



Image Analysis

Tim B. Dyrby
Rasmus R. Paulsen
DTU Compute




rapa@dtu.dk

<http://www.compute.dtu.dk/courses/02503>

Plenty of slides adapted from Thomas Moeslunds lectures

Lecture 4

Neighbourhood Processing (I) and Morphology (II)






Part I: What can you do after today?

- Describe the difference between point processing and neighbourhood processing
- Compute a rank filtered image using the min, max, median, and difference rank filters
- Compute a mean filtered image
- Decide if median or average filtering should be used for noise removal
- Choose the appropriate image border handling based on a given input image
- Implement and apply template matching
- Compute the normalised cross correlation and explain why it should be used
- Apply given image filter kernels to images
- Use edge filters on images
- Describe finite difference approximation of image gradients including the magnitude and the direction
- Compute the magnitude of the gradient
- Compute edge detection in images

Point processing



- The value of the output pixel is only dependent on the value of one input pixel
- A global operation – changes all pixels

Point processing

- Grey level enhancement
 - Process one pixel at a time independent of all other pixels
 - For example, used to correct Brightness and Contrast

Correct



Too high
brightness



Too low
brightness



Too high
contrast




Too low
contrast





Neighbourhood processing



- Several pixels in the input has an effect on the output



Use of filtering



Noise removal



Enhance edges



Smoothing

- Image processing
- Typically done before actual image analysis

Salt and pepper noise



- Pixel values that are very different from their neighbours
- Very bright or very dark spots
- Scratches in X-rays

What is that?

Salt and pepper noise



- Fake example
 - Let us take a closer look at noise pixels

169	169	173	170	170	172	171	171	169
172	173	172	172	169	171	168	171	170
168	171	169	168	0	169	170	169	255
173	175	170	172	173	168	170	169	171
169	175	170	172	170	255	169	255	169
173	172	255	171	170	172	169	169	170
176	175	172	173	172	171	169	168	173
173	172	169	168	166	0	170	165	166
170	172	172	170	169	169	169	168	172
174	172	172	166	167	168	168	170	172

They are all 0 or 255

Should we just remove all the 0's and 255's from the image?

What is so special about noise?

169	169	173	170	170	172	171	171	169
172	173	172	172	169	171	168	171	170
168	171	169	168	0	169	170	169	255
173	175	170	172	173	168	170	169	171
169	175	170	172	170	255	169	255	169
173	172	255	171	170	172	169	169	170
176	175	172	173	172	171	169	168	173
173	172	169	168	166	0	170	165	166
170	172	172	170	169	169	169	168	172
174	172	172	166	167	168	168	170	172

- What is the value of the pixel compared to the neighbours?
- Average of the neighbours
 - 170
- Can we compare to the average?
 - Difficult – should we remove all values bigger than average+1 ?
- It is difficult to detect noise!

172, 169, 171, 168, 0, 169, 172, 173, 168

Noise – go away!



- We cannot tell what pixels are noise
- One solution
 - Set all pixels to the average or mean of the neighbours (and the pixel itself)
- Oh no!
 - Problems!
 - The noise “pollutes” the good pixels

172, 169, 171, 168, 0, 169, 172, 173, 168

169, 168, 0, 170, 172, 173, 170, 172, 170

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

What is the median value of [169, 168, 0, 170, 172, 173, 170, 172, 170]?

170


173

169

171


172

What is the median value of [169, 168, 0, 170, 172, 173, 170, 172, 170]?



Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

What is the median value of [169, 168, 0, 170, 172, 173, 170, 172, 170]?



Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

The median value

- The values are sorted from low to high
- The middle number is picked
 - The median value


169, 168, 0, 170, 172, 173, 170, 172, 170

0, 168, 169, 170, 170, 170, 172, 172, 173

Median

Noise has no influence on the median!

Noise away – the median filter



- All pixels are set to the median of its neighbourhood (including the pixel itself)
- Noise pixels do not pollute good pixels

172, 169, 171, 168, 0, 169, 172, 173, 168

169, 168, 0, 170, 172, 173, 170, 172, 170

Noise removal – average filter



Scanned X-ray with salt and pepper noise



Average filter (3x3)

Noise removal – median filter



Scanned X-ray with salt and pepper noise



Median filter (3x3)

Image Filtering

169	169	173	170	170	172	171	171	169
172	173	172	172	169	171	168	171	170
168	171	169	168	0	169	170	169	255
173	175	170	172	173	168	170	169	171
169	175	170	172	170	255	169	255	169
173	172	255	171	170	172	169	169	170
176	175	172	173	172	171	169	168	173
173	172	169	168	166	0	170	165	166
170	172	172	170	169	169	169	168	172
174	172	172	166	167	168	168	170	172

- Creates a new *filtered* image
- Output pixel is computed based on a neighbourhood in the input image
- 3 x 3 neighbourhood
 - Filter size 3 x 3
 - Kernel size 3 x 3
 - Mask size 3 x 3
- Larger filters often used
 - Size
 - 7 x 7
 - Number of elements
 - 49

Median filter on image

The image is filtered with a 3 x 3 median filter. What is the result in the pixel marked with a circle?

66	222	102	230	199	147	166	175
204	148	19	241	99	15	187	47
110	140	61	125	62	60	165	94
232	37	31	125	103	90	115	160
46	218	47	86	25	209	139	199
67	159	61	230	34	4	76	21
37	89	106	94	240	11	190	237
35	131	13	28	244	43	48	198

25

90

198

86

103

Median filter on image

The image is filtered with a 3 x 3 median filter. What is the result in the pixel marked with a circle?

66	222	102	230	199	147	166	175
204	148	19	241	99	15	187	47
110	140	61	125	62	60	165	94
232	37	31	125	103	90	115	160
46	218	47	86	25	209	139	199
67	159	61	230	34	4	76	21
37	89	106	94	240	11	190	237
35	131	13	28	244	43	48	198

25

0

90 ✓

36

198

0

86

1

103

3

Median filter on image

The image is filtered with a 3 x 3 median filter. What is the result in the pixel marked with a circle?

66	222	102	230	199	147	166	175
204	148	19	241	99	15	187	47
110	140	61	125	62	60	165	94
232	37	31	125	103	90	115	160
46	218	47	86	25	209	139	199
67	159	61	230	34	4	76	21
37	89	106	94	240	11	190	237
35	131	13	28	244	43	48	198

25

0

90 ✓

36

198

0

86

1


103

3

Rank filters

169	169	173	170	170	172	171	171	169
172	173	172	172	169	171	168	171	170
168	171	169	168	0	169	170	169	255
173	175	170	172	173	168	170	169	171
169	175	170	172	170	255	169	255	169
173	172	255	171	170	172	169	169	170
176	175	172	173	172	171	169	168	173
173	172	169	168	166	0	170	165	166
170	172	172	170	169	169	169	168	172
174	172	172	166	167	168	168	170	172

- Based on sorting the pixel values in the neighbouring region as the median filter
- **Minimum** rank filter
 - Darker image. Noise problems.
- **Maximum** rank filter
 - Lighter image. Noise problems.
- **Difference** filter
 - Enhances changes (edges)



Rank filters on image

The image is filtered with a 3 x 3 median filter (medI). The image (the original) is also filtered with a 5 x 5 minimum rank filter (minI). The final image is made by subtracting minI from medI. What is the result in the marked pixel?

67	159	61	230	34	4	76	21
37	89	106	94	240	11	190	237
35	131	13	28	244	43	48	198
222	102	230	199	147	166	175	124
148	19	241	99	15	187	47	111
140	61	125	62	60	165	94	114
37	31	125	103	90	115	160	78
218	47	86	25	209	139	199	130

3

84

112

73

202

Rank filters on image

The image is filtered with a 3 x 3 median filter (medI). The image (the original) is also filtered with a 5 x 5 minimum rank filter (minI). The final image is made by subtracting minI from medI. What is the result in the marked pixel?

67	159	61	230	34	4	76	21
37	89	106	94	240	11	190	237
35	131	13	28	244	43	48	198
222	102	230	199	147	166	175	124
148	19	241	99	15	187	47	111
140	61	125	62	60	165	94	114
37	31	125	103	90	115	160	78
218	47	86	25	209	139	199	130

3

0

84 ✓

0

112

0

73

0

202

0

Rank filters on image

The image is filtered with a 3 x 3 median filter (medI). The image (the original) is also filtered with a 5 x 5 minimum rank filter (minI). The final image is made by subtracting minI from medI. What is the result in the marked pixel?

67	159	61	230	34	4	76	21
37	89	106	94	240	11	190	237
35	131	13	28	244	43	48	198
222	102	230	199	147	166	175	124
148	19	241	99	15	187	47	111
140	61	125	62	60	165	94	114
37	31	125	103	90	115	160	78
218	47	86	25	209	139	199	130

3

0

84 ✓

0

112

0


73

0


202

0

Correlation



Low




High

- What is it?
- Two measurements
 - Low correlation
 - High correlation
- High correlation means that there is a *relation* between the values
- They *look* the same
- Correlation is a *measure of similarity*

Why do we need similarity?

- Image analysis is also about recognition of patterns
- Often an example pattern is used
 - Often with some kind of meta data to apply to the targets
- We need something to tell us if there is a high match between our pattern and a part of the image




Find
matches




Example pattern
With meta information



Correlation (1D)




Correlation (1D)





Normalisation

- The sum of the kernel elements is used
- Keep the values in the same range as the input image



Normalisation

$$h(x) \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array}$$

- Normalisation factor
 - Sum of kernel coefficients

$$\sum_x h(x) = 1 + 2 + 1$$

2D Kernel Normalisation

Normalisation factor:

$$\sum_x \sum_y h(x, y)$$

$$1 + 2 + 1 + 1 + 3 + 1 + 1 + 2 + 1 = 13$$

1	2	1
1	3	1
1	2	1

h



Correlation on images

■ The filter is now 2D

Kernel coefficients

1	1	1
1	1	1
1	1	1

$\frac{1}{9}$

Kernel

Input

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

Output

	$\frac{12}{9}$				





Correlation on images

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

Input

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

Output


	$\frac{12}{9}$	$\frac{11}{9}$			





Correlation on images

The mask is moved row by row



Correlation on image - no kernel normalisation

1	2	1
1	3	1
1	2	1

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

4

7

16

23

25

Correlation on image - no kernel normalisation

1	2	1
1	3	1
1	2	1

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3



Correlation on image - no kernel normalisation

1	2	1
1	3	1
1	2	1

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3



Mathematics of 2D Correlation

$$g(x, y) = f(x, y) \circ h(x, y)$$

Correlation operator

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

f


1	2	1
1	3	1
1	2	1

h



Mathematics of 2D Correlation

$$g(x,y) = \sum_{j=-R}^R \sum_{i=-R}^R h(i,j) \cdot f(x+i,y+j)$$





Mathematics of 2D Correlation

$$g(x,y) = \sum_{j=-R}^R \sum_{i=-R}^R h(i,j) \cdot f(x+i,y+j)$$

$$g(x,y) = 1 \cdot 2 + 2 \cdot 0 + 1 \cdot 1 + 1 \cdot 1 + 3 \cdot 4 + 1 \cdot 2 + 1 \cdot 0 + 2 \cdot 1 + 1 \cdot 0$$


1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

f

1	2	1
1	3	1
1	2	1

h





$$g(x, y) = \sum_{j=-R}^R \sum_{i=-R}^R h(i, j) \cdot f(x + i, y + j)$$

$$g(2,2) = h(-1,-1) \cdot f(1,1) + h(0,-1) \cdot f(2,1) + h(1,-1) \cdot f(3,1) + h(-1,0) \cdot f(1,2) + h(0,0) \cdot f(2,2) + h(1,0) \cdot f(3,2) + h(-1,1) \cdot f(1,3) + h(0,1) \cdot f(2,3) + h(1,1) \cdot f(3,3)$$

Correlation is a dot-product between two vectors

$$g(x, y) = \sum_{j=-R}^R \sum_{i=-R}^R h(i, j) \cdot f(x + i, y + j)$$

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

1	2	1
1	3	1
1	2	1


h

- In correlation we multiply h and f
- Kernel h and 3x3 patch in figure f are two high dimensional vectors:

$$f(2,1) = [2, 0, 1, 1, 4, 2, 0, 1, 0]$$

$$h = [1, 2, 1, 1, 3, 1, 1, 2, 1]$$

- Multiplication is a dot product
 - $h \cdot f = \|h\| \|f\| \cos \theta$



Template match on image

A template match is done on the image to the left with the template seen to the right. To find the best match the correlation is computed. What is the correlation in the marked pixel?

227	208	90	97	145	42	58	27
245	62	212	145	120	154	233	245
140	237	149	19	3	67	39	1
35	89	140	14	86	167	211	198
38	50	234	135	41	176	137	208
66	64	73	199	203	191	254	222
214	157	193	238	79	115	20	22
65	121	192	33	135	21	113	102

66	232	37
204	46	35
110	67	222

50122

123001

11233

2550

90454

Template match on image

A template match is done on the image to the left with the template seen to the right. To find the best match the correlation is computed. What is the correlation in the marked pixel?

227	208	90	97	145	42	58	27
245	62	212	145	120	154	233	245
140	237	149	19	3	67	39	1
35	89	140	14	86	167	211	198
38	50	234	135	41	176	137	208
66	64	73	199	203	191	254	222
214	157	193	238	79	115	20	22
65	121	192	33	135	21	113	102

66	232	37
204	46	35
110	67	222

50122

0

123001

1

11233

1

2550

0

90454

20

Template match on image

A template match is done on the image to the left with the template seen to the right. To find the best match the correlation is computed. What is the correlation in the marked pixel?

227	208	90	97	145	42	58	27
245	62	212	145	120	154	233	245
140	237	149	19	3	67	39	1
35	89	140	14	86	167	211	198
38	50	234	135	41	176	137	208
66	64	73	199	203	191	254	222
214	157	193	238	79	115	20	22
65	121	192	33	135	21	113	102

66	232	37
204	46	35
110	67	222

50122

0

123001

1

11233

1

2550

0

90454

20

Smoothing filters

- Also know as
 - Smoothing kernel, Mean filter, Low pass filter, blurring
- Kernel normalisation?
 - Yes, secures the intensity range after filtering
- The simplest filter:
 - *Spatial low pass filter*
 - Removes high frequencies
- Another mask:
 - Gaussian filter

$$\frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$$\frac{1}{16}$$

1	2	1
2	4	2
1	2	1

Why Gaussian?

Use of smoothing




3x3



7x7



11x11



15x15

Use of smoothing



3x3



15x15

- Large kernels smooth more
- Removes high frequency information
- Good at enhancing *big structures*

Border handling

Input

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

Output

	○	○	○	○	
	○	○	○	○	
	○	○	○	○	
	○	○	○	○	


No values at the border!

Border handling – extend the input

Input

0	0	0	0	0	0	
0	1	2	0	1	3	1
0	2	1	4	2	2	2
0	1	0	1	0	1	3
0	1	2	1	0	2	4
	2	5	3	1	2	2
	2	1	3	1	6	3

- Zero padding – what happens?
- Zero is black – creates dark border around the image



Correlation on image with zero padding

1	2	1
1	3	1
1	2	1

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

5, 8

12, 16

13, 3


15, 21

17, 12

Correlation on image with zero padding

1	2	1
1	3	1
1	2	1


1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3



Correlation on image with zero padding

1	2	1
1	3	1
1	2	1

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3




Border handling – extend the input


Input

1	1	2	0	1	3	
1	1	2	0	1	3	1
2	2	1	4	2	2	2
1	1	0	1	0	1	3
1	1	2	1	0	2	4
	2	5	3	1	2	2
	2	1	3	1	6	3

- Reflection
- Normally better than zero padding




What is the connection to deep learning?



<https://se.mathworks.com/videos/introduction-to-deep-learning-what-are-convolutional-neural-networks--1489512765771.html>


Banks of filters

- The part of the network that touches the image consists of a bank of filters i.e., kernel coefficients
 - Organised in a multi-level hierarchy
- The weights of the filters are adapted to the problem



Template Matching


- Template
- Locates objects in images



Input

processing


Template



Output: Correlation image


Template Matching

- The correlation between the template and the input image is computed for each pixel




Template Matching

- The pixel with the highest value is found in the output image
 - Here is the highest correlation



Template Matching

- This corresponds to the found pattern in the input image



Problematic Correlation

- Correlation matching has problem with light areas – why?

$$g(x, y) = \sum_{j=-R}^R \sum_{i=-R}^R h(i, j) \cdot f(x + i, y + j)$$


Very High!



Input (f)



Template (h)



Fake max

Output: Correlation image

Real max

Normalised Cross Correlation

$$\text{NCC}(x, y) = \frac{\text{Correlation}}{\text{Length of image patch} \cdot \text{Length of template}}$$



Input (f)




Template (h)



Output: Correlation image

Real max



Length of template

- Vector length (or magnitude)
 - Put all pixel values into a vector
 - Compute the length of this vector
- Describes the intensity of the template
 - Bright template has a large length
 - Dark template has a small length


$$\text{Length of template} = \sqrt{\sum_{j=-R}^R \sum_{i=-R}^R h(i, j) \cdot h(i, j)}$$



Template (h)

Length of image patch

- Vector length based on pixel values in image patch
- Describes the intensity of the image patch



Input (f) with patch



Template (h)



Normalised Cross Correlation

- The length of the image patch ($\|f\|$) and the length of template ($\|h\|$) normalises the correlation
- Correlation is a dot-product i.e., $\text{correlation} = h \cdot f = \|h\| \|f\| \cos \theta$
- A normalised dot product is the *angle* between vector f and h

$$\text{NCC}(x, y) = \frac{\text{Correlation}}{\text{Length of image patch} \cdot \text{Length of template}}$$

$$\text{NCC}(x, y) = \frac{\|h\| \|f\| \cos \theta}{\|h\| \|f\|}$$

- Normalised dot product: $\text{NCC}(x, y) = \cos \theta$

Normalised Cross Correlation

■ NCC will be between

- 0 : No similarity between template and image patch
- 1 : Template and image patch are identical

$$\text{NCC}(x, y) = \frac{\text{Correlation}}{\text{Length of image patch} \cdot \text{Length of template}} = \cos \theta$$



Input (f)



Template (h)



Output: Correlation image

Real max

Normalised cross correlation on image

11

A template match using normalised cross correlation is performed. What is the resulting value in the marked pixel?

227	208	90	97	145	42	58	27
245	62	212	145	120	154	233	245
140	237	149	19	3	67	39	1
35	89	140	14	86	167	211	198
38	50	234	135	41	176	137	208
66	64	73	199	203	191	254	222
214	157	193	238	79	115	20	22
65	121	192	33	135	21	113	102

66	232	37
204	46	35
110	67	222

0.10

0.33

0.83

0.62

0.98


Normalised cross correlation on image

11

A template match using normalised cross correlation is performed. What is the resulting value in the marked pixel?

227	208	90	97	145	42	58	27
245	62	212	145	120	154	233	245
140	237	149	19	3	67	39	1
35	89	140	14	86	167	211	198
38	50	234	135	41	176	137	208
66	64	73	199	203	191	254	222
214	157	193	238	79	115	20	22
65	121	192	33	135	21	113	102

66	232	37
204	46	35
110	67	222




Normalised cross correlation on image

11


A template match using normalised cross correlation is performed. What is the resulting value in the marked pixel?

227	208	90	97	145	42	58	27
245	62	212	145	120	154	233	245
140	237	149	19	3	67	39	1
35	89	140	14	86	167	211	198
38	50	234	135	41	176	137	208
66	64	73	199	203	191	254	222
214	157	193	238	79	115	20	22
65	121	192	33	135	21	113	102


66	232	37
204	46	35
110	67	222



Edges




Gray level profile




- An edge is where there is a high change in gray level values
- Objects are often separated from the background by edges

Edges



$$f'(d)$$



- The profile as a function $f(d)$
- What value is high when there is an edge?
 - The slope of f
 - The slope of the tangent at d



Finite Difference

■ Definition of slope

$$f'(d) = \lim_{h \rightarrow 0} \frac{f(d+h) - f(d)}{h}$$

■ Approximation


$$f'(d) \approx \frac{f(d+h) - f(d)}{h}$$

■ Simpler approximation

$$h = 1$$


$$f'(d) \approx f(d+1) - f(d)$$

Edges




- Discrete approximation of $f'(d)$
- Central differences is used: $d = +/-1$
- Can be implemented as a filter

$$f'(d) \approx f(d \pm 1) - f(d)$$




Edges in 2D



- Changes in gray level values
 - Image gradient
 - Gradient is the 2D derivative of a 2D function $f(x,y)$
 - Equal to the *slope* of the image
 - A steep slope is equal to an edge

$$\nabla f(x, y) = \vec{G}(g_x, g_y)$$

2D Gradient



$$\text{magnitude} = \sqrt{g_x^2 + g_y^2}$$

Edge filter kernel

-1	0	1
-1	0	1
-1	0	1

Vertical Prewitt filter
(gradient: $g(x)$)

- The Prewitt filter is a typical edge filter
- Output image has high values where there are edges


Prewitt filter (vertical)




Original




Prewitt



Prewitt
Hot colormap




Smooth
15x15






Smooth 15x15
Prewitt

Edge detection



- Edge filter
 - Prewitt for example
- Find magnitude of 2D gradient
- Thresholding
 - Separate edges from non-edges
- Output is binary image
 - Edges are white

Lecture 4b – Morphology



0	0	1	1	1	0	0
0	1	1	1	1	1	0
1	1	1	1	1	1	1
1	1	1	1	1	1	1
1	1	1	1	1	1	1
0	1	1	1	1	1	0
0	0	1	1	1	0	0




Part II: What can you *also* do after today?

- Describe the similarity between filtering and morphology
- Describe a structuring element
- Compute the dilation of a binary image
- Compute the erosion of a binary image
- Compute the opening of a binary image
- Compute the closing of a binary image
- Apply compound morphological operations to binary images
- Describe typical examples where morphology is suitable
- Remove unwanted elements from binary images using morphology
- Choose appropriate structuring elements and morphological operations based on image content

Morphology

- The science of *form*, *shape* and *structure*
- In biology: The form and structure of animals and plants



Common leaf morphologies



Mathematical morphology

- Developed in 1964
- Theoretical work done in Paris
- Used for classification of minerals in cut stone
- Initially used for binary images

Theorem 4.10

$$\begin{cases} \psi_m = \tilde{\varphi} \tilde{\gamma} = \tilde{\gamma} \tilde{\varphi} \tilde{\gamma} = \psi \tilde{\gamma} & , \\ \psi_M = \tilde{\gamma} \tilde{\varphi} = \tilde{\varphi} \tilde{\gamma} \tilde{\varphi} = \psi \tilde{\varphi} & , \\ \psi = \tilde{\gamma} \psi = \tilde{\varphi} \psi, & \\ \tilde{\gamma} \leq \psi_m \leq \psi \leq \psi_M \leq \tilde{\varphi} & . \end{cases}$$

The same theorem may be restated in another way. If $\mathcal{Jd}(\mathcal{B}) \neq \emptyset$ then let B_i be a family of elements of $\mathcal{B}$. We have $\vee B_i \in \sim B$, and thus $\tilde{\gamma}(\vee B_i) = \vee B_i$. From the first relation above, it follows for any $\psi \in \mathcal{Jd}(\mathcal{B})$, that

$$\psi(\vee B_i) = \psi \tilde{\gamma}(\vee B_i) = \tilde{\varphi} \tilde{\gamma}(\vee B_i).$$

But $\tilde{\gamma}(\vee B_i) = \vee B_i$, so that

$$\tilde{\varphi}(\vee B_i) = \psi(\vee B_i) \in \mathcal{B}.$$

In the same way, we also obtain

$$\tilde{\gamma} \tilde{\varphi}(\wedge B_i) = \tilde{\gamma}(\wedge B_i) = \psi(\wedge B_i) \in \mathcal{B}.$$

In other words, $\mathcal{B}$ is a *complete lattice* with respect to the ordering on $\mathcal{B}$ induced by $\leq$, i.e. any family B_i in $\mathcal{B}$ has a smallest upper bound $\tilde{\varphi}(\vee B_i) \in \mathcal{B}$ and a greatest lower bound $\tilde{\gamma}(\wedge B_i) \in \mathcal{B}$.


Conversely, let us assume that $\mathcal{B}$ is a complete lattice. Thus, for any $A \in \mathcal{L}$, the family $\{B : B \in \mathcal{B}, B \geq A\}$ has in $\mathcal{B}$ a greatest lower bound, which is

$$\tilde{\gamma}(\wedge \{B : B \in \mathcal{B}, B \geq A\}) = \tilde{\gamma} \tilde{\varphi}(A) \in \mathcal{B}.$$

But this implies $\mathcal{B}_{\psi_M} \subseteq \mathcal{B}$ for the filter $\psi_M = \tilde{\gamma} \tilde{\varphi}$. Conversely, for any



Do not worry! We use a much less theoretical approach!

Relevance?




- Point wise operations
- Filtering
- Thresholding
 - Gives us objects that are separated by the background
- Morphology
 - Manipulate and enhance binary objects


What can it be used for?



- Remove noise
 - Small objects
 - Fill holes
- Isolate objects
- Customized to specific shapes



How does it work?



- Grayscale image
- Preprocessing
 - Inversion
- Threshold => Binary image
- Morphology

Filtering and morphology

1	2	0	1	3	1
2	1	4	2	2	2
1	0	1	0	1	3
1	2	1	0	2	4
2	5	3	1	2	2
2	1	3	1	6	3

■ Filtering

- Gray level images
- Kernel
- Moves it over the input image
- Creates a new output image

Filtering and morphology

0	1	0
1	1	1
0	1	0

Disk

1	1	1
1	1	1
1	1	1

Box

■ Filtering

- Gray level images
- Kernel
- Moves it over the input image
- Creates a new output image

■ Morphology

- Binary images
- Structuring element (SE)
- Moves the SE over the input image
- Creates a new binary output image



1D Morphology

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element
(SE)

1	1	1
---	---	---



Output Image

			?						
--	--	--	---	--	--	--	--	--	--




1D Morphology : The hit operation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

Structuring Element
(SE)



1	1	1
---	---	---

Output Image

			1						
--	--	--	---	--	--	--	--	--	--

- If just one 1 in the SE match with the input
 - output 1
- else
 - output 0

1D Morphology : The fit operation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

Structuring Element
(SE)

1	1	1
---	---	---


Output Image

			0						
--	--	--	---	--	--	--	--	--	--

- If all 1 in the SE match with the input
 - output 1
- else
 - output 0

1D Morphology : Dilation

- Dilate : To make wider or larger
- Based on the *hit* operation



1D Dilation example

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1								
--	---	--	--	--	--	--	--	--	--

$$g(x) = f(x) \oplus SE$$

to make bigger

Hit

- If just one 1 in the SE match with the input
 - output 1
- else
 - output 0



Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1	0							
--	---	---	--	--	--	--	--	--	--





Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1	0	1						
--	---	---	---	--	--	--	--	--	--





Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1	0	1	1					
--	---	---	---	---	--	--	--	--	--



Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1	0	1	1	1				
--	---	---	---	---	---	--	--	--	--

Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1	0	1	1	1	1			
--	---	---	---	---	---	---	--	--	--

Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	1	0	1	1	1	1	1		
--	---	---	---	---	---	---	---	--	--



Example for Dilation

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

Structuring Element

1	1	1
---	---	---

Output Image


	1	0	1	1	1	1	1	1	
--	---	---	---	---	---	---	---	---	--

The object gets bigger, and holes are filled!



1D Morphology : Erosion

- Erode : To wear down (*Waves eroded the shore*)
- Based on the *fit* operation





Example for Erosion

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element

1	1	1
---	---	---



Output Image

	0								
--	---	--	--	--	--	--	--	--	--

$g(x) = f(x) \ominus SE$


to make smaller

Fit

- If all 1 in the SE match with the input
 - output 1
- else
 - output 0



1D Erosion



01001100


00101000

00001000

00100001

01000100

1D Erosion



01001100

0%

00101000

0%

00001000 ✓

100%


00100001

0%

01000100

0%

1D Erosion



01001100

0%

00101000

0%

00001000 ✓

100%

00100001

0%

01000100

0%



Example for Erosion

Input image

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

Structuring Element

1	1	1
---	---	---

Output Image

	0	0	0	0	1	0	0	0	
--	---	---	---	---	---	---	---	---	--

The object gets smaller



Structuring Element (Kernel)

0	1	0
1	1	1
0	1	0

Disk

1	1	1
1	1	1
1	1	1

Box

- Structuring Elements can have varying sizes
- Usually, element values are 0 or 1, but other values are possible (including none!)
- Structural Elements have an **origin**
- Empty spots in the Structuring Elements are *don't cares*!

		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	1	1	1	1
	1	1	1	1	1	
		1	1	1		

Structuring Element Origin

0	1	0
1	1	1
0	1	0

- The origin is not always the center of the SE

1	1	1
1	1	1
1	1	1

Special structuring elements

- Structuring elements can be customized to a specific problem


0	0	0	1	0	0	0
0	0	1	1	1	0	0
0	1	1	1	1	1	0
1	1	1	1	1	1	1
0	1	1	1	1	1	0
0	0	1	1	1	0	0
0	0	0	1	0	0	0

Diamond

0	0	0	0	0	1	1
0	0	1	1	1	0	0
1	1	0	0	0	0	0



Line

Dilation on images - disk



$$g(x,y) = f(x,y) \oplus SE$$

Dilation on image - box



SE

$$g(x, y) = f(x, y) \oplus SE$$

1 0 1 1



0 1 0 0

0 1 1 1

0 1 0 1

1 1 0 1

Dilation on image - box



SE


$$g(x, y) = f(x, y) \oplus SE$$

1 0 1 1 ✓




100%

0 1 0 0




0%

0 1 1 1




0%

0 1 0 1





0%

1 1 0 1



0%

Dilation on image - box



SE

$$g(x, y) = f(x, y) \oplus SE$$

1 0 1 1 ✓

100%

0 1 0 0

0%

0 1 1 1

0%





0 1 0 1

0%




1 1 0 1

0%

Dilation – the effect of the SE



Dilation Example




- Round structuring element (disk)
- Creates round corners

0	0	1	1	1	0	0
0	1	1	1	1	1	0
1	1	1	1	1	1	1
1	1	1	1	1	1	1
1	1	1	1	1	1	1
0	1	1	1	1	1	0
0	0	1	1	1	0	0

Threshold and Dilation

A threshold of 200 is applied to the image and the result is a binary image. Now a dilation is performed with the structuring element below. How many foreground pixels are there in the resulting image?

145	56	86	42	191
19	33	41	255	115
14	240	203	234	21
135	120	209	167	58
199	3	135	176	116



14

17

6


3

12

Threshold and Dilation

A threshold of 200 is applied to the image and the result is a binary image. Now a dilation is performed with the structuring element below. How many foreground pixels are there in the resulting image?

145	56	86	42	191
19	33	41	255	115
14	240	203	234	21
135	120	209	167	58
199	3	135	176	116



14 ✓

0%

17

0%

6

0%

3

0%


12

0%

Threshold and Dilation

A threshold of 200 is applied to the image and the result is a binary image. Now a dilation is performed with the structuring element below. How many foreground pixels are there in the resulting image?

145	56	86	42	191
19	33	41	255	115
14	240	203	234	21
135	120	209	167	58
199	3	135	176	116



14 ✓

0%

17

0%

6

0%


3

0%

12

0%


Erosion on images - disk



$$g(x,y) = f(x,y) \ominus SE$$

Object is smaller

Erosion on images



0010


1010

0110



0100

1000


Erosion on images



Erosion on images




Erosion example




Counting Coins

- Counting these coins is difficult because they touch each other!
- **Solution:** Threshold and Erosion separates them!
- More on counting later in the course!



Compound operations




- Compound
 - *made of two or more separate parts or elements*
- Combining Erosion and Dilation into more advanced operations
 - Finding the outline
 - Opening
 - Isolate objects and remove small objects (better than Erosion)
 - Closing
 - Fill holes (better than Dilation)

Finding the outline

1. Dilate input image (object gets bigger)
2. Subtract input image from dilated image
3. The outline remains!

$$g(x, y) = (f(x, y) \oplus SE) - f(x, y)$$





Opening

- Motivation: Remove small objects BUT keep original size (and shape)
- Opening = Erosion + Dilation
 - Use the same structuring element!
 - Similar to erosion but less destructive
- Math:

$$g(x, y) = f(x, y) \circ SE = (f(x, y) \ominus SE) \oplus SE$$


- Opening is **idempotent**: Repeated operations have no further effects!

$$f(x, y) \circ SE = (f(x, y) \circ SE) \circ SE$$




Opening


$$g(x,y) = (f(x,y) \ominus SE) \oplus SE$$



Original




Eroded



Dilated

Opening = erosion+dilation





SE




Opening Example

■ 9x3 and 3x9 Structuring Elements



Opening example

- Size of structuring element should fit into the smallest object to keep
- Structuring Element: 11 pixels disc





Closing

- Motivation: Fill holes BUT keep original size (and shape)
- Closing = Dilation + Erosion
 - Use the same structuring element!
 - Similar to Dilation but less destructive
- Math:


$$g(x, y) = f(x, y) \bullet SE = (f(x, y) \oplus SE) \ominus SE$$

- Closing is **idempotent**: Repeated operations have no further effects!


$$f(x, y) \circ SE = (f(x, y) \circ SE) \circ SE$$

Closing


$$g(x,y) = (f(x,y) \oplus SE) \ominus SE$$



Original




Dilated



Eroded


Closing = dilation + erosion



SE

Closing Example




- Closing operation with a 22 pixels disc
- Closes small holes





What did we learn today

Part I: Neighbourhood Processing



Part II: Morphology of binary images

