Spatial statistics and image analysis (TMS016/MSA301)

Introduction

2021-03-21

Practical information: Lectures and computer exercises

Lectures: Konstantinos Konstantinou and Aila Särkkä (examiner)

Monday 8:00-9:45, Euler

Wednesday 10:00-11:45, Euler

Computer exercises: Konstantinos Konstantinou

Monday 13:15-15:00, MVF24 and MVF25

Wednesday 13:15-15:00, MVF24 and MVF25

(Needed to be able to do the projects.)

Practical information: course literature

The course is mainly based on the lecture notes "Statistics of Imaging" by Mats Rudemo.

Other useful reading material:

- Gelfand, A.E., Diggle, P., Fuentes. M. and Guttorp, P. (2010). Handbook of Spatial Statistics, Chapman and Hall/CRC.
- Hastie, T., Tibshirani, R. and Friedman, J. (2009). The Elements of Statistical Learning, Springer.
- Efron, B. and and Hastie, T. (2016). Computer Age Statistical Inference, Cambridge University Press.
- Glasbey, C.A. and Horgan, G.W. (1995). Image Analysis for the Biological Sciences, Wiley.

The three first books are available as eBooks and the fourth one can be found chapter-wise in Canvas in the directory Files/GlasbeyHorgan.

Practical information: examination

Examination consists of two components

- written exam at the end of the course,
- project assignments

The project

- is performed in groups of 1-3 students
- consists of three parts
  - two problems introduced in the computer exercises and graded by Konstantinos
  - one problem (major part, see Files/Project in Canvas) that you choose on your own (with approval from Aila), presented at a seminar and as a written report at the end of the course.

Practical information

How to find a group?

- As soon as you have a group, please, register it in Canvas in People/Group (or inform me by email).
- If you want to be assigned to a group, please, send an email to me.

Programming language

- For project parts 1 and 2, Matlab is recommended since it has all the code you will need.
- For project part 3, you can use any program.

Practical information: deadlines for the projects

For Parts 1 and 2 (see the course PM)

- Submit a PDF file containing these two parts in Canvas (Assignments) at the latest May 11. Include the Matlab code as a zip-file.
- Konstantinos may ask you to complement your answers and you should submit your revised versions in Canvas (Assignments) at the latest May 25.

Practical information: deadlines for the projects

For Part 3 (see the course PM)

- Submit a planning report in Canvas (Assignments) at the latest April 25
  - Preliminary title
  - Names of the people in the group
  - Purpose for the study
  - Short description of the data
- Submit a PDF file containing a preliminary version of the report (one submission for each group) at the latest May 18. It does not need to be complete.
- Submit the final version of the report at the latest May 27, together with the Matlab code (or other program code) as a zip-file.

Course content

- Classical image analysis methods
- Statistical and machine learning methods for image data
- Model-based spatial statistics
- Spatial point processes
- Applications
  - Pattern recognition
  - Remote sensing
  - Microscopy

Examples of images: aerial photographs

A thinning experiment of Norway spruce trees in northern Sealand in Denmark. The airplane was 560m above "Nadir".

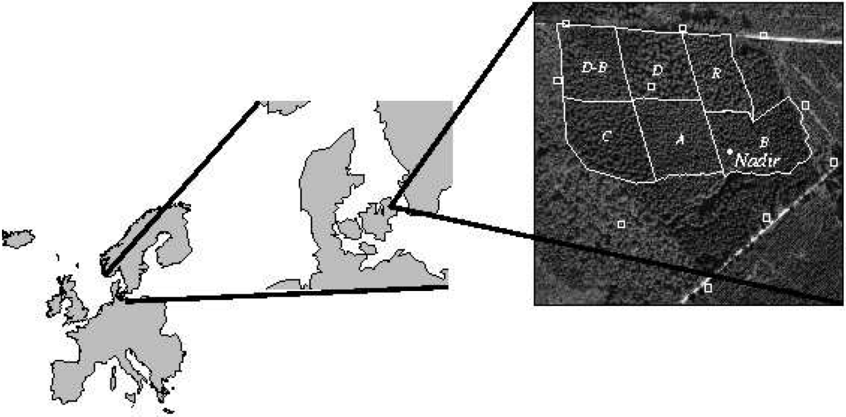

Examples of images: aerial photographs

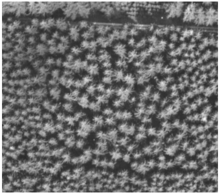

Subplot D with very heavy thinning

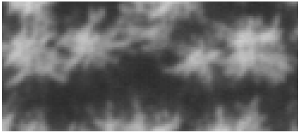

The southeastern corner of subplot D

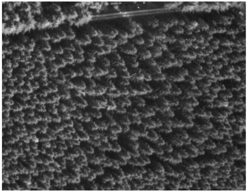

Subplot D from another angle (the airplane located northwest from the experimental area)

Images of plants (carrots)

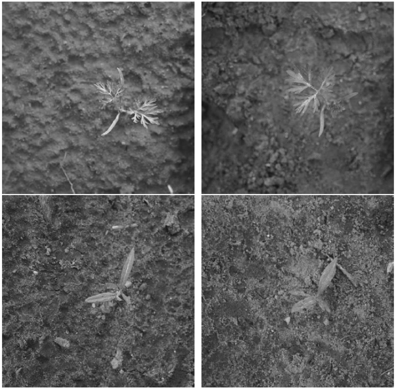

Microscopy images

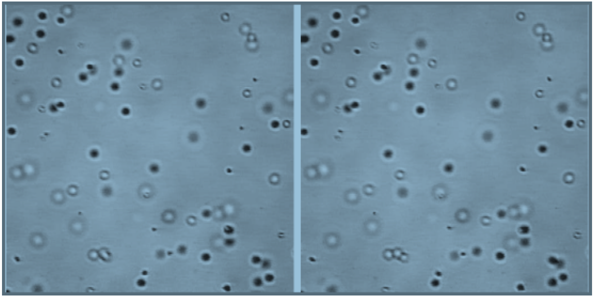

Particles in focus are shown as small distinct black objects, while particles out of focus are extended and shown either white (above the focal plane) or black (below the focal plane).

Point patterns

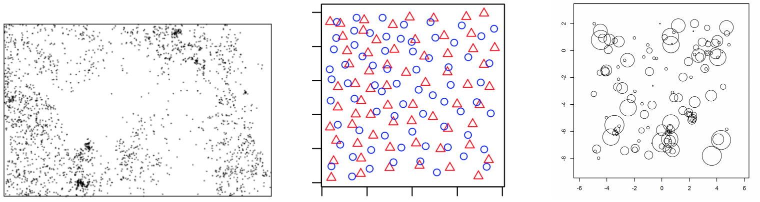

Locations of trees (Beilschmiedia Pendula) in a tropical rain forest (left), locations of on (red) and off (blue) Beta-type retina cells in the retina of a cat (middle), and locations and sizes of Finnish pine saplings (right).

Methods: Kriging and geostatistics

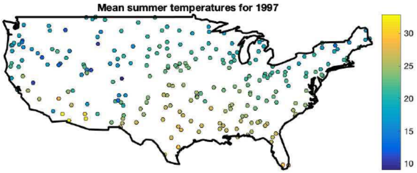

Mean summer temperatures in the US in 1997

- Predict an observation (temperature) at an unobserved location
- Create a temperature map for the whole US

Methods: Segmentation

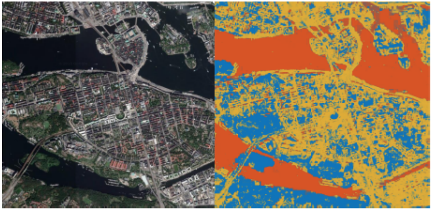

- Segmentation into three types: water (red), park (blue) and buildings/streets (yellow)
- Methods
  - thresholding
  - k-means
  - neural networks

Methods: Pattern recognition/classification

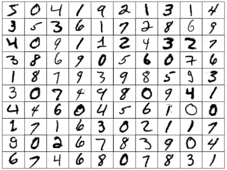

(twistedsifter.com/2016/03/puppy-or-bagel-meme-gallery/)

- Can you recognize the different digits? A puppy or a bagel?
- Methods
  - classification methods
  - clustering methods
  - regression methods

Digital images

A digital image is a matrix of pixels

$$f = (f_{ij}) = (f_{ij}, i = i, ..., m, j = 1, ..., n)$$

$$f_{ij} \in V,$$

where V is, for example,

$$\begin{aligned} V &= \{0, 1\} \text{ (binary image)} \\ V &= \{0, ..., 255\} \text{ (grey level image)} \\ V &= \{0, ..., 255\}^3 \text{ (colour image)} \end{aligned}$$

A pixel (voxel) is specified by a location (i, j) and a pixel value f_{ij}.

Image filtering

Filtering is used to modify or improve image properties, e.g. remove noise, and/or to emphasize information, such as edges, of the images.

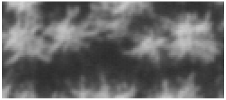

A new image g can be constructed from a given image f by linear filtering,

$$g_{ij} = \sum_{k=-p}^{p} \sum_{l=-p}^{p} w_{k,l} f_{i+k,j+l},$$

where $w = \{w_{k,l}\} = (w_{k,l}, k = -p, -p+1, ..., p, l = -p, ..., p)$ is a matrix of real numbers, a filter.

Simple filters

A 3×3 averaging filter ($p = 1$)

$$w = \begin{bmatrix} w_{-1,-1} & w_{-1,0} & w_{-1,1} \\ w_{0,-1} & w_{0,0} & w_{0,1} \\ w_{1,-1} & w_{1,0} & w_{1,1} \end{bmatrix} = \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

and a median filter

$$g_{ij} = \text{median}\{f_{i+k,j+l} : |k| \leq p, |l| \leq p\}$$

Drawbacks

1. not isotropic (i.e. circularly symmetric), smooths further away from the diagonals than along rows and columns
2. weights have a sudden cut-off rather than decaying gradually to zero which leaves discontinuities in the smoothed image

Gaussian filter

Circular 2D Gaussian filter (smooth)

$$w_{k,l} = c \cdot \exp\left(-\frac{k^2 + l^2}{2\sigma^2}\right),$$

where σ^2 is the variance, and c is chosen such that

$$\sum_{k=-p}^{p} \sum_{l=-p}^{p} w_{k,l} = 1.$$

$$\left(w = \begin{bmatrix} w_{-1,-1} & w_{-1,0} & w_{-1,1} \\ w_{0,-1} & w_{0,0} & w_{0,1} \\ w_{1,-1} & w_{1,0} & w_{1,1} \end{bmatrix} \right)$$

Edge detection

To emphasize edges in an image, some weights in the filter have to be negative. A linear filter with

$$w = \frac{1}{6}\begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$$

tends to emphasize vertical edges and a filter with

$$w = \frac{1}{6}\begin{bmatrix} -1 & -1 & -1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}$$

horisontal edges.

These filters give the difference in pixel values in columns (or rows) on either side of the location in the input image.

The output from the filter is large in magnitude if there is a large difference in pixel values to the left and right (above and below) of the pixel location.

Example: aerial photo

From left to right: Original image, smoothed version after Gaussian filter with $\sigma = 4.5$ pixels (approximately 0.7m), a 3D illustration where light intensity is the vertical coordinate, and the extracted tree locations.

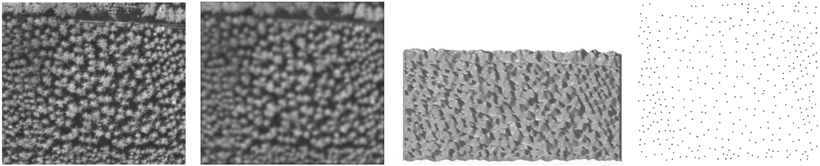

The locations of the trees can be estimated as whiteness maxima in the smoothed image.

Segmentation

- Division of an image into regions or categories which correspond to different objects (or parts of objects). Every pixel in an image is allocated to one of these categories.
- In a good sementation
  - pixels in the same category have a similar greyscale (or multivariate values) and form a connected region
  - neighbouring pixels that are in different categories have dissimilar values
- A good segmentation makes the remaining image analysis steps easier
- Can be applied either to the original images or the filtered images

Segmentation: thresholding

The simplest and most commonly used segmentation method is thresholding, where pixels are allocated to categories according to the range of values in which pixel lies.

For example, for a grey scale image

$$(f_{ij}) = (f_{ij}, i = 1, ...m, j = 1, ..., n)$$

where $f_{ij} \in V$ and V is a set of real numbers, we can define a histogram

$$h_k = \text{card}(\{(i,j) : f_{ij} \in I_k\}), k = 1, ..., K$$

where $\text{card}(A)$ denotes the number of elements in the set A and $\{I_1, ..., I_K\}$ is a set of disjoint intervals with V as the union.

Segmentation: thresholding

If the image consists of two parts, you can use a single threshold t and allocate the pixel f_{ij} in category 1 if

$$f_{ij} \leq t$$

and to category 2 otherwise. The threshold t can be chosen by using the histogram.

Segmentation: weed seed

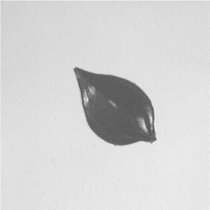

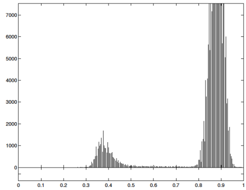

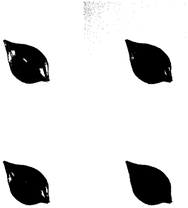

- An image of a weed seed and the corresponding histogram of greyvalues.
- Threshold t should be somewhere between 0.5 and 0.8.

- $t = 0.5$ (upper left)
- $t = 0.8$ (upper right)
- $t = 0.65$ (lower left)
- lower left with wholes filled (lower right)