Contextual Analysis of Image Elements

Daniele Cerra, German Aerospace Center (DLR)

DLR

Marr-Hildrith Edge detector

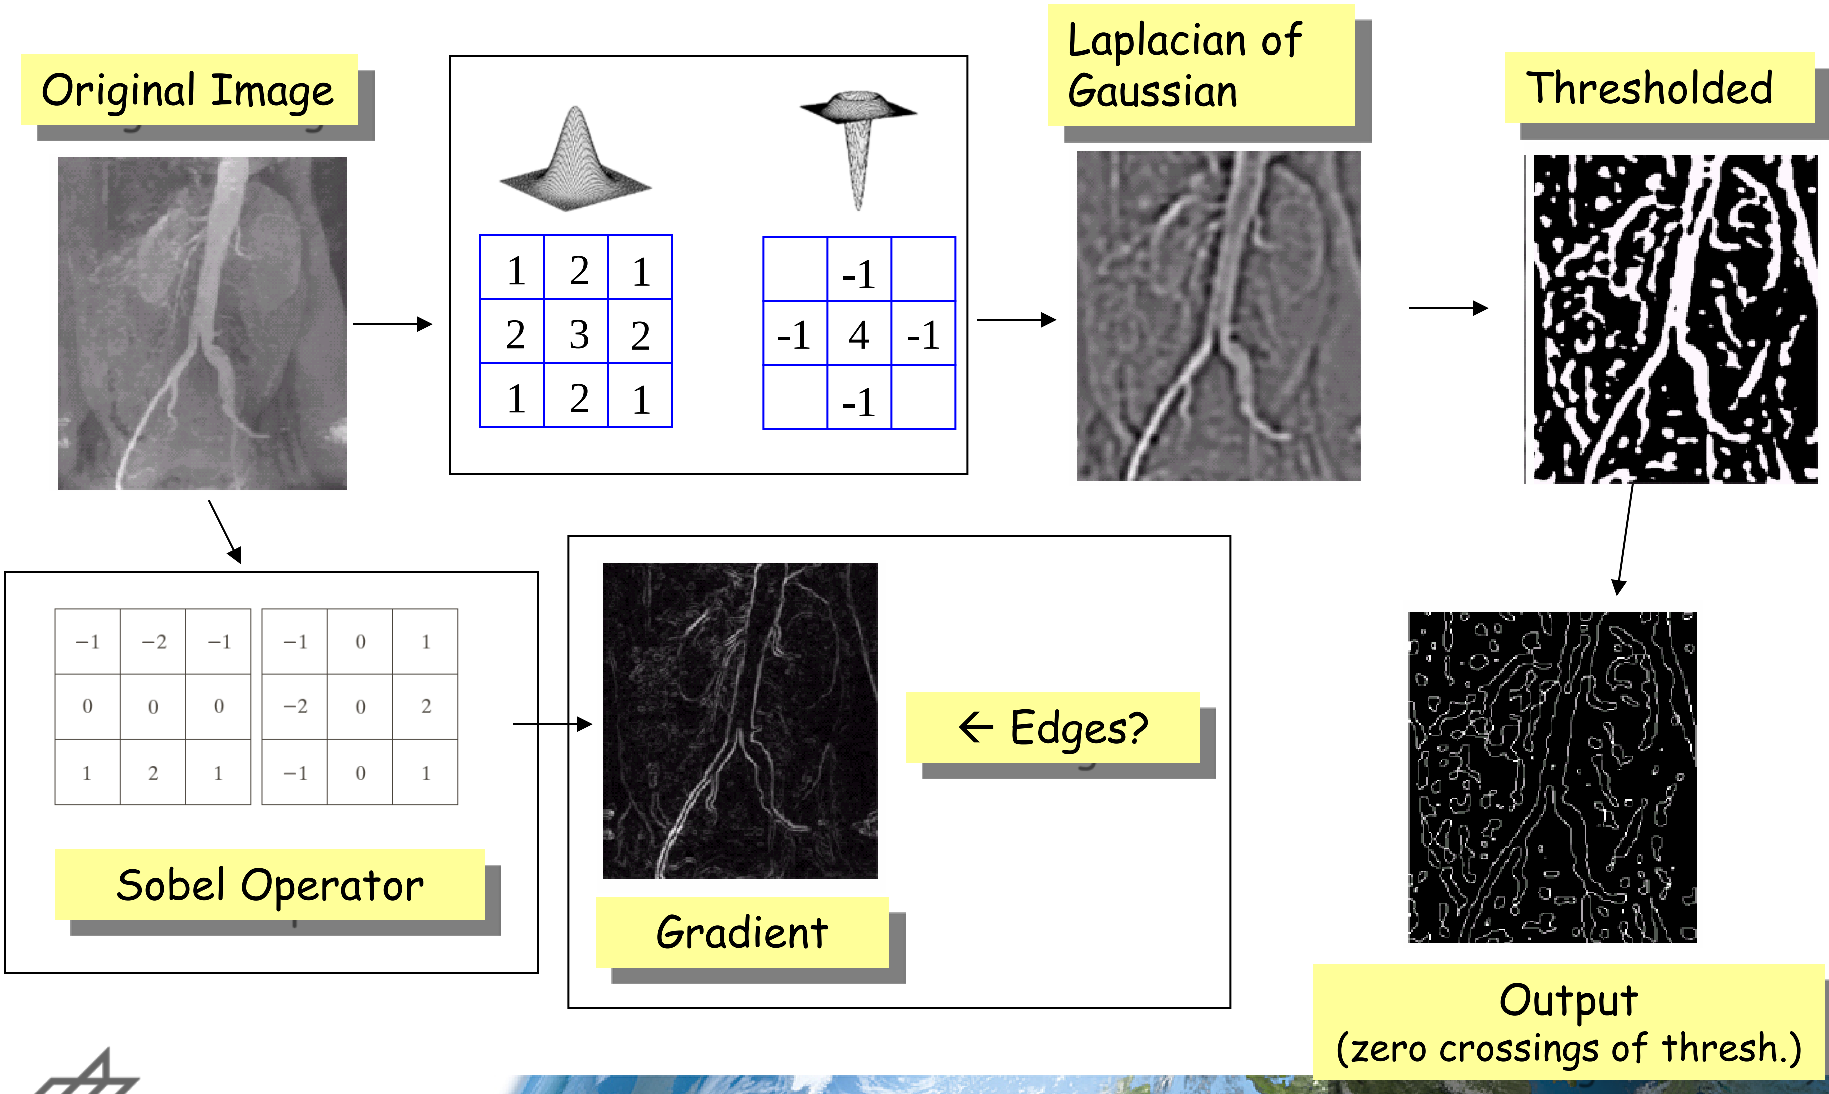

Context Analysis:

Edge Extraction

Canny edge detector

John Canny

- Canny defined three objectives for edge detection:
  1. Low error rate: All edges should be found and there should be no spurious responses.
  2. Edge points should be well localized: The edges located must be as close as possible to the true edges.
  3. Single edge point response: The detector should return only one point for each true edge point. That is, the number of local maxima around the true edge should be minimum.

Canny is the most widely used edge operator today; no one has done better since it came out in the late 80s.

Canny edge detector – first steps

1. Smooth the image with a Gaussian filter with spread σ.
   - Decreases the noise in the image: remember what happens if you apply a high-pass filter to a noisy image??
2. Compute gradient magnitude and direction at each pixel of the smoothed image.
   - Which means: apply a high-pass filter to the image
3. Threshold gradient image
   - Keep only meaningful edges

Canny edge detector

Original image

Magnitude of the gradient

Canny edge detector

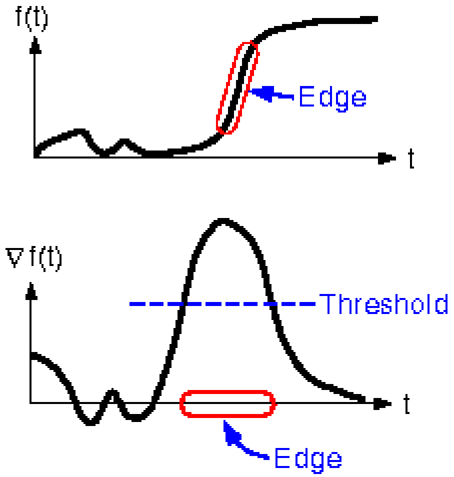

How to turn these thick regions of the gradient into curves?

Magnitude of the gradient

Canny edge detector

- Non-maxima suppression:
  - Check if pixel is local maximum along gradient direction.
  - Select single max across width of the edge.
  - Requires checking interpolated pixels p and r.
  - This operation can be used with any edge operator when thin boundaries are wanted.

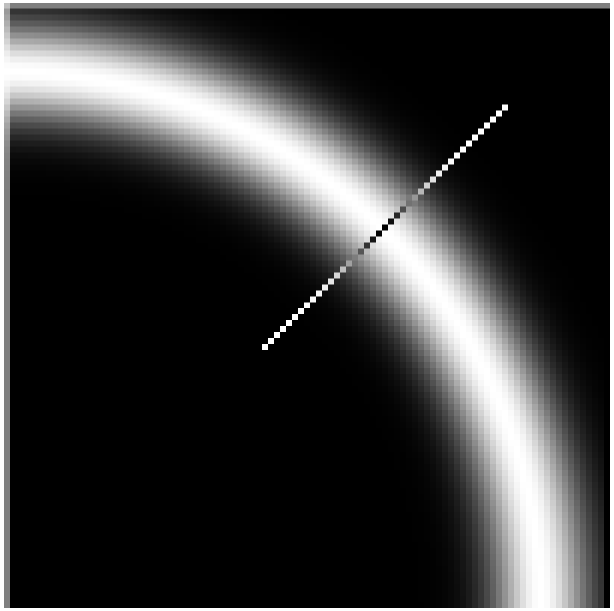

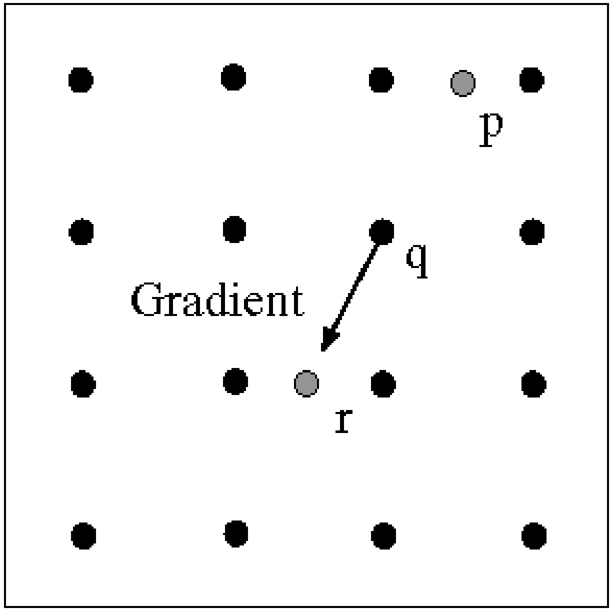

Canny edge detector

Original image

Gradient magnitude

courtesy of G. Loy

Non-maxima suppressed

Canny edge detector

Problem: pixels along this edge did not "survive" the thresholding

Original image

courtesy of G. Loy

Strong edges

Canny edge detector

- Hysteresis thresholding:
  - Use a high threshold Th to start edge curves, and a low threshold Tl to continue them.

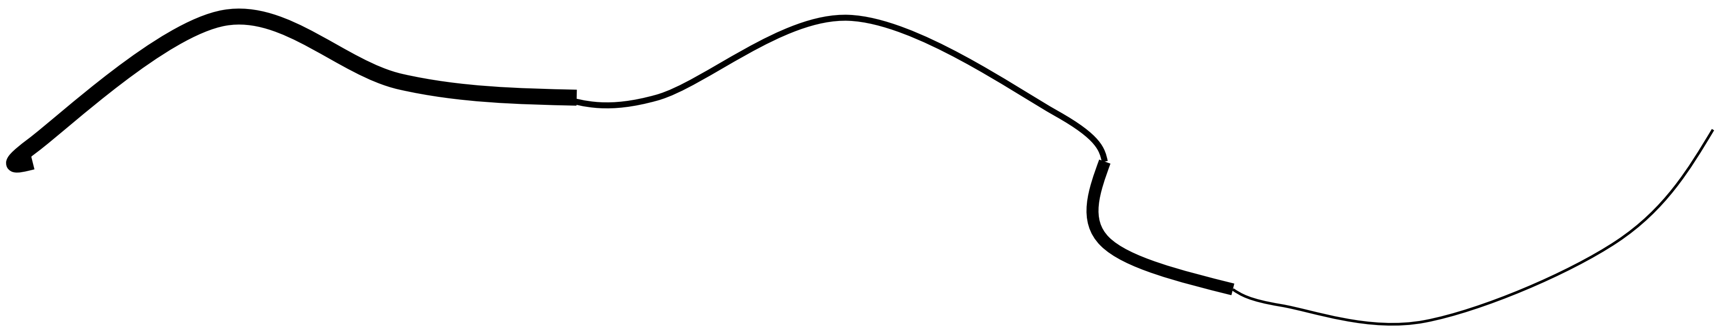

- Select the pixels with value v > Th
- Then collect the pixels with value v > Tl that are connected to selected pixels

Canny Edge Detector: Final Result

gap is gone

Original image

Strong + connected weak edges

Strong Edges (gradient > Thigh)

Weak Edges (gradient > Tlow)

courtesy of G. Loy

Canny Edge detector

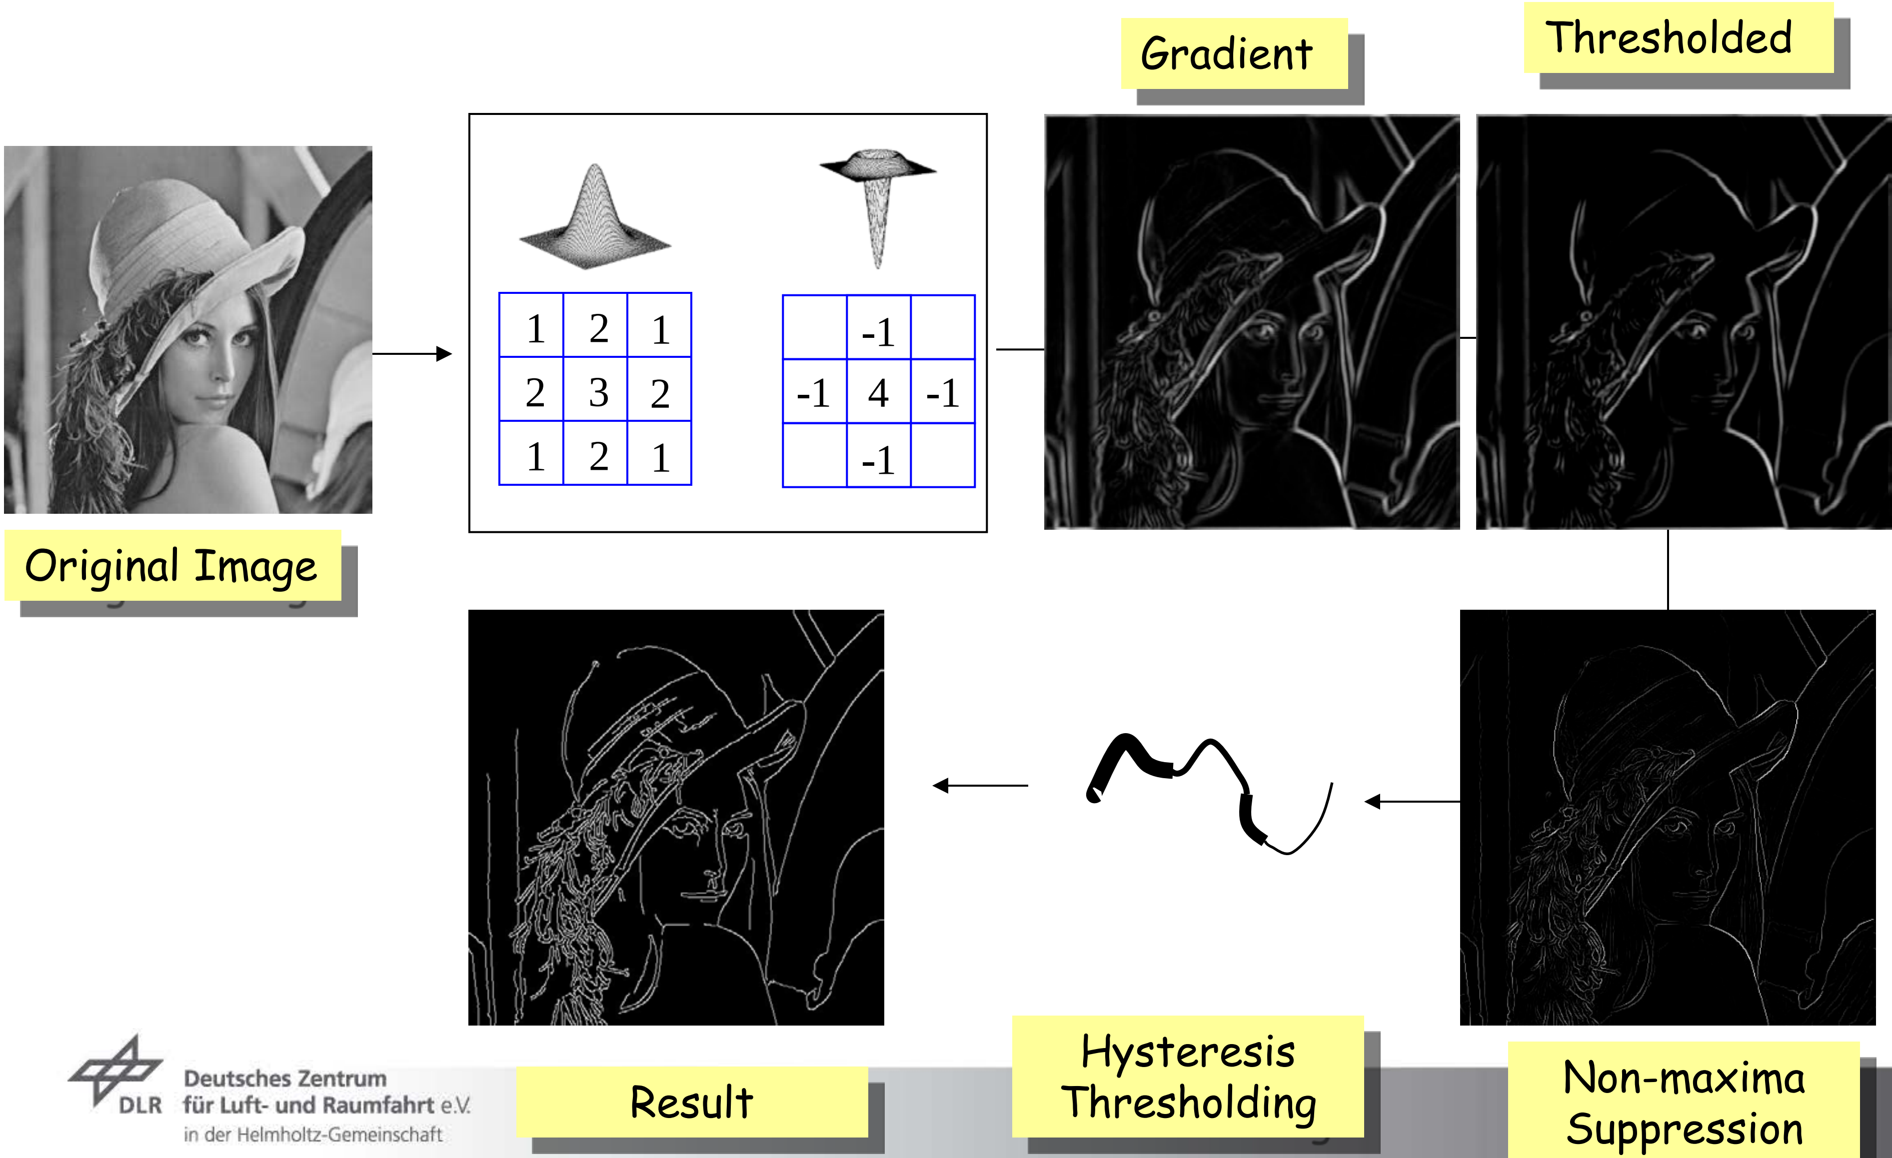

Canny Edge Detector

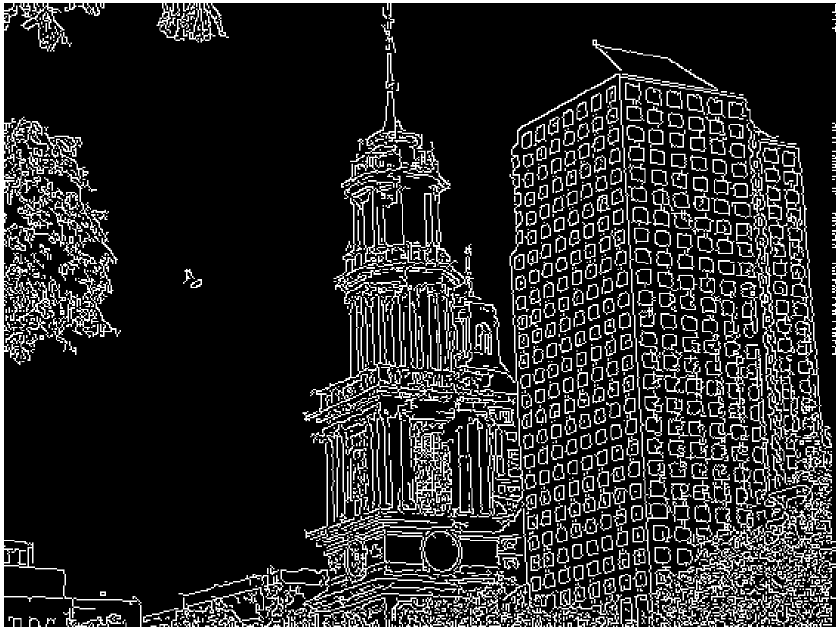

- Canny algorithm is very sensitive to its parameters, which need to be adjusted for different application domains.
    - Smoothing parameter σ
    - Threshold for strong edges
    - Threshold for weak edges

original | Canny with $\sigma = 1$ | Canny with $\sigma = 2$

Canny Results $\sigma = 1.5$ $t_{low} = 0.5$

Canny results: varying 1 threshold

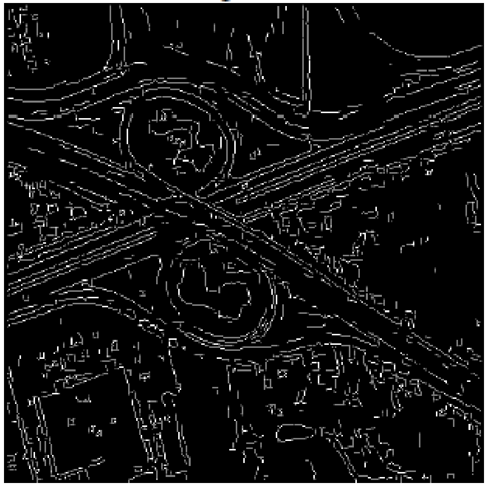

$t_{high} = 0.7$

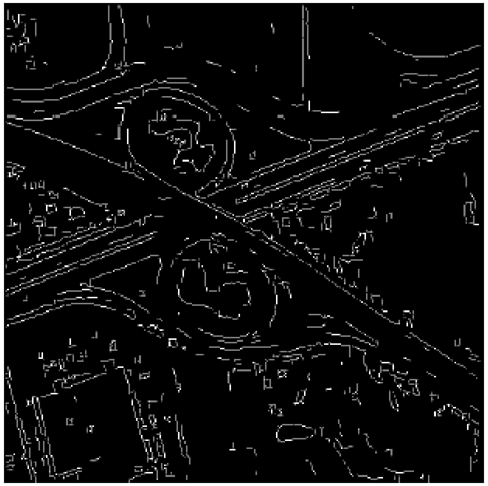

$t_{high} = 0.8$

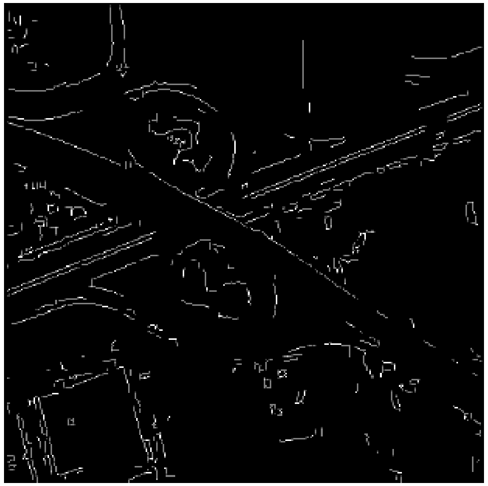

$t_{high} = 0.9$

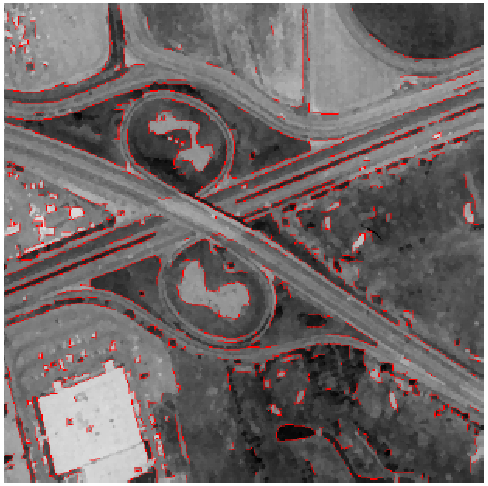

$t_{high} = 0.8$

Try your own combinations at:

http://matlabserver.cs.rug.nl/

Canny Detector

- Quality of results depends on the desired level of detail

original

σ =2.5, T_h = 0.0214, T_l = 5e-5

σ =0.5, T_h = 0.17, T_l = 0.002

Canny Edge Detector in Remote Sensing

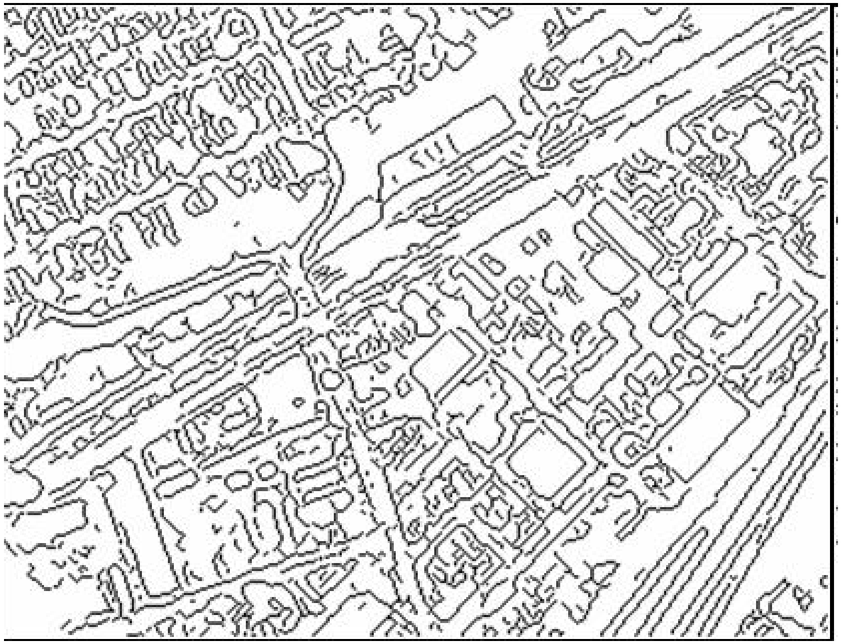

- Image Acquisition & Correction
  - Raw Data→ Raw Image → Image

03	29	38	48
59	96	94	04
05	06	96	97
87	76	75	45

- Low-level Analysis
  - Image → Image
    - Time domain
    - Frequency domain

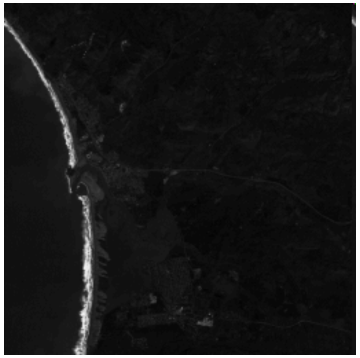

- **Mid-level Analysis**
  - **Image → Features / Attributes**
    - **Feature Extraction**
    - Clustering / Segmentation

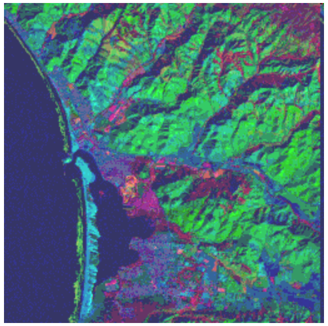

Beach Bar
Wave Breakers
Vegetation1
Vegetation2
Golf Course
Urban Area
Shadows
Sea
Mountains (bright slopes)

- High-level Analysis
  - Features → Recognition

Summary

- Image Acquisition
- Image enhancement
- Sampling & Aliasing
- Image Features
  - Spectral Features
    - We will see more about this in the lectures on hyperspectral remote sensing
  - Features based on relations between pixels
- Image Clustering
- Image Classification

Which features can we extract from an image?

Pixel Value for each band

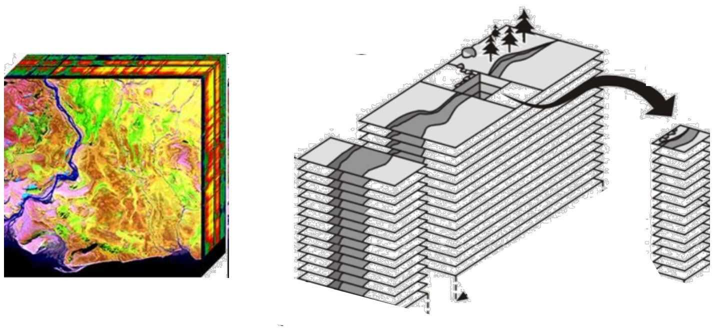

Texture Parameters

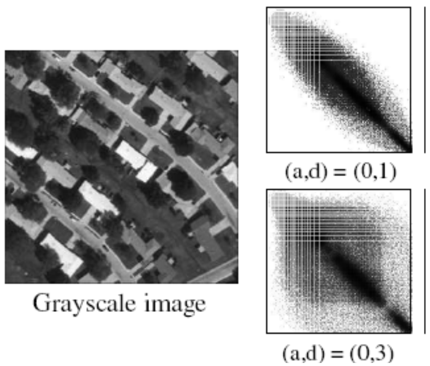

„Interesting“ Points

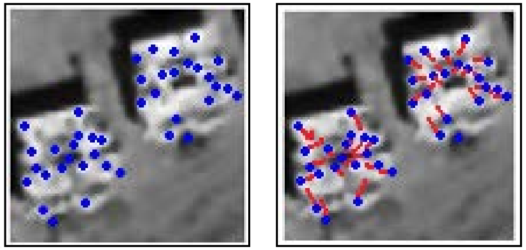

Power Spectrum Features

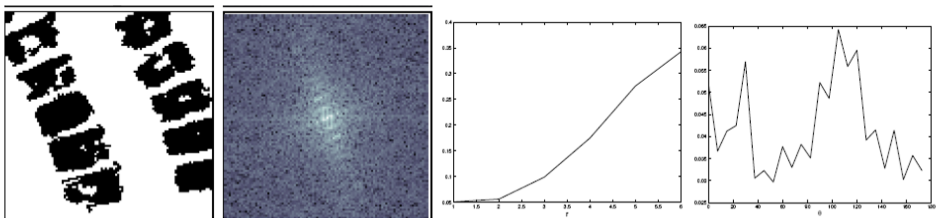

…and more!

Texture

- An important approach to image description is to quantify its texture content.
- Texture gives us information about the spatial arrangement of the colors or intensities in an image.

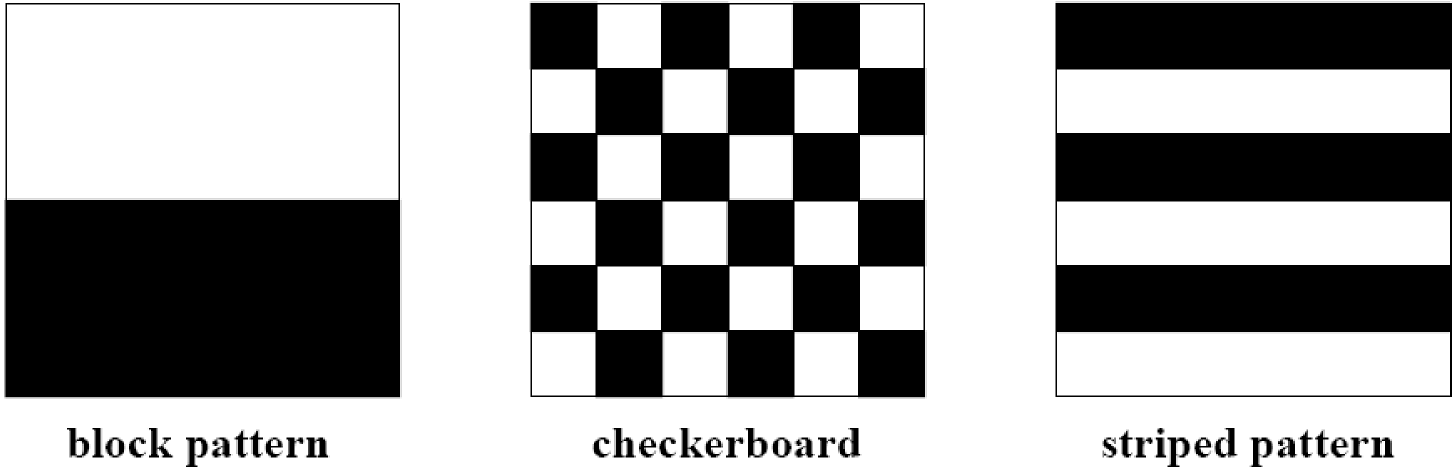

Figure 7.2: Three different textures with the same distribution of black and white.

Texture


Bark

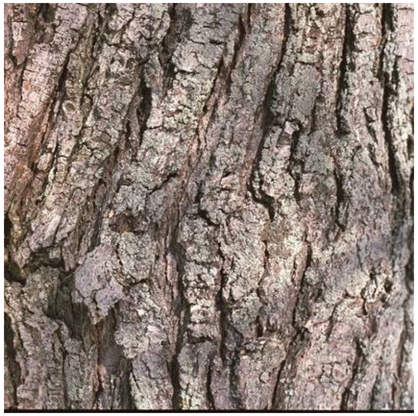
Bark

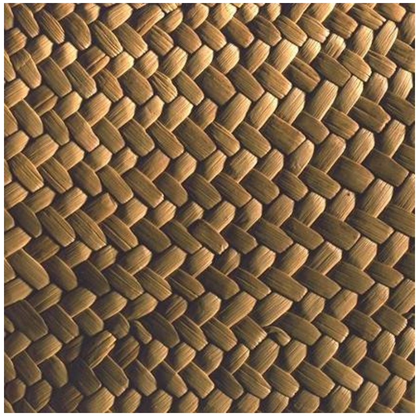
Fabric


Fabric

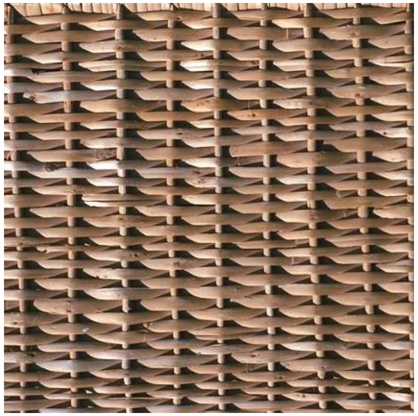
Fabric


Flowers


Flowers


Flowers

Texture Analysis → Local Analysis of Pixels Distribution

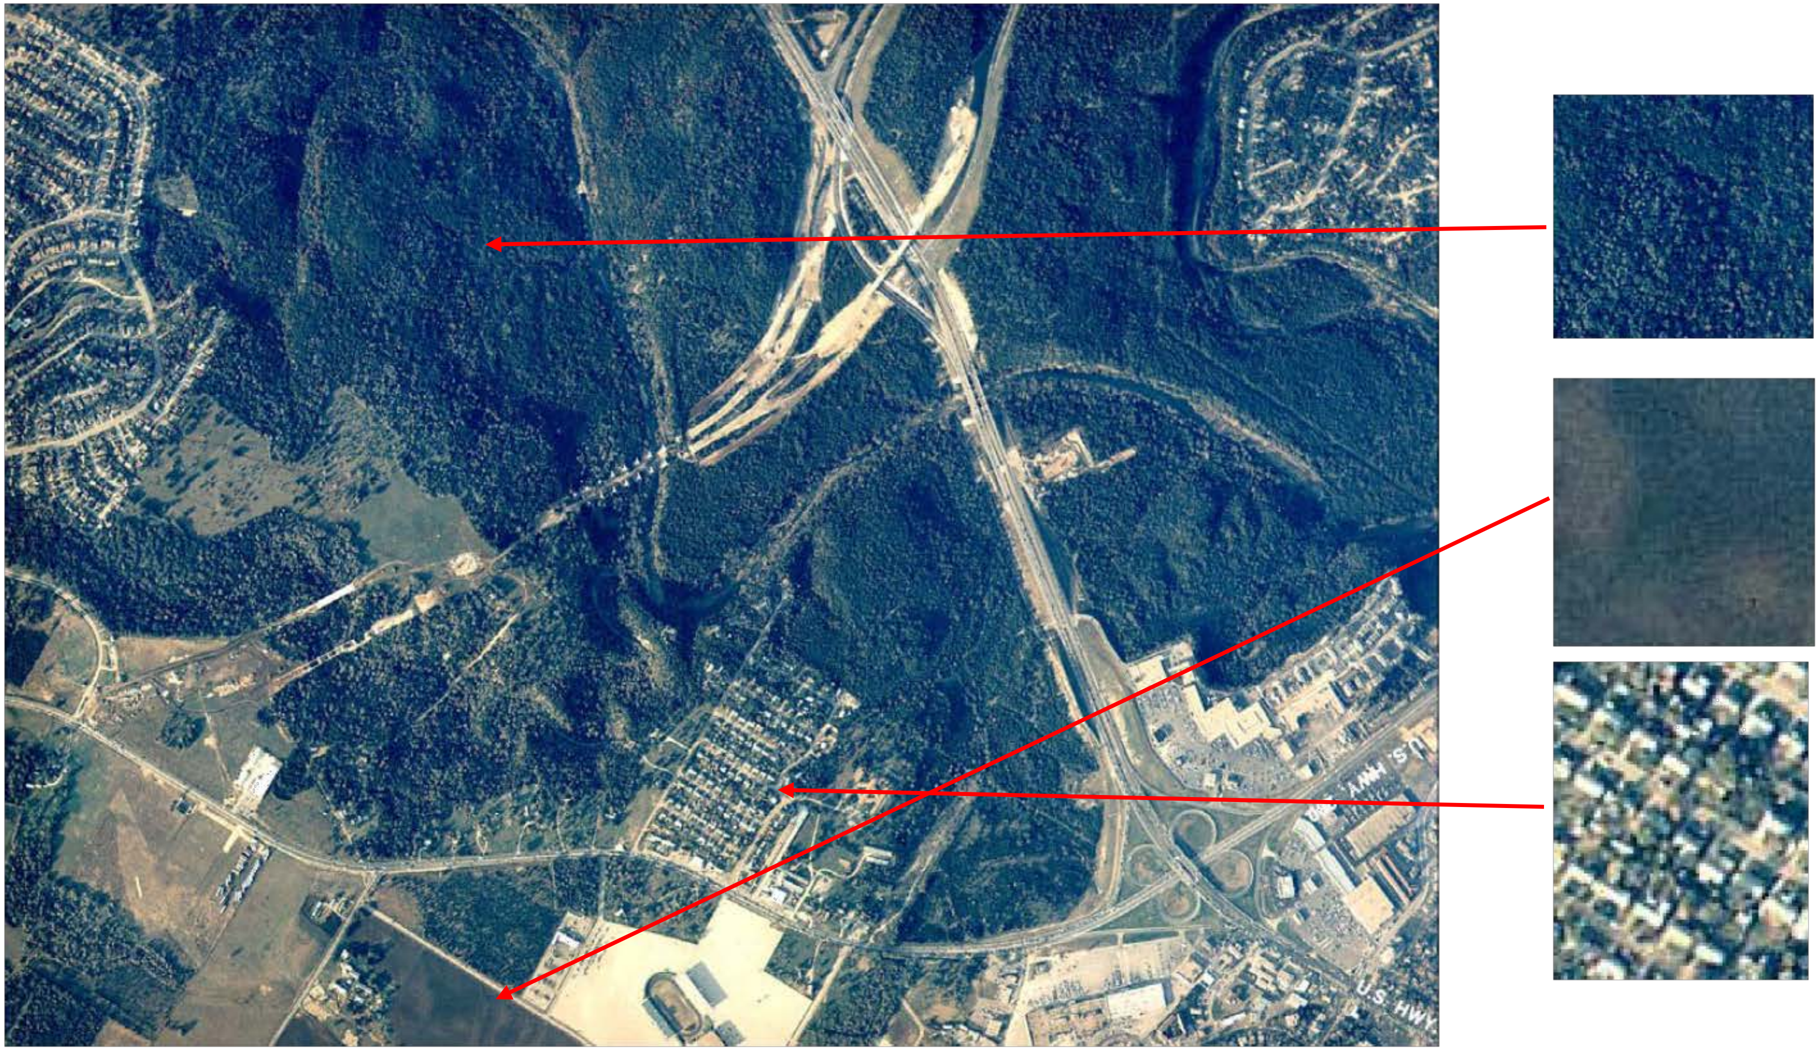

Texture

- Whether an effect is a texture or not depends on the scale at which it is viewed.

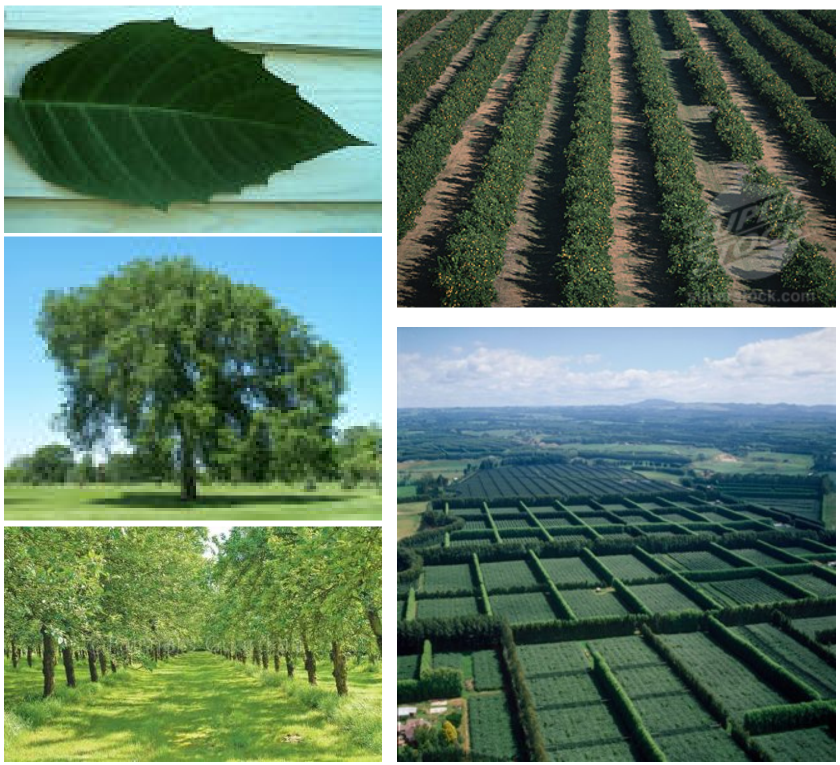

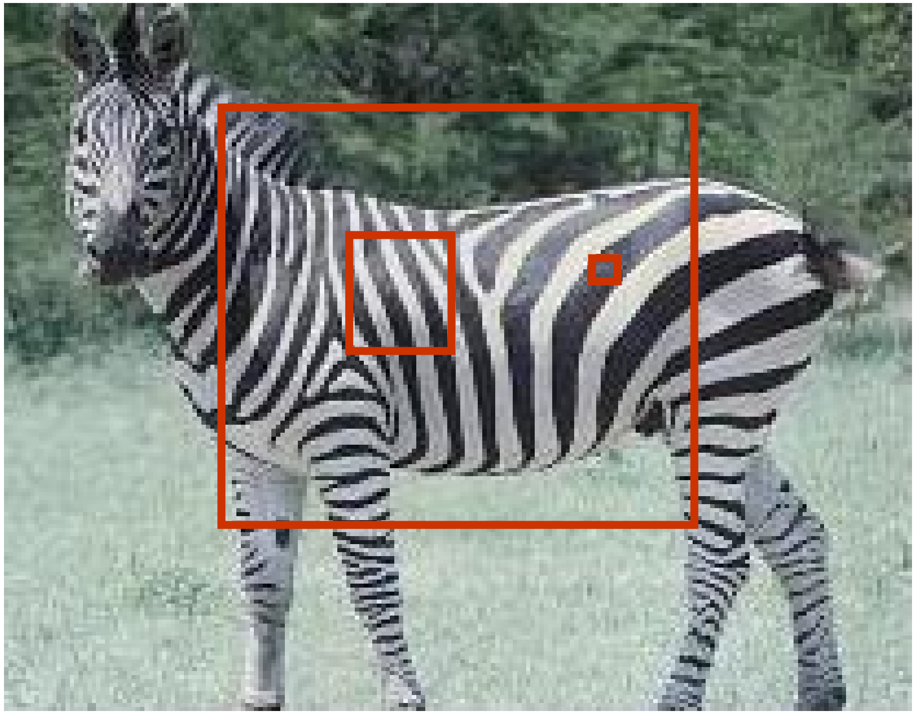

Statistical moments

- The easiest thing we can do is to check the statistics of the histogram of a small window in the image
  - Mean
  - Standard deviation
  - Variance
  - Kurtosis
  - Skewness…

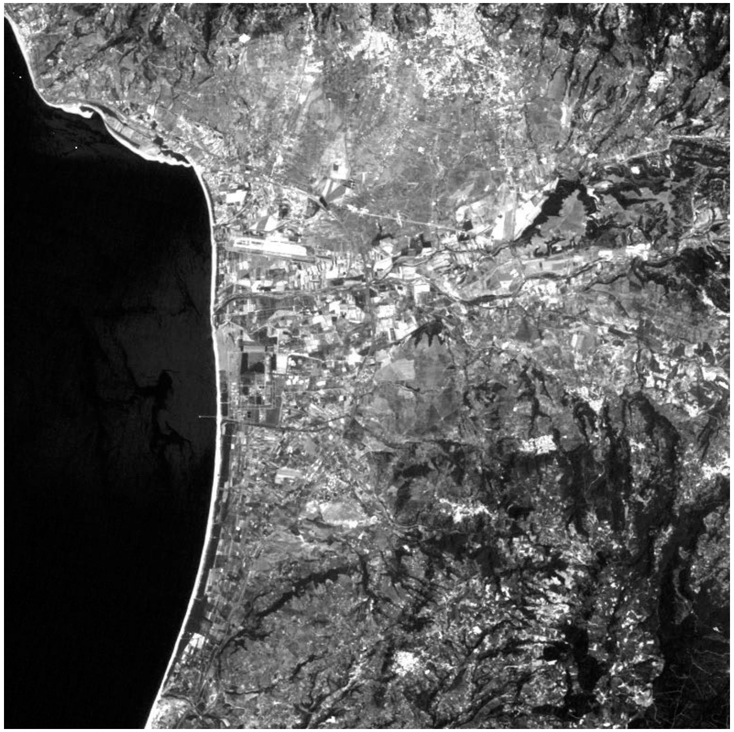

Example

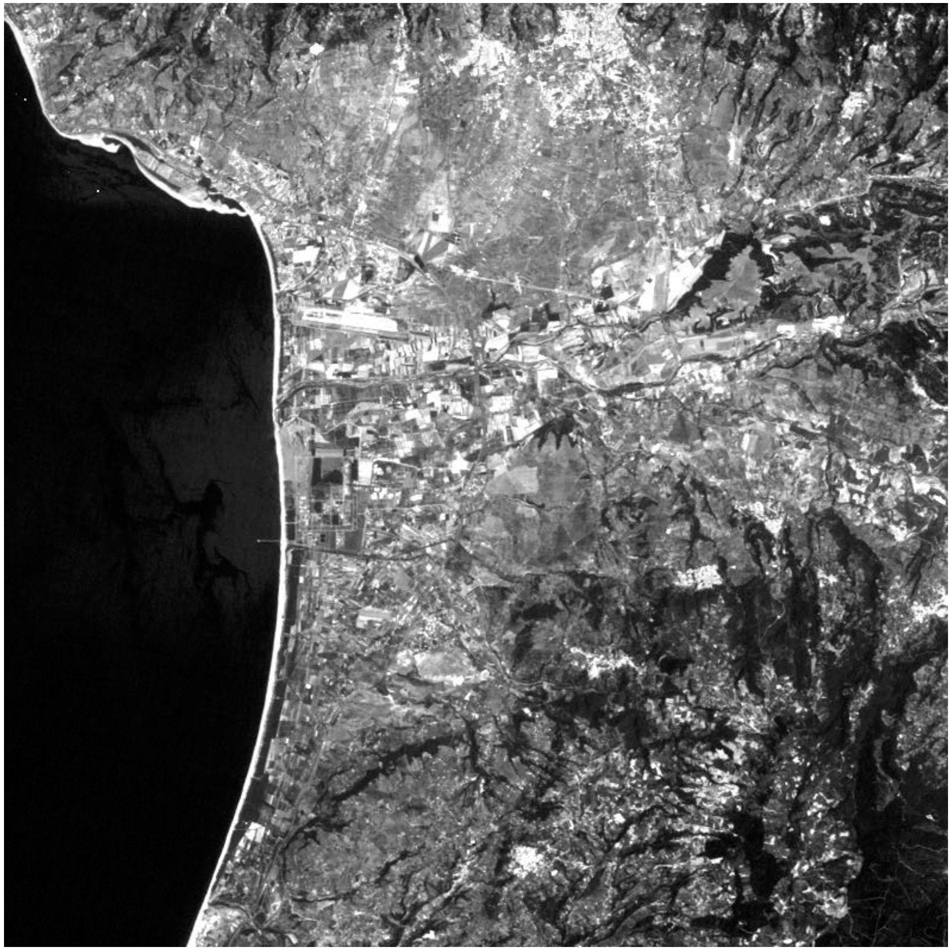

Variance

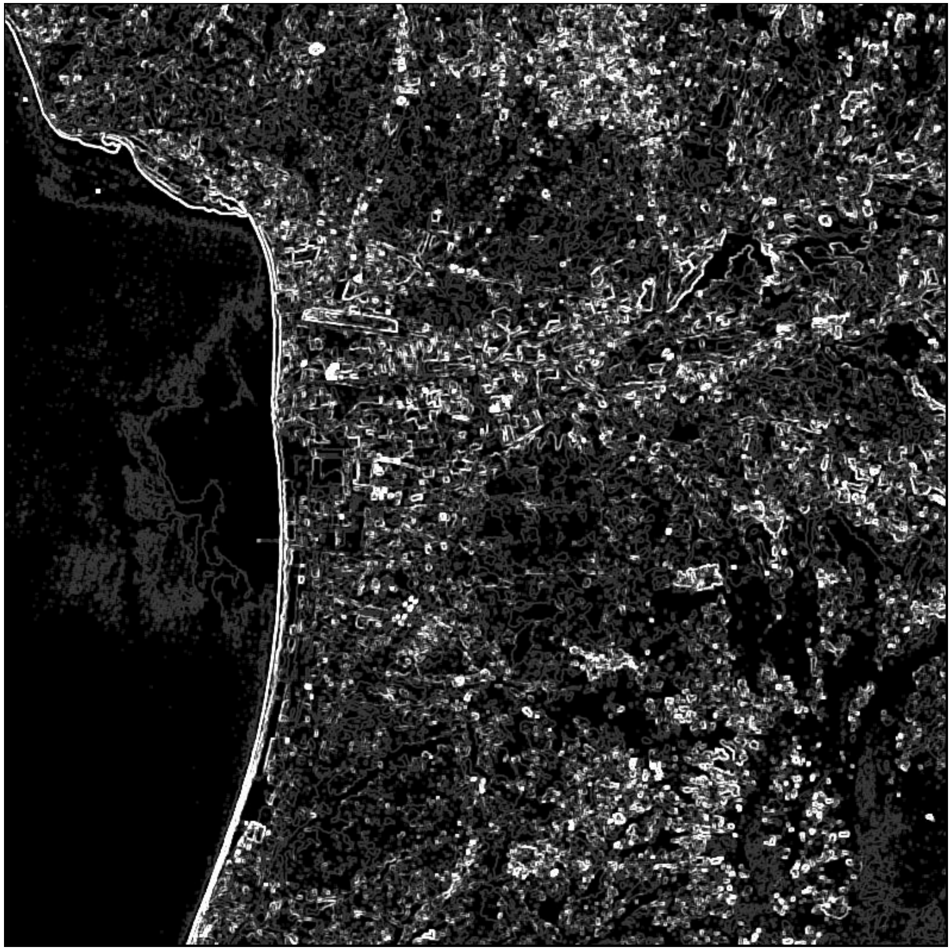

Co-occurrence matrices

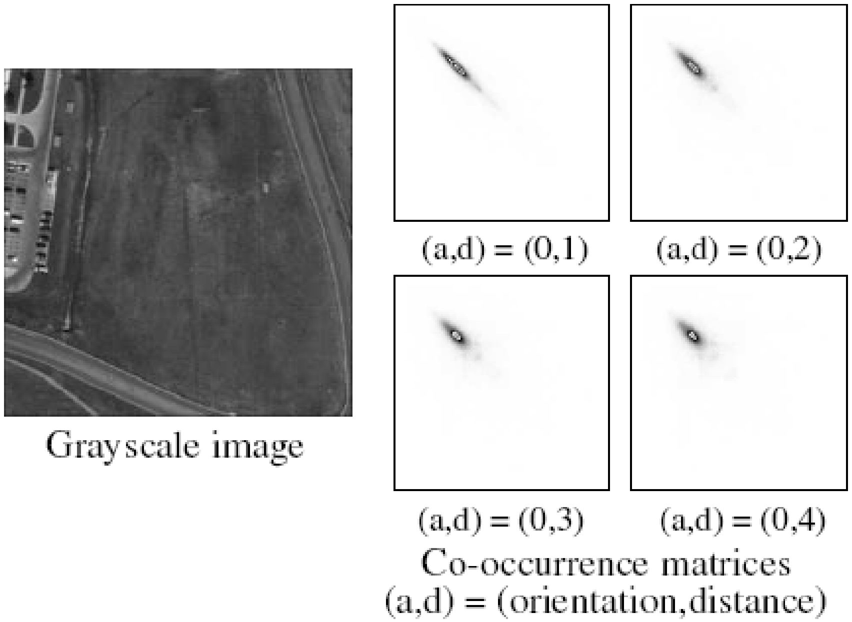

(a) Co-occurrence matrices for an image with a small amount of local spatial variations.

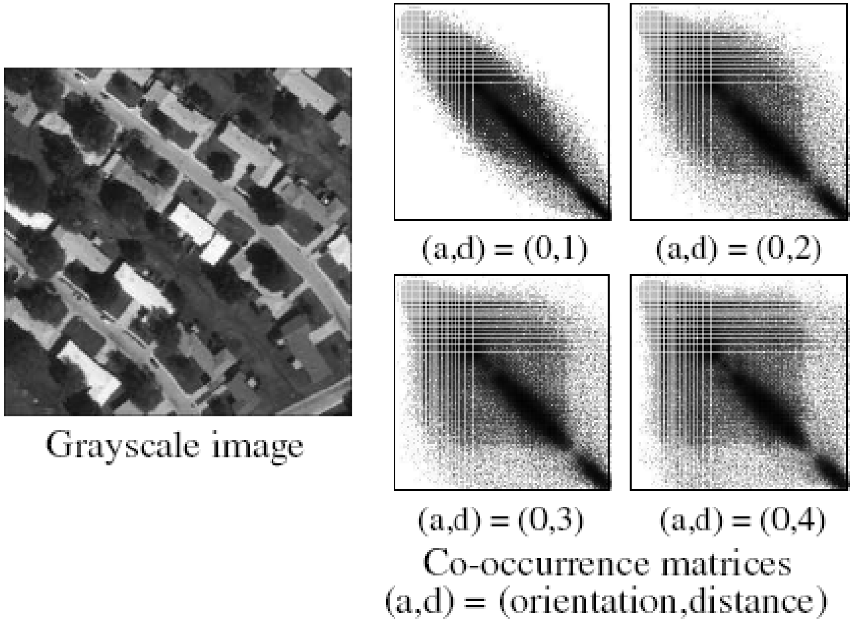

(b) Co-occurrence matrices for an image with a large amount of local spatial variations.

Fourier power spectrum

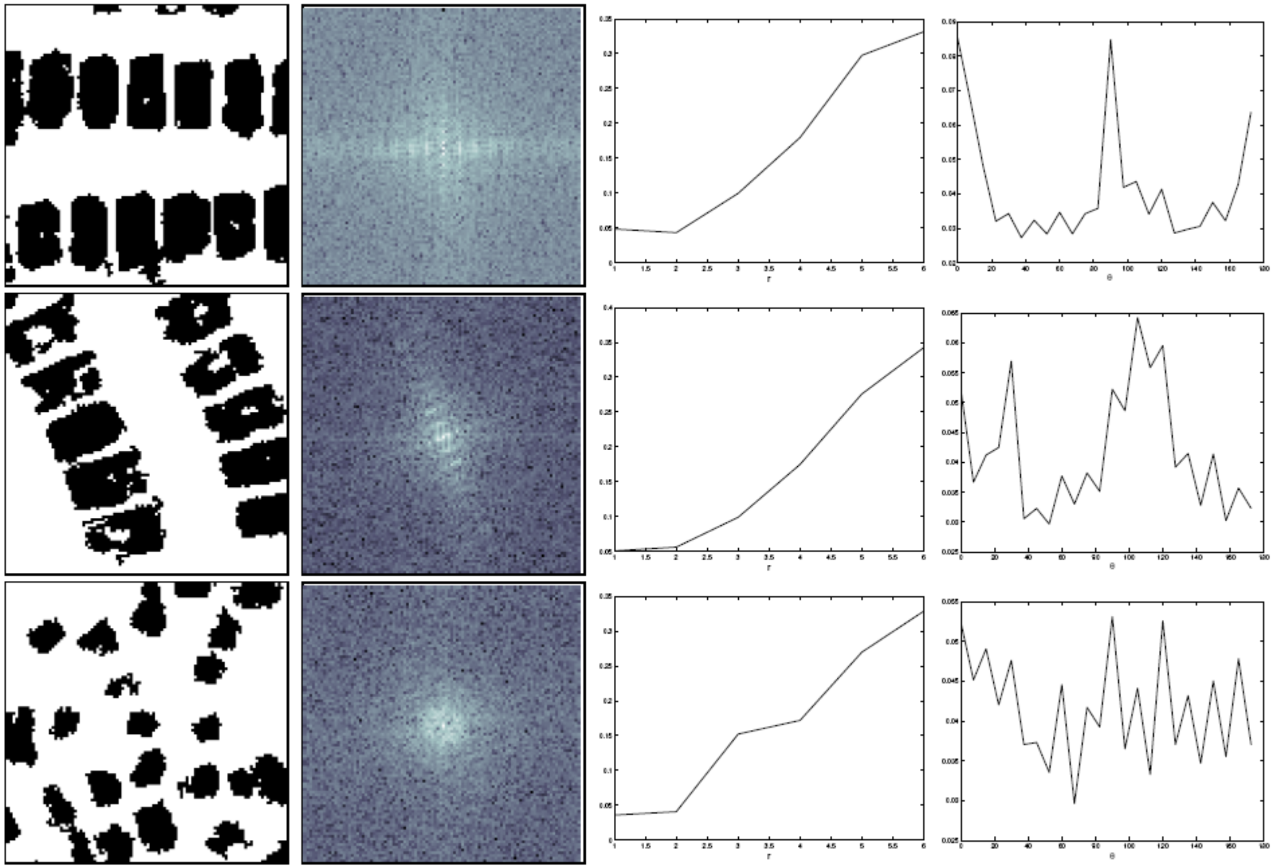

Example building groups (first column), Fourier spectrum of these images (second column), and the corresponding ring- and wedge-based features (third and fourth columns). X-axes represent the rings in the third column and the wedges in the fourth column plots. The peaks in the features correspond to the periodicity and directionality of the buildings, whereas no dominant peaks can be found when there is no regular building pattern.

Texture Classification

Bam, Iran, suffered an earthquake in 2003

IKONOS image acquired in the aftermath of the earthquake

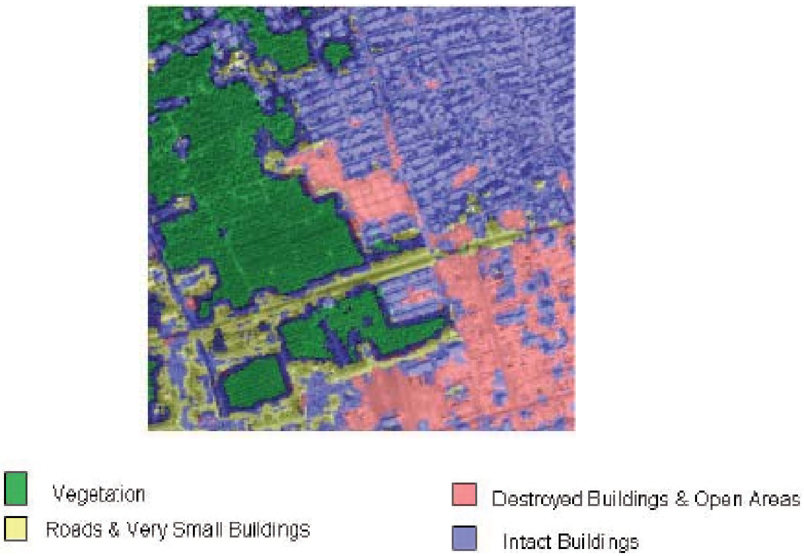

Classification obtained on the basis of the texture parameters only

Edge texture

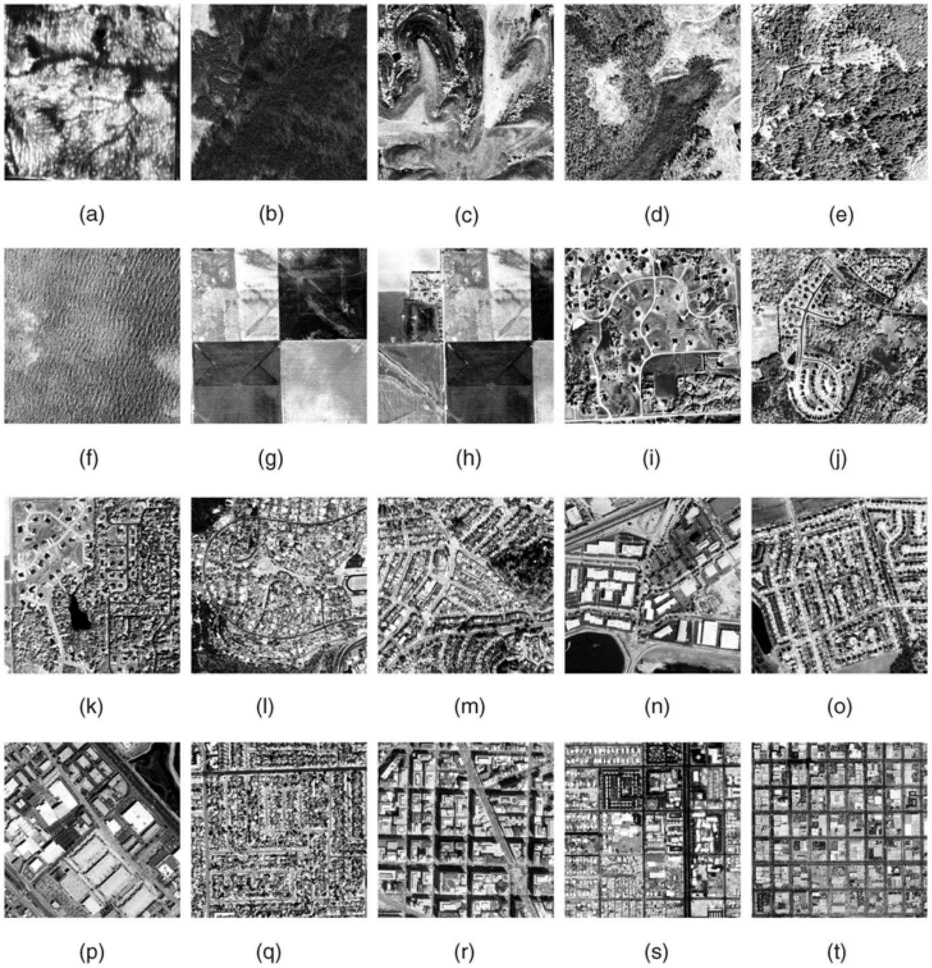

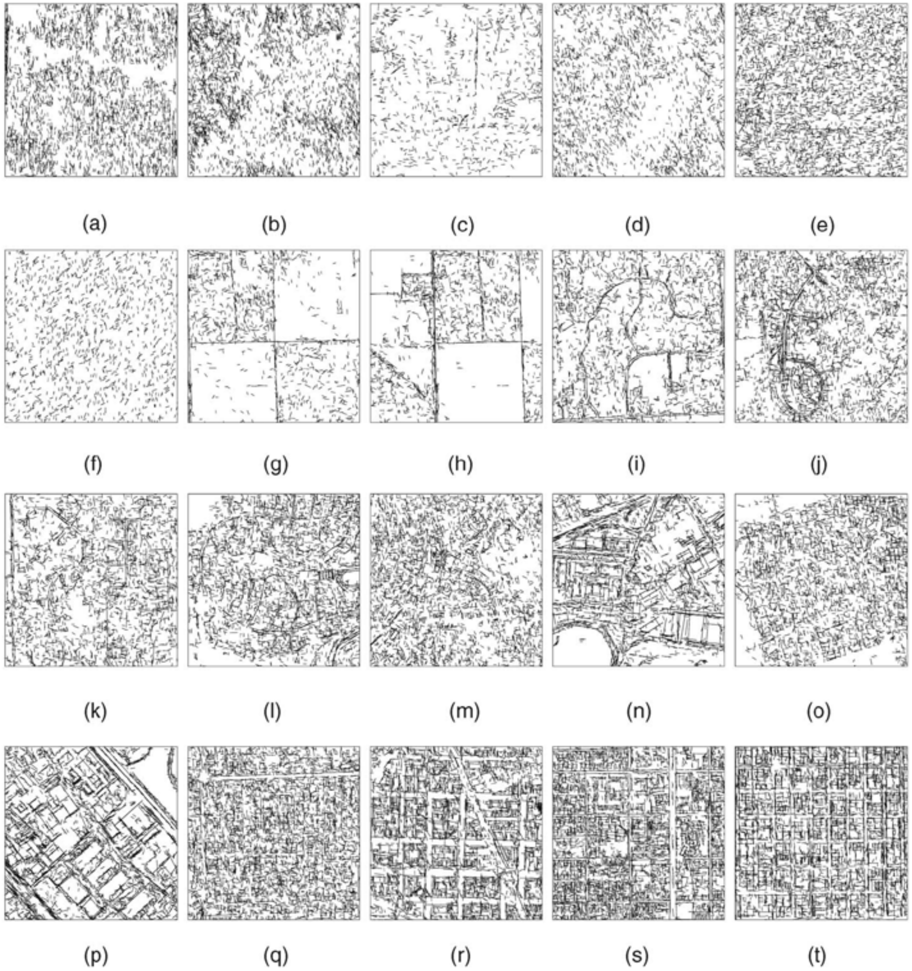

Satellite images sorted according to the amount of land development (left). Properties of the arrangements of line segments can be used to model the organization in an area (right).

Matching based on “Interesting” Points

Object recognition: Find correspondences between feature points in training and test images.

Local Features Detectors for Image Matching

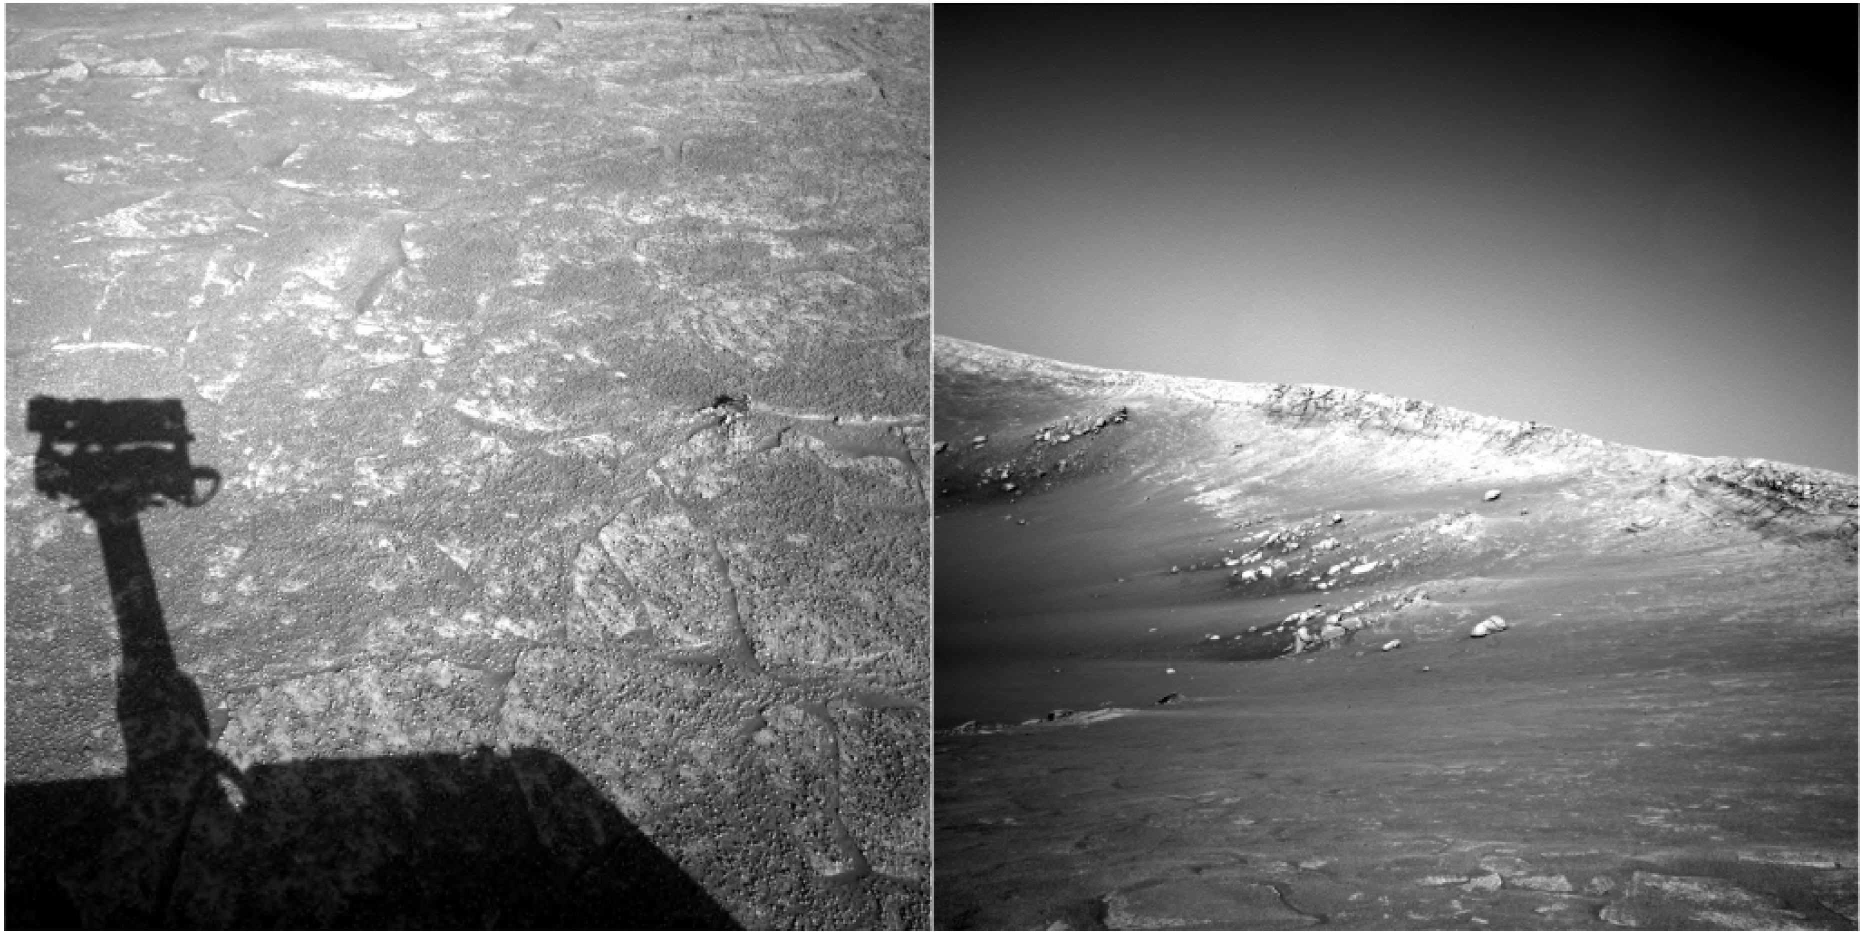

Two images from NASA Mars Rover: very hard matching case

Local Features Detectors for Image Matching

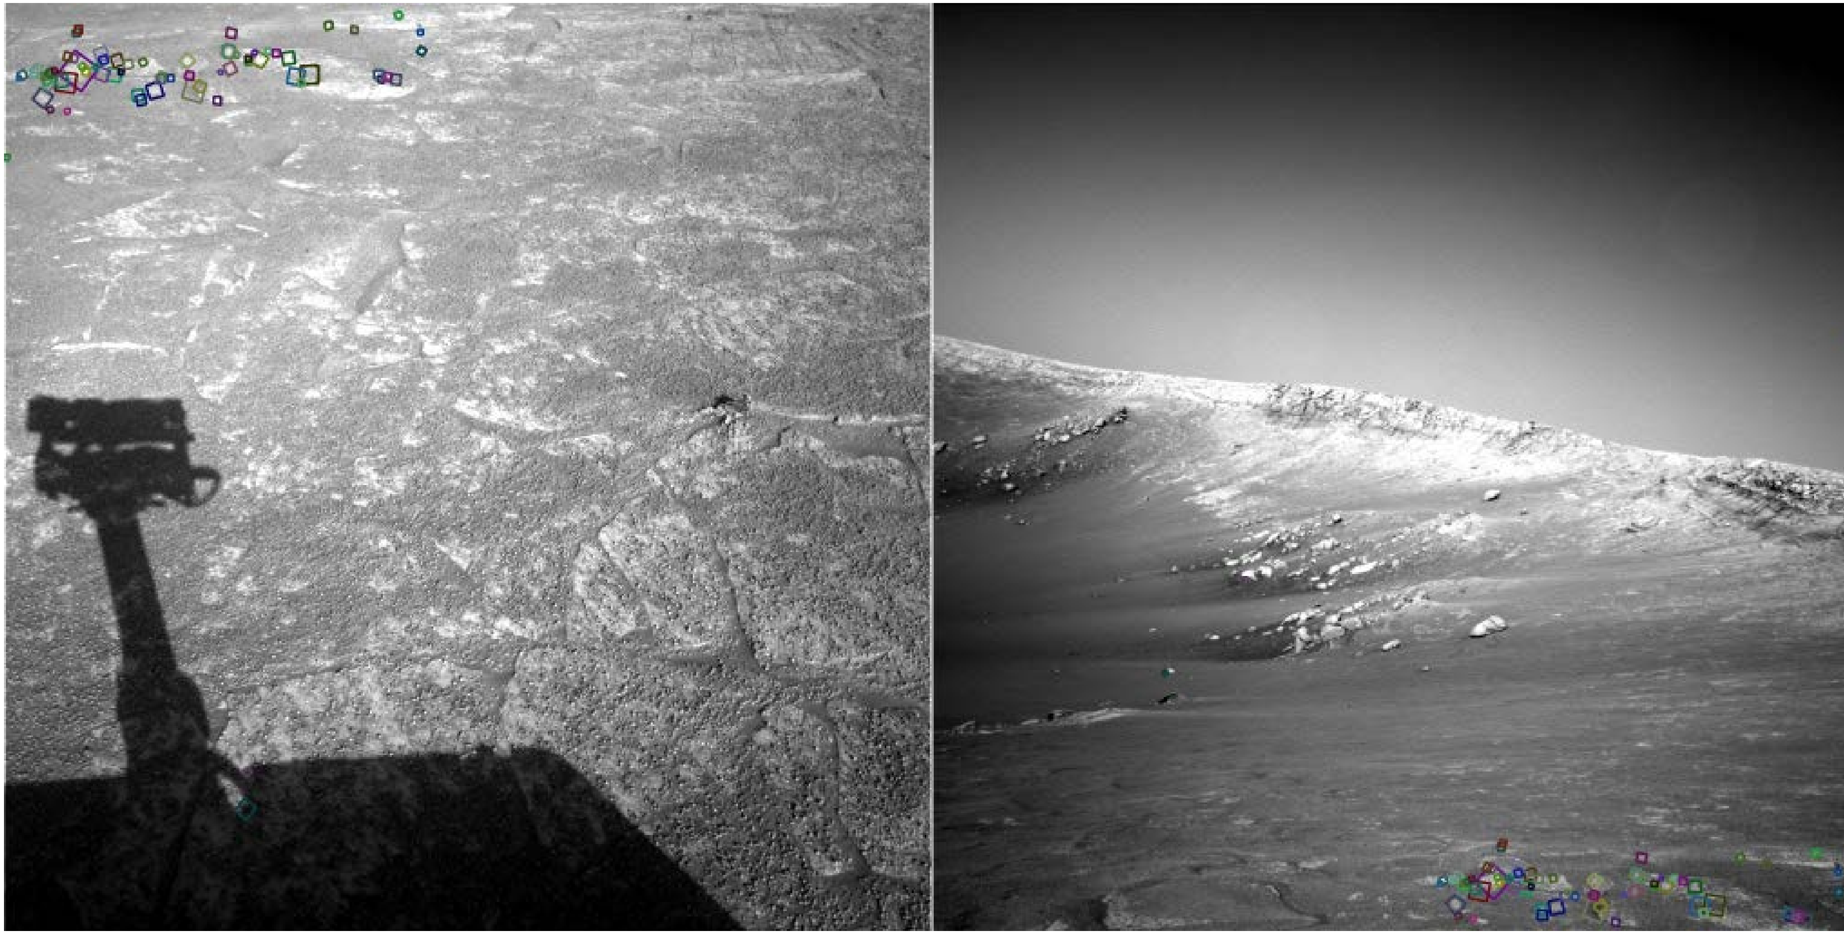

Two images from NASA Mars Rover: matching using local features

Remote Sensing: Building Detection

Model Building Database

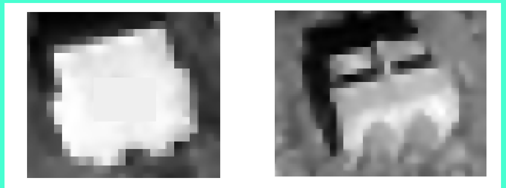

(a) Bright building (b) Dark building

Each feature is matched with the most similar features

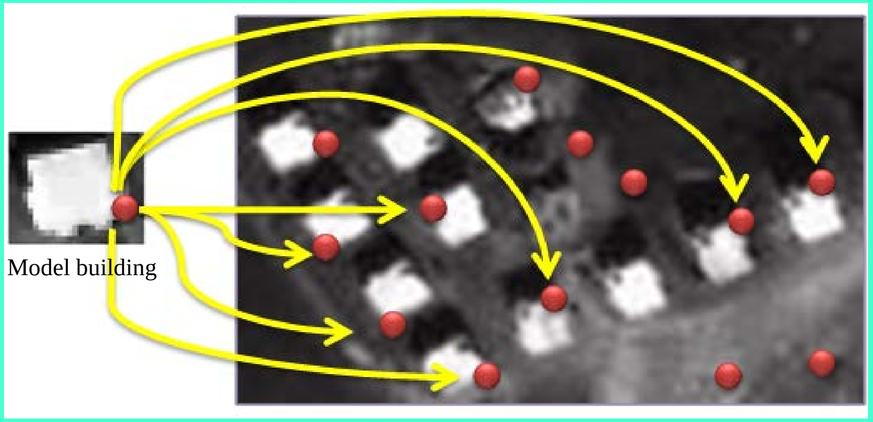

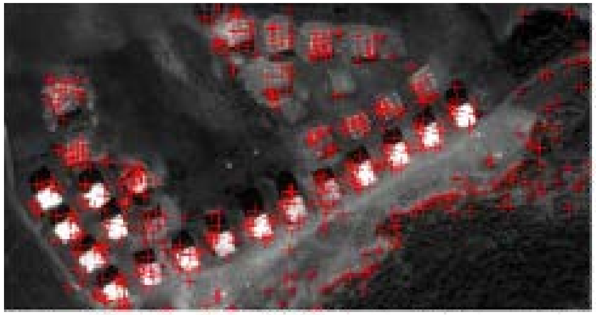

Detected „Interesting Points“

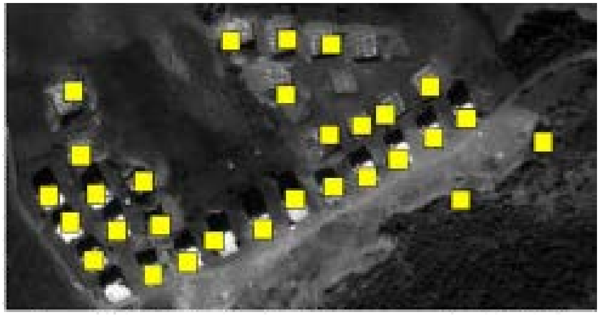

Detected buildings

Adapted from B. Sirmacek, DLR

What do you do when you choose Ground Control Points (GCP)?

A.S. Kiseleva

Name	Type	L	R	Corr.	Par.
*16	Tie	+	+	0.958	-0.001
*17	Tie	+	+	0.961	-0.002
*18	Tie	+	+	0.932	0.002
*19	Tie	+	+	0.961	0.003

Example: Image coregistration

(I had to remove this)

Local features

- What makes a good feature?
- We want uniqueness.
  - Look for image regions that are unusual.
  - Lead to unambiguous matches in other images.
- How to define "unusual"?

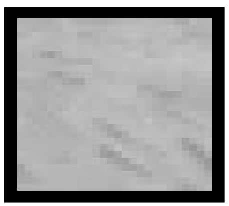

0D structure

not useful for matching

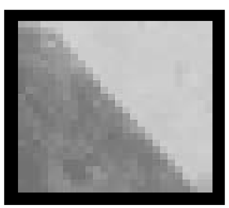

1D structure

edge, can be localized in 1D, not so good for matching

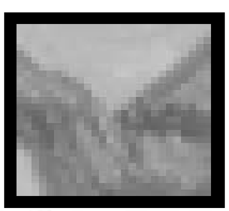

2D structure

corner, can be localized in 2D, good for matching

Which points make good features?

Good candidates are points with strong variations in all directions

Full image

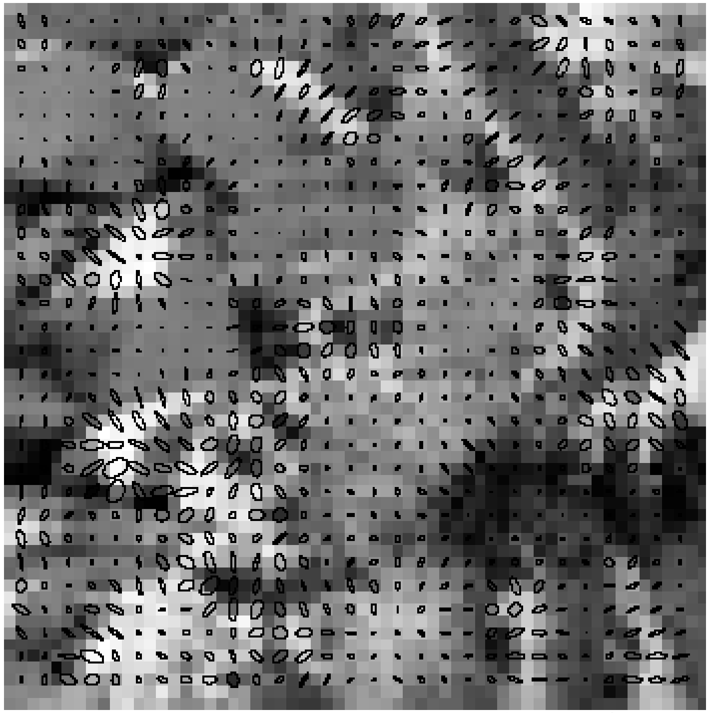

from Forsyth & Ponce

Detail of image with gradient covar-
iance ellipses for 3 x 3 windows

Sample Output of a Corner Detector

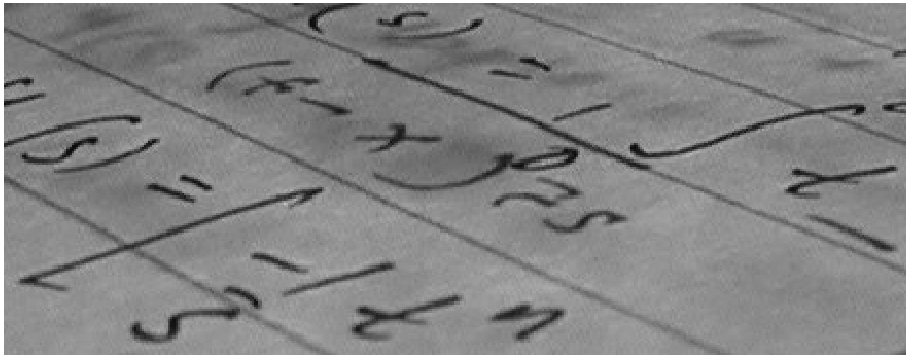

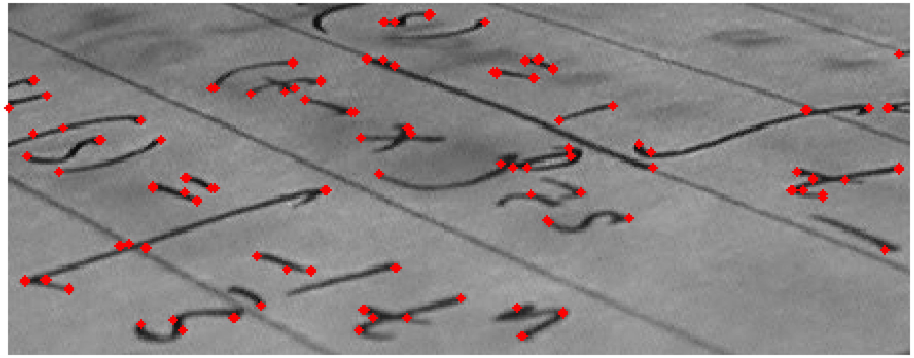

Scale Invariant Point of Interest Detection

- The problem: “corners” are dependent on the scale of the image
- How do we choose corresponding circles ***independently*** in each image?

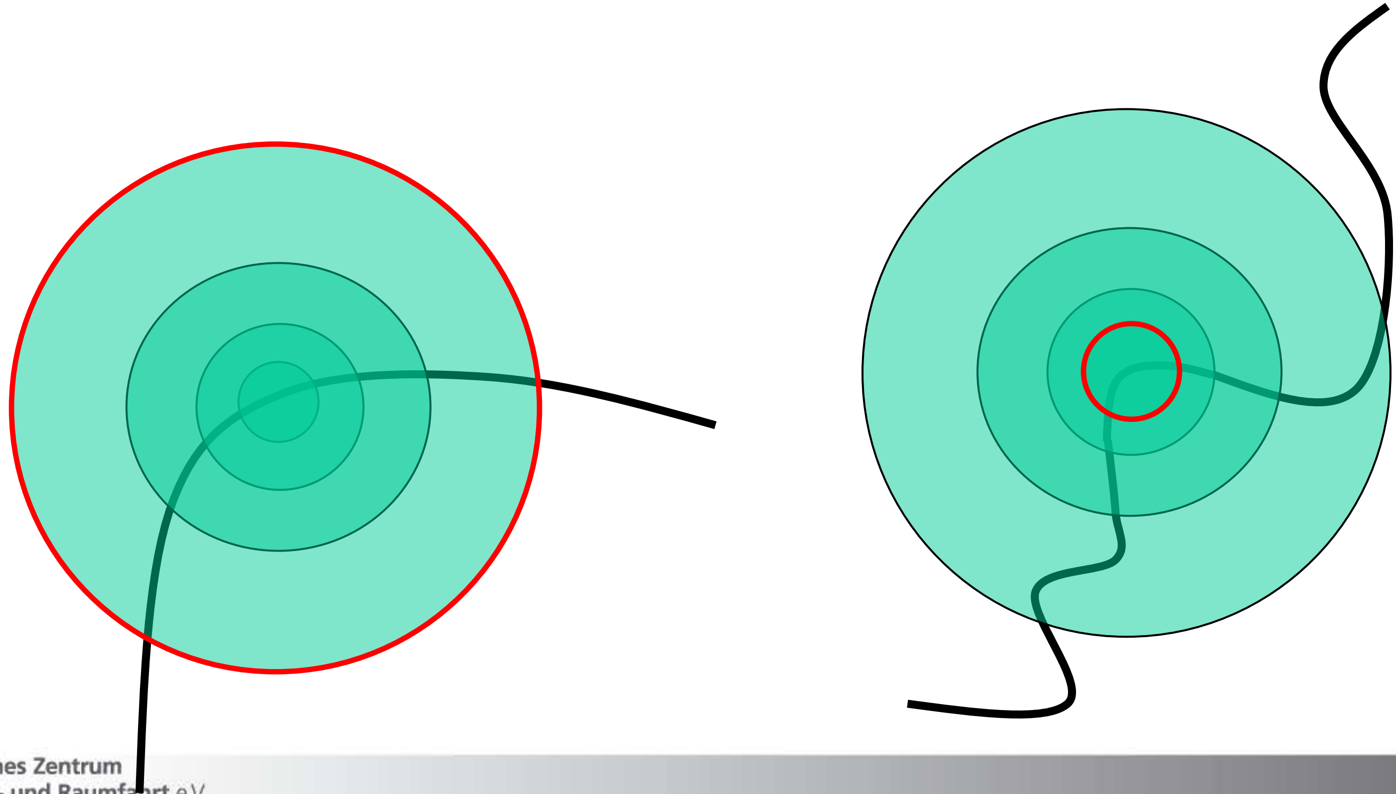

Look for “Corners” in Difference-of-Gaussians

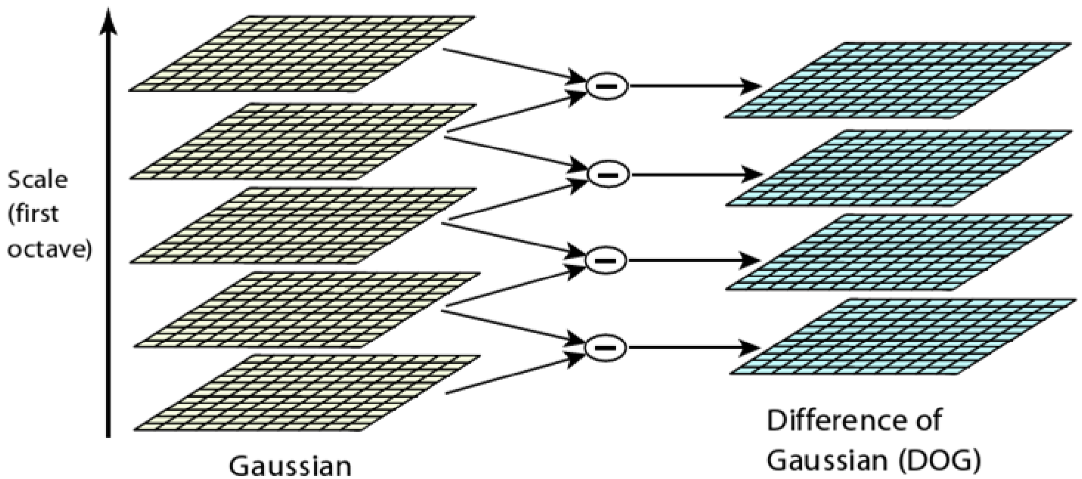

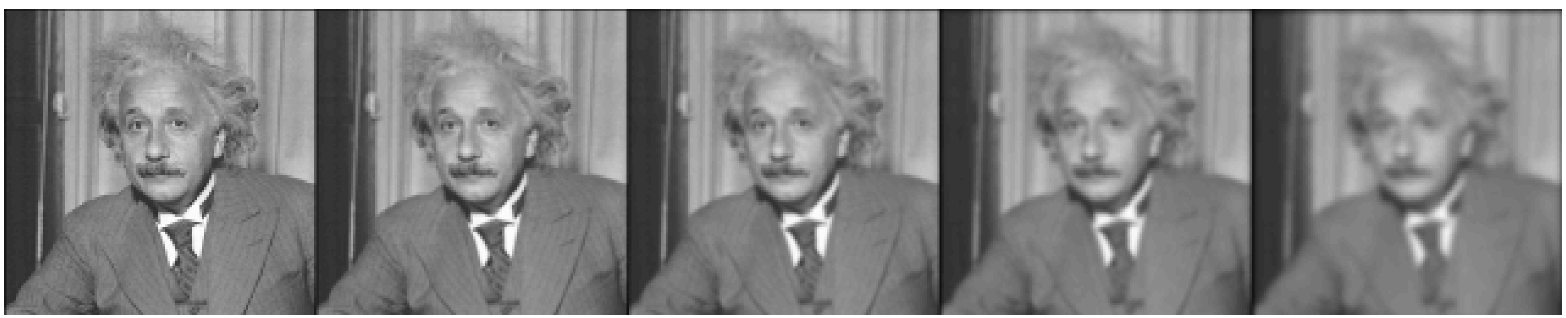

Scale Invariant Feature Transform (SIFT) descriptors

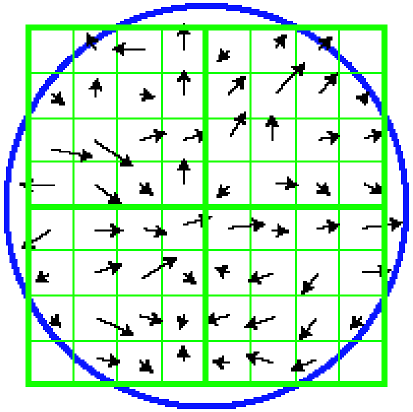

Image gradients

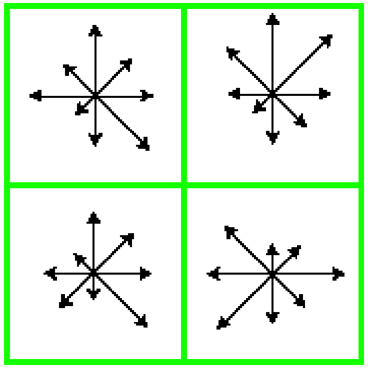

Keypoint descriptor

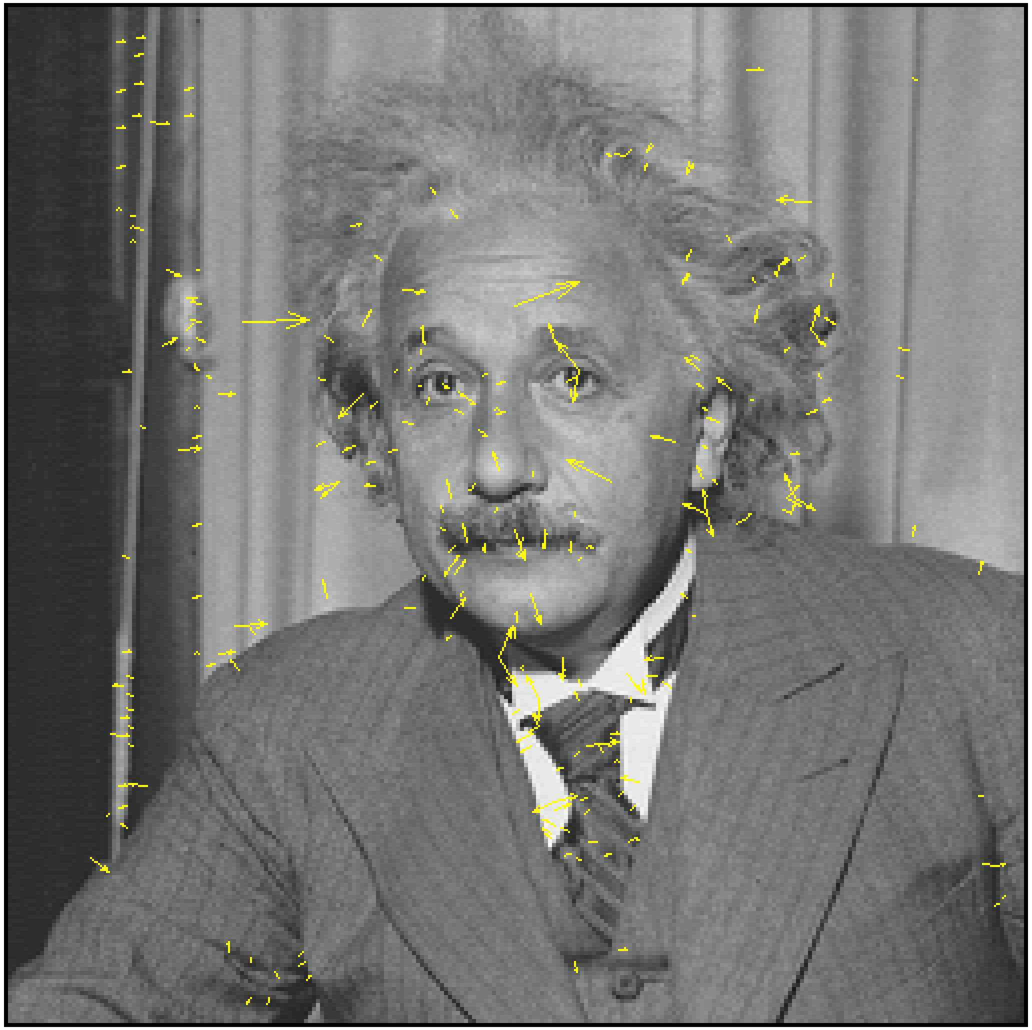

Matching examples

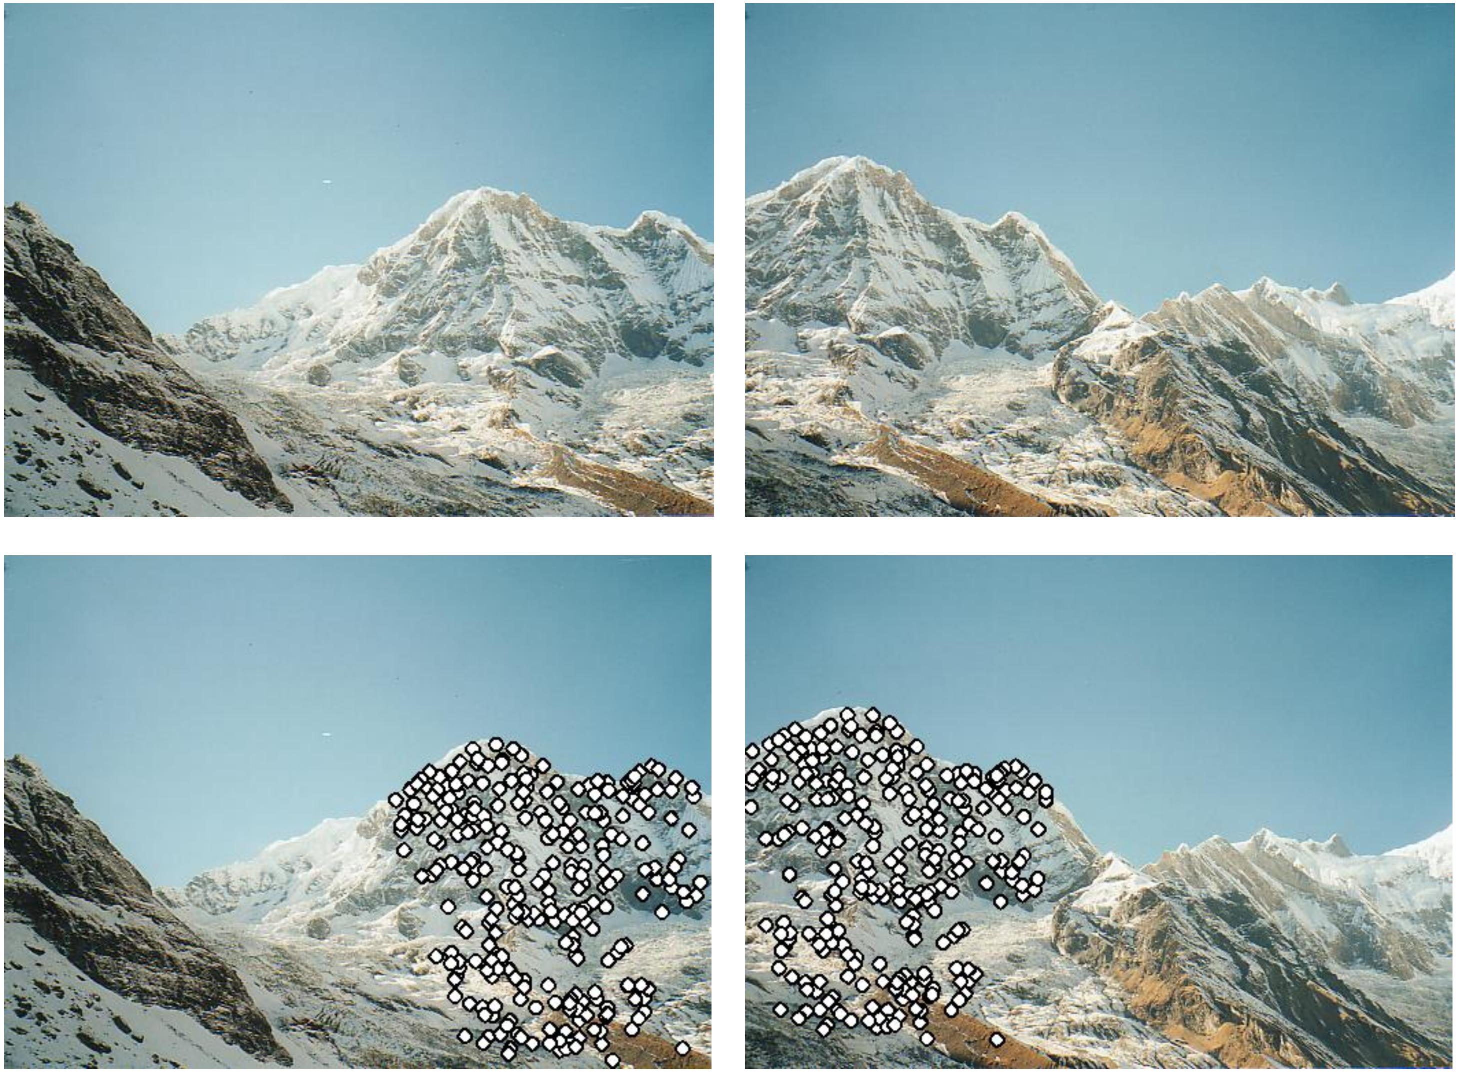

Matching examples

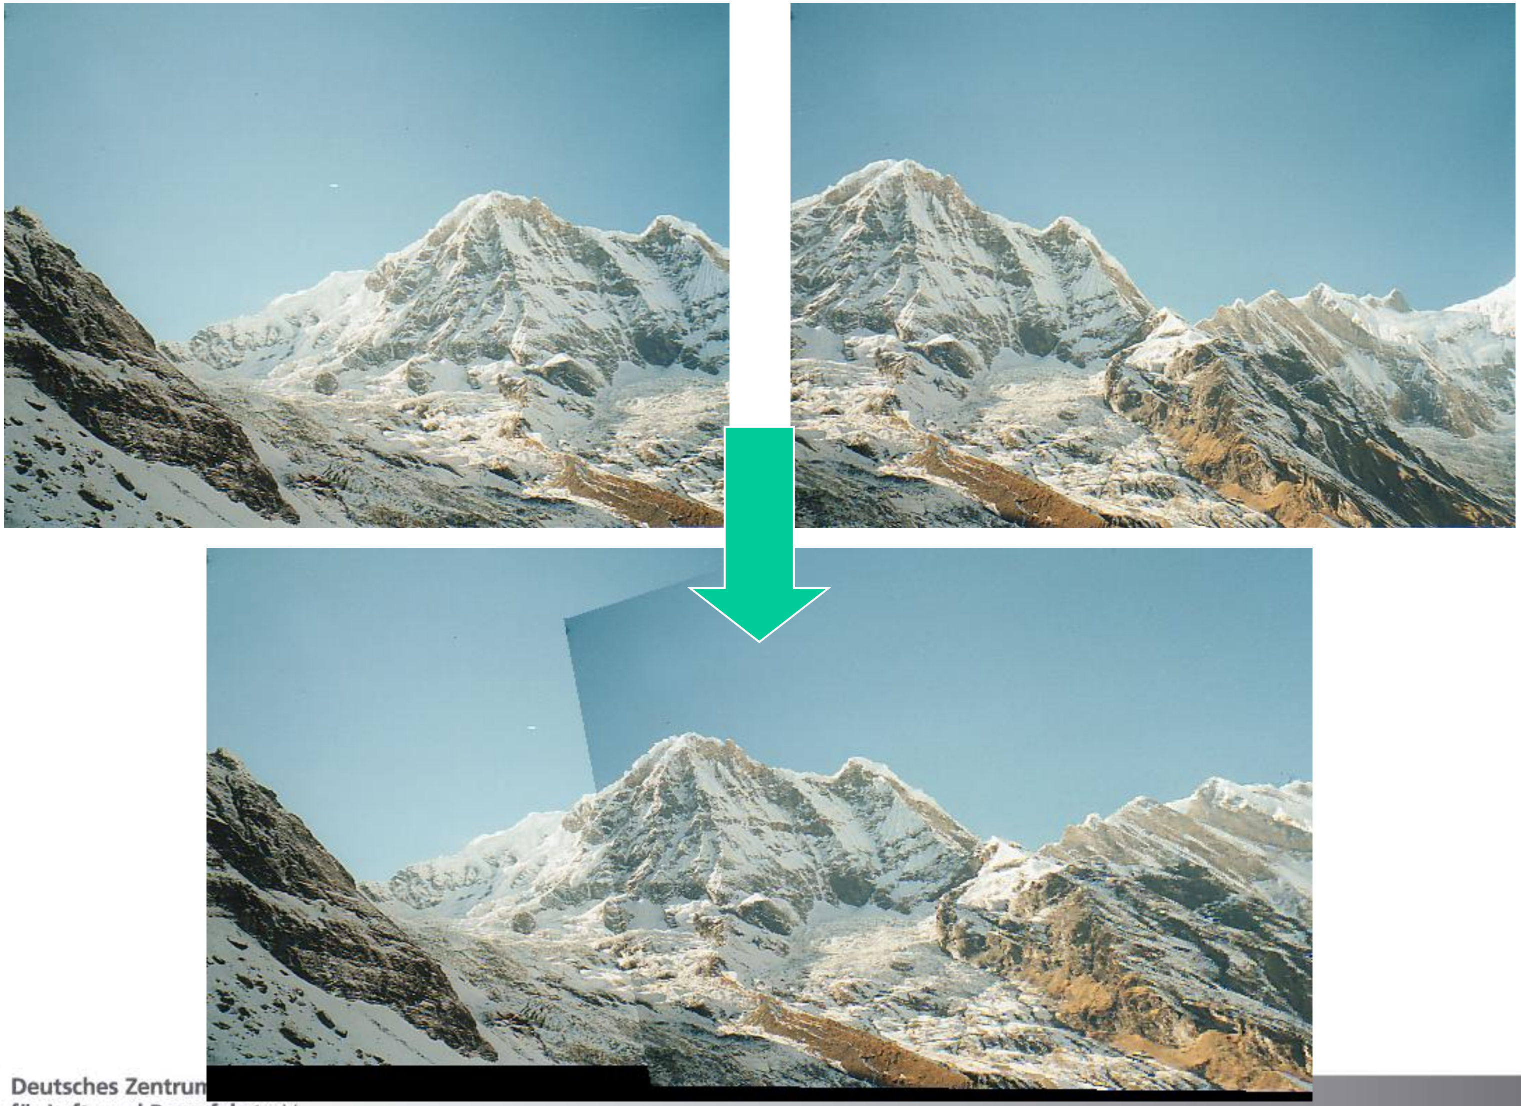

Examples: panoramas

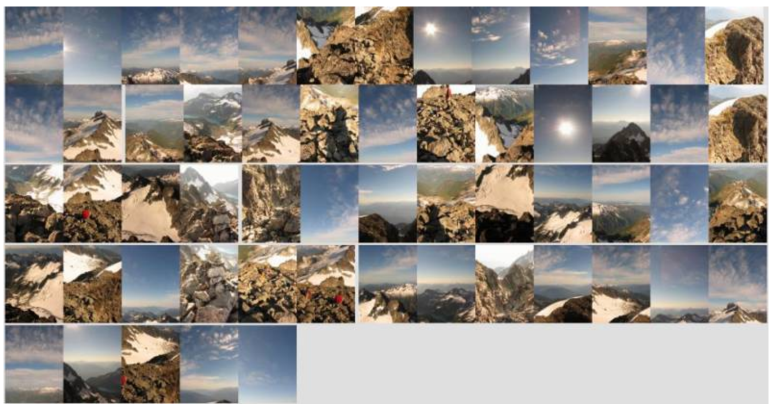

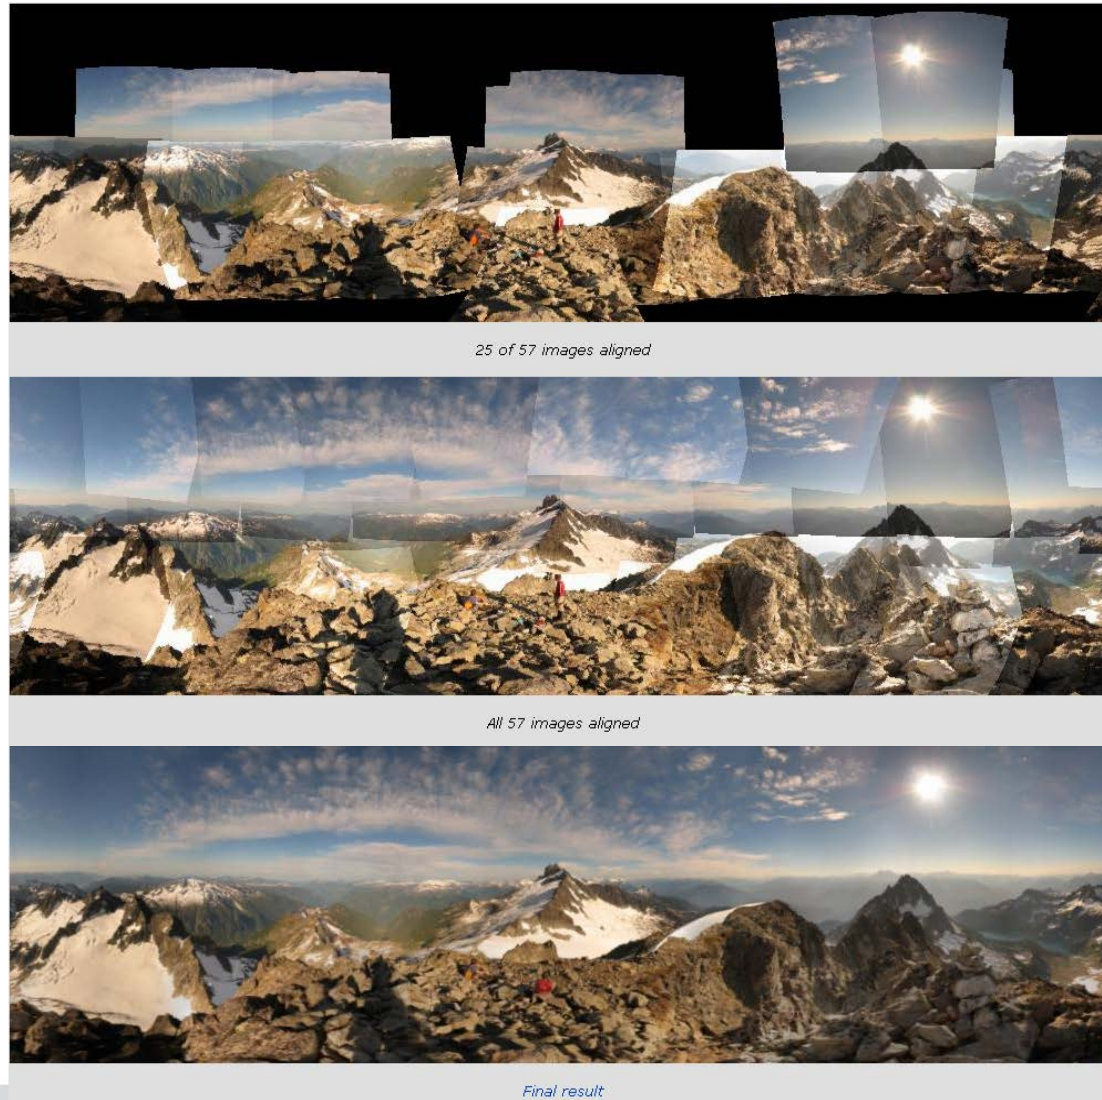

25 of 57 images aligned

All 57 images aligned

Final result

Sony Aibo

- SIFT usage:
  - Recognize charging station
  - Communicate with visual cards
  - Teach object recognition

Photo tourism: exploring photo collections

- Joint work by University of Washington and Microsoft Research
  - http://phototour.cs.washington.edu/
  - http://research.microsoft.com/IVM/PhotoTours/
- Photosynth Technology Preview at Microsoft Live Labs
  - http://labs.live.com/photosynth/

Photo tourism: exploring photo collections

- Detect features using SIFT.

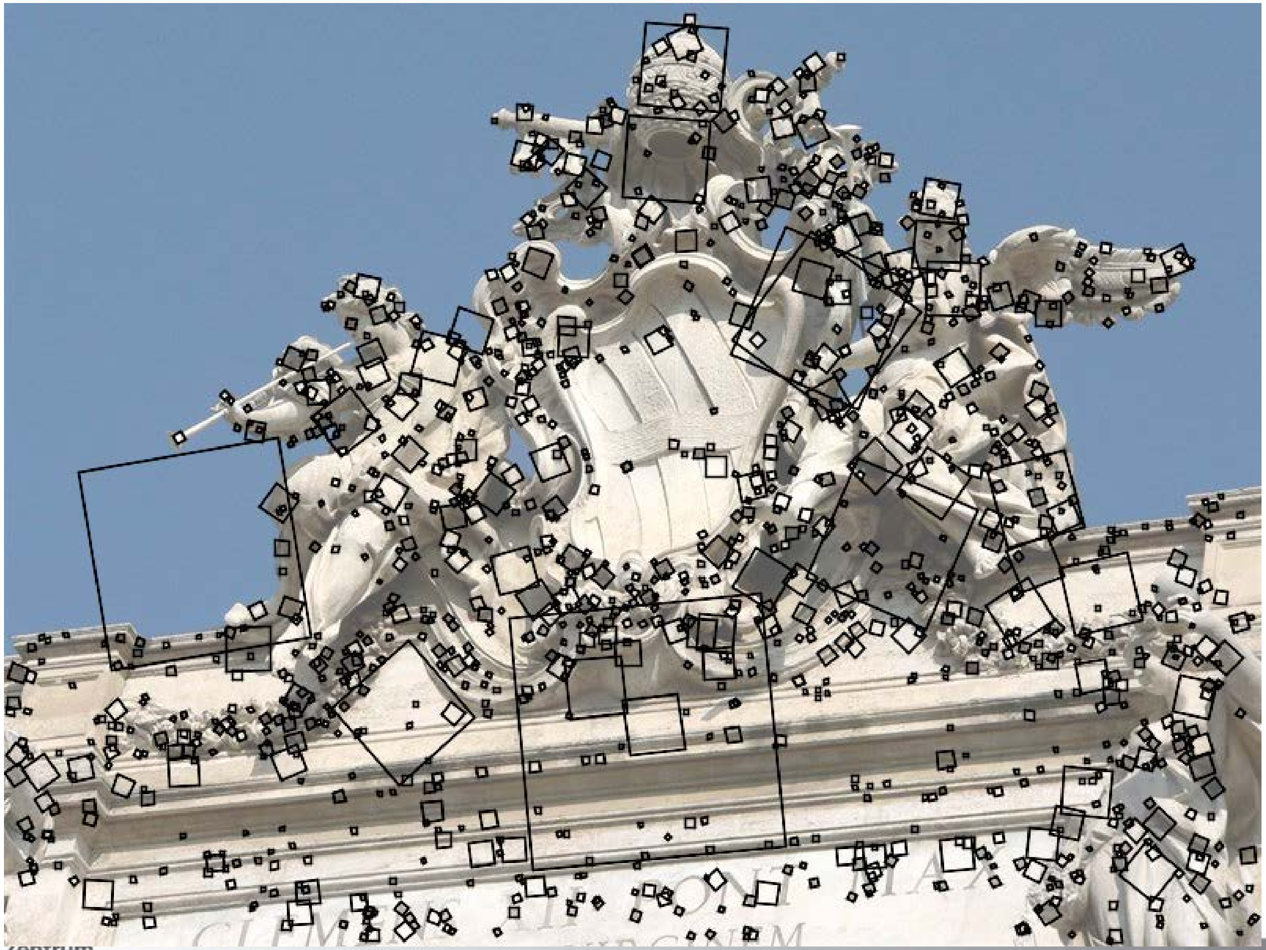

Photo tourism: exploring photo collections

- Detect features using SIFT.

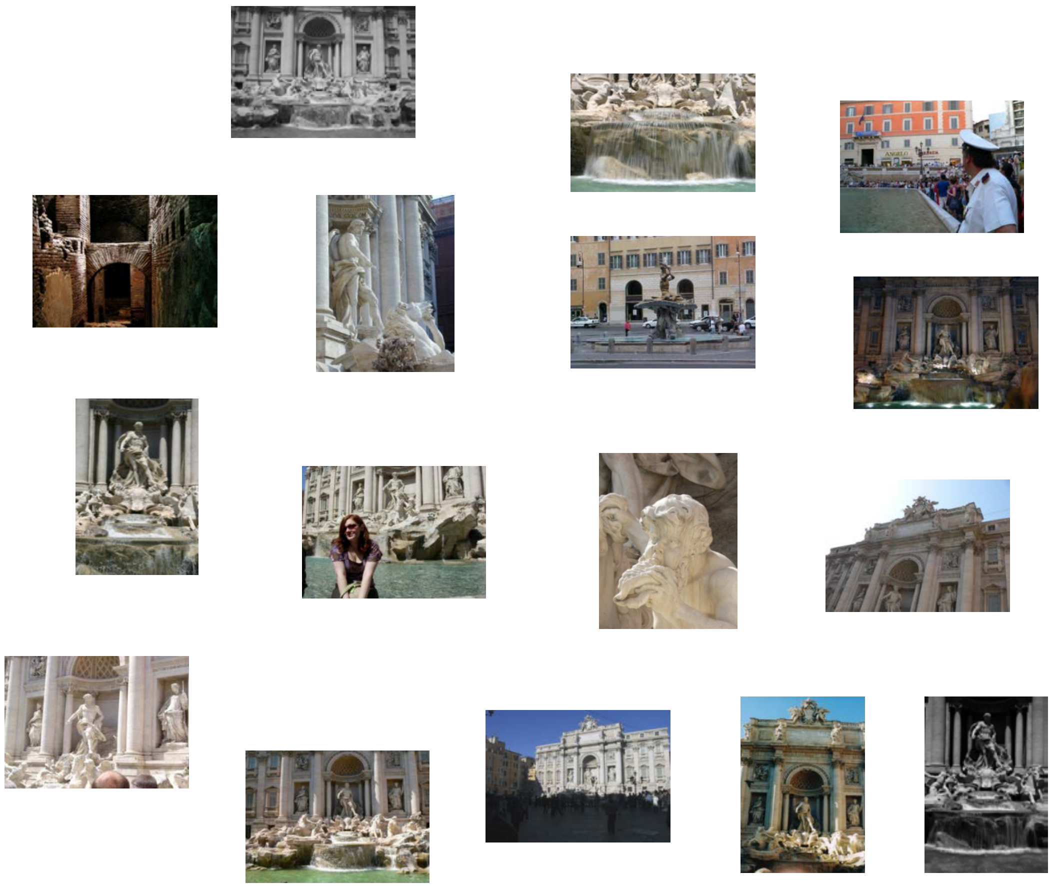

Photo tourism: exploring photo collections

- Detect features using SIFT.

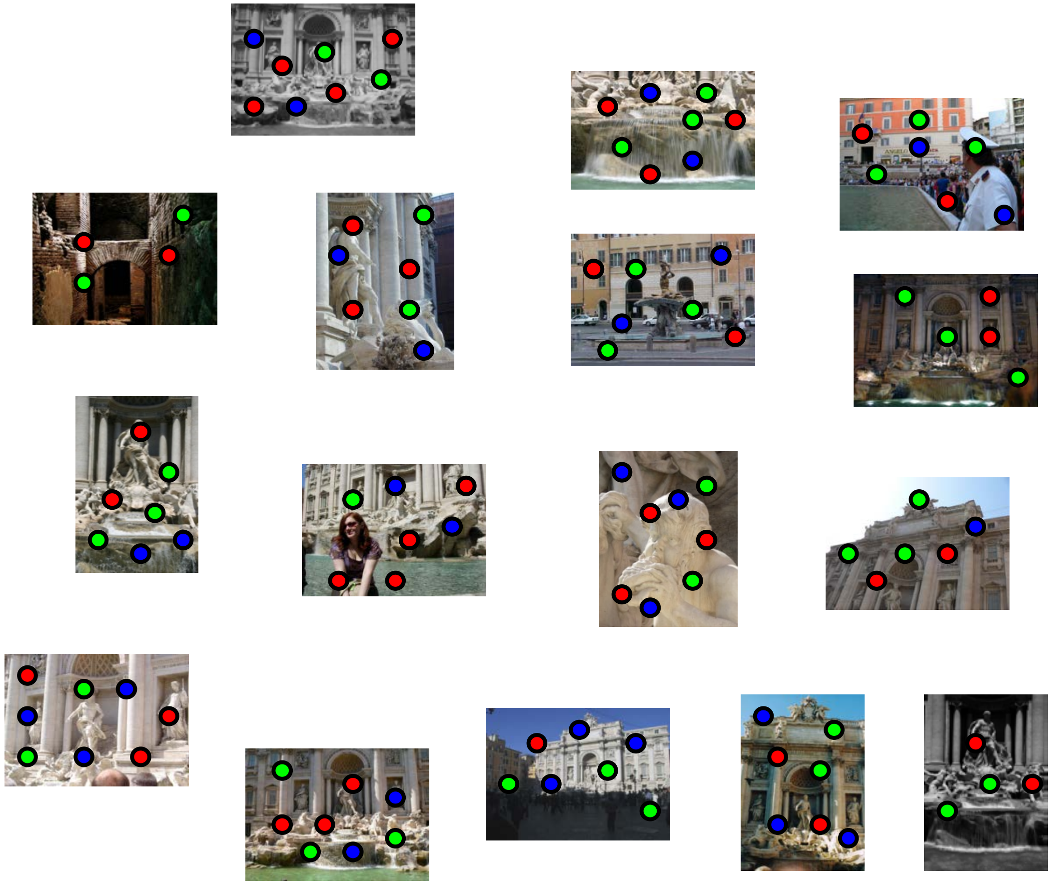

Photo tourism: exploring photo collections

- Match features between each pair of images.

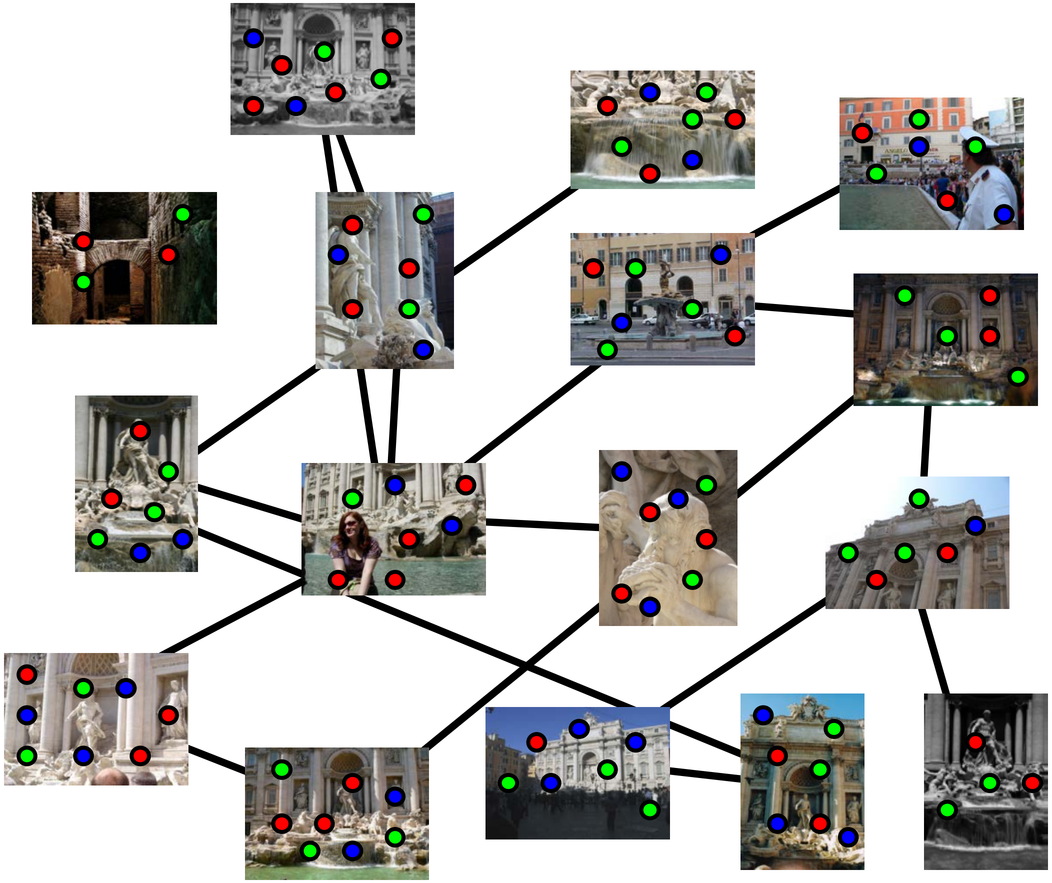

Photo tourism: exploring photo collections

- Link up pairwise matches to form connected components of matches across several images.

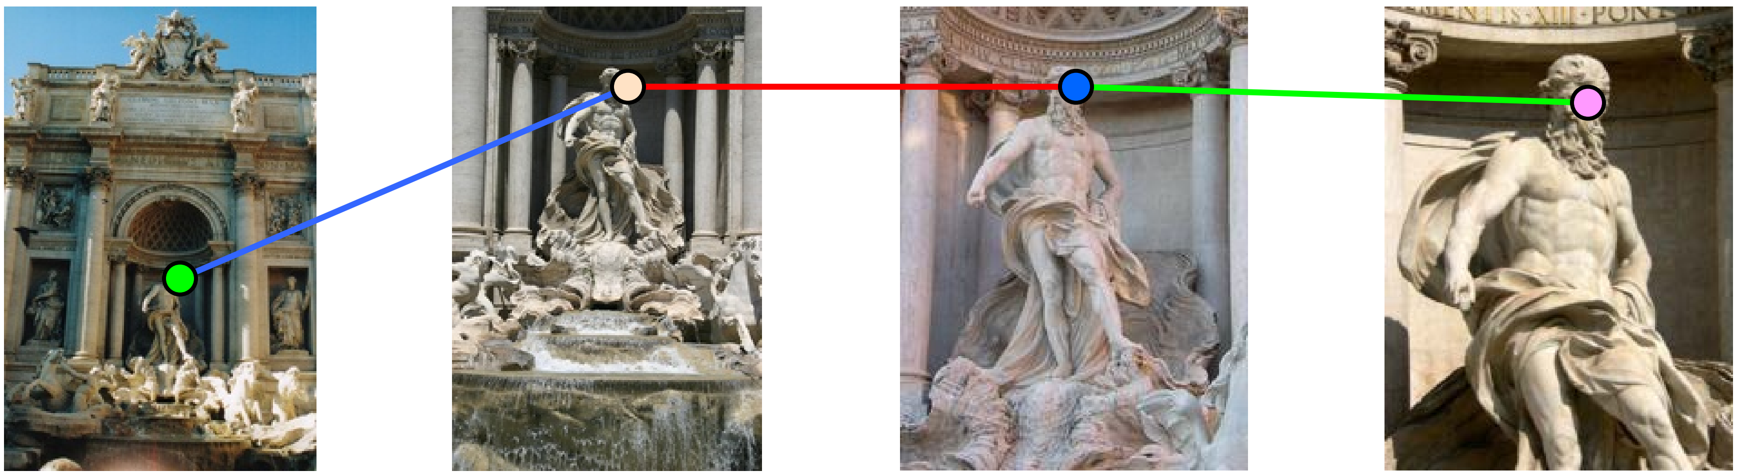

Image 1 Image 2 Image 3 Image 4

Photo tourism: exploring photo collections