SINTEF

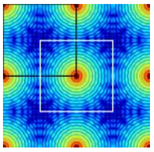

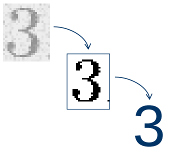

INF4300
Introduction and preliminaries

Asbjørn Berge 25-08-2010

SINTEF

Plan for today

- Practical information
- What will you learn in this course?
- Examples of applications of digital image analysis
- Repetition of key material from INF2310.

SINTEF

Practical information - Lecturers

- **Asbjørn Berge**
  - SINTEF ICT, IFI/UiO wednesdays
  - Telephone: *22067694*
  - Email: *asbjorb@ifi.uio.no*
- **Fritz Albregtsen**
  - IFI/UiO (Fourth floor, IFI building)
  - Telephone: *22852463*
  - Email: fritz@ifi.uio.no
- **NN**

SINTEF

Practical information - Schedule

- Lectures
  - **Fritz Albregtsen and Asbjørn Berge**
  - When: Wednesday 1215-1400.
  - Where: Lille Auditorium, Informatikkbygget
- Exercise lectures
  - **NN**
  - Group 1:
  - When: Friday 1015-1200.
  - Where: C205, Vilhelm Bjerknes hus
  - Group 2:
  - When: Tuesday 1215-1400.
  - Where: C207, Vilhelm Bjerknes hus

SINTEF

Practical information – Web page

- http://www.uio.no/studier/emner/matnat/ifi/INF4300/
  - Information about the course
  - Lecture plan
  - Lecture notes
  - Exercise material
  - Course requisite description
  - Exam information
  - Messages

SINTEF

Practical information – Course material

- All foils will be made available on the course web site.
- The foils define the course requisites.
- Exercises will be introduced as we go along.
- No books defining all course requisites
  - Gonzalez & Woods: Digital Image Processing, 3rd

SINTEF

Practical information – Exercises

- The ordinary exercises are NOT obligatory.
  - Probably a good idea to do them anyway ☺
  - The ordinary exercises can be solved in any programming language, solutions will be provided in Matlab.
- One large exercise (term project)
  - Individual work
  - Counts for 30% percent of the final grade

SINTEF

Practical information - Exam

- Written or oral exam depending on the number of students
  - Final exam counts for 70% of the final grade
  - The term project counts for the remaining 30%
- No written sources of information available at exam
- Follow the web page for updates on term paper and exam

SINTEF

Practical information – Term project

- Sadly, plagiarism and cheating on term papers is very common
- Therefore you are obliged to read the document at the following web address, and attach a copy to your submissions
  - Norwegian: http://www.ifi.uio.no/studinf/skjemaer/erklaring.pdf
  - English: http://www.ifi.uio.no/studinf/skjemaer/declaration.pdf
- However: Using available source code and applications is **perfectly ok** and will be **credited** as long as the origin is cited
- The term project is **individual** work, and the handed in result should clearly be your own work

SINTEF

Practical information - Lecture plan

	23	24	25	26	27	28	29	Introduction and preliminaries	Asbjørn
september	30	31	01	02	03	04	05	Features from images I	Fritz
	06	07	08	09	10	11	12	Features from images II	Fritz
	13	14	15	16	17	18	19	Region and edge based segmentation	Fritz
	20	21	22	23	24	25	26	The Hough transform	Fritz
oktober	27	28	29	30	01	02	03	No lecture	
	04	05	06	07	08	09	10	Object descriptors	Fritz
	11	12	13	14	15	16	17	Learning from data I	Asbjørn
	18	19	20	21	22	23	24	Learning from data II	Asbjørn
	25	26	27	28	29	30	31	Regularization	Asbjørn
november	01	02	03	04	05	06	07	Tracking and flow	Asbjørn
	08	09	10	11	12	13	14	No lecture	
	15	16	17	18	19	20	21	Course summary	Fritz/Asbjørn

SINTEF

What is image analysis?

- **Image analysis** is the art and science whose ultimate goal is to give computers "*vision*"
  - Read handwritten documents
  - Recognize people
  - Find objects
  - Measure the world in 3D
  - Guide robots
- **Image processing** is often used in the more limited sense of *simple image manipulations*:
  - Removing noise
  - Changing contrast
  - Improving edges
  - Joining different data sources

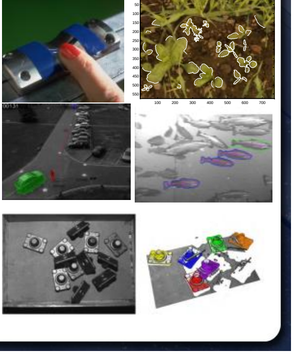

SINTEF

Image analysis:
Industrial research and examples of applications from SINTEF

Asbjørn Berge

Commercial break

- Want to work with state-of-the-art applied image analysis?
  - SINTEF ICT is leading in applied research for industrial vision applications
  - Prototyping and innovative solutions based on OTS hardware
- We are currently looking for Masters- and PhD-students!
- Accelerating image analysis using GPUs
  - Video analysis, background modeling, object tracking
  - 3d object search and recognition

Tom Kavli, Chief Scientist, tka@sintef.no
Asbjørn Berge, Research Scientist, asbe@sintef.no

From pixels to features to class

- Objects often correspond to regions
  We need the spatial relationship between the pixels
- For text recognition: the information is in shape, not grey levels
- Classification: learn features that are common for one type of objects

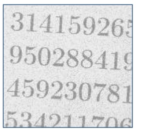

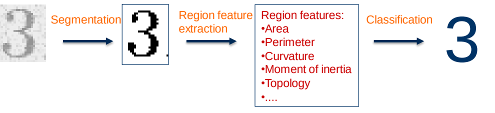

INF 2310 repetition

- See http://www.uio.no/studier/emner/matnat/ifi/INF2310/v09/undervisningsplan.xml
- Topics covered in the course:
  - Image representation, sampling and quantization.
  - Compression and coding
  - Color imaging
  - Grey-level mapping
  - Geometrical operations

  Assumed known

  - Filtering and convolution in the image domain
  - Fourier transform
  - Segmentation by thresholding
  - Edge detection

  Good understanding needed

INF 2310 repetition
2-D convolution

- The result image g(x,y) is given by

$$g(x,y) = \sum_{j=-w_1}^{w_1} \sum_{k=-w_2}^{w_2} h(j,k) f(x-j, y-k)$$

$$= \sum_{j=x-w_1}^{x+w_1} \sum_{k=y-w_2}^{y+w_2} h(x-j, y-k) f(j,k)$$

- h is a m n filter with size $m=2w_1+1$, $n=2w_2+1$
- The result is a weighed sum of the input pixels surrounding pixel (x,y). The weights are given by *h(j,k).*
- The pixel value of the next pixel in the out image is found by moving the filter one position and computing again.

INF 2310 repetition
Separable filters

- Geometrical shapes: rectangles and quadrats
- Rectangular mean filters are separable.

$$h(i,j) = \frac{1}{25}\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix} = \frac{1}{25}\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \end{bmatrix} * \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

- Advantage: fast filter

INF 2310 repetition
Non-uniform lowpass filters

- 2D Gauss-filter:

$$h(x,y) = \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

- Parameter σ is standard deviation (width)
- Filter size must be set using σ

SINTEF

INF 2310 repetition
Digital gradient operators

- The gradient of *f(x)* is

$$\lim_{h\to 0}\frac{f(x+h)-f(x)}{h}$$

- The gradient of an image:

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

- The gradient points in the direction of most rapid change in intensity

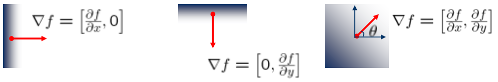

SINTEF

INF 2310 repetition
Gradient operators

- Prewitt-operator

$$H_x(i,j)=\begin{bmatrix}1&0&-1\\1&0&-1\\1&0&-1\end{bmatrix} H_y(i,j)=\begin{bmatrix}-1&-1&-1\\0&0&0\\1&1&1\end{bmatrix}$$

- Sobel-operator

$$H_x(i,j)=\begin{bmatrix}1&0&-1\\2&0&-2\\1&0&-1\end{bmatrix} H_y(i,j)=\begin{bmatrix}-1&-2&-1\\0&0&0\\1&2&1\end{bmatrix}$$

- Frei-Chen-operator

$$H_x(i,j)=\begin{bmatrix}1&0&-1\\\sqrt{2}&0&-\sqrt{2}\\1&0&-1\end{bmatrix} H_y(i,j)=\begin{bmatrix}-1&-\sqrt{2}&-1\\0&0&0\\1&\sqrt{2}&1\end{bmatrix}$$

SINTEF

INF 2310 repetition
G_x, G_y and the gradient

- Horisontal edges:
  - Compute $g_x(x,y)=H_x*f(x,y)$
  - Convolve with the horisontal filter kernel H_x
- Vertical edges:
  - Compute $g_y(x,y)=H_y*f(x,y)$
- Compute the gradient:

The gradient direction is given by:

$$\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} \Big/ \frac{\partial f}{\partial x}\right)$$

- how does this relate to the direction of the edge?

The *edge strength* is given by the gradient magnitude

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

SINTEF

INF 2310 repetition – Edge extraction

- Several basic edge extraction techniques were taught in INF2310
- In this context edges are both edges in intensity, color and texture
- Edges are important for many reasons:
  - Much of the information in an image is contained in the edges. In many cases semantic objects are delineated by edges
  - We know that biological visual systems are highly dependent on edges

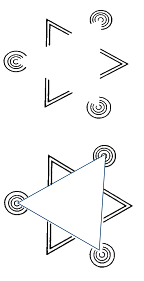

SINTEF

INF 2310 repetition
Edge extraction

- The edge detection operators are operators that produce strong responses in image regions where pixel values (in intensity images) change rapidly.
- In such, they are digital approximations to the gradient operator:

$$\nabla f(x)$$

- …which is a vector quantity given by:

$$\nabla \mathbf{f}(x,y) = \begin{bmatrix} G_x \\ G_y \end{bmatrix} = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix}$$

SINTEF

INF 2310 repetition
Edge extraction

- The gradient is a measure of how the function f(x,y) changes as a function of changes in the arguments x and y.
- The gradient vector points in the direction of maximum change.
- The length of this vector indicates the size of the gradient:

$$\nabla f = |\nabla \mathbf{f}| = \sqrt{G_x^2 + G_y^2}$$

INF 2310 repetition
Edge extraction

- The standard operator is the so called Sobel operator.
- In order to apply Sobel on an image you convolve the two x- and y-direction masks with the image:

-1	-2	-1
0	0	0
1	2	1

-1	0	1
-2	0	2
-1	0	1

INF 2310 repetition
Edge extraction

- This will give you two images, one representing the horizontal components of the gradient, one representing the vertical component of the gradient
- Thus using Sobel you can derive both the local gradient magnitude and direction

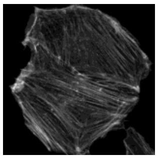
Grayscale image


Horizontal edges

INF 2310 repetition
Edge extraction

- Another frequently used technique for edge detection is based on the use of discrete approximations to the *second derivative*
- The *Laplace operator* is given by

$$\nabla^2(f(x,y)) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

- This operator changes sign where f(x,y) has an inflection point, it is equal to zero at the exact edge position

INF 2310 repetition
Edge extraction

- Approximating second derivatives on images as finite differences gives the following mask

$$\nabla^2(f(x,y)) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$
$$\approx -f(i-1,j) + 2f(i,j) - f(i+1-j)$$
$$-f(i,j-1) + 2f(i,j) - f(i,j+1)$$

0	-1	0
-1	4	-1
0	-1	0

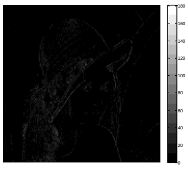

INF 2310 repetition
Edge extraction

- Since this operator is based on second derivatives it is extremely sensitive to noise.
- To counter this it is often combined with Gaussian prefiltering in order to reduce noise.
- This gives rise to the so called Laplacian-of-Gaussian (LoG) operator.

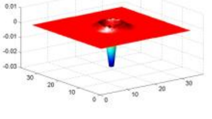

INF 2310 repetition
Sinusoids in images

$$f(x,y) = 128 + A\cos\left(\frac{2\pi(ux+vy)}{N} + \phi\right)$$

A - amplitude
u - horisontal frequency
v - vertical frekvency
ϕ - phase

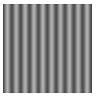
A=50, u=10, v=0

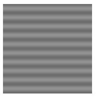
A=20, u=0, v=10

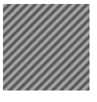
A=50, u=10, v=10

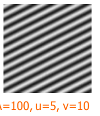
A=100, u=5, v=10

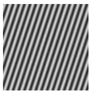
A=100, u=15, v=5

Note: u and v are the number of cycles (horisontally and vertically) in the image

SINTEF

INF 2310 repetition
2-D Discrete Fourier transform (DFT)

f(x,y) is a pixel in a *N*×*M* image

Definition: $$F(u,v)=\frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}f(x,y)e^{-j2\pi(ux/M+vy/N)}$$

$$e^{j\theta}=\cos\theta+j\sin\theta$$

This can also be written:

$$F(u,v)=\frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}f(x,y)\left[\cos(2\pi(ux/M+vy/N))-j\sin(2\pi(ux/M+vy/N))\right]$$

Inverse transform:

$$f(x,y)=\sum_{u=0}^{M-1}\sum_{v=0}^{N-1}F(u,v)e^{j2\pi(ux/M+vy/N)}$$

SINTEF

INF 2310 repetition
Example – oriented structure

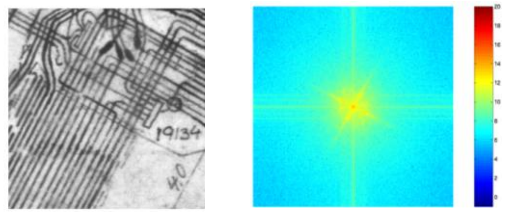

SINTEF

INF 2310 repetition
The convolution theorem

$$f(x,y)*h(x,y)\Leftrightarrow F(u,v)\cdot H(u,v)$$

Convolution in the image domain
⇔
Multiplication in the frequency domain

$$f(x,y)\cdot h(x,y)\Leftrightarrow F(u,v)*H(u,v)$$

Multiplication in the image domain
⇔
Convolution in the frequency domain

SINTEF

INF 2310 repetition
How do we filter out this effect?

SINTEF

INF 2310 repetition
Filtering in the frequency domain

1. Multiply the image by $(-1)^{x+y}$ to center the transform
2. Compute *F(u,v)* using the 2-D DFT
3. Multiply *F(u,v)* by a filter *H(u,v)*
4. Compute the inverse *FFT* of the result from 3
5. Obtain the real part from 4
6. Multiply the result by $(-1)^{x+y}$

SINTEF

INF 2310 repetition
The "ideal" low pass filter

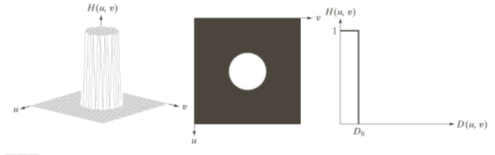

FIGURE 4.40 (a) Perspective plot of an ideal lowpass-filter transfer function. (b) Filter displayed as an image. (c) Filter radial cross section.

INF 2310 repetition
Example - ideal low pass

Original

r_0=0.2

r_0=0.3

Look at these image in high resolution. You should see ringing effects in the two rightmost images.

INF 2310 repetition
What causes the ringing effect?

Ideal lowpass in the image domain

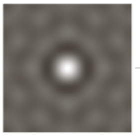

fft of H(u,v) for ideal lowpass

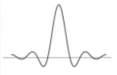

1D profile for ideal lowpass

- Note that the filter profile has negative coefficients
- It has similar profile to a Mexican-hat filter (Laplace-of-Gaussian)
- The radius of the circle and the number of circles per unit is inversely proportional to the cutoff frequency
  - Low cutoff gives large radius in image domain

INF 2310 repetition
Butterworth low pass filter

- Window-functions are used to reduce the ringing effect.
- Butterworth low pass filter of order *n*:

$$H(u,v) = \frac{1}{1 + \left[D(u,v)/D_0\right]^{2n}}$$

- D_0 describes the point where *H(u,v)* has decreased to the half of its maximum)
  - Low filter order (*n* small): *H(u,v)* decreases slowly: Little ringing
  - High filter order (*n* large): *H(u,v)* decreases fast: More ringing
- Other filters can also be used, e.g. Gaussian, Bartlett, Blackman, Hamming, Hanning

INF 2310 repetition
Gaussian lowpass filter

$$H(u,v) = e^{-D^2(u,v)/2\sigma^2}$$

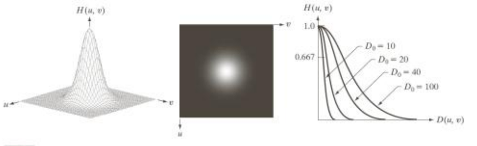

FIGURE 4.47 (a) Perspective plot of a GLPF transfer function. (b) Filter displayed as an image. (c) Filter radial cross sections for various values of D_0.

INF 2310 repetition
High pass filtering

- Simple ("Ideal") high pass filter:

$$H_{hp}(u,v) = \begin{cases} 0, & D(u,v) \le D_0, \\ 1, & D(u,v) > D_0. \end{cases}$$

or

$$H_{hp}(u,v) = 1 - H_{lp}(u,v)$$

- Butterworth high pass filter:

$$H_{hpB}(u,v) = \frac{1}{1 + \left[D_0/D(u,v)\right]^{2n}}$$

- Gaussian high pass filter:

$$H_{hpG}(u,v) = 1 - e^{-D^2(u,v)/2D_0^2}$$

INF 2310 repetition
Ideal, Butterworth and Gaussian highpass

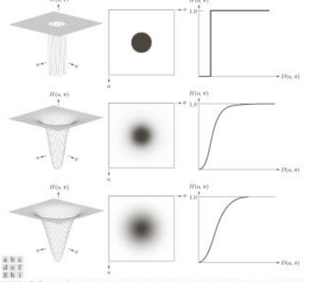

FIGURE 4.52 Top row: Perspective plot, image representation, and cross section of a typical ideal highpass filter. Middle and bottom rows: The same sequence for typical Butterworth and Gaussian highpass filters.

SINTEF

INF 2310 repetition
Example – Butterworth highpass

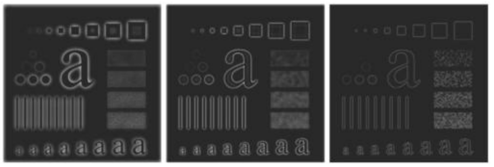

a b c

FIGURE 4.55 Results of highpass filtering the image in Fig. 4.41(a) using a BHPF of order 2 with $D_0 = 30, 60,$ and 160, corresponding to the circles in Fig. 4.41(b). These results are much smoother than those obtained with an IHPF.

SINTEF

INF 2310 repetition
Bandpass and bandstop filters

- Bandpass filter: Keeps only the energy in a given frequency band $<D_{low},D_{high}>$ (or $<D_0\text{-}W/2, D_0+W/2>$)
- W is the width of the band
- D_0 is its radial center.
- Bandstop filter: Removes all energy in a given frequency band $<D_{low},D_{high}>$

SINTEF

INF 2310 repetition
Bandstop/bandreject filters

- Ideal

$$H_{bs}(u,v) = \begin{cases} 1 & \text{if } D(u,v) < D_0 - \frac{W}{2} \\ 0 & \text{if } D_0 - \frac{W}{2} \le D(u,v) \le D_0 + \frac{W}{2} \\ 1 & \text{if } D(u,v) > D_0 + \frac{W}{2} \end{cases}$$

- Butterworth

$$H_{bsB}(u,v) = \frac{1}{1 + \left[\frac{D(u,v)W}{D^2(u,v) - D_0^2}\right]^{2n}}$$

- Gaussian

$$H_{bsG}(u,v) = 1 - e^{-\frac{1}{2}\left[\frac{D^2(u,v) - D_0^2}{D(u,v)W}\right]^2}$$

SINTEF

INF 2310 repetition
An example of bandstop filtering

a b
c d

FIGURE 5.16 (a) Image corrupted by sinusoidal noise. (b) Spectrum of (a). (c) Butterworth bandreject filter (white represents 1). (d) Result of filtering. (Original image courtesy of NASA.)

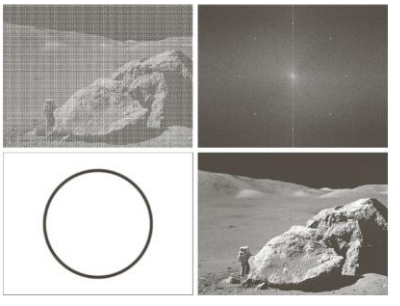

SINTEF

INF 2310 repetition
Bandpass filters

- Are defined by

$$H_{bp}(u,v) = 1 - H_{bs}(u,v)$$

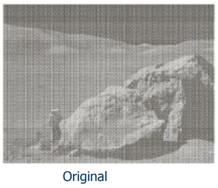

Original

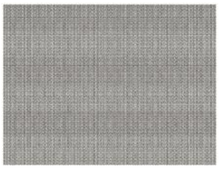

Result after bandpass filtering

SINTEF

INF 2310 repetition
Segmentation and thresholding

- Segmentation
  - Function that labels each pixel in input image with a group label
  - Usually "foreground" and "background"
  - Each group shares some common properties
    - Similar color
    - Similar texture
    - Surrounded by the same edge
- Thresholding
  - One way of segmentation is by defining a threshold on pixel intensity

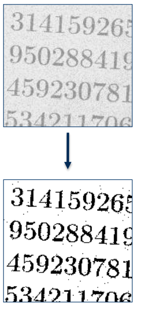

SINTEF

INF 2310 repetition
Segmentation and thresholding

Remember, regions that have semantic importance do not always have any particular visual distinction.

SINTEF

INF 2310 repetition
Segmentation and thresholding

- The only segmentation method taught in INF2310 was thresholding.
- Thresholding is a transformation of the input image f to an output (segmented) image g as follows:

$$g(i,j) = \begin{cases} 1, & f(i,j) \geq T \\ 0, & f(i,j) < T \end{cases}$$

SINTEF

INF 2310 repetition
Segmentation and thresholding

Original image

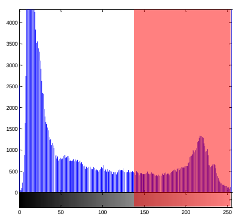

Intensity histogram

Image thresholded at T=140

SINTEF

INF 2310 repetition
Segmentation and thresholding

- Many variants of the basic definition:

$$g(i,j) = \begin{cases} 1, & f(i,j) \in D \\ 0, & otherwise \end{cases}$$

SINTEF

INF 2310 repetition
Segmentation and thresholding

- Many variants of the basic definition:

$$g(i,j) = \begin{cases} 1, & f(i,j) \in D_1 \\ 2, & f(i,j) \in D_2 \\ \ldots & \\ n, & f(i,j) \in D_n \\ 0, & otherwise \end{cases}$$

SINTEF

INF 2310 repetition
Segmentation and thresholding

- Many variants of the basic definition (semithresholding):

$$g(i,j) = \begin{cases} f(i,j), & f(i,j) \geq T \\ 0, & otherwise \end{cases}$$

SINTEF

INF 2310 repetition
Segmentation and thresholding

- This seemingly simple method must be considered with some care:
  - How do you select the threshold, manually or automatically?
  - Do you set a threshold that is global or local (on a sliding window or blockwise)?
  - Purely local method, no contextual considerations are taken
- Automatic threshold selection will be covered later
  - Otsu's method
  - Ridler-Calvard's method
- Local thresholding methods will also be covered
  - Local applications of Otsu and Ridler-Calvard
  - Niblack's method

SINTEF

INF 2310 repetition
Segmentation and thresholding

- Remember that you normally make an error performing a segmentation using thresholding:

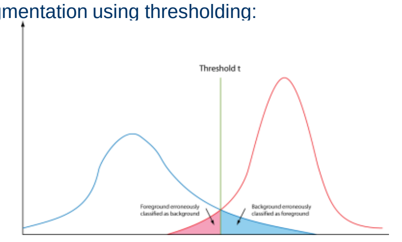

SINTEF

INF 2310 repetition
Segmentation and thresholding

- Assume that the histogram is the sum of two distributions b(z) and f(z), b and f are the normalized background and foreground distributions respectively, and z is the gray level.
- Let B and F be the prior probabilities for the background and foreground (B+F=1).
- In this case the histogram can be written p(z)=Bb(z)+Ff(z).

SINTEF

INF 2310 repetition
Segmentation and thresholding

- In this case the probabilities of erroneously classifying a pixel, given a threshold t, is given by:

$$E_B(t) = \int_{-\infty}^{t} f(z)dz$$

$$E_F(t) = \int_{t}^{\infty} b(z)dz$$

SINTEF

INF 2310 repetition
Segmentation and thresholding

- The total error will be:

$$E(t) = F\int_{-\infty}^{t} f(z)dz + B\int_{t}^{\infty} b(z)dz$$

- Using Leibnitz's rule for derivation of integrals and by setting the derivative equal to zero you can find the optimal value for t:

$$\frac{E(t)}{dt} = 0 \Rightarrow Ff(T) = Bb(T)$$

SINTEF

INF 2310 repetition
Segmentation and thresholding

$$\frac{E(t)}{dt} = 0 \Rightarrow Ff(T) = Bb(T)$$

- This is a general solution that does not depend on the type of distribution.
- Remember that in the case of f and b being Gaussian distributions, it is possible to solve the above equation explicitly.

SINTEF

INF 2310 repetition
Segmentation and thresholding

- In INF2310 we briefly introduced two methods (Ridler-Calvard and Otsu) for determining the thresholds automatically
- These and other methods will be covered in much more detail in the INF4300 lectures