ILMATIETEEN LAITOS
METEOROLOGISKA INSTITUTET
FINNISH METEOROLOGICAL INSTITUTE

Automatic Image Analysis – Day 2

Mikko Syrjäsuo

Earth Observation

Finnish Meteorological Institute

Segmentation of images

Segmentation = dividing the image into (meaningful) regions

Object vs. background

Segmentation is typically based on either

- finding boundaries (edges) that separate regions
- combining similar subregions into larger regions

Application: recognition of targets (humans, galaxies, stars) partial targets (human face) or regions (land areas) for further analysis.

Segmentation of images

In most applications, a full segmentation of an image is not necessary!!

In some applications, no segmentation is needed at all!!

Partial or rough segmentation can be used

- to detect/locate targets
- to decide whether further processing is needed
- in constructing a rough understanding of image contents

Finding boundaries or discontinuities

The boundaries between regions are discontinuities of region content (types). Traditionally, these are called “edges” in images.

Edge detection is probably the oldest not yet satisfactorily solved problem in image processing!

What is a discontinuity in a discrete world, by the way?

And in discrete world corrupted with noise?

No noise, original data, first derivative and second derivative

Greyscale image

Greyscale values

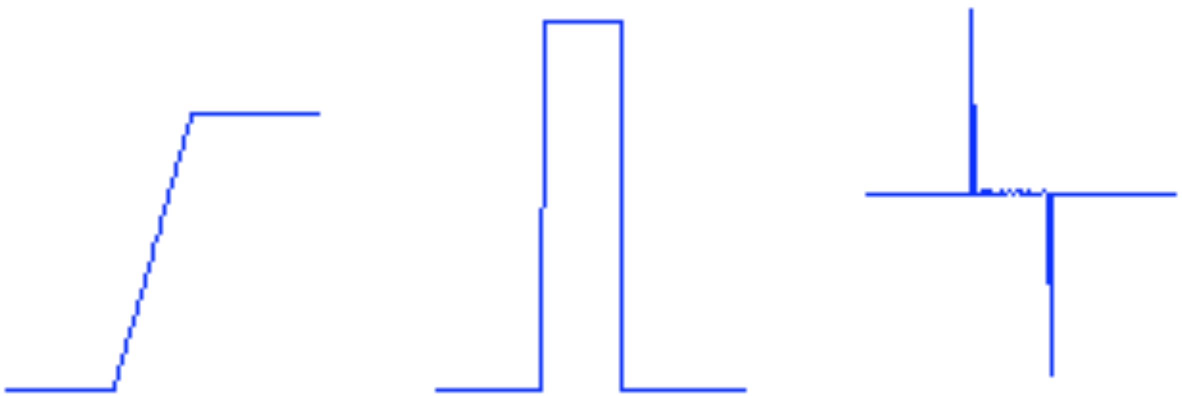

0.5% noise, original data, first derivative and second derivative

Greyscale image

Greyscale values

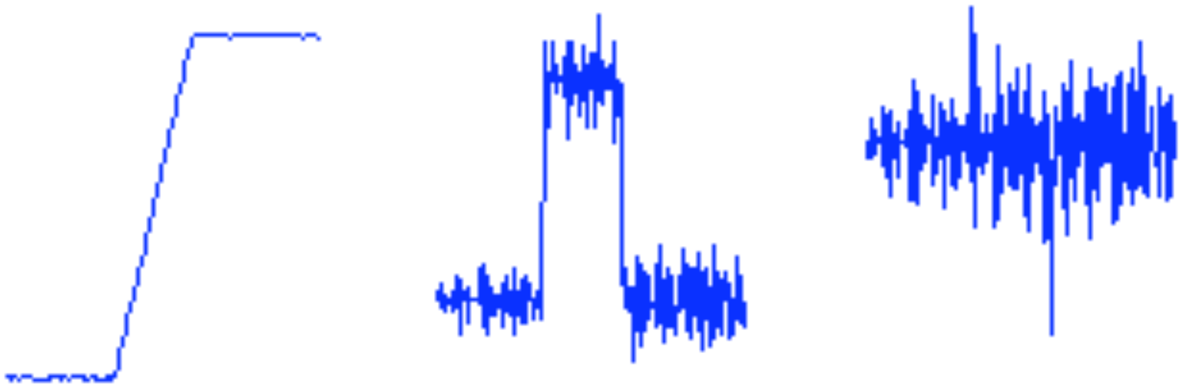

5% noise, original data, first derivative and second derivative

Greyscale image

Greyscale values

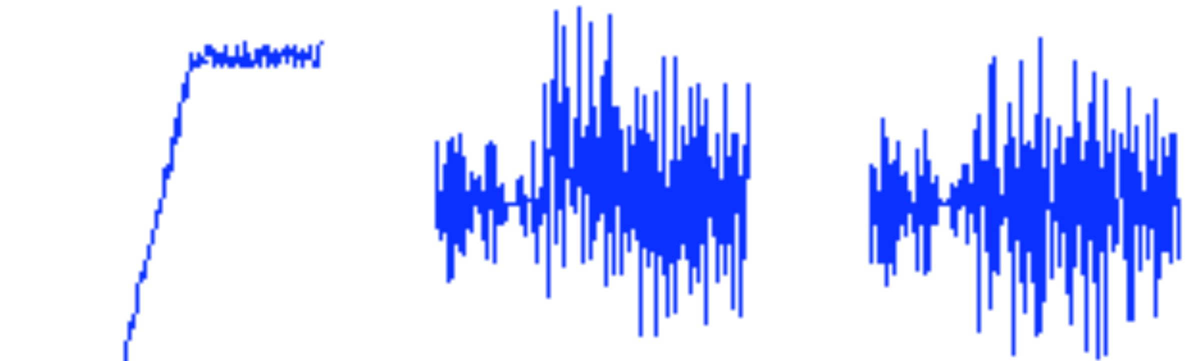

Finding boundaries in 2D, 3D etc.

In short: MORE DIFFICULT than in 1D...

- No general algorithm that works all the time exists
- Some boundary points will be detected but not all, combining partial results to a complete result is non-trivial
- Smoothing the data reduces the effects of noise but also blurs the boundaries...
- Several approaches are proposed in the literature, no clear winner yet, but if you need to start somewhere use the “Canny edge detector”

- MUCH easier if the target shape is known in advance

Finding boundaries in 2D, 3D etc.

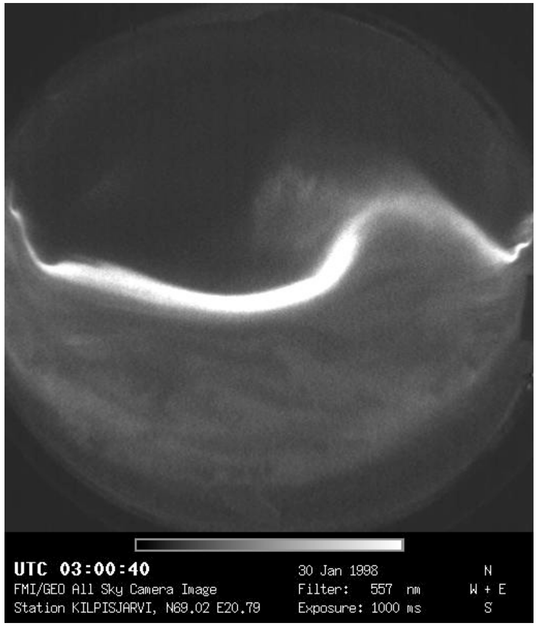

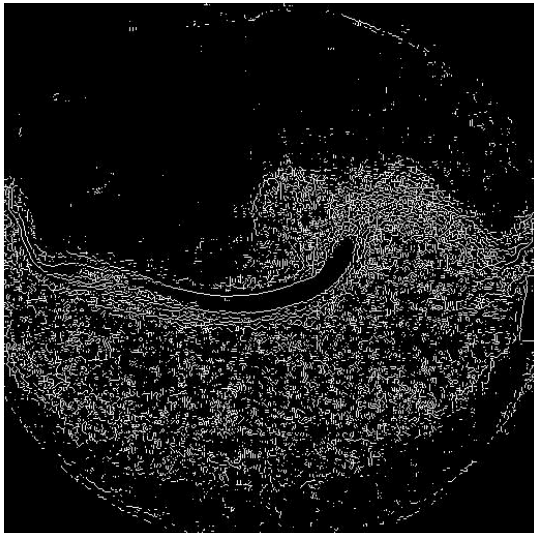

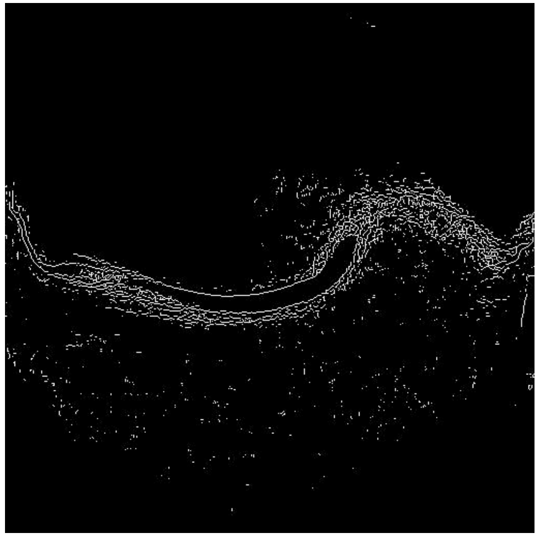

Edge detection output depends on the selected threshold!

ILMATIETEEN LAITOS
METEOROLOGISKA INSTITUTET
FINNISH METEOROLOGICAL INSTITUTE

Finding boundaries: Hough transform

If we know the target shape(s), the segmentation becomes easier. There are many methods in the literature, but most are based on **parametric representation** of the target(s) and **selecting** parameters to match the data.

Hough transform: line detection

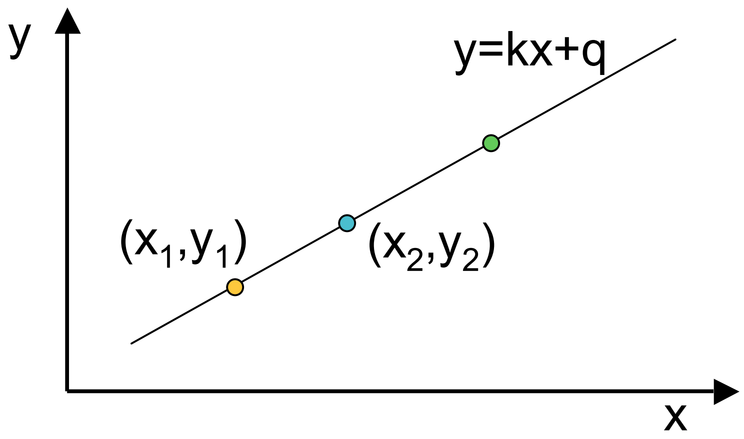

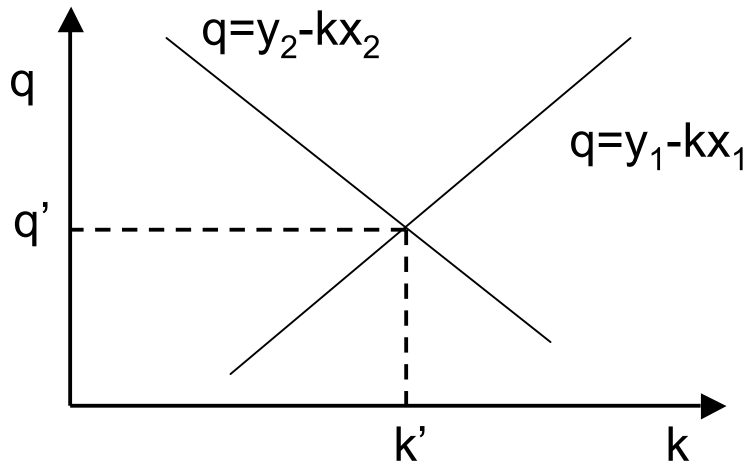

Finding boundaries: Hough transform

Hough transform: line detection

Use edge detection to see line candidates!

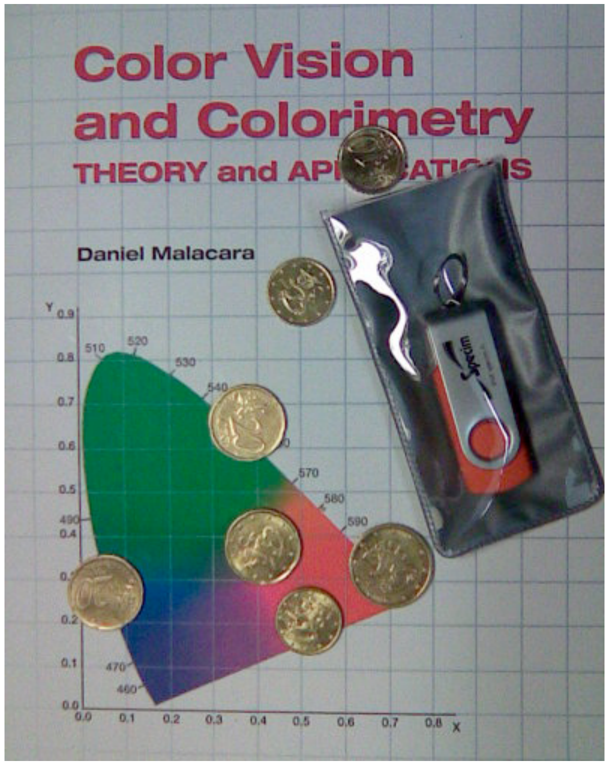

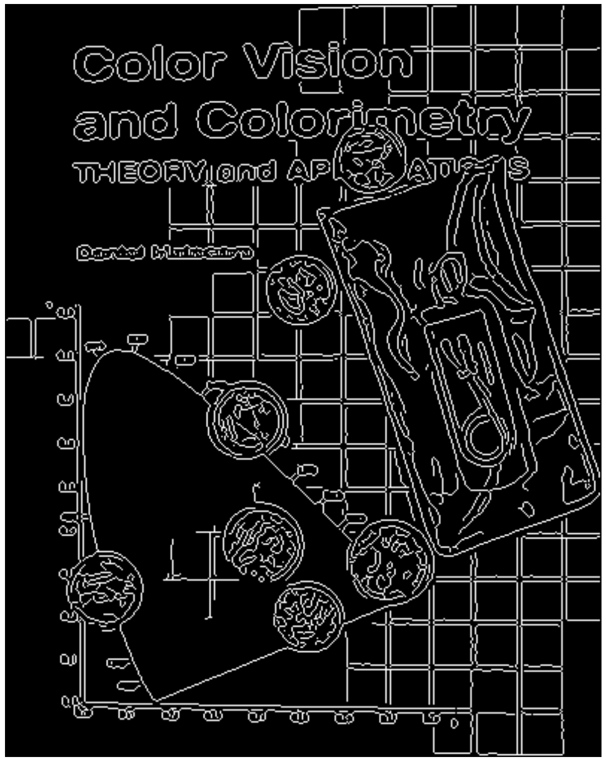

Finding boundaries: Hough transform

Accumulated evidence of lines!

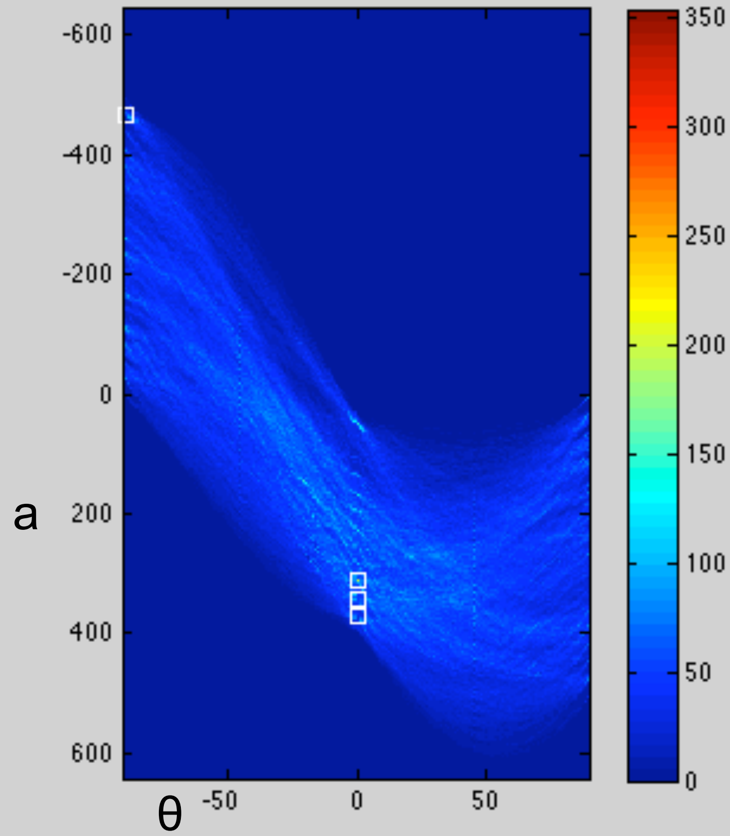

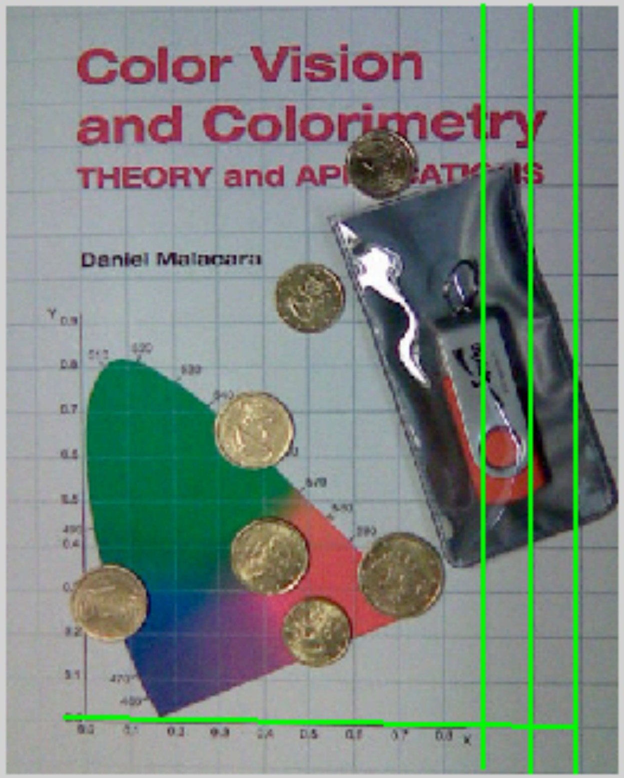

Finding boundaries: active shape models

More complex shapes can be represented as a B-spline or another computationally "convenient" form.

Shotton, J., Blake, A. and Cipolla, R., "*Contour-based learning for object detection*", Proc. IEEE Int. Conf. on Computer Vision (ICCV-05), 2005.

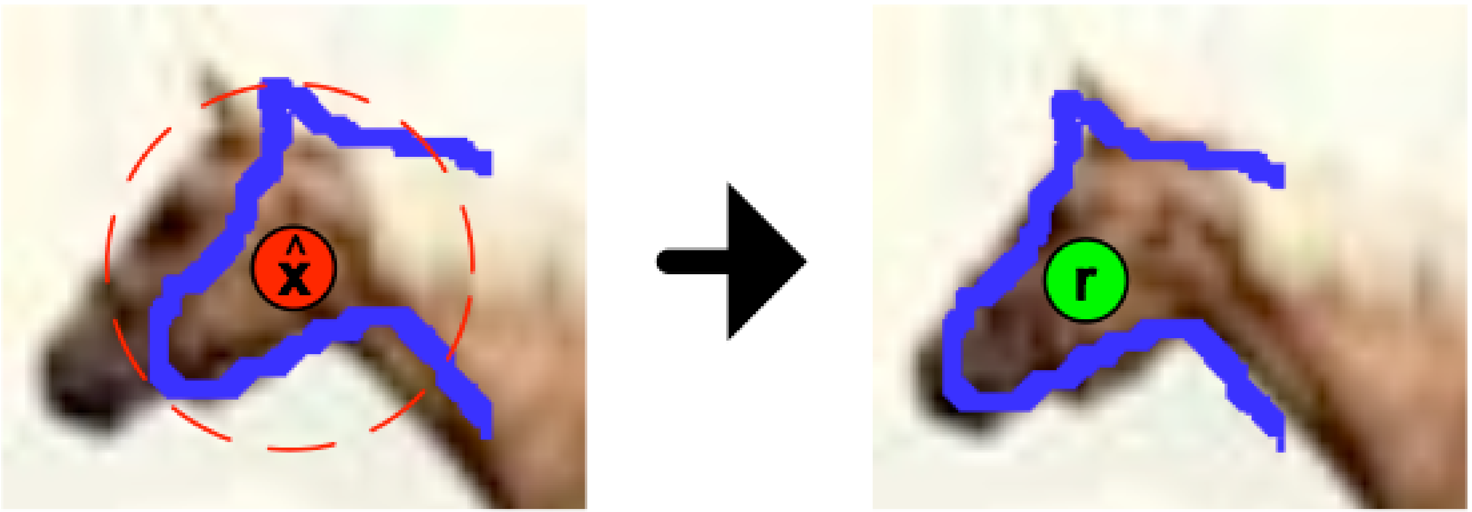

Finding boundaries: appearance models

Active appearance models learn a statistical model of the target shape(s).

Example: face detection

W. Zhao, R. Chellappa, A. Rosenfeld, P.J. Phillips, "Face Recognition: A Literature Survey", *ACM Computing Surveys*, 2003, pp. 399-458

Boundary (shape) descriptions

- The exact boundaries (object/region outlines) are usually approximated
  - **first few coefficients of Taylor/Fourier etc. series**
  - **splines**
  - **polygons**

- Region properties are commonly expressed using statistics
  - **moments (area, centroid, eccentricity)**

Auroral shapes

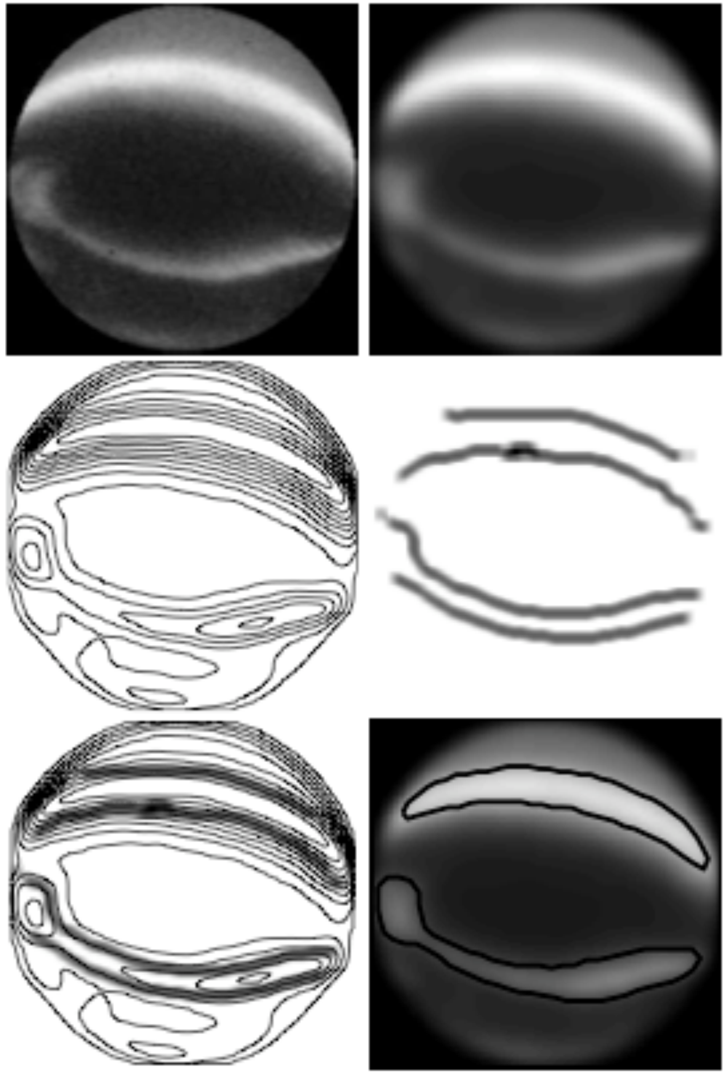

Syrjäsuo et al., “Content-based retrieval of auroral images – thousands of irregular shapes”, *Proc. Int. Conf. on Visualization, Imaging and Image Processing*, 224-228, 2004.

Extracting salient shapes (brighter spots)

Selected auroral shapes

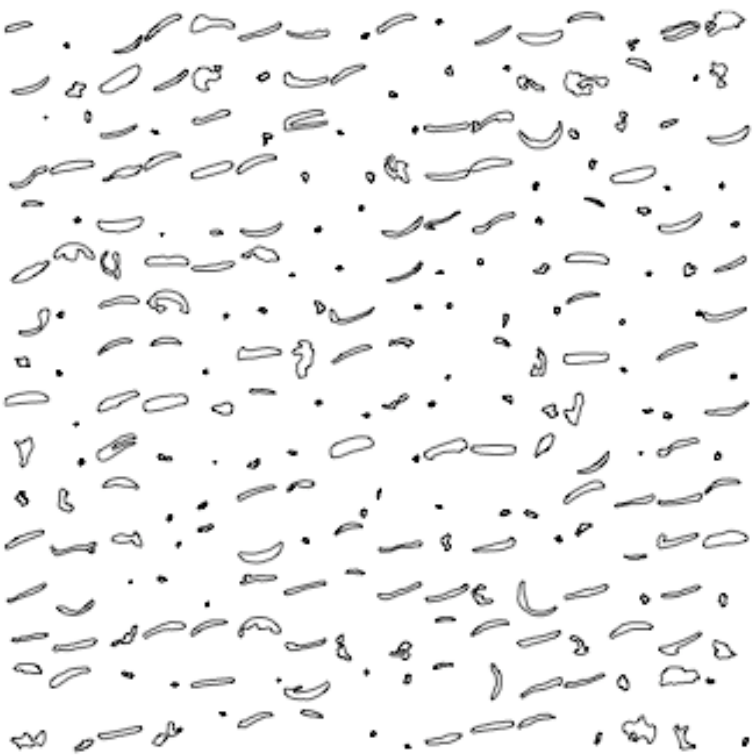

Various extracted shapes

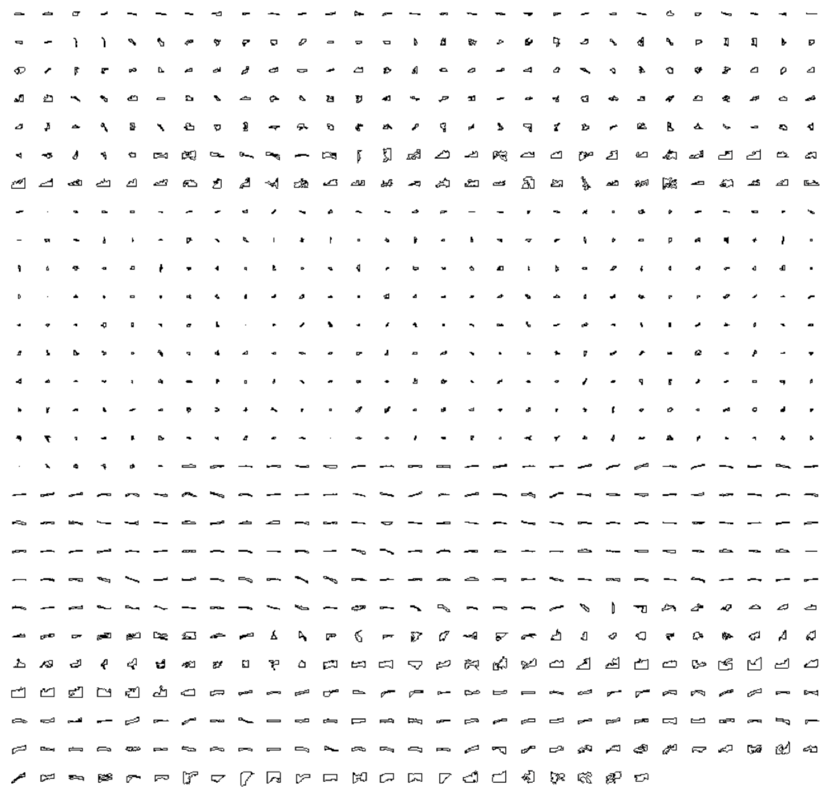

Shapes sorted using K-means

Similarity within a region

Note: a region is normally a continuous area of an image (volume in 3D etc.)

The region "appearances" differ between regions

- similarity is usually based on statistical measures
- relate to multidimensional classification/clustering problems!

Some commonly used properties:

- texture
- colour
- labels from classification/clustering based on (multidimensional) pixel data values

Texture

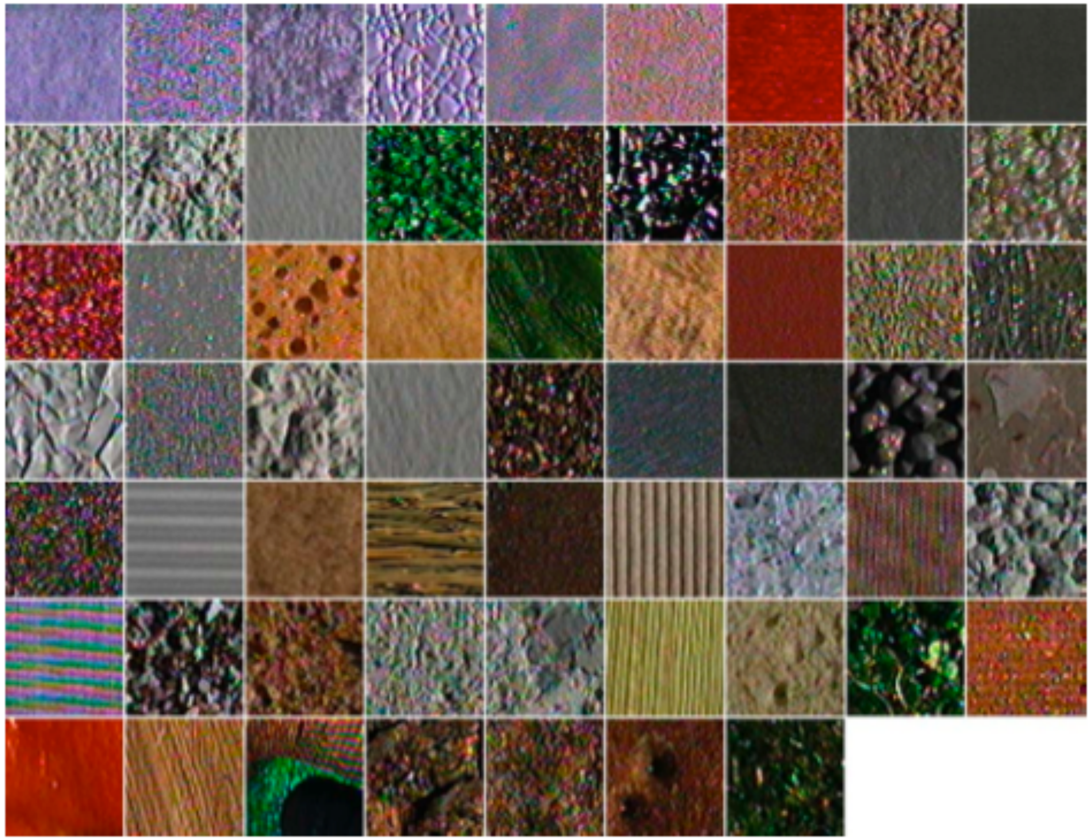

"Texture operators"
- a region of image as input
- the output is a feature vector

Example textures from the Columbia-Utrecht texture set.

Dana et al., "Reflectance and texture of real world surfaces", *ACM Transactions on Graphics*, 18(1):1-34, 1999.

Varma & Zisserman, "A statistical approach to texture classification from single images", *Int. Journal of Computer Vision*, 62(1/2), 61-81, 2005.

Segmentation based on similar regions

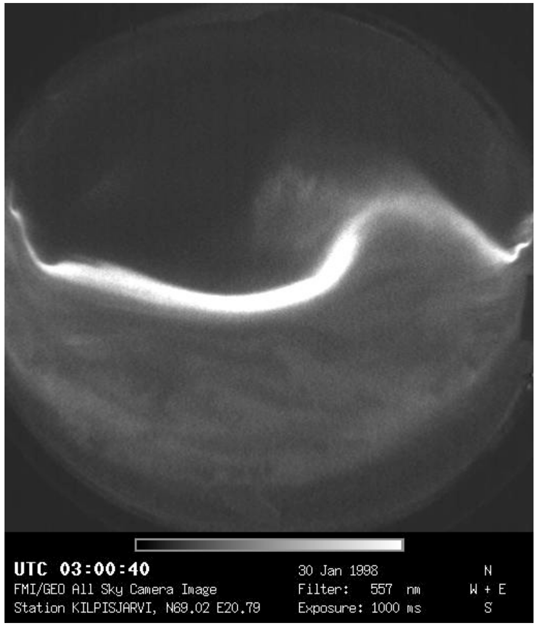

MIRACLE all-sky camera image, Kilpisjärvi, 1998-01-30 03:00:40UT

Threshold

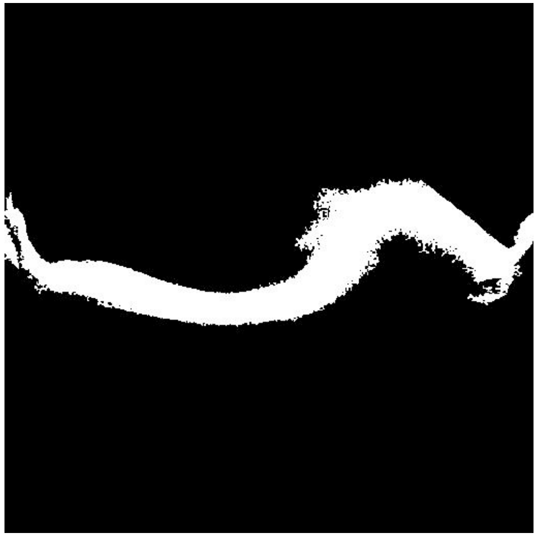

Largest contiguous region selected

Segmentation based on region colour

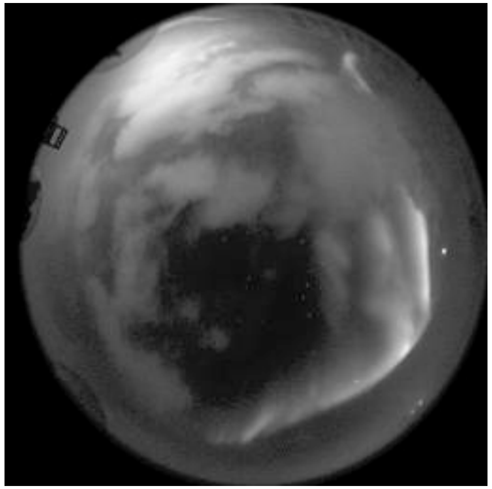

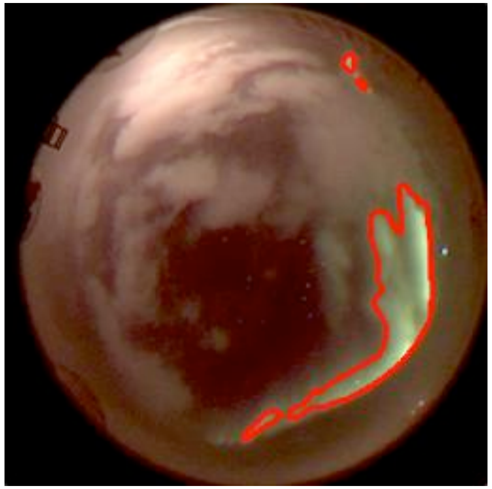

Typical “white light” auroral image:
• usually enough for a human expert

Syrjäsuo M.T., Jackel B.J., Donovan E.F., Trondsen T.S. and Greffen M., "Low-cost multi-band ground-based imaging of the aurora", *Proc. of SPIE Volume 5901 Solar Physics and Space Weather Instrumentation* (eds. Silvano Fineschi, Rodney A. Viereck), 2005.

The same scene captured with a camera designed with automatic analysis in mind:
• auroral outlines automatically added
• also better for human expert...

New approaches

A new auroral imager designed with computer vision applications in mind
— a surprisingly good spectral recording device, too!

Partamies, N. et al., “Using colour in auroral images”, *Canadian Journal of Physics*, Vol. 85, No. 2, pp. 101-109, 2007.

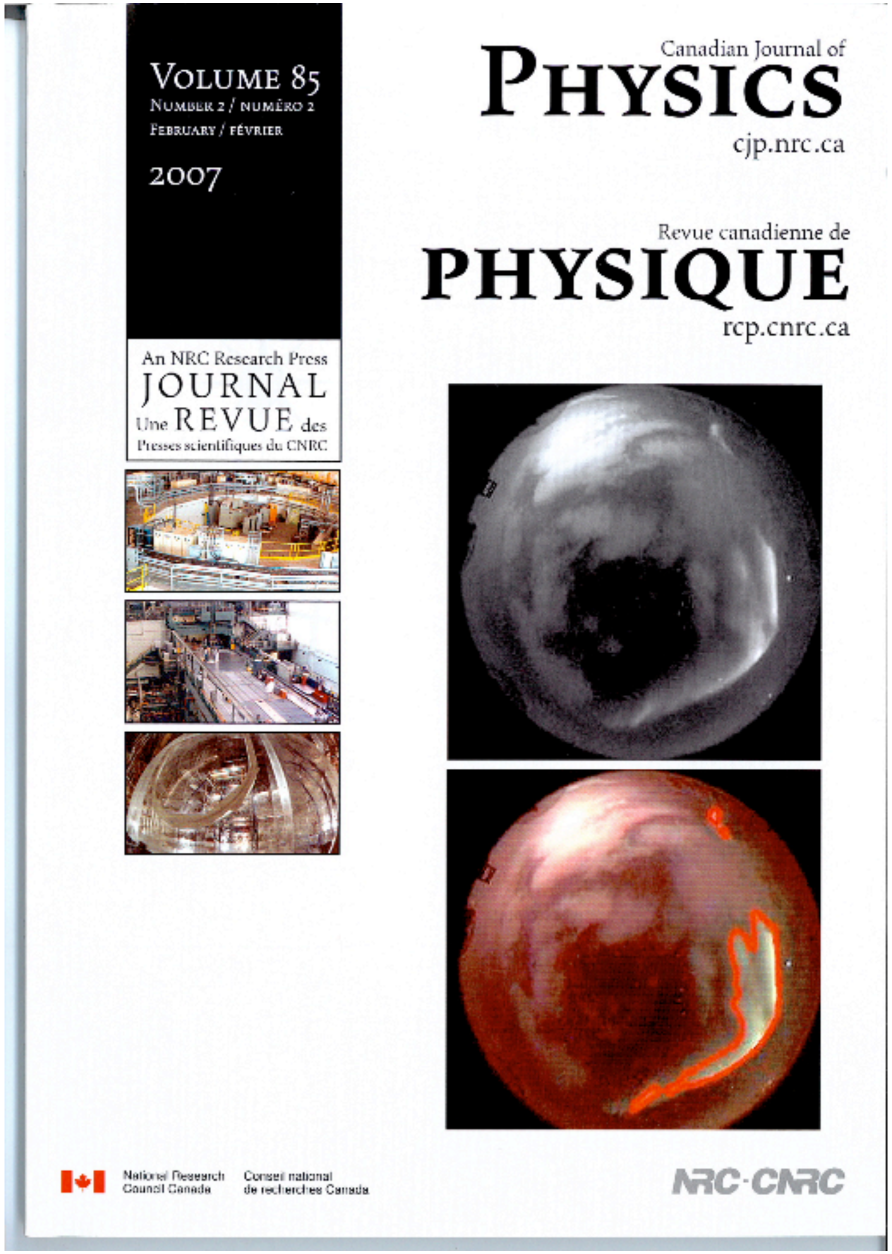

Segmentation demonstration

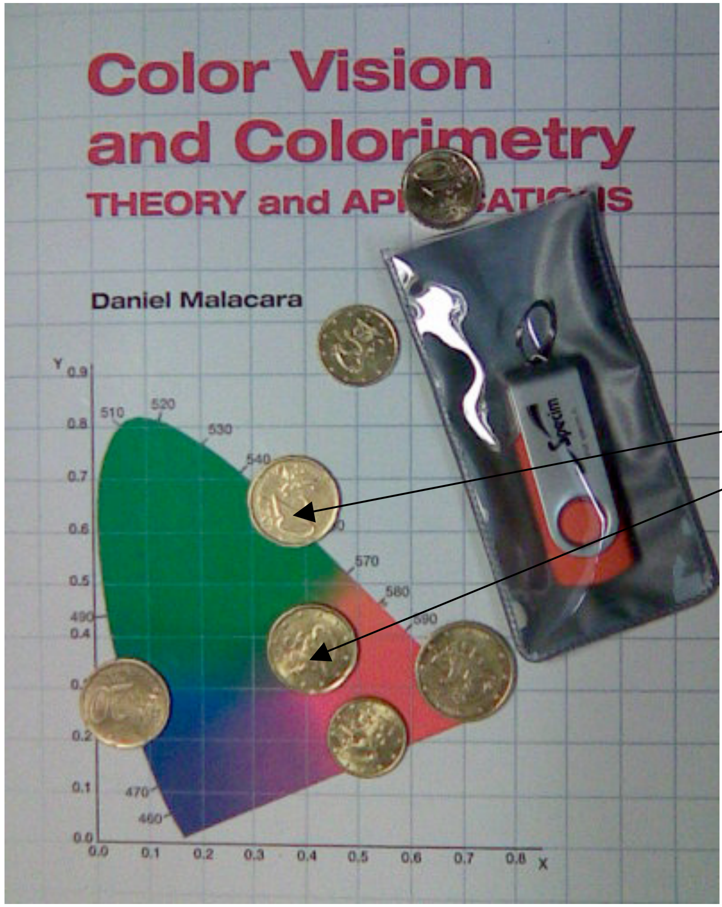

Sample the coin colours to obtain a colour range

Segmentation demonstration (cont.)

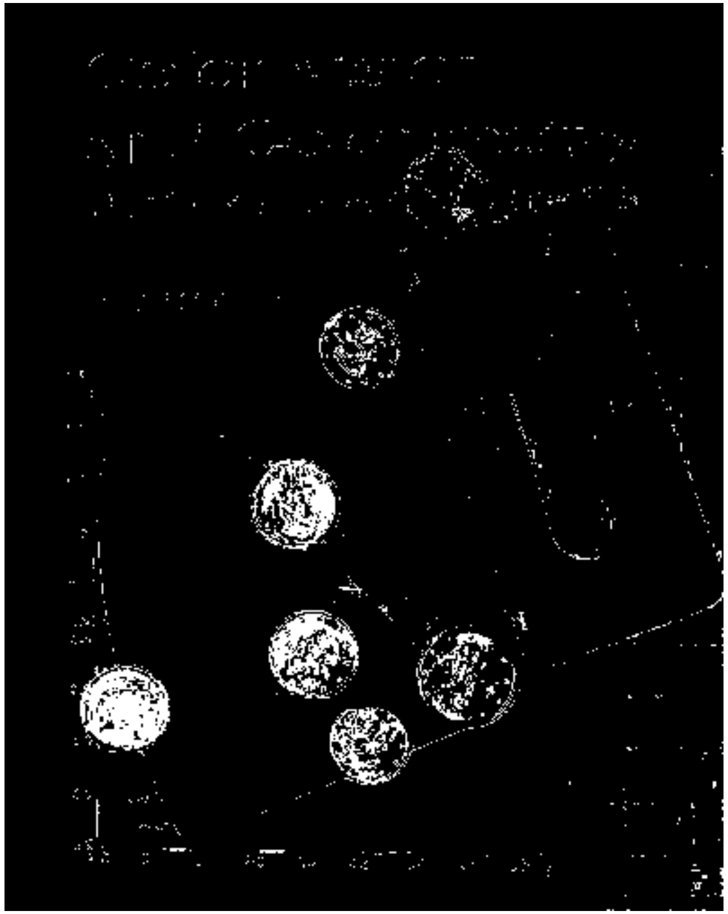

All pixels that have a RGB colour value in the 3D-space

R=98...163
G=85...148
B=57...104

are shown in white.

Segmentation demonstration (cont.)

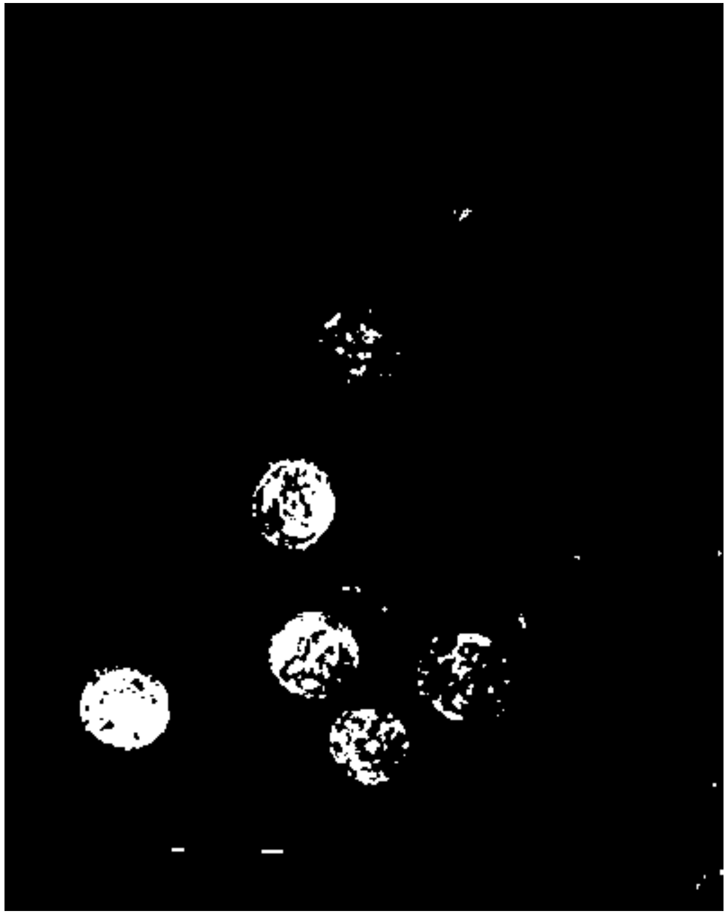

Noise suppressed by using a 3x3 median filter: individual white pixels have been removed.

Smoothing (averaging) is not as effective as median for “salt & pepper” noise (why?).

Segmentation demonstration (cont.)

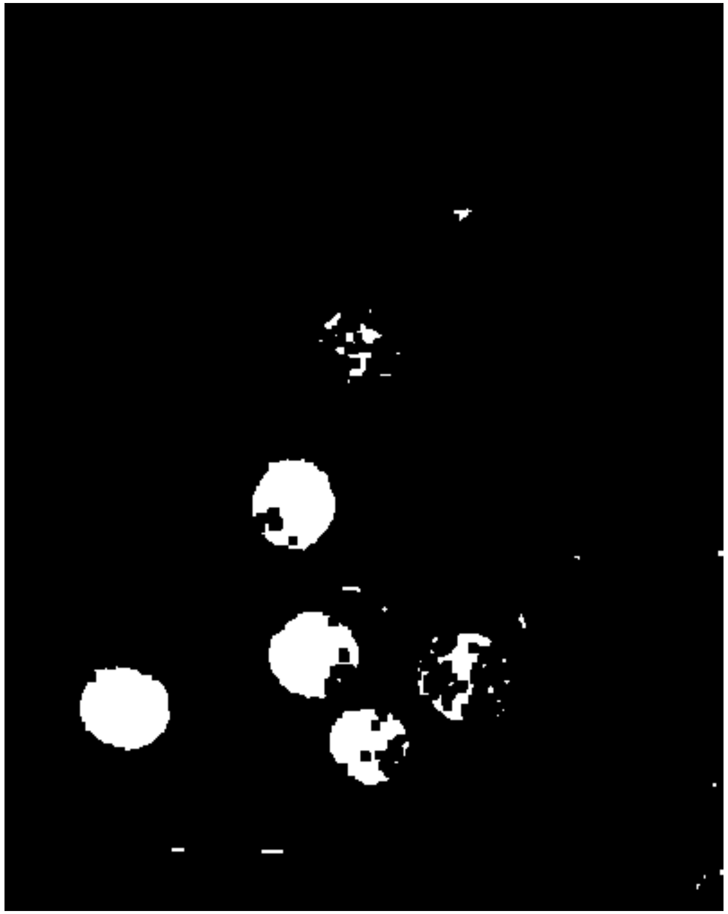

Using morphological operators to "close" small holes in the image.

Use "region labelling" to assign a label to each contiguous blob in the image and compute the areas...

Segmentation demonstration (cont.)

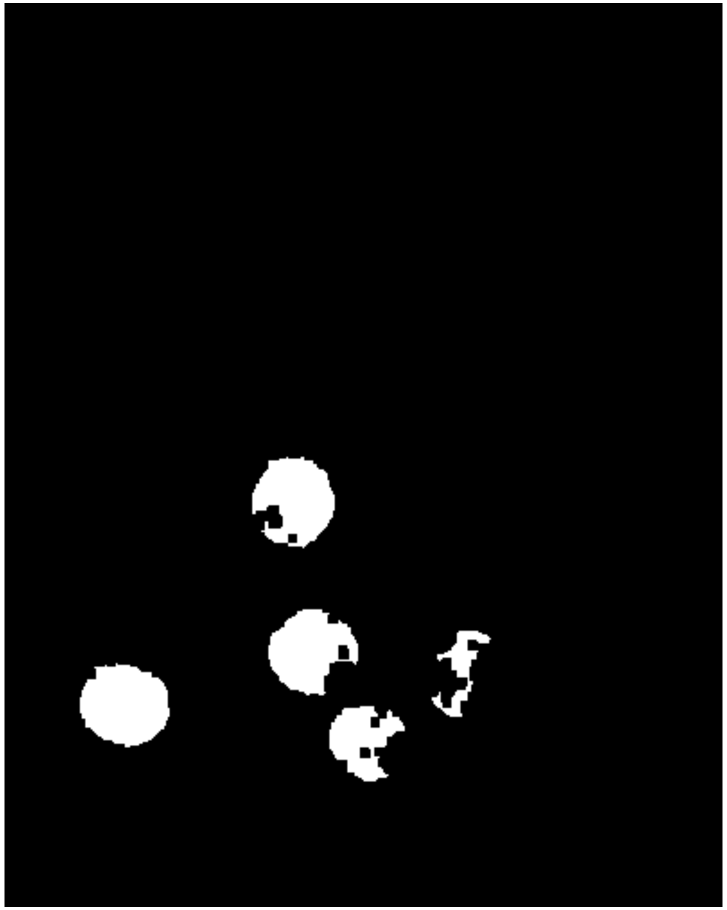

Select regions that large enough

Segmentation demonstration (cont.)

Final segmentation

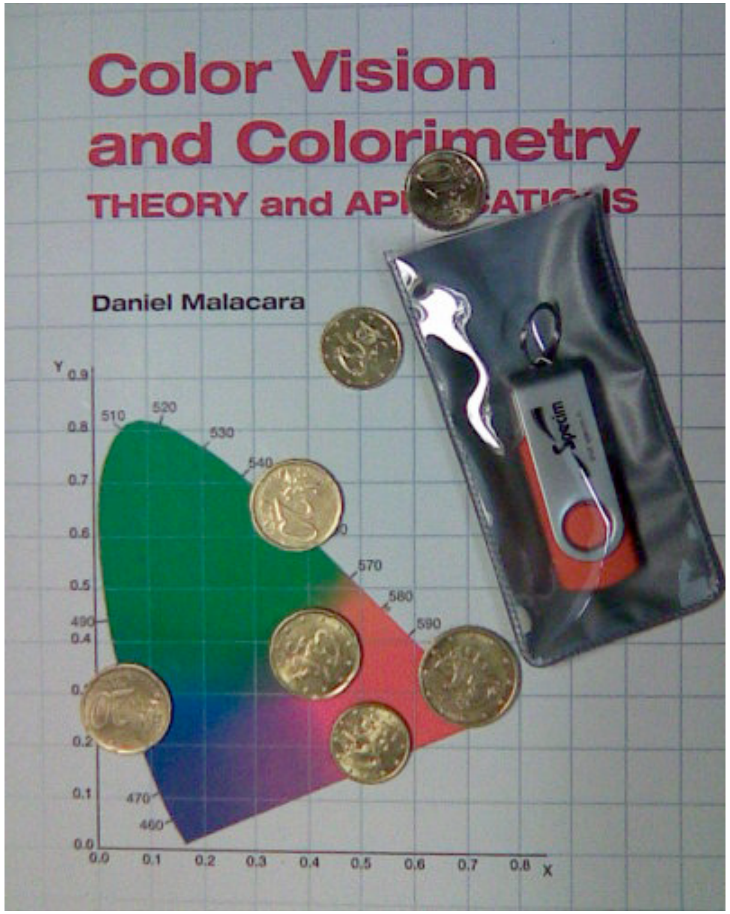

Detection, classification, adaptation

- in many cases, the detection of a target is implemented as a classification problem
  - **known targets ⇒ matching model shapes**
  - **“blobs” ⇒ classification, clustering**
  - **“is this the same object?” (⇒ 1-class problems)**
- multidimensional feature vectors
  - **similarity metrics**
  - **distributions (curse of dimensionality!)**
  - **probabilistic (statistical) methods**

Modern segmentation approach

- the state-of-art segmentation methods are complex
  - **probabilistic approaches**
  - **multiple “cues” in the image are utilised**
  - **interest regions**
  - **motion (if available)**
  - **hierarchical ordering**

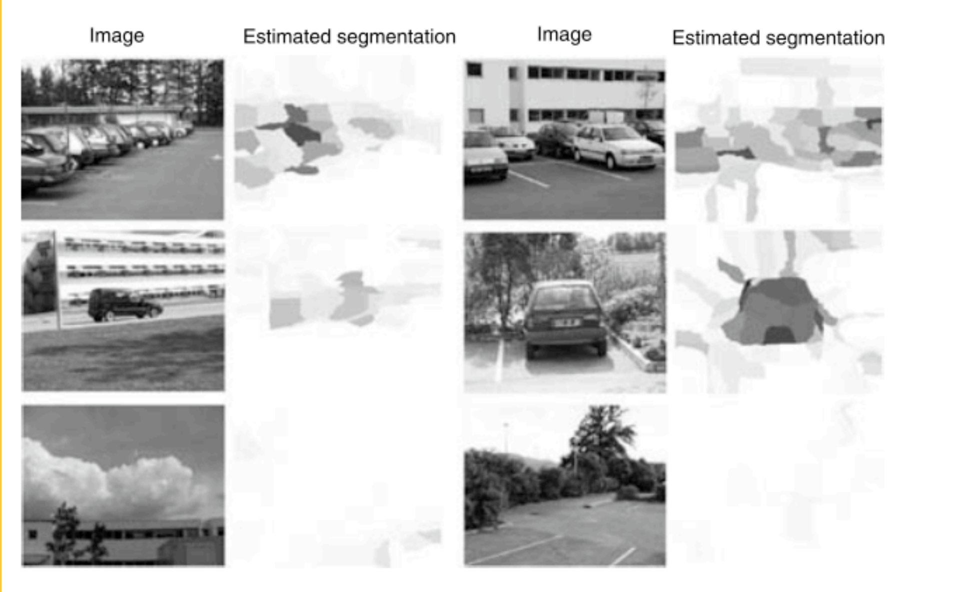

Darker segments show higher likelyhoods of car locations

Carbonetto et al.,”Learning to recognize objects with little supervision”, *Int. J. Computer Vision*, 77:219-237, 2008.

Image understanding

- the “ultimate goal” of automatic image analysis
  - **explaining the image contents**
- again, general solutions are more difficult (impossible) than particular solutions
- fortunately:
  - **not usually required in space physics or astronomy applications!**

Image understanding strategies

Given a random image, understanding the contents is more or less impossible! (even for humans!)

- **segmentation driven by a priori information**
  - camera/telescope point combined with date / time
- **region labeling**
  - merging/splitting regions
  - relation of region positions
  - hierarchical structure
  - ignoring noise (impossible regions)
- **mapping the regions to a human concepts**

Image understanding – real world case (1)

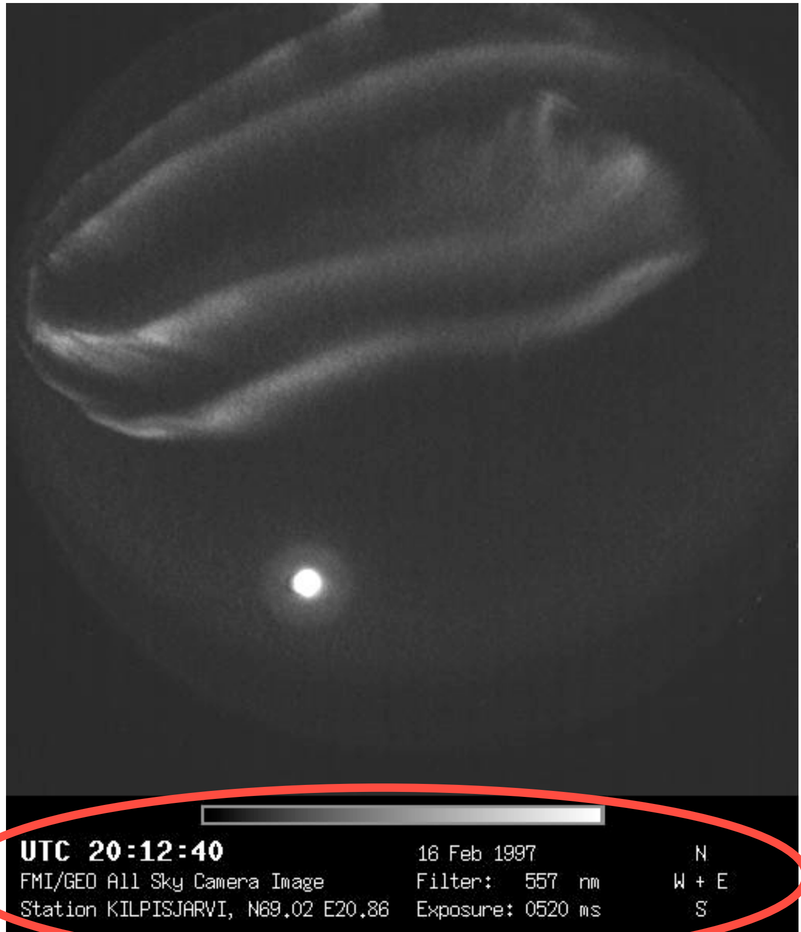

Type of image
Time and location
Approximate orientation

Image understanding – real world case (2)

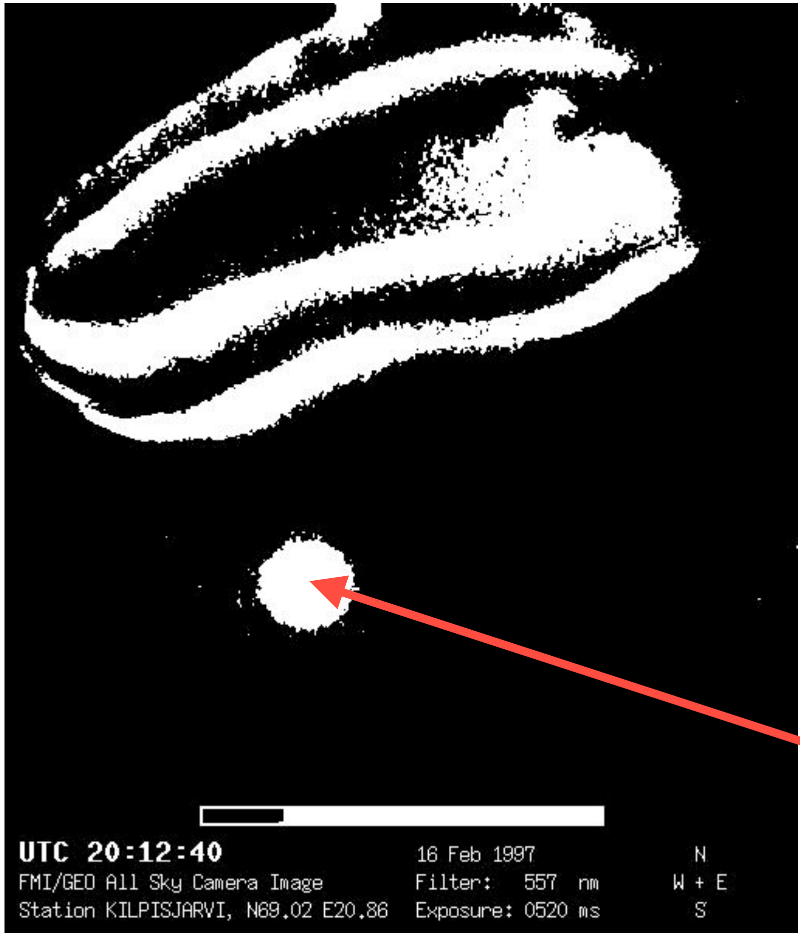

Segmentation (threshold)
+
region properties
+
knowledge of the Moon position in the sky
=
a round object where the moon is expected to be

⇒ Moon identified

Image understanding – real world case (3)

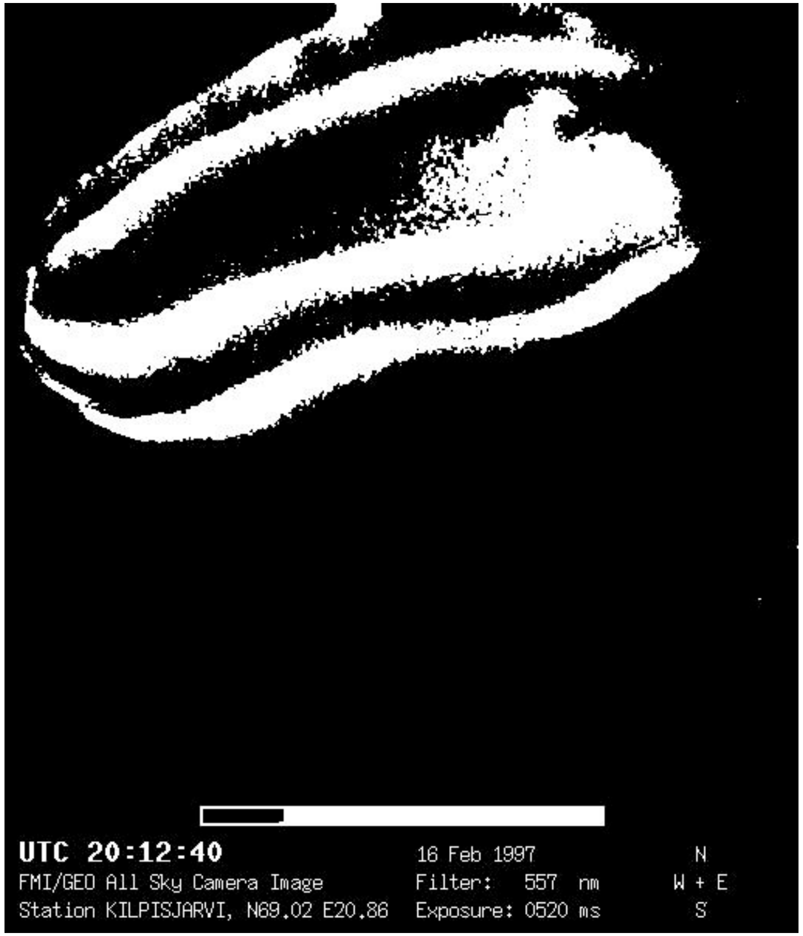

Moon “removed” from the remaining segments

Remaining regions are classified as auroras because of their brightness, shape, size and alignment (*auroral arcs*)

⇒

Scene: bright auroras in the sky north of Kilpisjärvi (with the Moon in the southern sky)

Case study: auroral occurrence (1)

- Auroral appearance depends on the time of day
- Early studies have been performed manually

Syrjäsuo and Donovan, “Diurnal auroral occurrence statistics obtained via machine vision”, *Annales Geophysicae*, 22, 1103-1113, 2004.

- CANOPUS all-sky imager data from Gillam, Manitoba, Canada
- 350,000 image subset
- KNN classifier using a manually classified training set (supervised learning)
- several image features concatenated to form one multidimensional feature vector for each image

Case study: auroral occurrence (2)

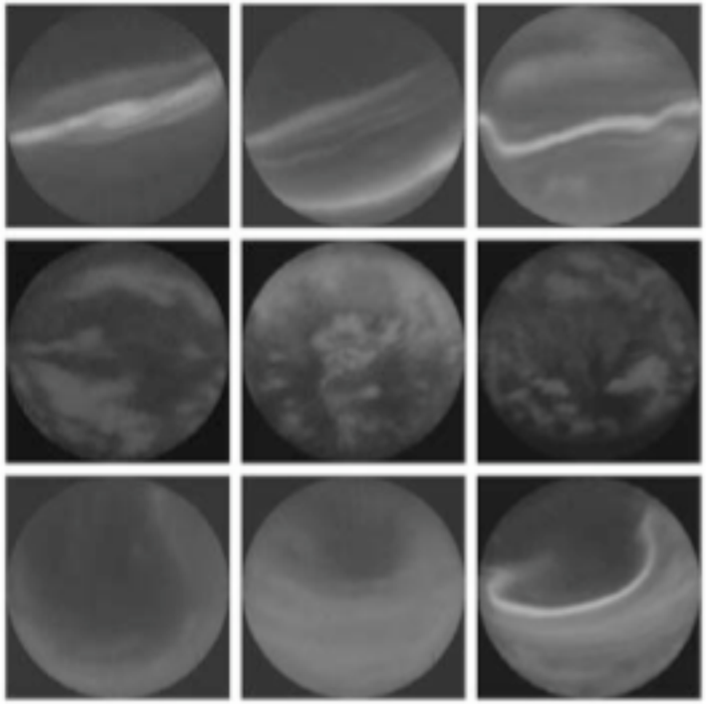

Auroral appearances: arcs, patchy auroras and Omega-bands

Several features:

1. auroral brightness
2. north-south brightness distribution
3. east-west brightness distribution
4. texture (i.e. overall appearance)

Case study: auroral occurrence (3)

- Classification into four categories
  1. **No aurora**
  2. **Arcs (one or more)**
  3. **Patchy aurora**
  4. **Omega-bands**
  5. **Other/reject**

- 300 carefully selected sample images from categories 2,3 and 4

ILMATIETEEN LAITOS
METEOROLOGISKA INSTITUTET
FINNISH METEOROLOGICAL INSTITUTE

Case study: auroral occurrence (5)

- **The sample images were used as a dictionary: a previously unseen image (feature vector) is compared to known images (feature vectors)**
  - K-nearest neighbours classifier
  - training set classification accuracy 90±3%

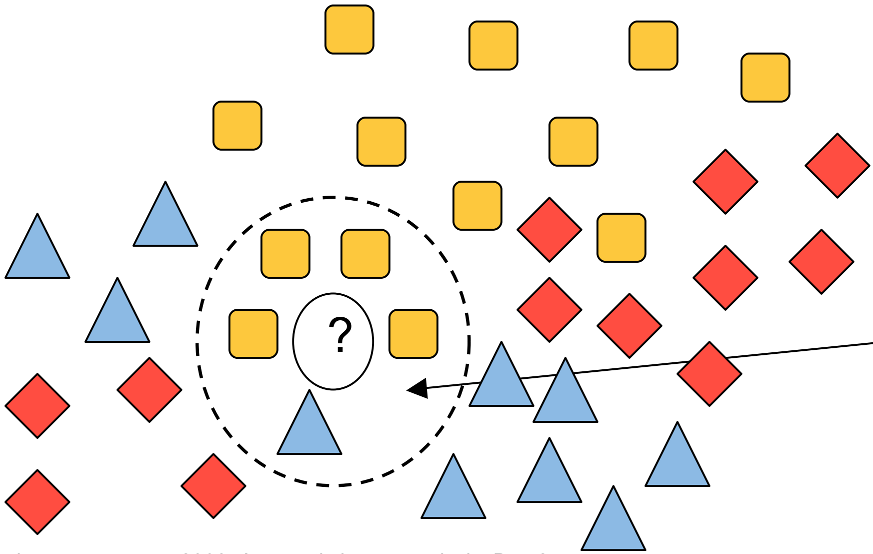

The previously unseen sample is classified to the class of the majority of the neighbours: here

Case study: auroral occurrence (6)

- 350,000 auroral images ⇒ 220,000 images with aurora ⇒ 30,000 images with confident classification
  - **the auroral appearance is not well-defined even from the human observer point of view and this is reflected in this study as well!**

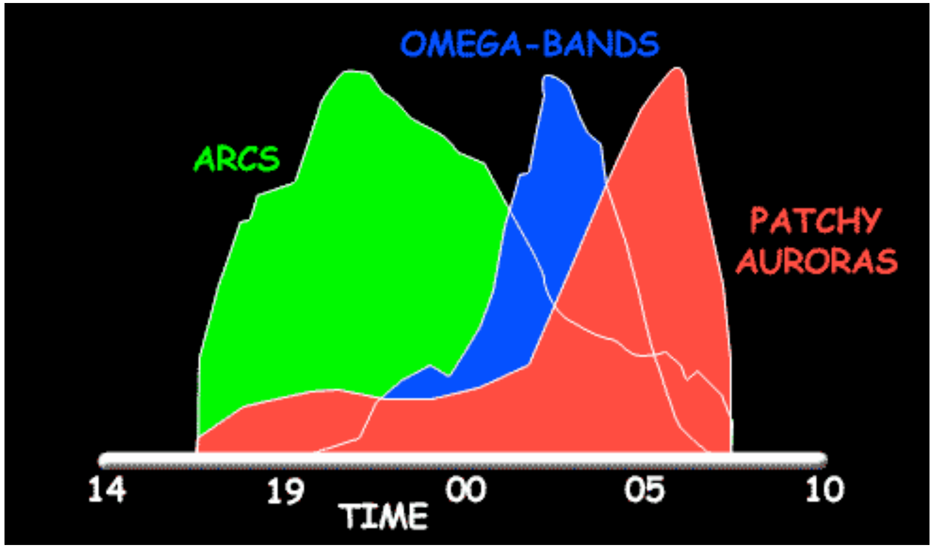

Sharp cutoffs in evening and morning are instrument effects.

Occurrences are normalised within each class, there were 17,000 arcs, 9700 patchy auroras and 600 Omega-bands

Reality check: could objects be detected?

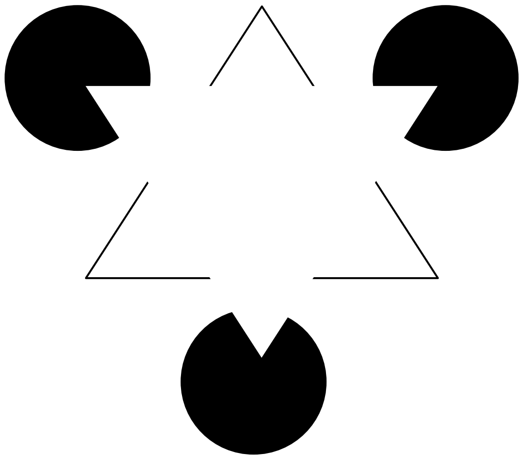

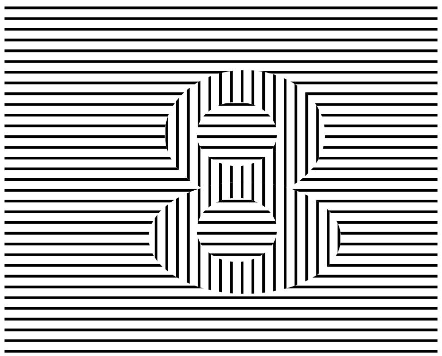

Exercise: Finding all circular objects

In this case study, we use incomplete information of edges (boundaries) to locate circular objects (coins) in an image.

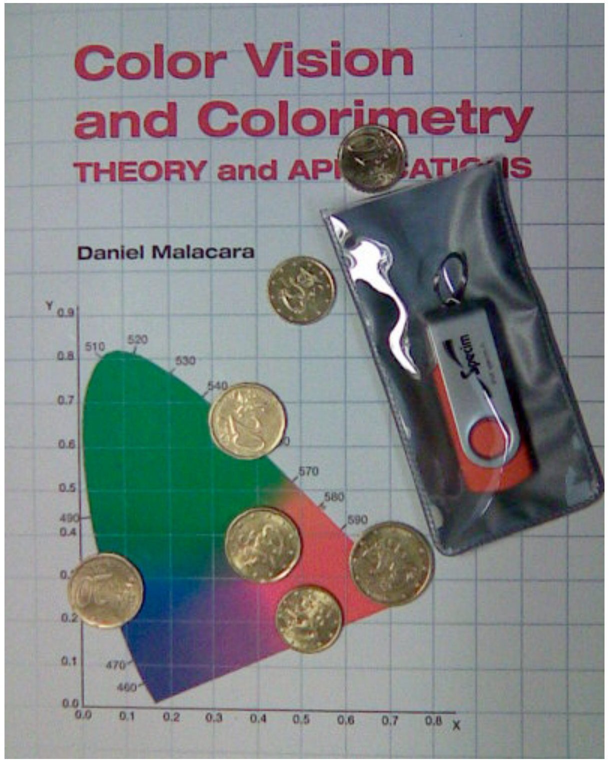

Exercise: Finding all circular objects

- The input image shows detected edges (pixel value is 1)
- Use Hough transform to find circles $(x-x_0)^2+(y-y_0)^2=r^2$
- Search for circles with r=23 pixels
- Accept those circle locations where at least a third of the circle is supported by detected edges