepted in Bridges 2011. Congratulations!
Ricardo Sousa is in an internship at UFC, Fortaleza, Brazil. Enjoy!
We are co-organizing . Take a look!
Site Stats
Visits Total 63
Visits Today 3
Online 1
Counted days 10
01.Apr.2011

Visual Computing and Machine Intelligence http://vcmi.inescporto.pt/

- Our Team
  - **Jaime S. Cardoso, PhD, Assistant Professor DEEC/FEUP**
  - Pedro Carvalho, PhD Std.
  - Ana Rebelo, PhD Std.
  - Ricardo Sousa, PhD Std.
  - Hélder Oliveira, PhD Std.
  - Inês Domingues, PhD Std.
  - Samaneh Khoshrou, PhD Std.
  - Ana F. Sequeira, PhD Std.
  - Inês Moreira, Researcher
  - João Moreira, Researcher
  - etc.

- Our Projects
  - 3d BCT
  - Semantic PACS
    - Picture Archiving and Communication System with Semantic Search Engine
  - BCCT
    - Advanced Objective Method for the Evaluation of the Aesthetical Result of Breast Interventions
  - OMR
    - Optical Recognition System for Handwritten Music Scores
  - NeTS
    - Next Generation Network Operations and Management
  - INCT-MACC
  - SINPATCO
  - etc

Pattern Recognition Concepts

Goal of computer vision

- Provide computers with human-like perception capabilities so that they can sense the environment, understand the sensed data, take appropriate actions (make decisions), learn from this experience in order to enhance future performance
  - **Understand the visual information** with no accompanying structural, administrative or descriptive text information

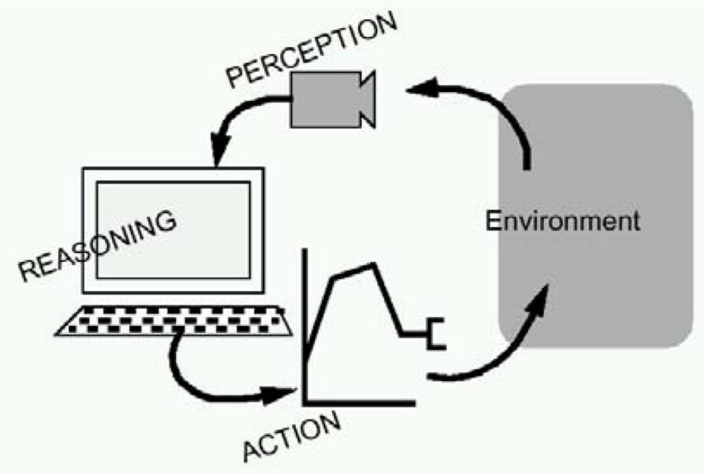

- Sources of difficulties:
  - Sensory gap
  - Semantic gap

From Pixels to Perception

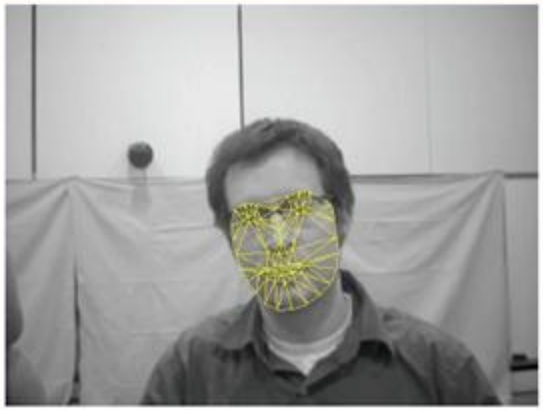
"Face Recognition"

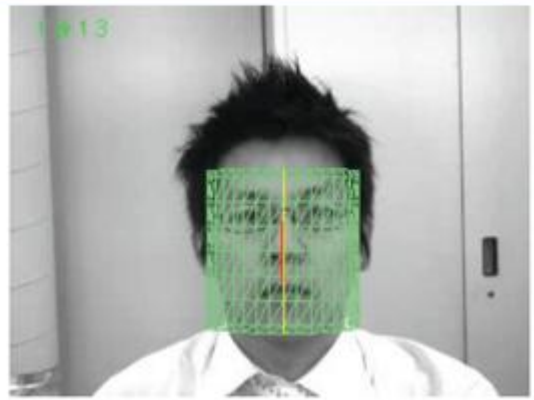
"Pose Estimation"


"Body Tracking"


"Speech Reading"

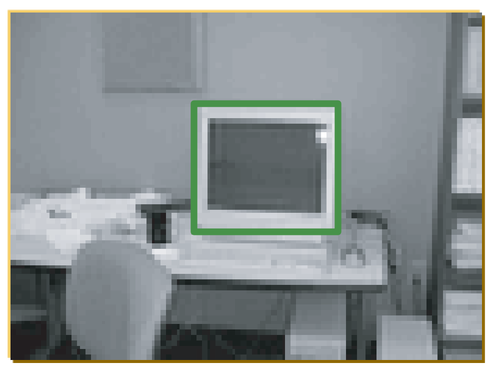
"Object detection"


"Car Tracking"

Object Recognition

- Perception of function: We can perceive the 3D shape, texture, material properties, without knowing about objects. But, the concept of category encapsulates also information about what can we do with those objects.

Object recognition: Is it really so hard?

This is a chair

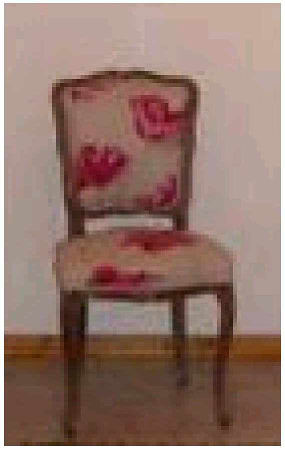

Find the chair in this image

Output of normalized correlation

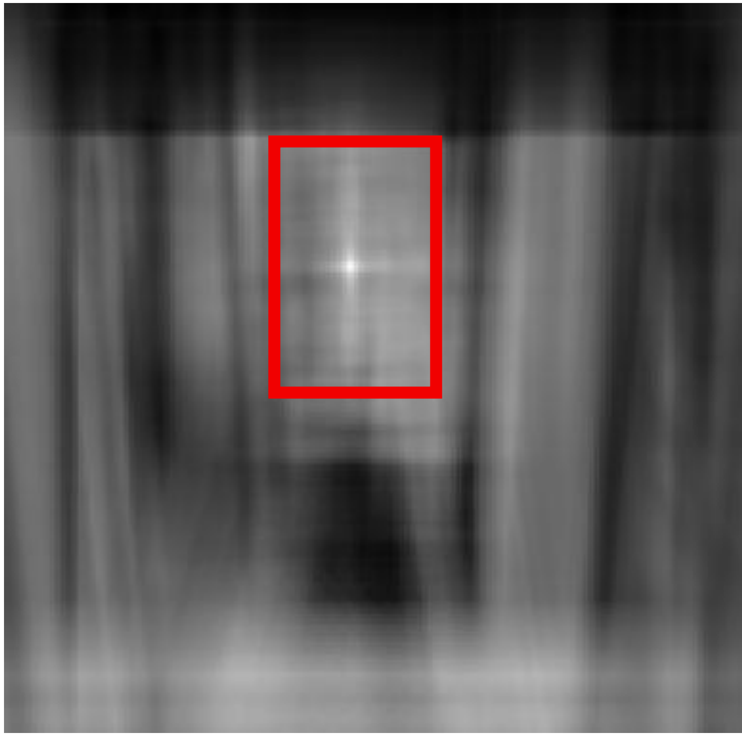

Object recognition: Is it really so hard?

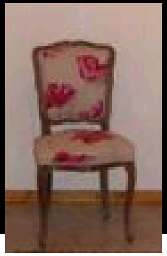

Find the chair in this image

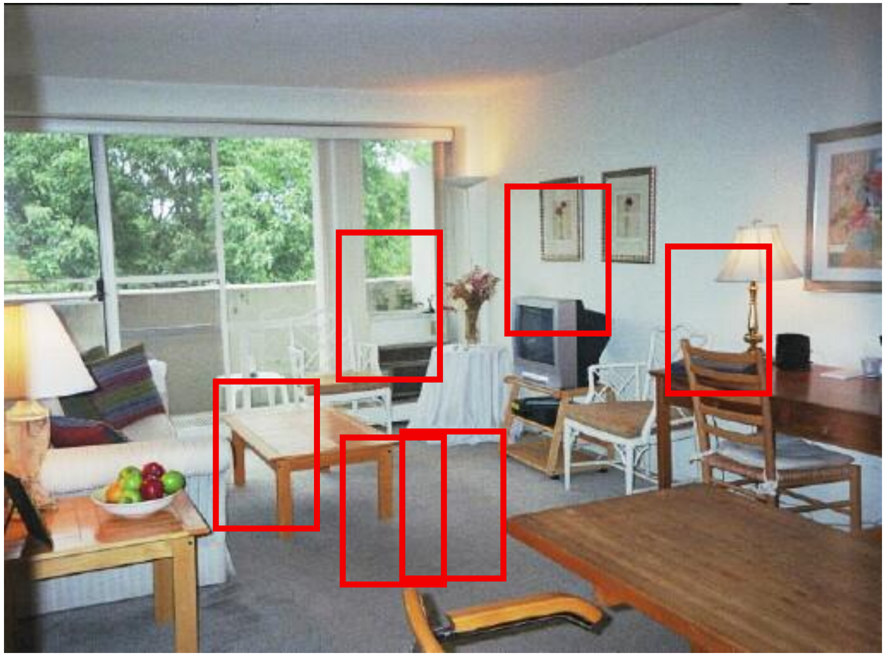

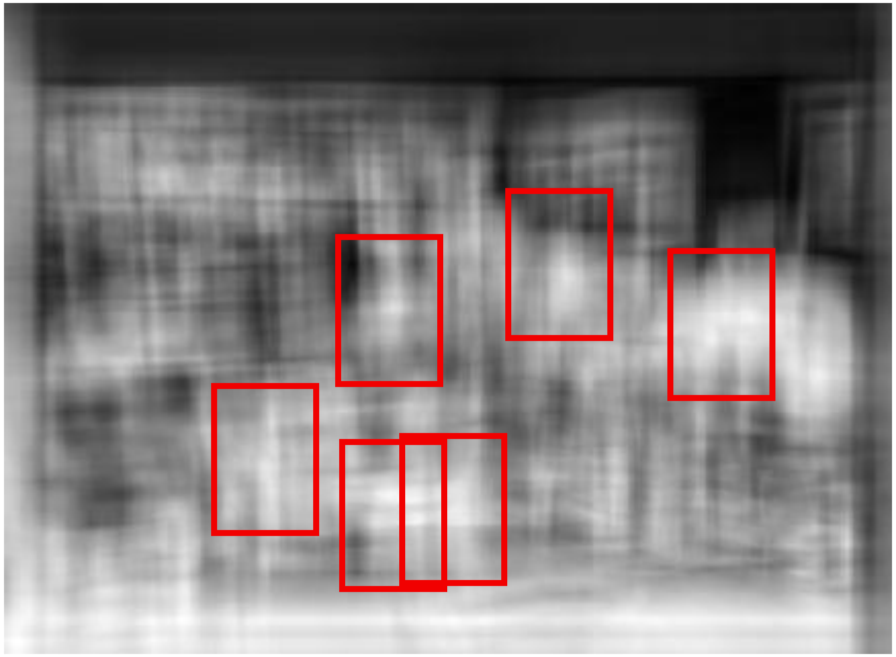

Pretty much garbage
Simple template matching is not going to make it

Object recognition: Is it really so hard?

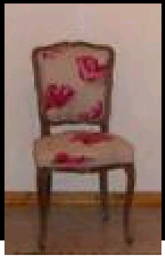

Find the chair in this image

A “popular method is that of template matching, by point to point correlation of a model pattern with the image pattern. These techniques are inadequate for three-dimensional scene analysis for many reasons, such as occlusion, changes in viewing angle, and articulation of parts.” Nivatia & Binford, 1977.

Challenges

- Why is vision hard?

Grayscale Image

y \ x =	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72
y = 41	210	209	204	202	197	247	143	71	64	80	84	54	54	57	58
42	206	196	203	197	195	210	207	56	63	58	53	53	61	62	51
43	201	207	192	201	198	213	156	69	65	57	55	52	53	60	58
44	216	206	211	193	202	207	208	57	69	60	55	77	49	62	51
45	221	206	211	194	196	197	220	56	63	60	55	46	97	58	106
46	209	214	224	199	194	193	204	173	64	60	59	51	62	56	48
47	204	212	213	208	191	190	191	214	60	62	66	76	51	49	55
48	214	215	215	207	208	180	172	188	69	72	55	49	56	52	56
49	209	205	214	205	204	196	187	196	86	62	66	87	57	60	48
50	208	209	205	203	202	186	174	185	149	71	63	55	55	45	56
51	207	210	211	199	217	194	183	177	209	90	62	64	52	93	52
52	208	205	209	209	197	194	183	187	187	239	58	68	61	51	56
53	204	206	203	209	195	203	188	185	183	221	75	61	58	60	60
54	200	203	199	236	188	197	183	190	183	196	122	63	58	64	66
55	205	210	202	203	199	197	196	181	173	186	105	62	57	64	63

How do we go from an array of numbers to recognizing fruit?

Challenges

- viewpoint variation

Challenges

- Illumination

Challenges

- Occlusion

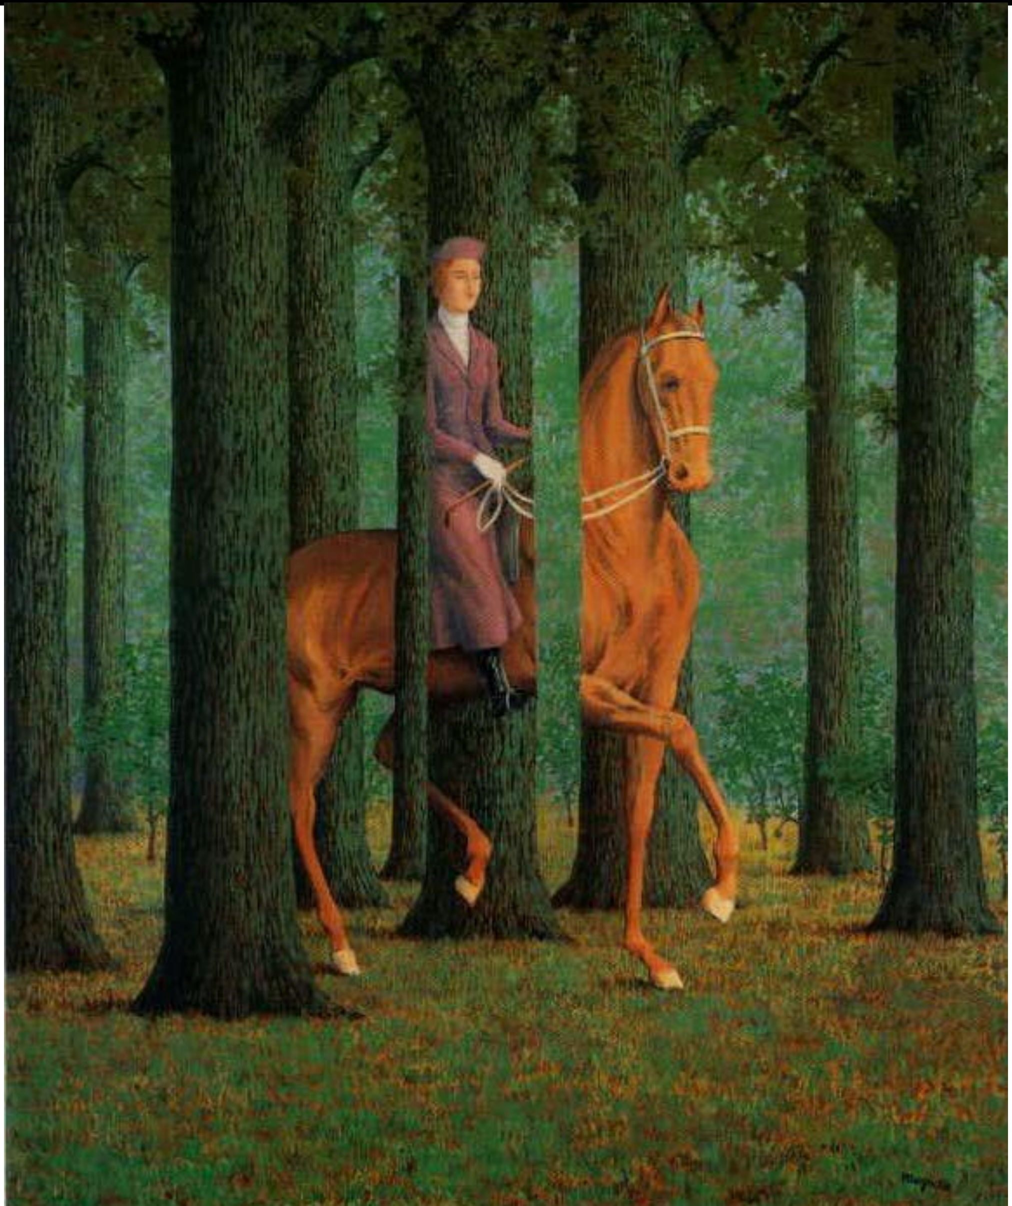

Challenges

- scale

Challenges

- deformation

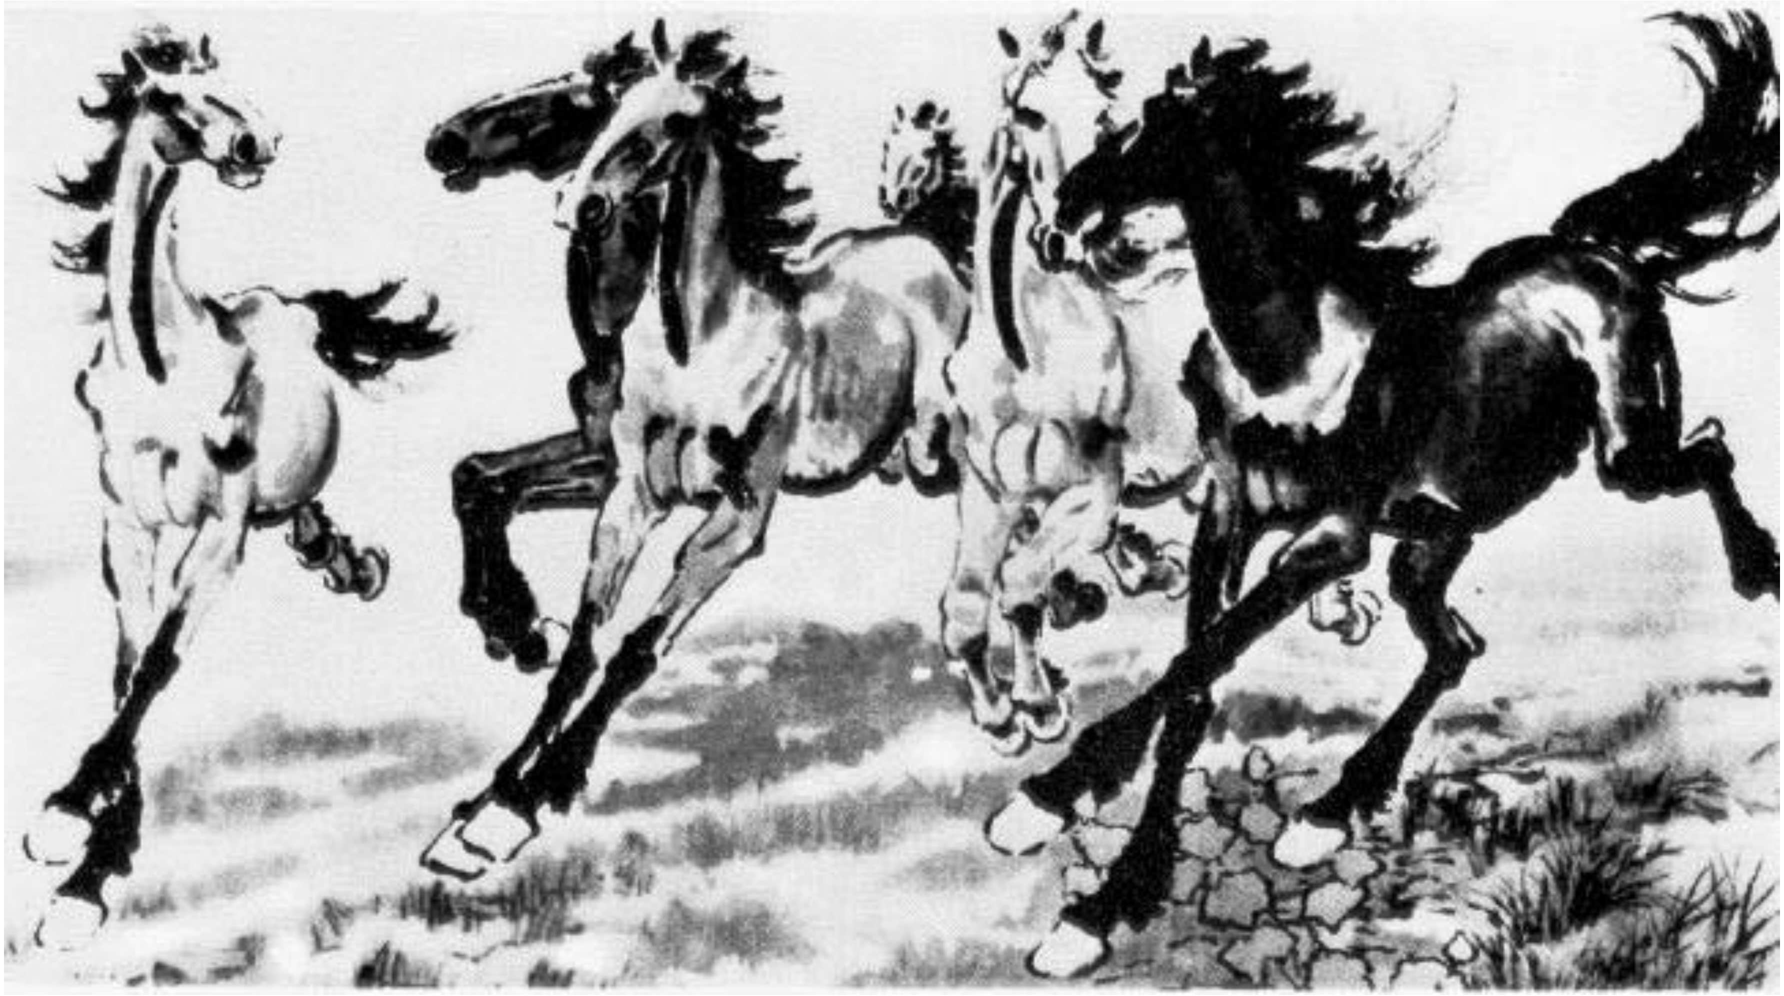

Challenges

- background clutter

An Example

- Problem: sorting incoming fish on a conveyor belt according to species
- Assume that we have only two kinds of fish:
  - Salmon
  - Sea bass

Picture taken with a camera

An Example: the problem

What ***we*** see

0	3	2	5	4	7	6	9	8
3	0	1	2	3	4	5	6	7
2	1	0	3	2	5	4	7	6
5	2	3	0	1	2	3	4	5
4	3	2	1	0	3	2	5	4
7	4	5	2	3	0	1	2	3
6	5	4	3	2	1	0	3	2
9	6	7	4	5	2	3	0	1
8	7	6	5	4	3	2	1	0

What a ***computer*** sees

An Example: Decision Process

- What kind of information can distinguish one species from the other?
  - Length, width, weight, number and shape of fins, tail shape, etc.
- What can cause problems during sensing?
  - Lighting conditions, position of fish on the conveyor belt, camera noise, etc.
- What are the steps in the process?
  - Capture image -> isolate fish -> take measurements -> make decision

An Example: Selecting Features

- Assume a fisherman told us that a sea bass is generally longer than a salmon.
  - We can use *length* as a feature and decide between sea bass and salmon according to a threshold on length.
  - I How can we choose this threshold?

An Example: Selecting Features

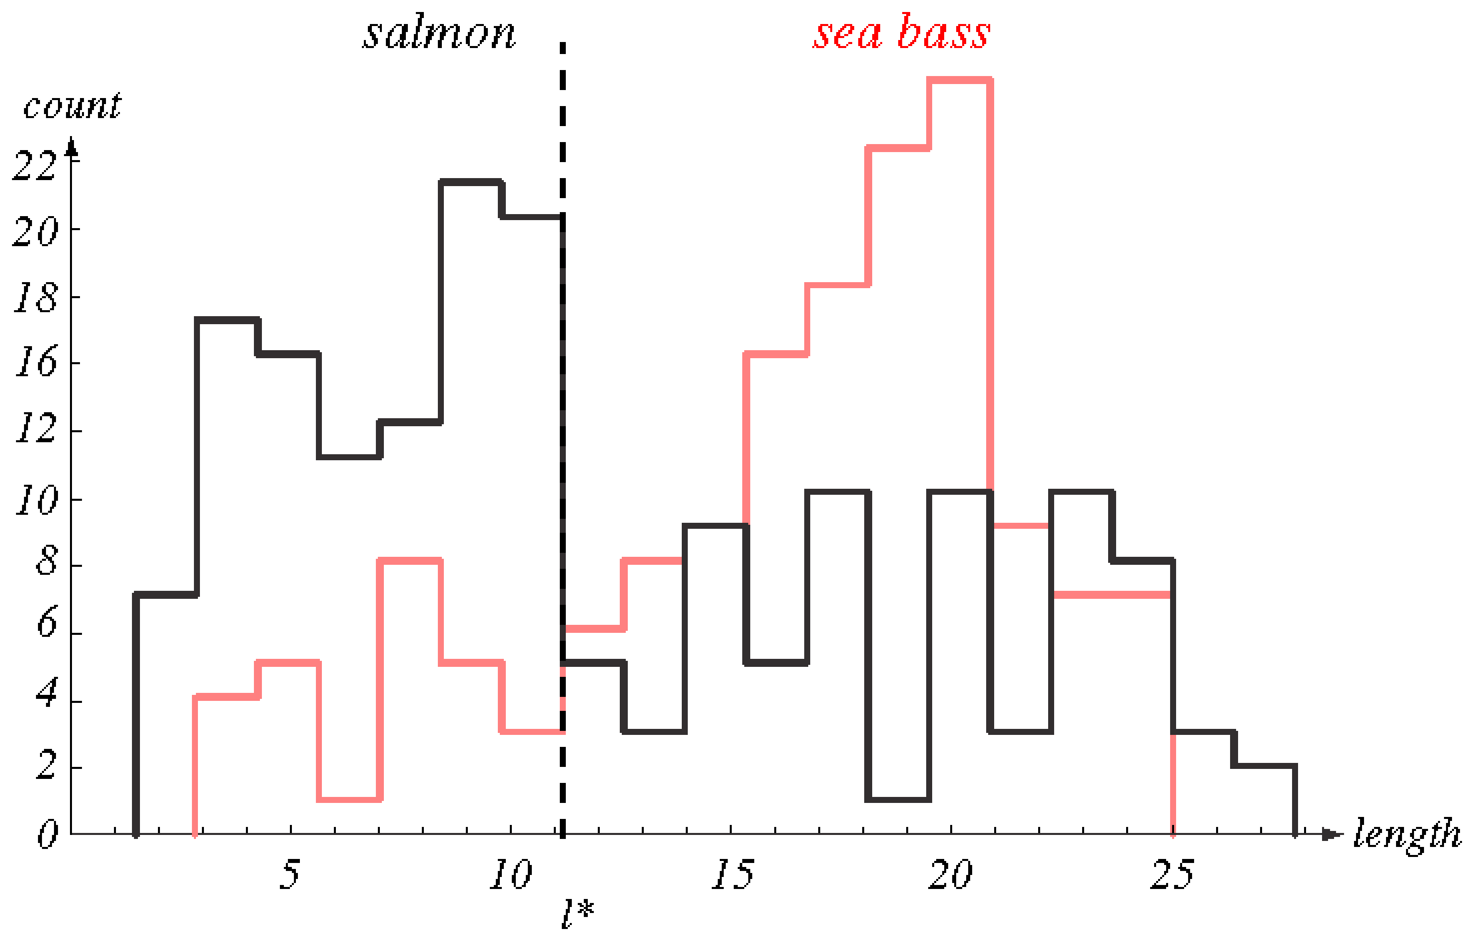

Histogram of the length feature for two types of fish in training samples. How can we choose the threshold ℓ^* to make a reliable decision?

An Example: Selecting Features

- Even though sea bass is longer than salmon on the average, there are many examples of fish where this observation does not hold.
- Try another feature: average lightness of the fish scales.

An Example: Selecting Features

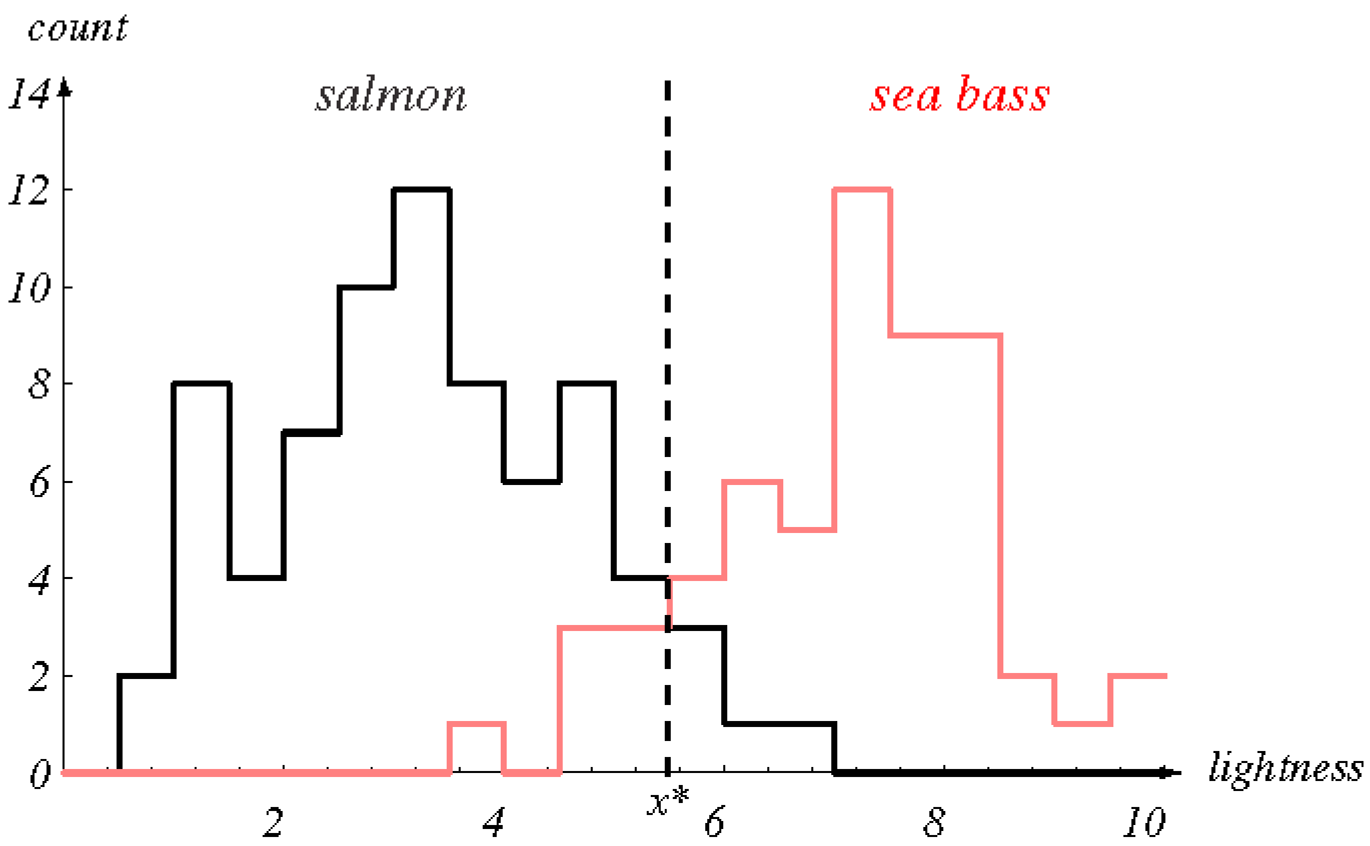

Histogram of the lightness feature for two types of fish in training samples. It looks easier to choose the threshold x^* but we still cannot make a perfect decision.

An Example: Multiple Features

- We can use two features in our decision:
  - lightness: $\boldsymbol{x}_1$
  - length: $\boldsymbol{x}_2$
- Each fish image is now represented as a point (feature vector) $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$ in a two-dimensional feature space.

An Example: Multiple Features

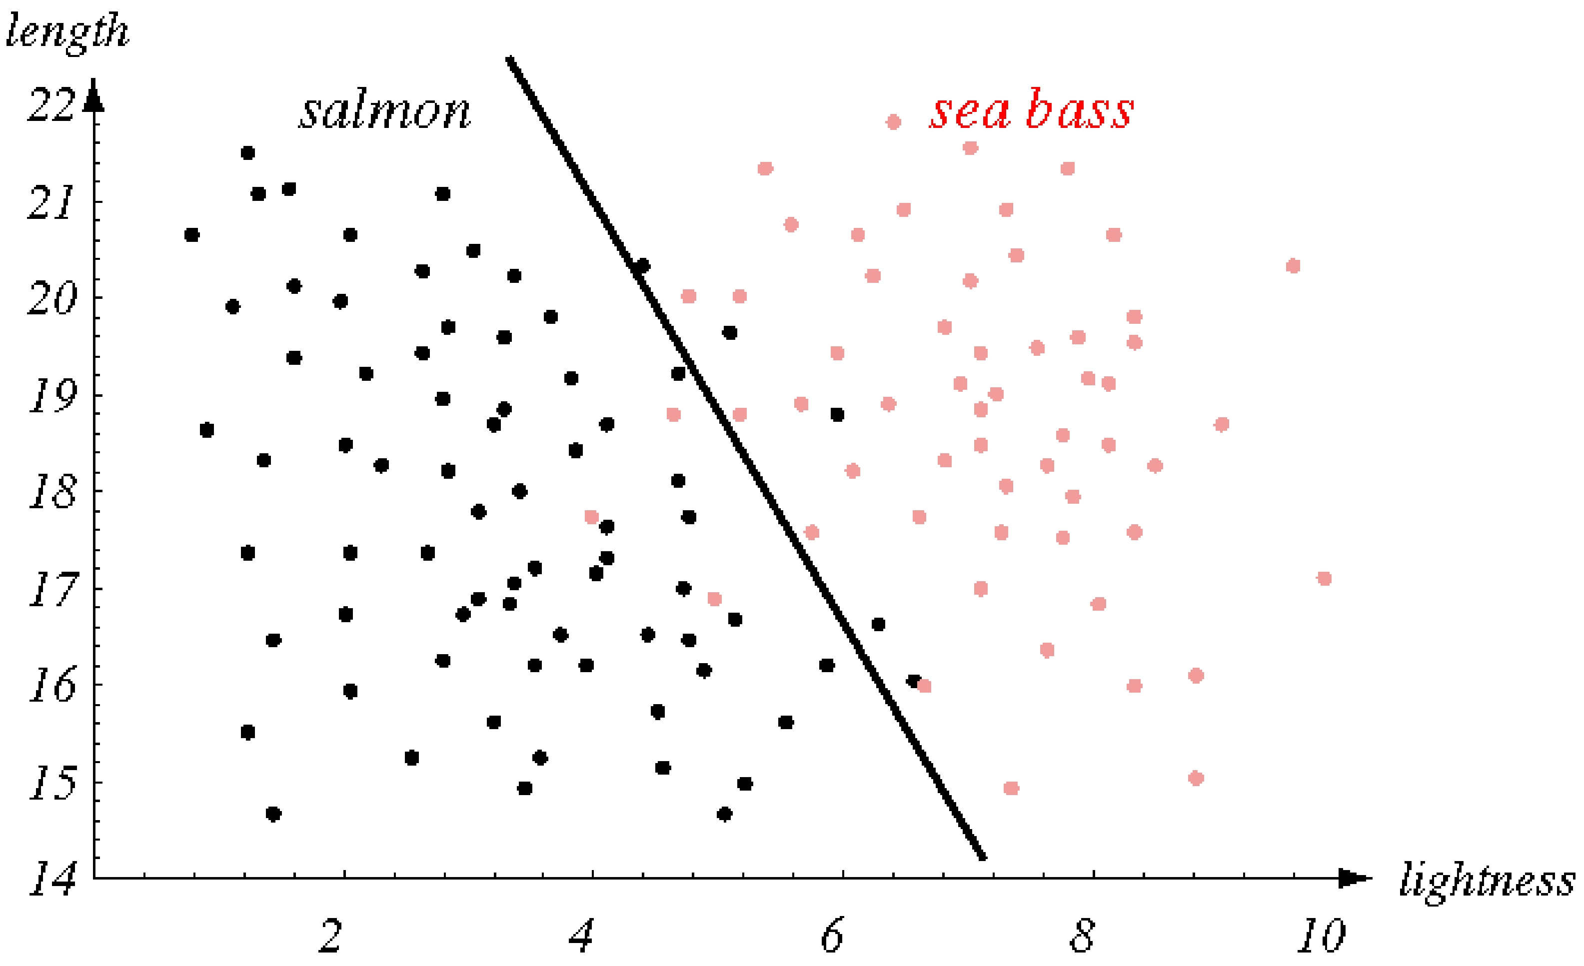

Scatter plot of lightness and length features for training samples. We can draw a decision boundary to divide the feature space into two regions.

An Example: Cost of Error

- We should also consider costs of different errors we make in our decisions.
- For example, if the fish packing company knows that:
  - Customers who buy salmon will object vigorously if they see sea bass in their cans.
  - Customers who buy sea bass will not be unhappy if they occasionally see some expensive salmon in their cans.
- How does this knowledge affect our decision?

An Example: Multiple Features

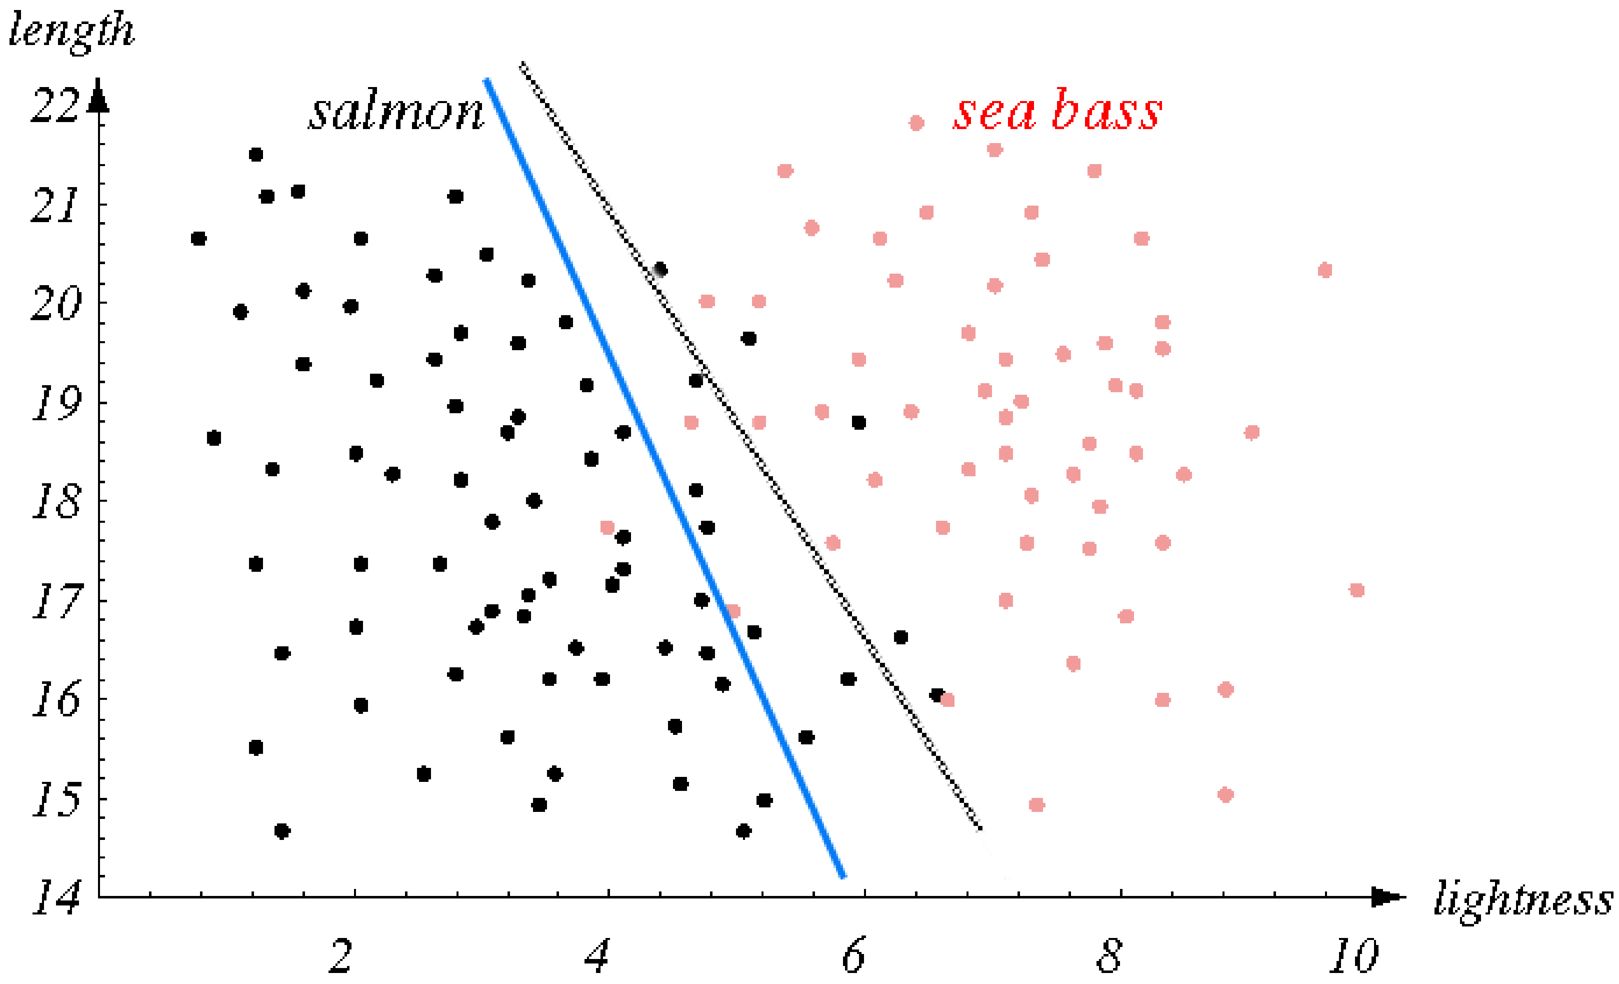

Scatter plot of lightness and length features for training samples with distinct costs.

An Example: Decision Boundaries

- Can we do better with another decision rule?
- More complex models result in more complex boundaries.

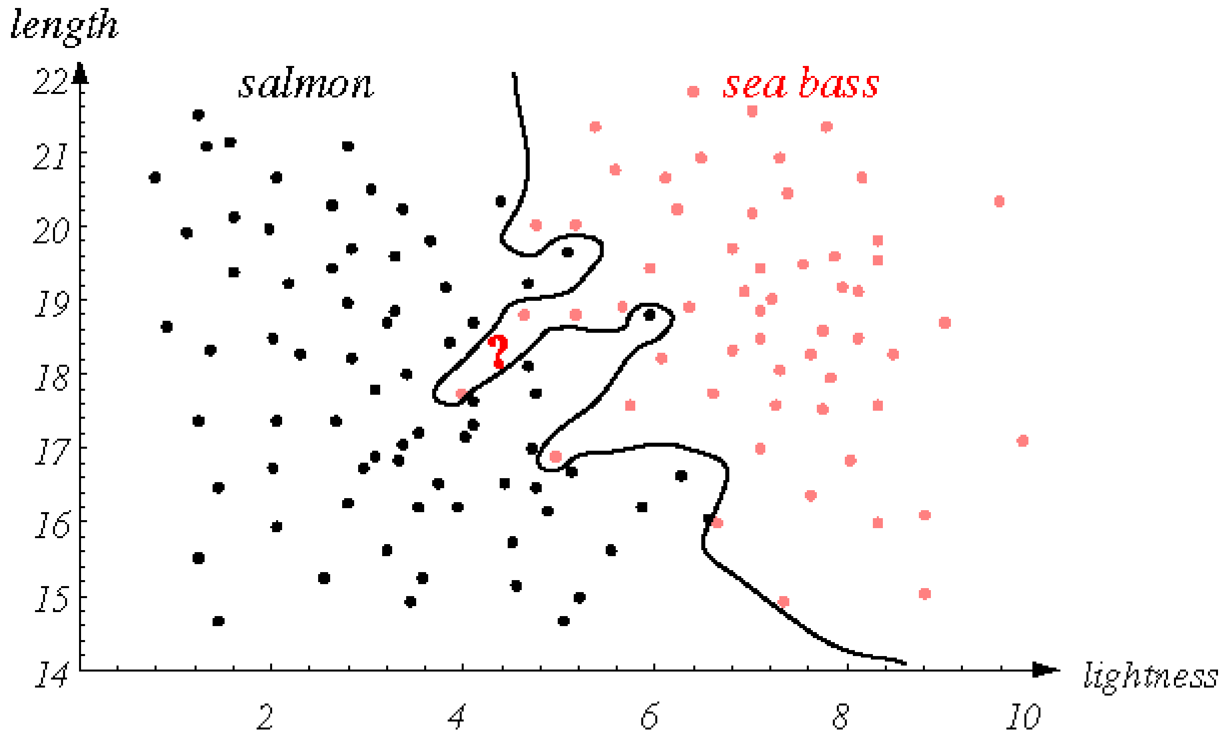

We may distinguish training samples perfectly but how can we predict how well we can generalize to unknown samples?

An Example: Decision Boundaries

- How can we manage the tradeoff between complexity of decision rules and their performance to unknown samples?

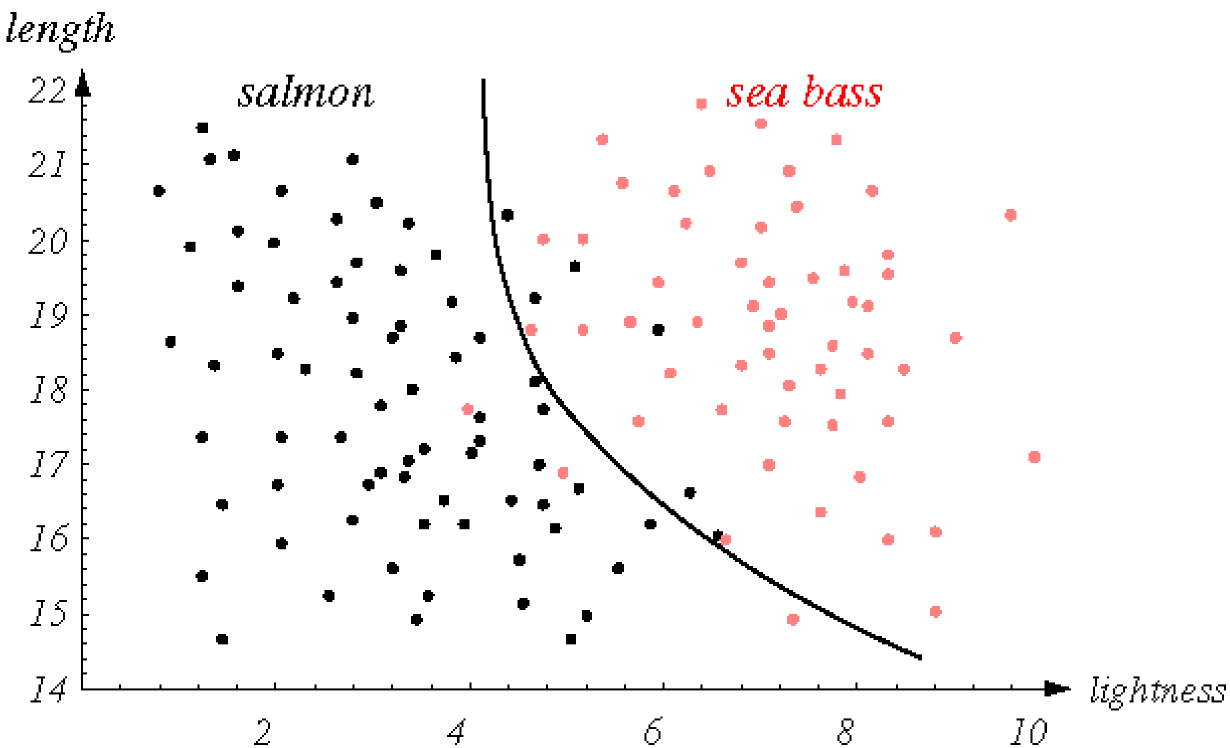

Different criteria lead to different decision boundaries.

The Design Cycle

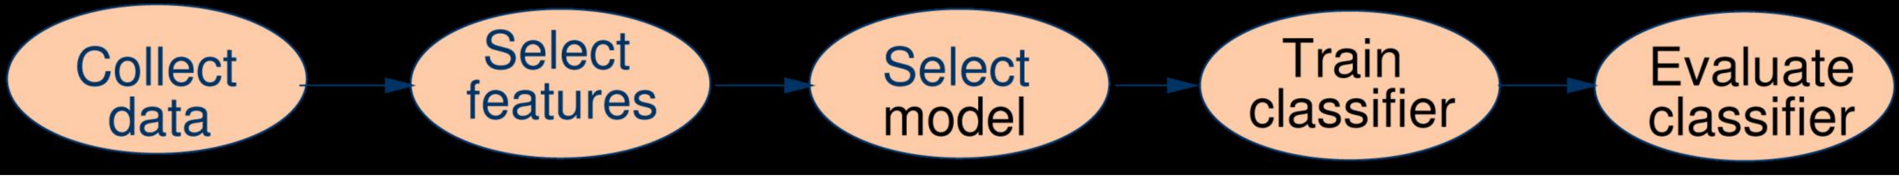

- Data collection:
  - Data acquisition and sensing
    - Measurements of physical variables
    - Important issues: bandwidth, resolution, sensitivity, distortion, SNR, latency, etc.
  - Collecting training and testing data
  - How can we know when we have adequately large and representative set of samples?

The Design Cycle

- Feature extraction and selection:
  - Finding a new **representation** in terms of features
  - Domain dependence and prior information
  - Computational cost and feasibility
  - Discriminative features
    - Similar values for similar patterns
    - Different values for different patterns
  - Invariant features with respect to translation, rotation and scale
  - Robust features with respect to occlusion, distortion, deformation, and variations in environment

The Design Cycle

- Model learning and estimation
  - Learning a mapping between features and pattern groups and categories
- Model selection & training:
  - Domain dependence and prior information
  - Definition of design criteria
  - Parametric vs. non-parametric models
  - Computational complexity
  - Types of models: templates, decision-theoretic or statistical, syntactic or structural, neural, and hybrid
  - How can we know how close we are to the true model underlying the patterns?
  - How can we learn the rule from data?

The Design Cycle

- Predicting:
  - Using features and learned models to assign a pattern to a category
- Evaluation:
  - How can we estimate the performance with training samples?
  - How can we predict the performance with future data?
  - Problems of overfitting and generalization

References

- **Selim Aksoy, Introduction to Pattern Recognition, Part I, http://retina.cs.bilkent.edu.tr/papers/patrec_tutorial1.pdf**
- Christopher M. Bishop, Pattern recognition and machine learning, Springer, 2006.
- Richard O. Duda, Peter E. Hart, David G. Stork, Pattern Classification, John Wiley & Sons, 2001