Tutorial on Image-Based Measurement

Cris Luengo
Centre for Image Analysis
Uppsala University
Sweden

Tutorial on Image-Based Measurement

- Part I: Stereology
  - On statistics, count-based measurement, and slices
- Part II: Measuring binarized objects
  - Objects represented as a set of pixels
- Part III: Taking grey-values into account
  - Sub-pixel precision through integration
- Part IV: Granulometries
  - Size and length distributions without segmentation

Cris Luengo
luengo@ieee.org

Measurement issues

- Sampling invariance: the choice of sampling grid should not influence the measurement result
- Accuracy (bias)
- Precision
- Preprocessing !?

Sampling invariance

- Translation invariance
  - distances are the same everywhere in the image
  - but also: sub-pixel shifts of sampling grid
- Rotation invariance
  - the rectangular sampling grid makes this difficult
- Scaling invariance
  - a denser grid gives higher precision
  - ...unless we use sub-pixel techniques
- For many grey-value measurements:
  band limit gives maximum attainable precision

Cris Luengo
luengo@ieee.org

Accuracy vs. precision

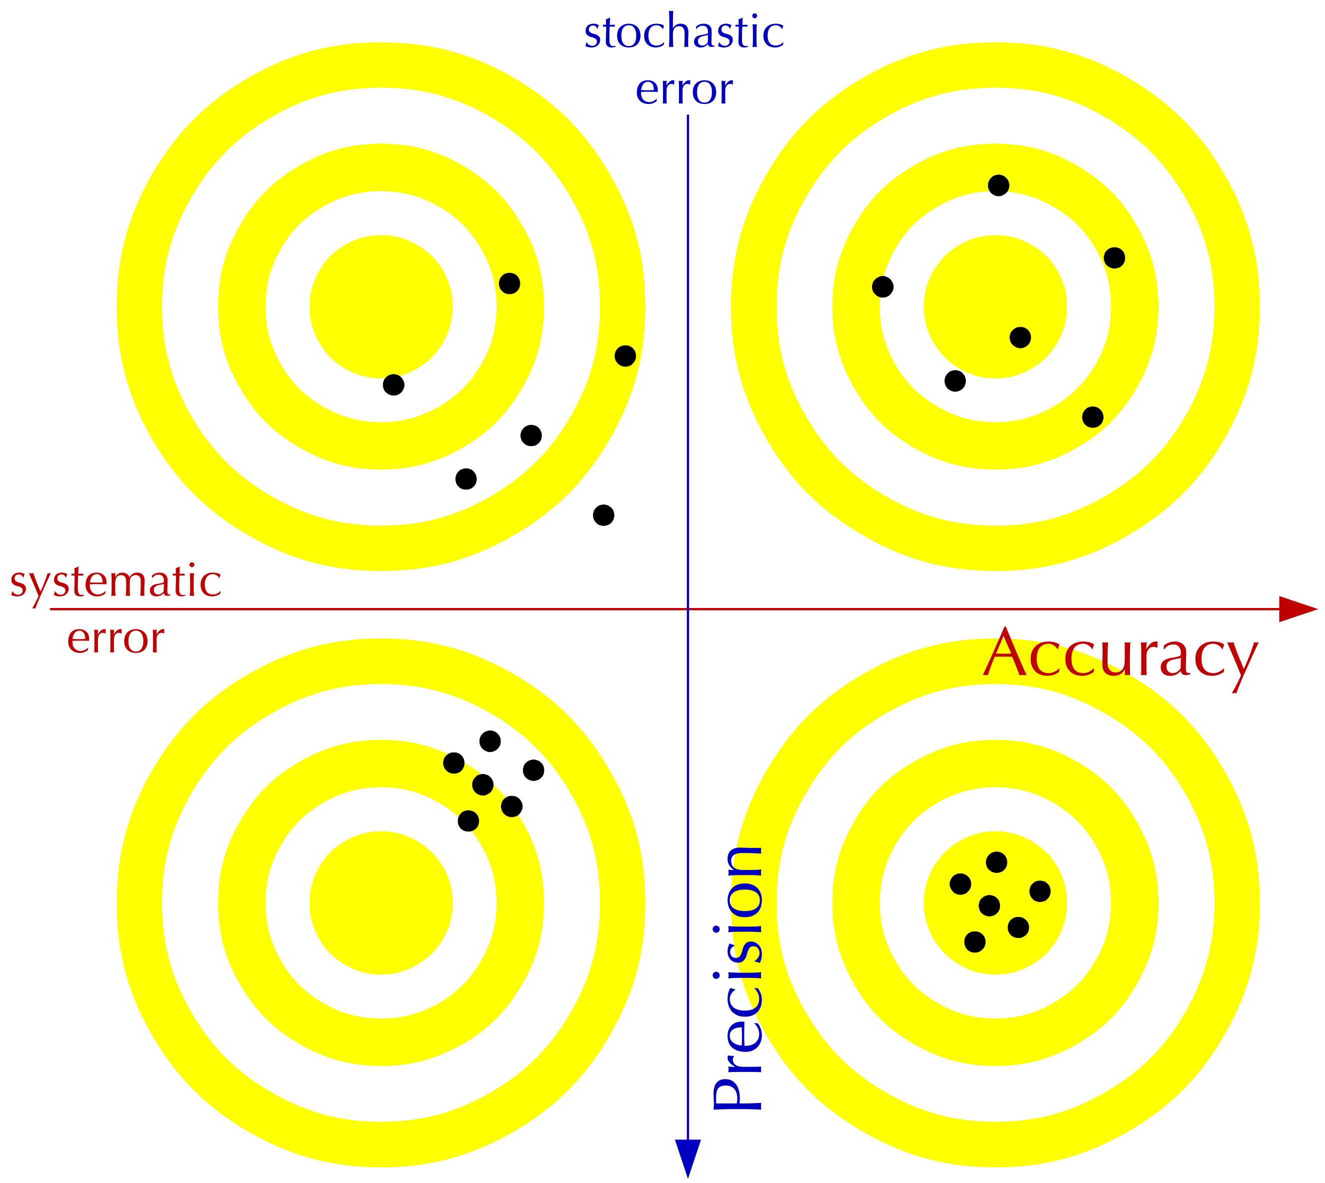

bias = lack of accuracy

Cris Luengo
luengo@ieee.org

Sources of bias

- Uncalibrated equipment
- Incorrect assumptions
  - When the model does not fit reality, conclusions drawn from that model are biased
  - Modern stereology is assumption-free
- Improper sampling
  - Selection bias
  - Biased sectioning of 3D sample
  - Imaging only the nice-looking cells
  - etc.

Cris Luengo
luengo@ieee.org

Selection bias example

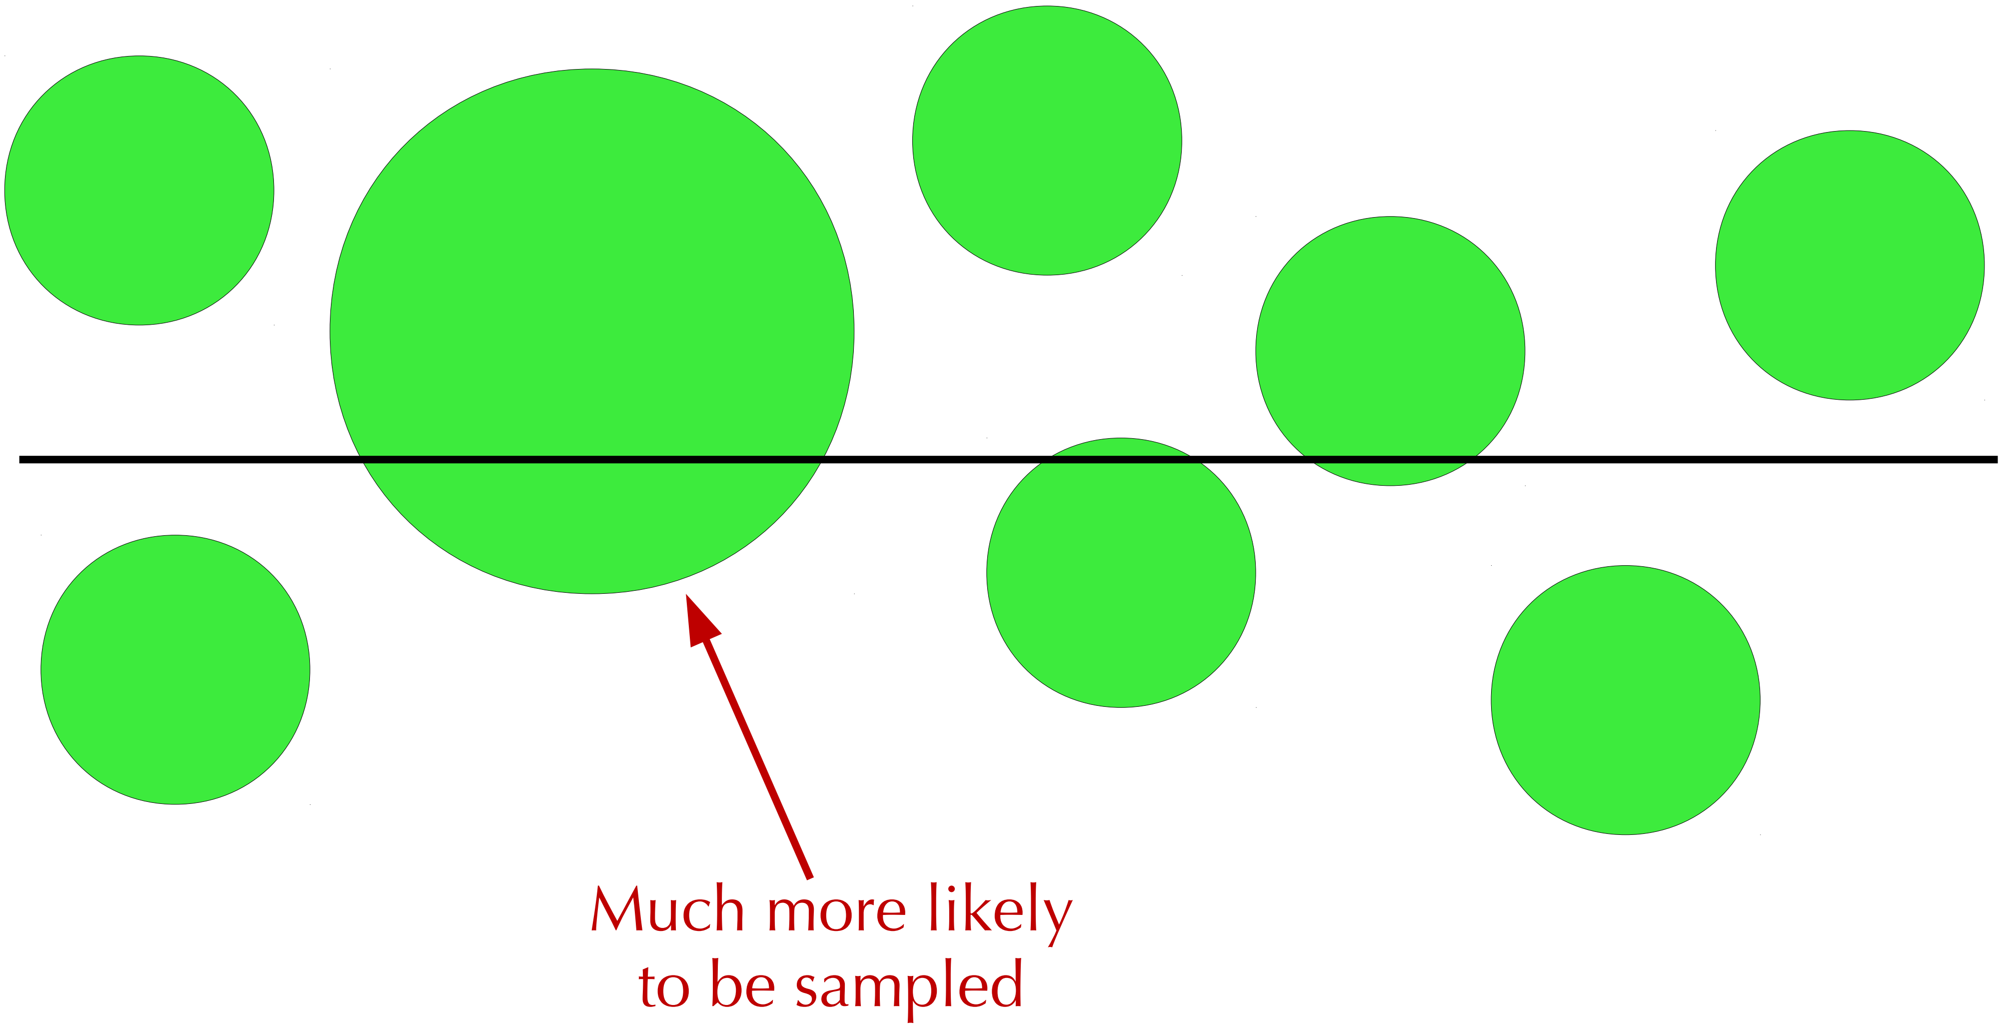

Precision

- We are measuring a property of a population
  - the population has a distribution $N(\mu,\sigma_{bio})$
  - the measurement has an error $N(0,\sigma_{meas})$
  - what we measure is $N(\mu,\sigma_{bio}) + N(0,\sigma_{meas})$

unbiased!

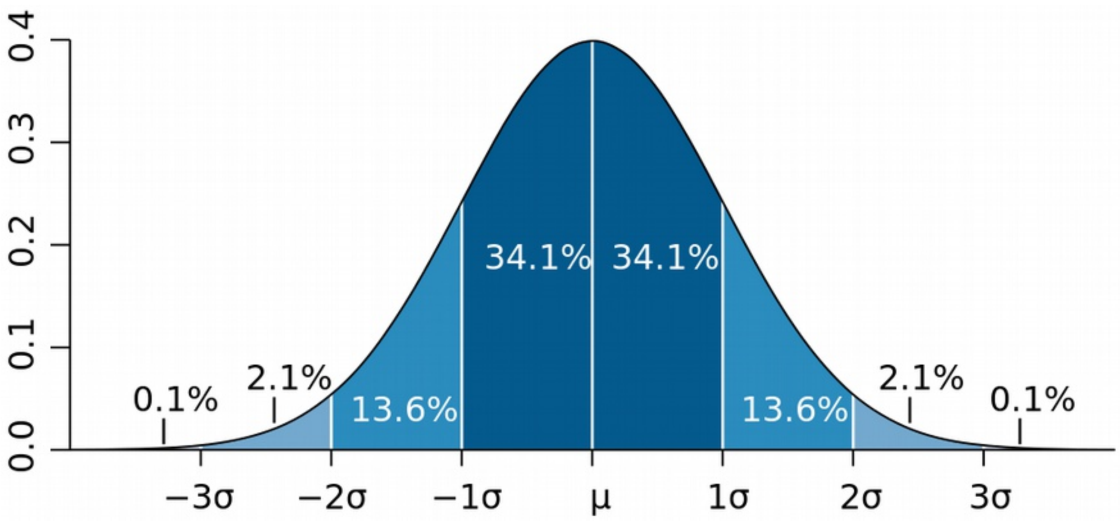

Precision

- In terms of coefficient of variation:
  - $CV_{\text{bio}} = \sigma_{\text{bio}} / \mu$
  - $CE = \sigma_{\text{meas}} / \mu$
  - we measure $CV_{\text{meas}}^{\;2} = CV_{\text{bio}}^{\;2} + CE^{2}$
- Optimally, CE should be somewhat smaller than CV_{bio}, but not by much!

$$0.1 < \frac{CE^2}{CV_{\text{bio}}^2} < 0.5$$

Cris Luengo
luengo@ieee.org

Filtering affects measurements!

- Low-pass filtering always moves the edges inwards
  - (Inwards = in the direction of curvature)
- Edge-preserving smoothing filters sometimes also move edges

Gauss, σ=10 threshold 0.5

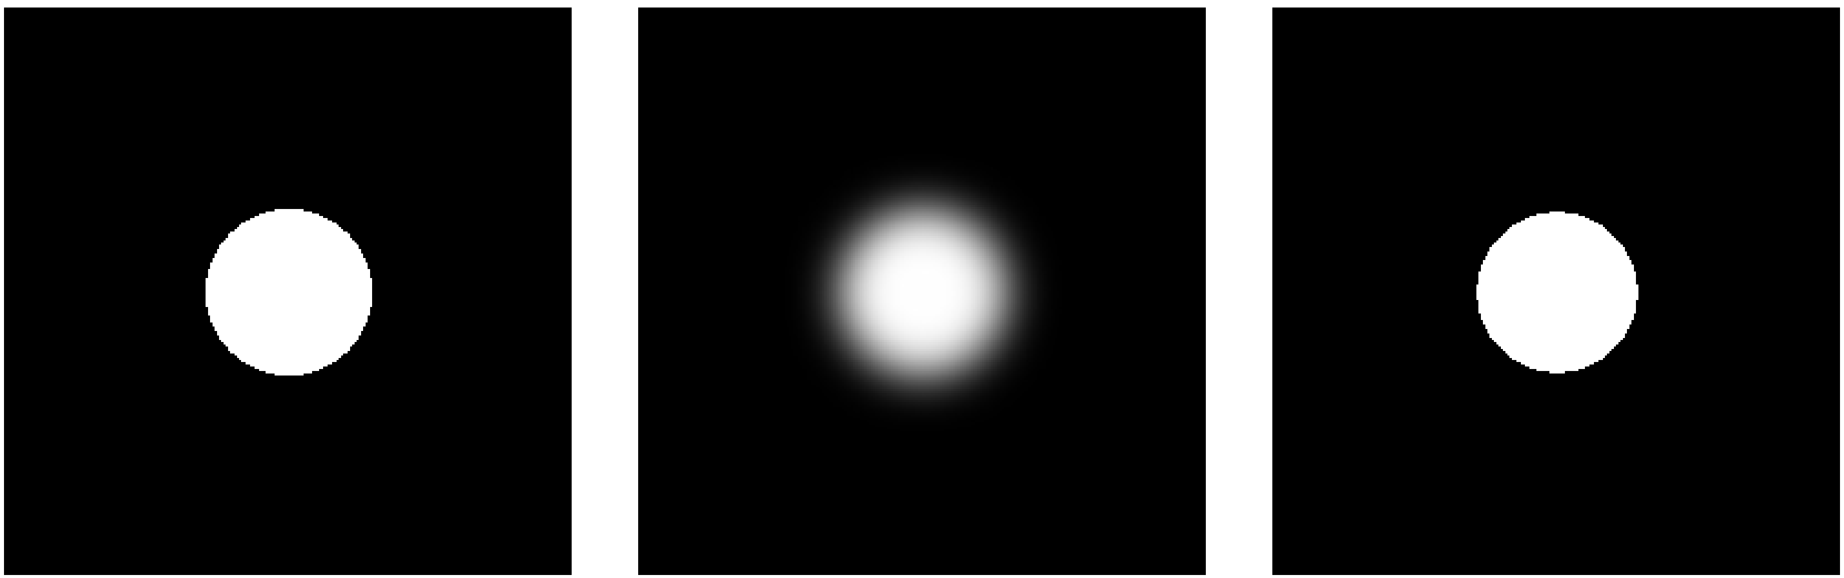

count = 4421 count = 4101

Cris Luengo
luengo@ieee.org

Tutorial on Image-Based Measurement

Part I: Stereology

Cris Luengo
luengo@ieee.org

Overview

- Sampling
- Stereological approach to measurement
- Estimating volume
- Estimating surface area
- Counting

- Delesse, Buffon, Cavalieri, Mandelbrot

Cris Luengo
luengo@ieee.org

Independent random sampling

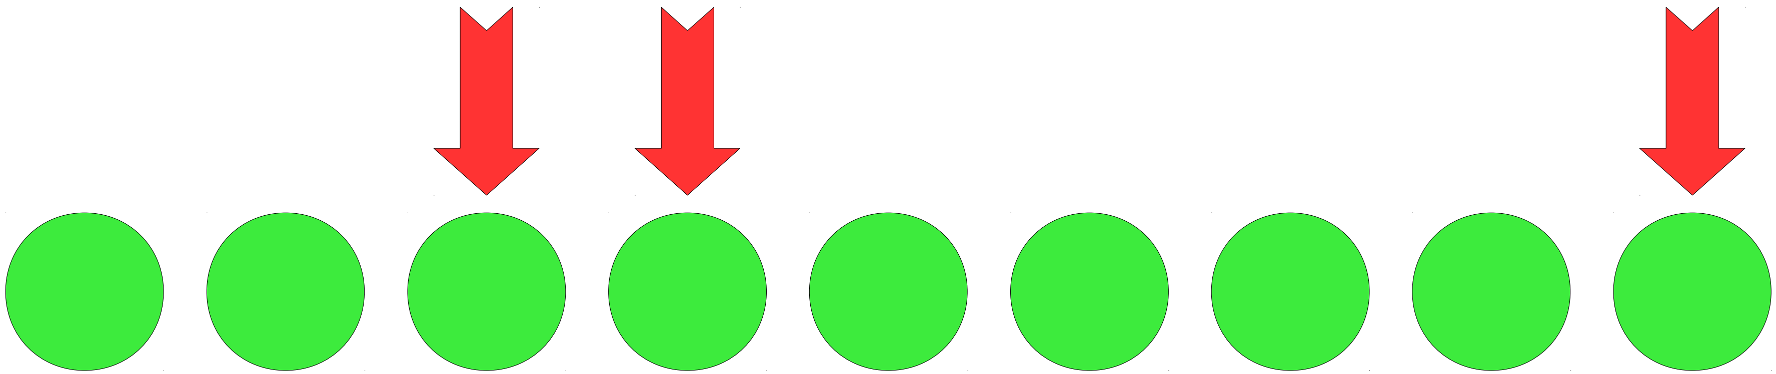

Pick three random numbers between 1 and 9

Cris Luengo
luengo@ieee.org

Systematic uniform random sampling

- Avoids having to generate so many random numbers
- More efficient than independent random sampling
  - it has been shown that fewer samples are needed to obtain the same precision

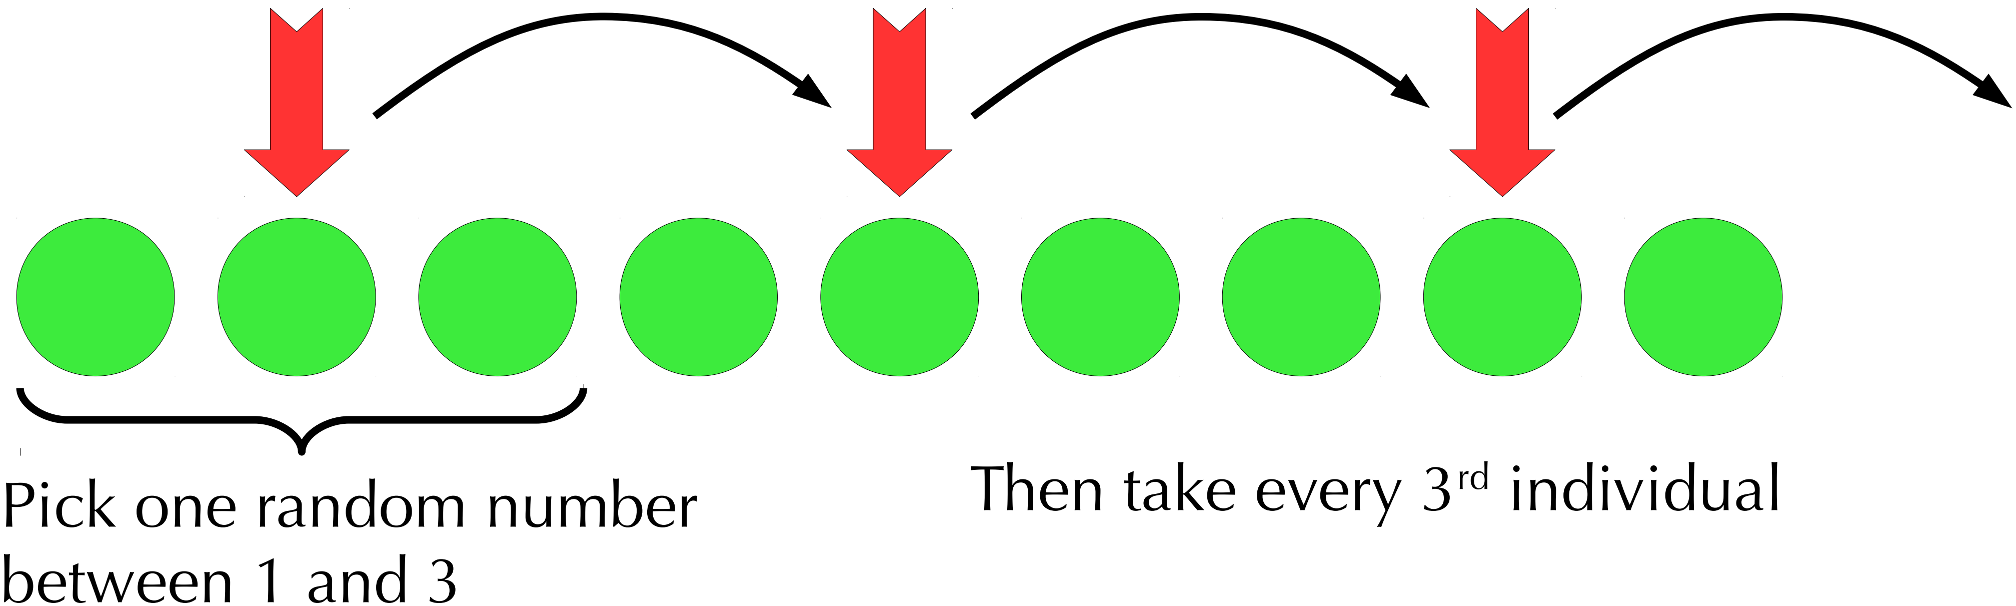

Cris Luengo
luengo@ieee.org

The fractionator principle

- Do the estimate in a fraction *f* of the population
- Compute the total by dividing by *f*

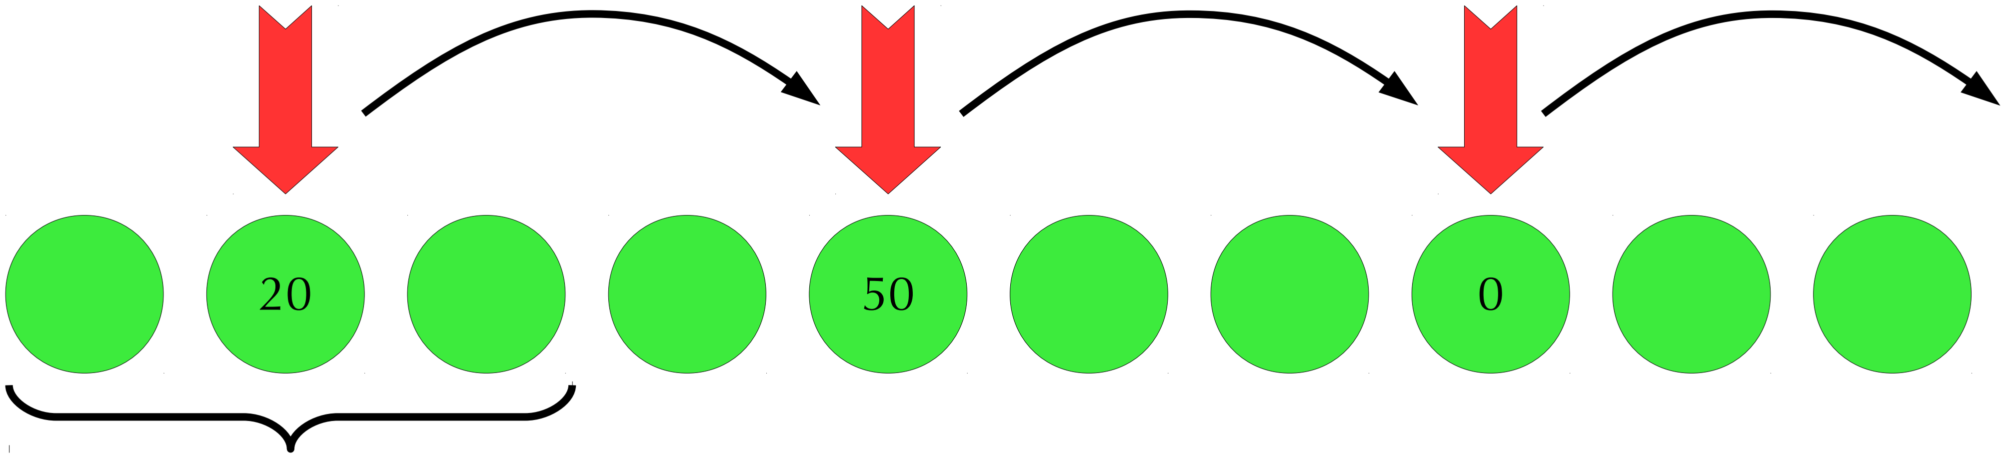

Pick one random number between 1 and 3

Then take every 3rd individual ($f = 1/3$)

Total counted = 70
Estimate for whole population = $70 \cdot 3 = 210$

Cris Luengo
luengo@ieee.org

Measuring in sections

Cris Luengo
luengo@ieee.org

Measuring in sections

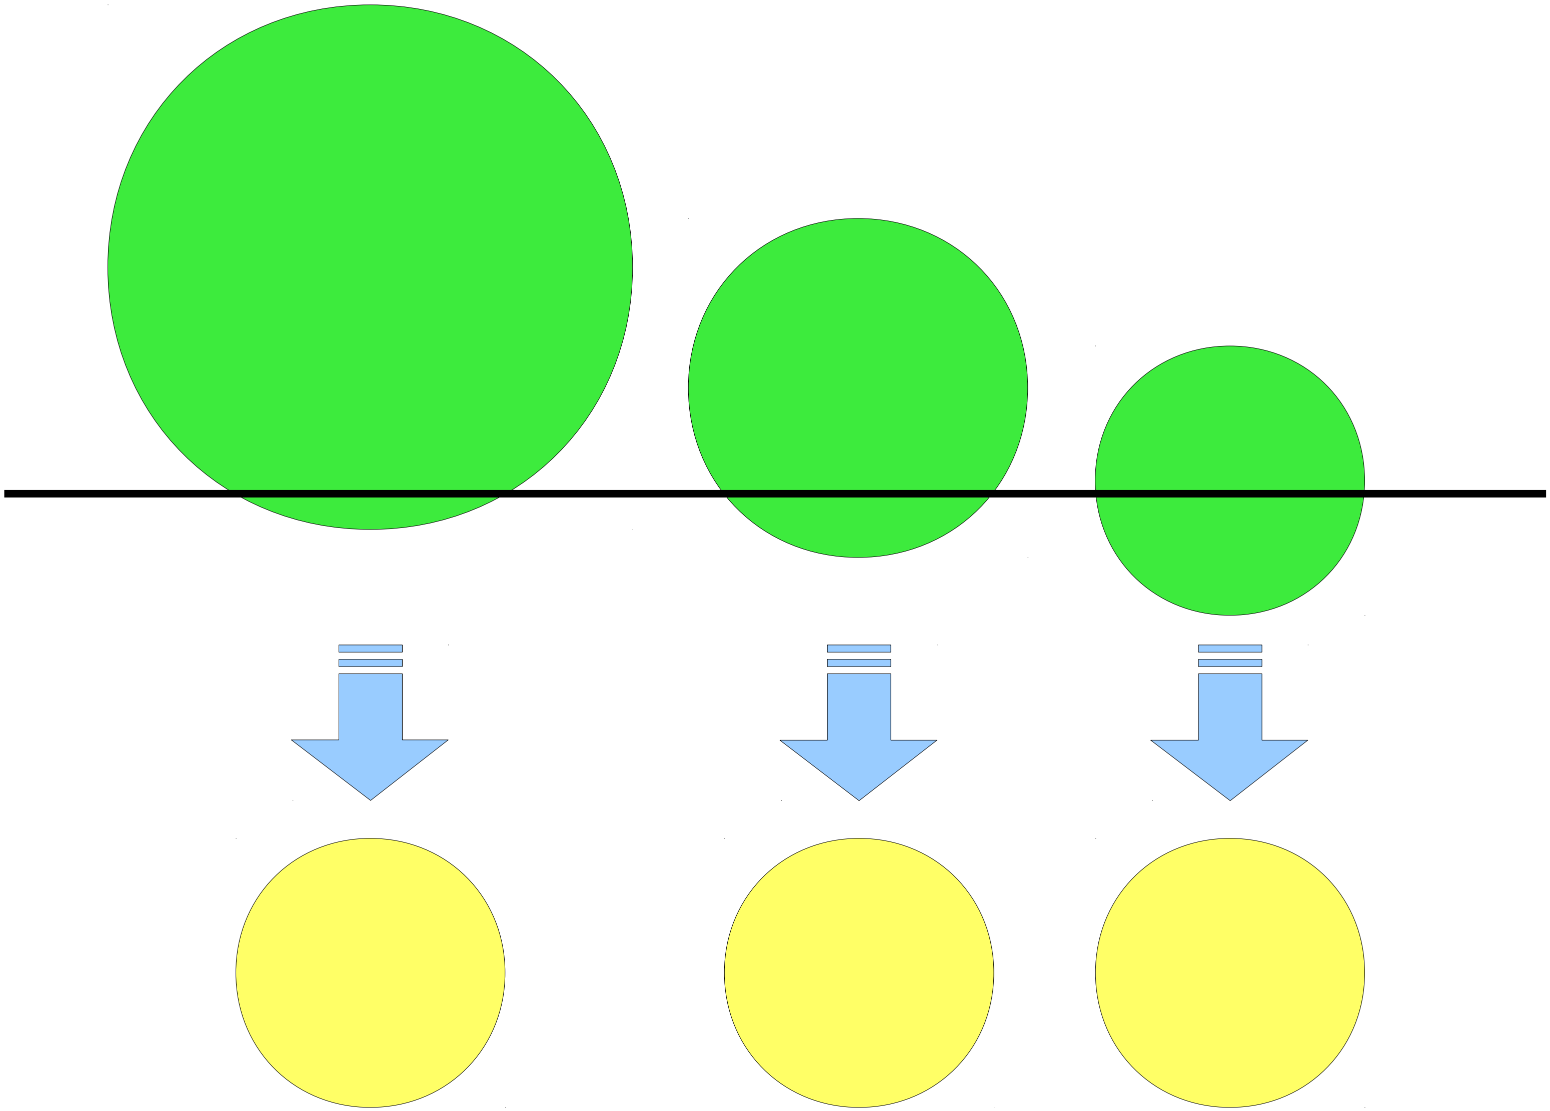

Measuring in sections

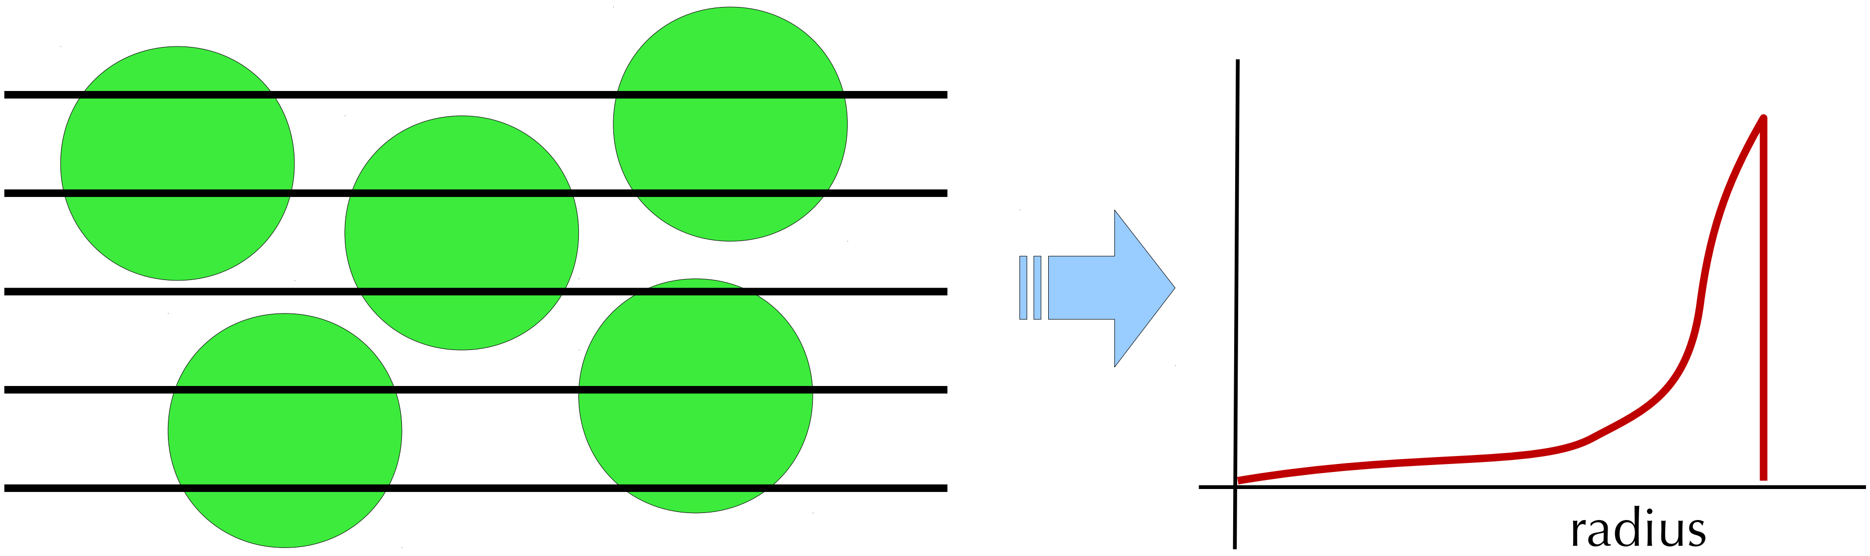

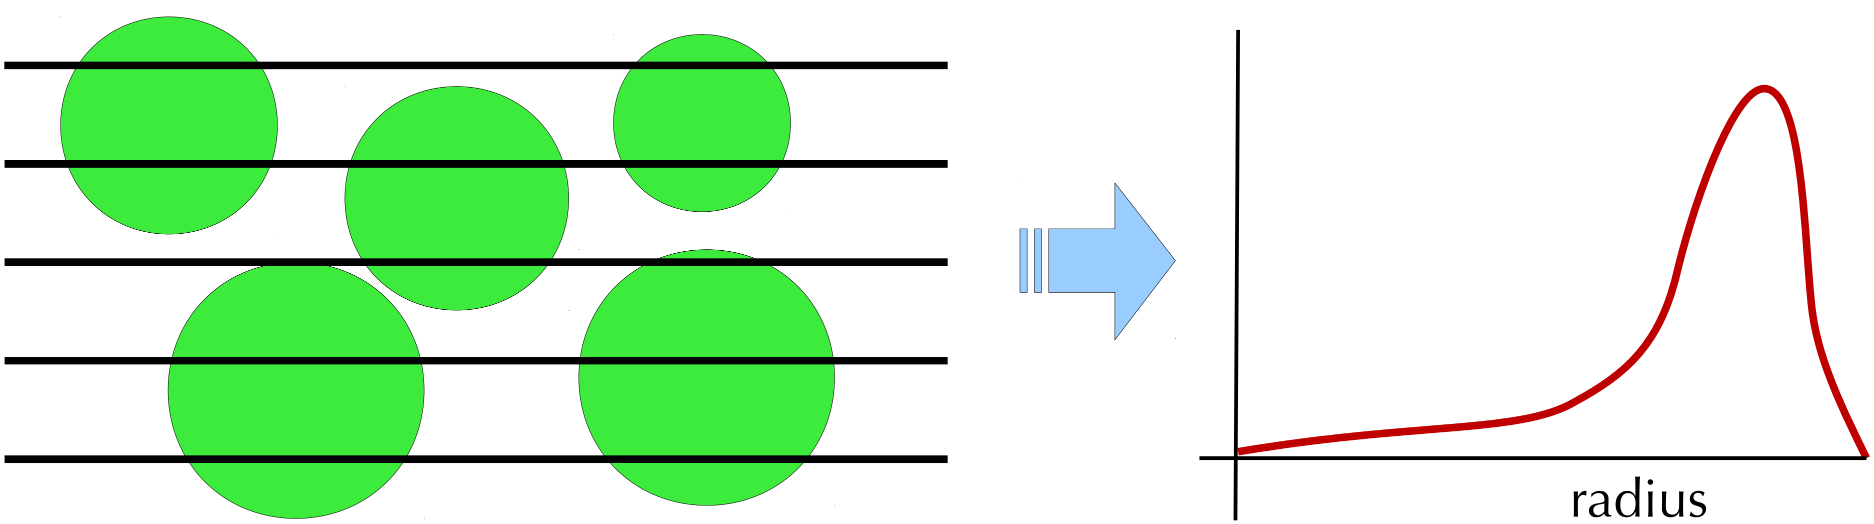

Measuring in sections

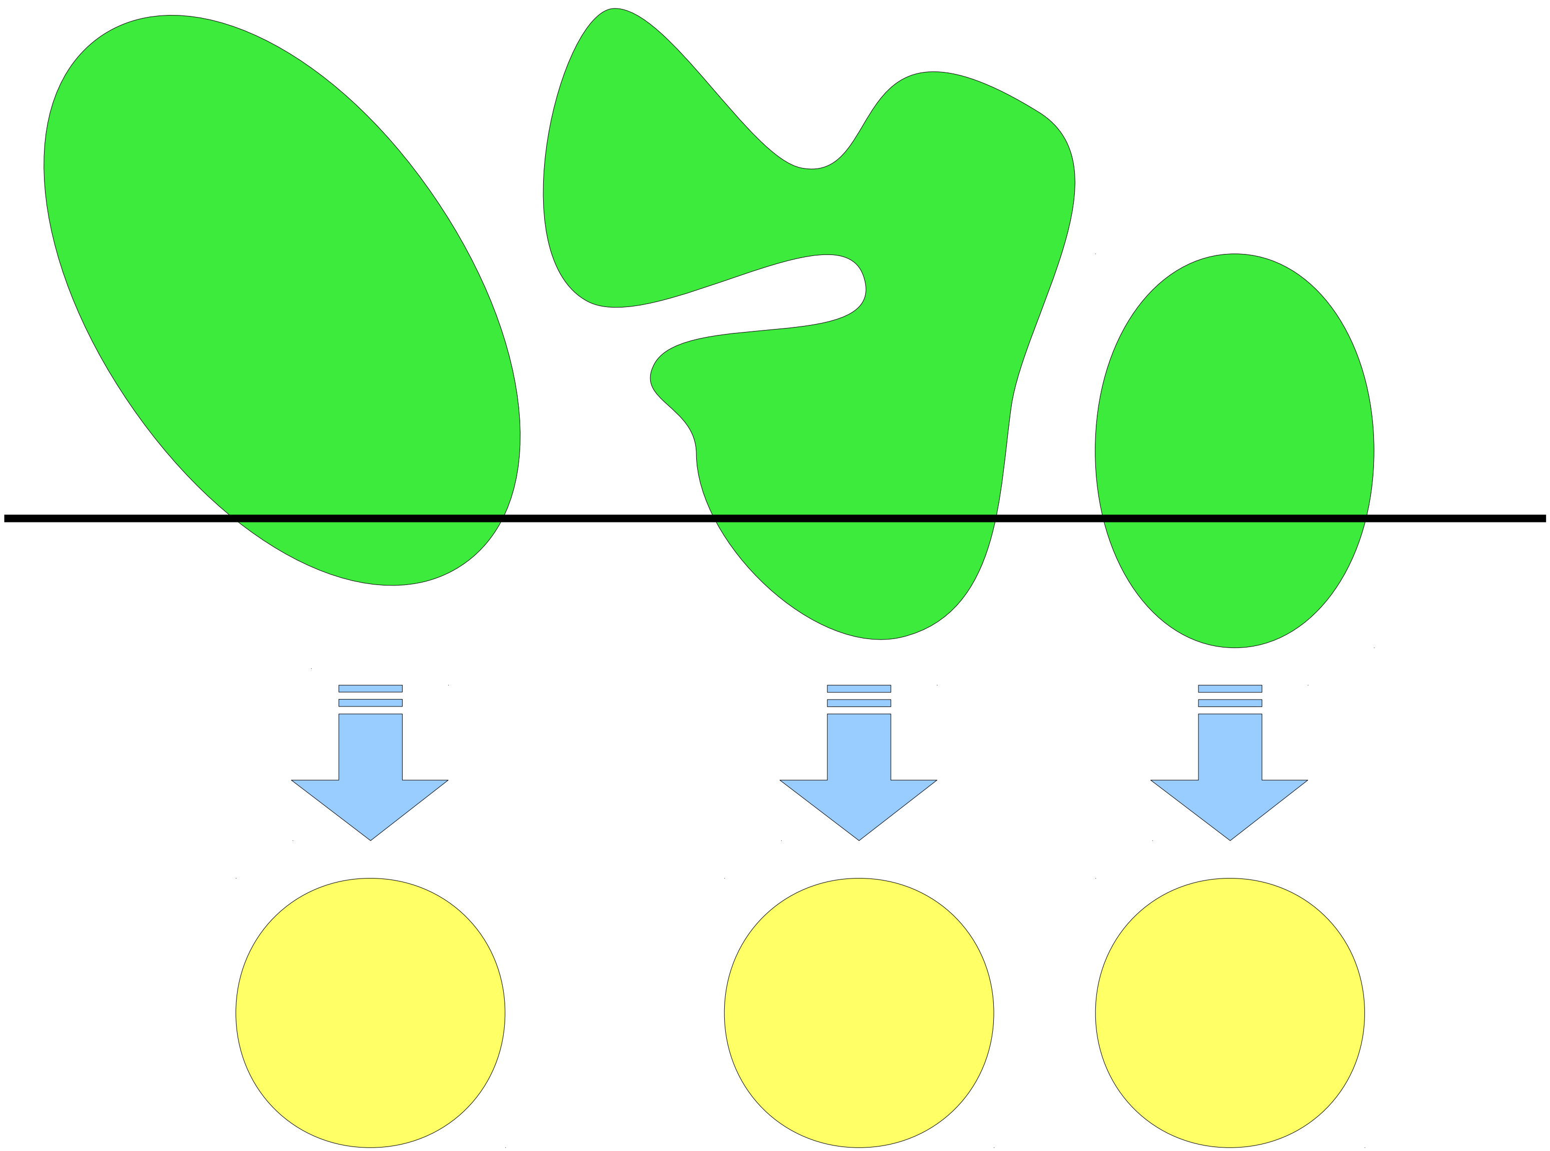

The Delesse principle

- Achille Delesse showed that the area fraction of oil in a 2D cross-section of rock is equal to the volume fraction of oil in the whole core sample (1847)

$$A_{\text{phase}} \,/\, A_{\text{ref}} = V_{\text{phase}} \,/\, V_{\text{ref}}$$

- Allows studying volume properties in 2D cross-sections

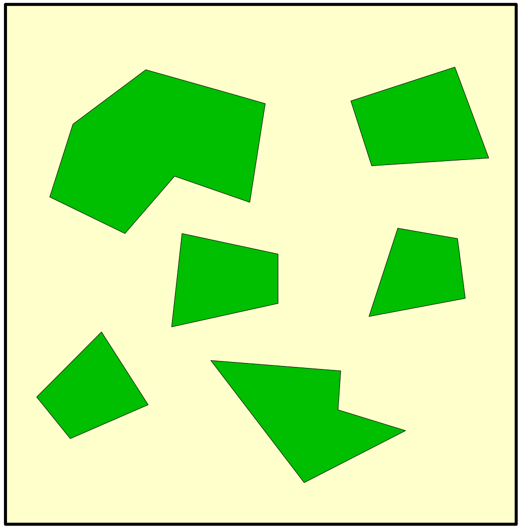

Cris Luengo
luengo@ieee.org

The Delesse principle

- He also showed that the number of objects in a 2D cross section is unrelated to the number in 3D
- The Delesse principle scales to lines and point counting:

$$V_{\mathrm{phase}} / V_{\mathrm{ref}} = A_{\mathrm{phase}} / A_{\mathrm{ref}} = L_{\mathrm{phase}} / L_{\mathrm{ref}} = N_{\mathrm{phase}} / N_{\mathrm{ref}}$$

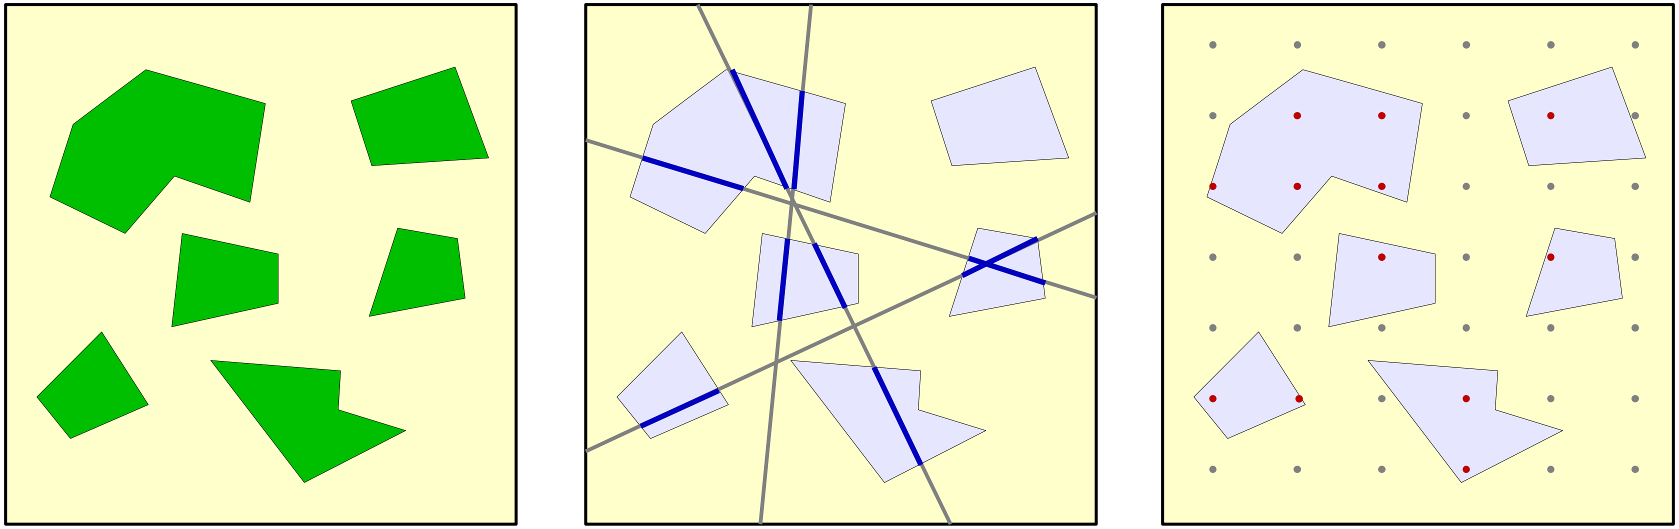

Delesse (1847) Rosival (1898) Glagolev (1933)

Stereological approach to measurement

Buffon’s needle problem

- Georges-Louis Leclerc, Comte de Buffon (1733):
  - What is the chance that a needle randomly tossed into the air will fall across the lines on a parquet floor?
  - His solution: The needle will intersect with a probability that is directly proportional to the length of the needle and inversely proportional to the distance between the lines on the floor.
  - Full solution: P = $(2/\pi)$ length / distance
- By tossing a needle of known length 200 times, you can estimate the distance between lines
  - distance = $(2/\pi)$ 4 cm / (50/200) = 0.63 · 4 · 4 cm = 10 cm

Buffon's needle problem

- The interesting bits:
  - This relation holds whether you toss 1 needle 200 times, or 200 needles 1 time.
  - This relation makes no assumption about the shape of the needle or the parquet lines

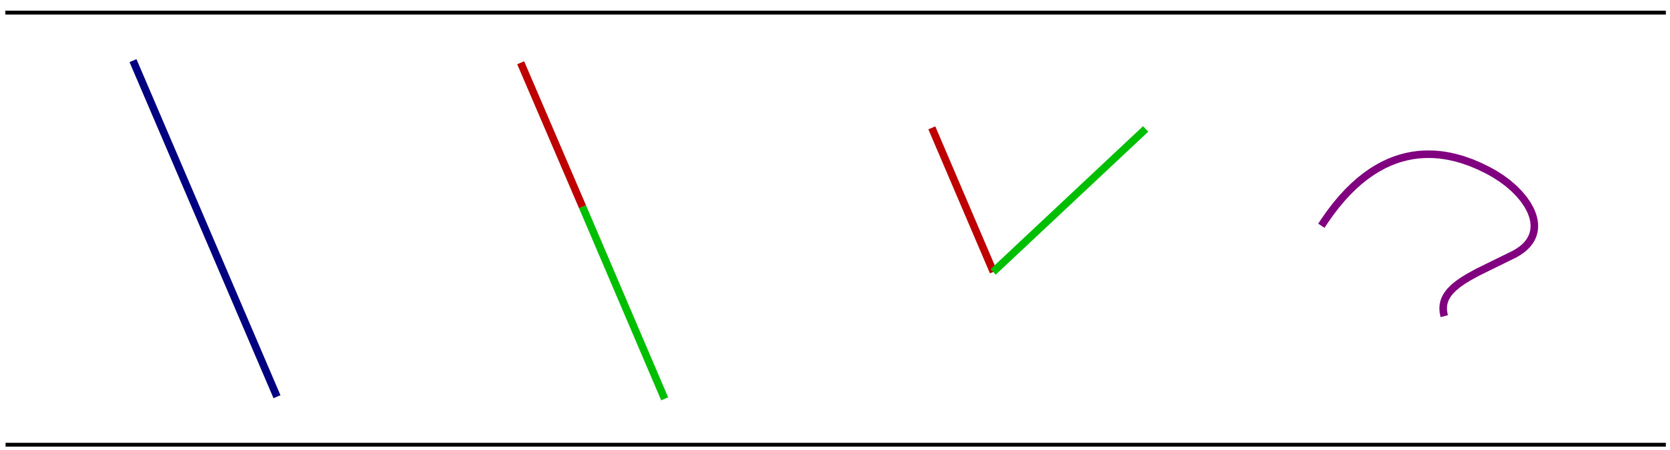

P $\quad$ $P = P + P$ $\quad$ $P = P + P$ $\quad$ $P = P$

Cris Luengo
luengo@ieee.org

Buffon's needle problem

- Thus: this relation can be used to estimate the perimeter of any object
  - # of intersections = $(2/\pi)$ total boundary length / distance between probe lines
- Randomness is important for the method to be unbiased
  - probe orientation needs to be random
    (from a uniform orientation distribution)

Cris Luengo
luengo@ieee.org

The point-counting method

- Estimate first order stereological parameters by counting
  - volume
  - surface area
  - length
  - number
- A probe intersects the feature to be measured at a point
- Every point on the feature must be equally likely to intersect the probe
- Count 100-200 events for each individual

Cris Luengo
luengo@ieee.org

Probe-based measurement

- The intersection is a point when
  - Probe dimensions + feature dimensions = 3
- When the intersection is a line or an area or a volume, other measurement techniques are needed
  - Probe dimensions + feature dimensions $\geq$ 3
- When the inequality is not satisfied: bias!

probe	**2D reference**	**3D reference**
point (0)	area (2)	volume (3)
line (1)	length (1)	surface area (2)
area / plane (2)	number (0)	length (1)
volume (3)		number (0)

Cris Luengo
luengo@ieee.org

Estimating volume (or area in a 2D section)

Cris Luengo
luengo@ieee.org

Estimating volume

- The probability that a random point within the reference volume hits the object is equal to the volume of the object divided by the reference volume

$$P = V_{obj} \,/\, V_{ref}$$

- Throw many points at the volume, count the number of points that hit the object

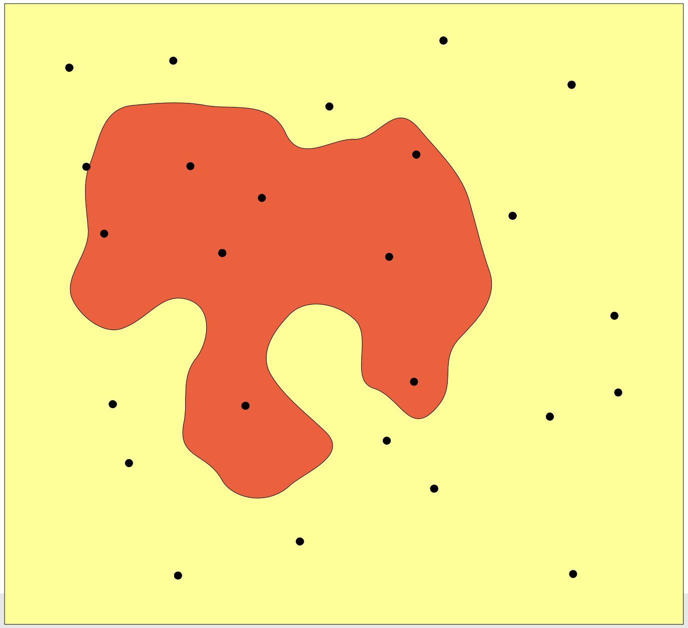

$$P \approx 8 \,/\, 25 = 0.32$$
$$V_{obj} = P\, V_{ref} \approx 0.32 \cdot 1\ \text{mm}^2$$
$$V_{obj} \approx 0.32\ \text{mm}^2$$

Cris Luengo
luengo@ieee.org

Estimating volume

- Instead use systematic uniform random sampling
- This leads to sampling grids, paced *randomly* over image: isotropic-uniform-random probes

$P \approx 27 / 119 = 0.227$

$V_{obj} = P\, V_{ref} \approx 0.227 \cdot 1\ \text{mm}^2$

$V_{obj} \approx 0.227\ \text{mm}^2$

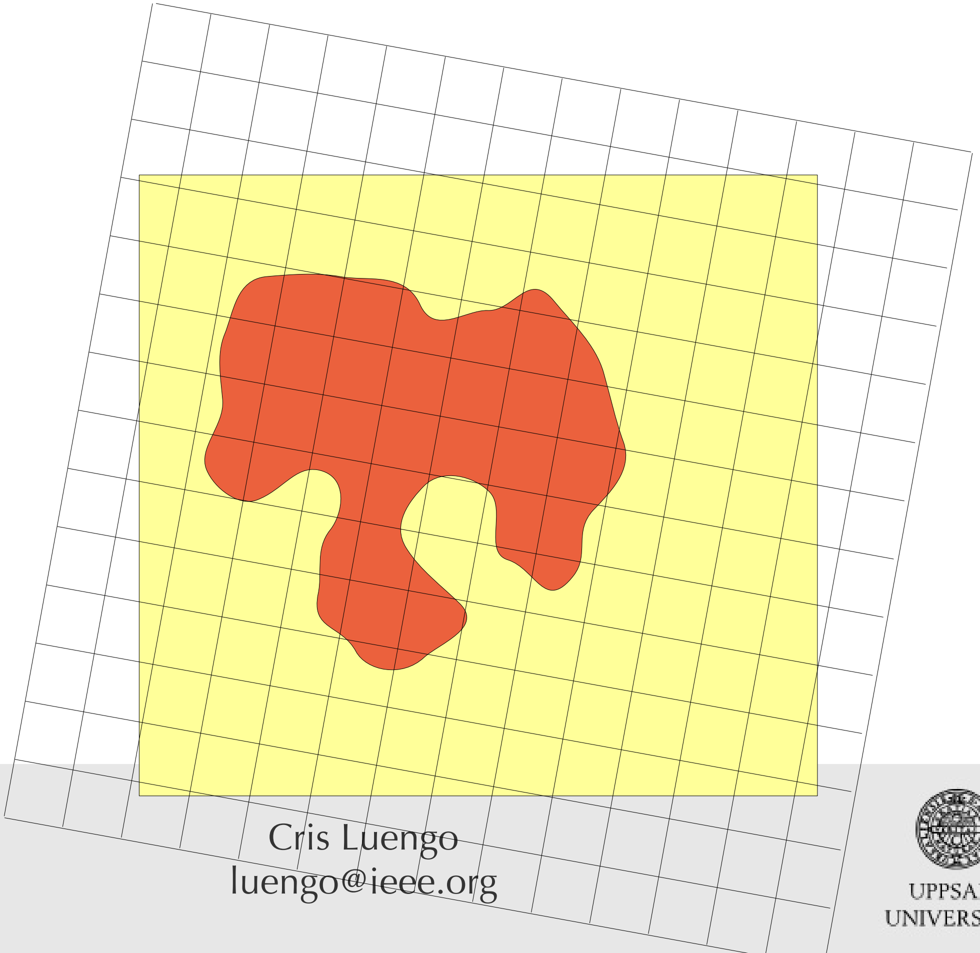

The Cavalieri principle

- Bonaventura Cavalieri showed that the volume of an arbitrarily shaped object can be estimated by serial sectioning (1635):
  - the first section must be random
  - subsequent sections at consistent intervals (distance T)
  - estimate the total area in each section (A_i)
  - volume is given by $V = T \, \Sigma \, A_i$

systematic uniform random sampling!

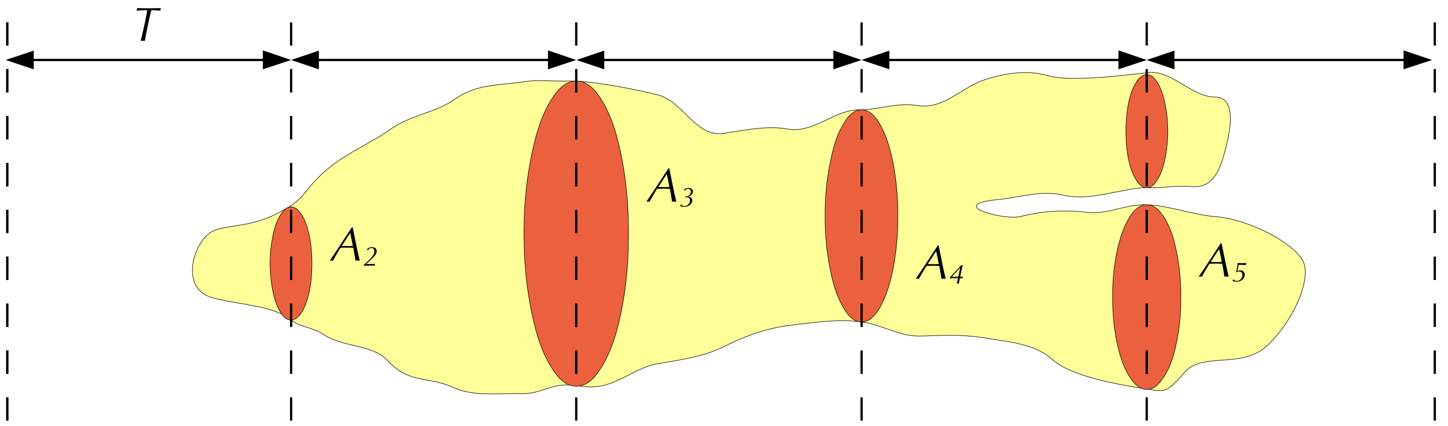

The Cavalieri principle

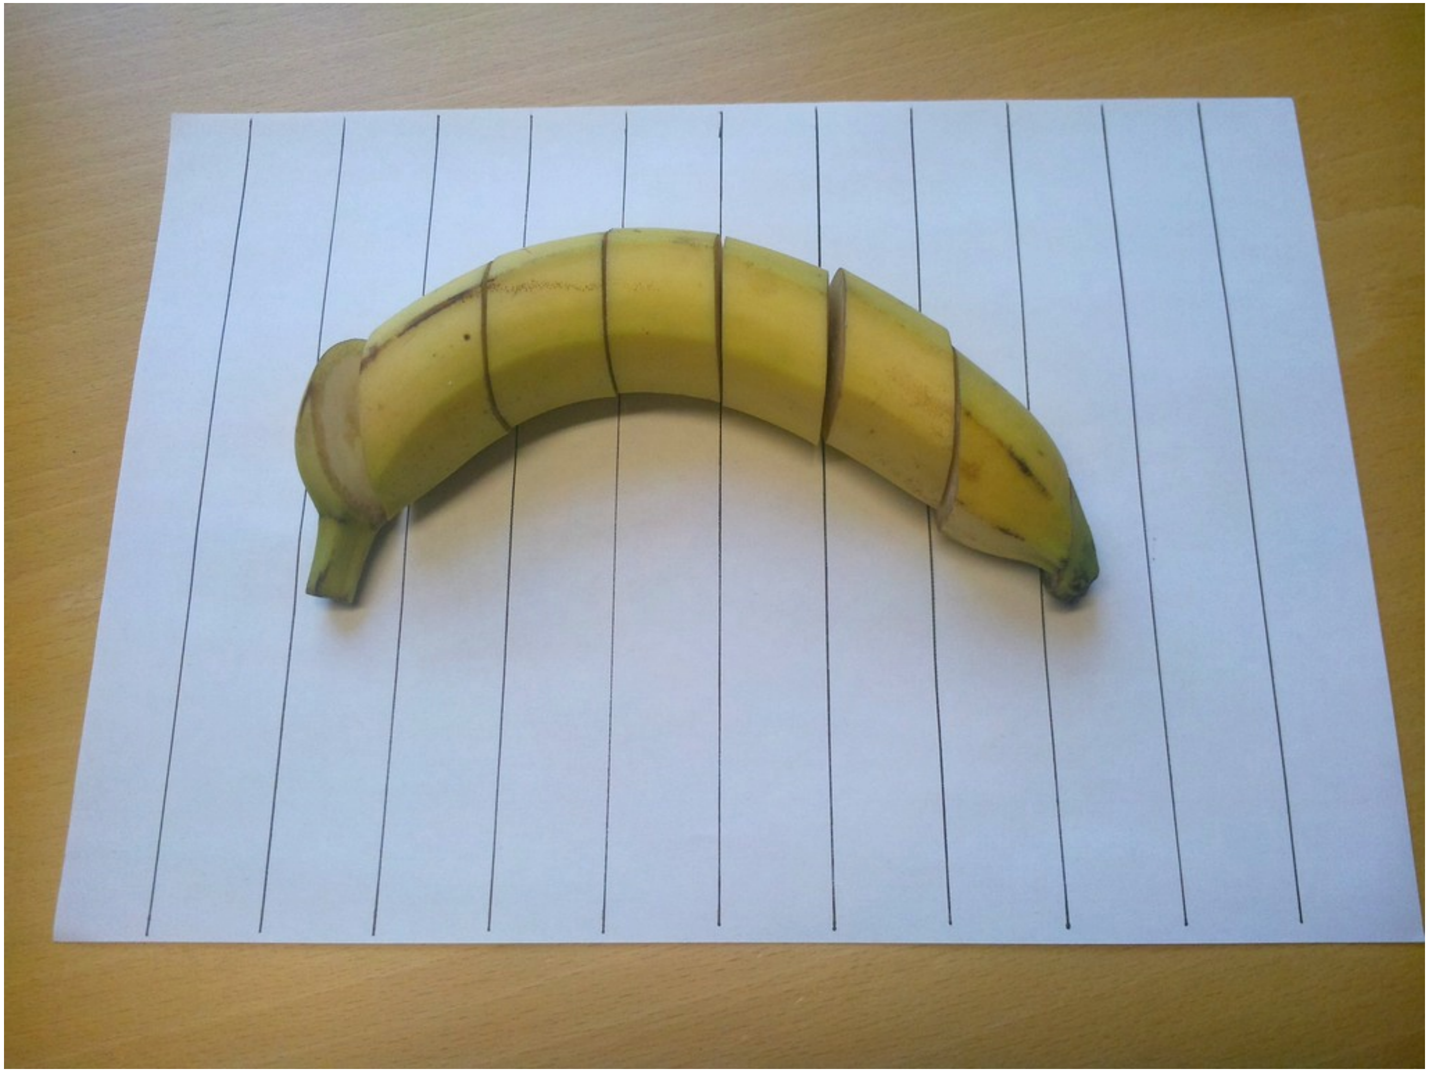

The Cavalieri principle

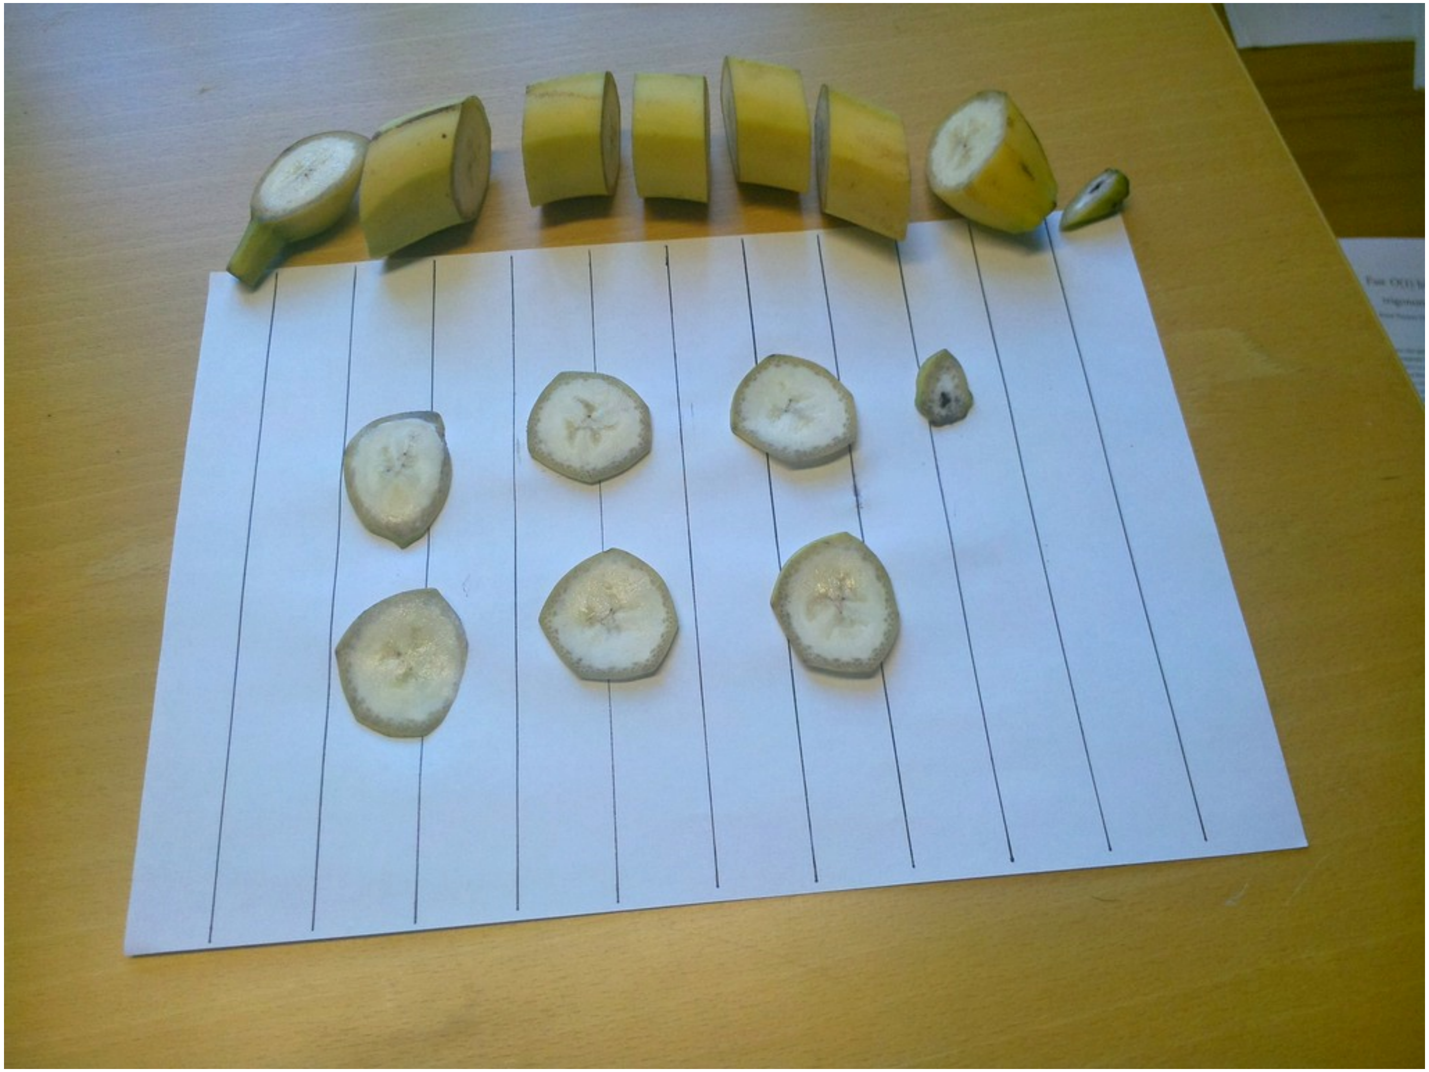

The Cavalieri principle

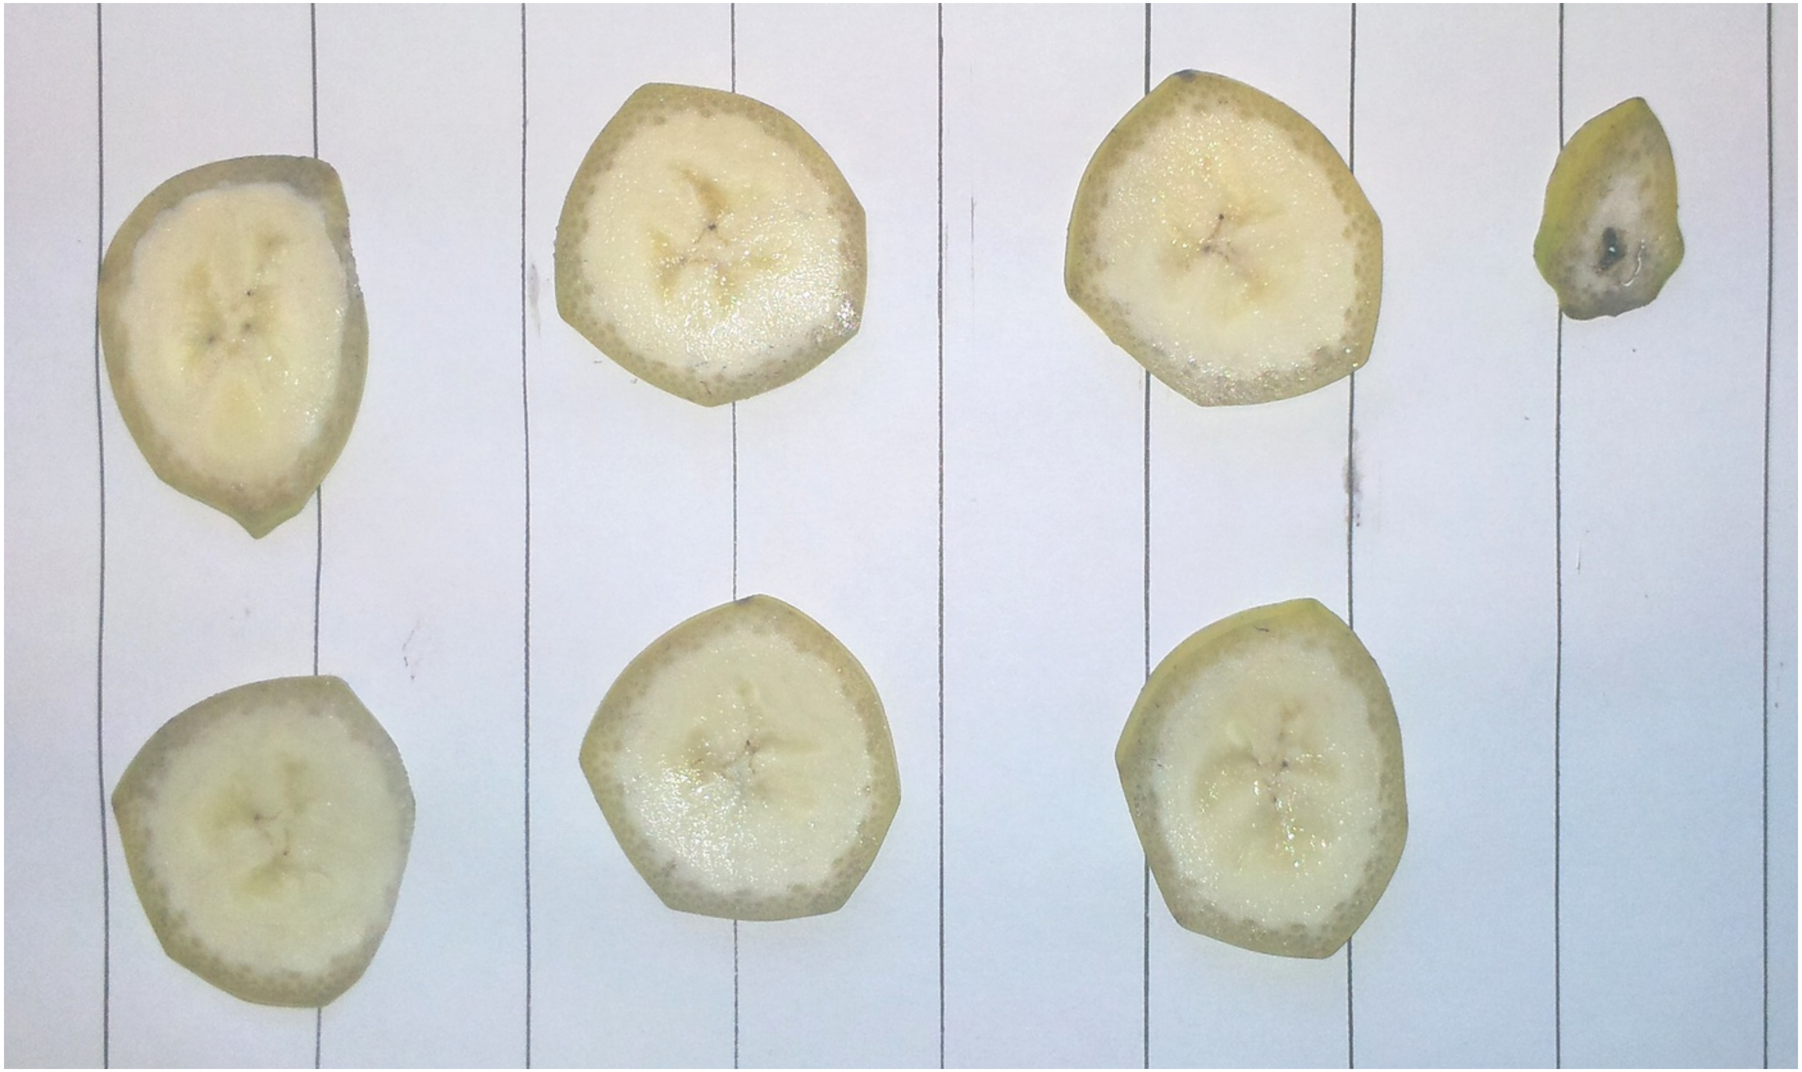

The Cavalieri principle

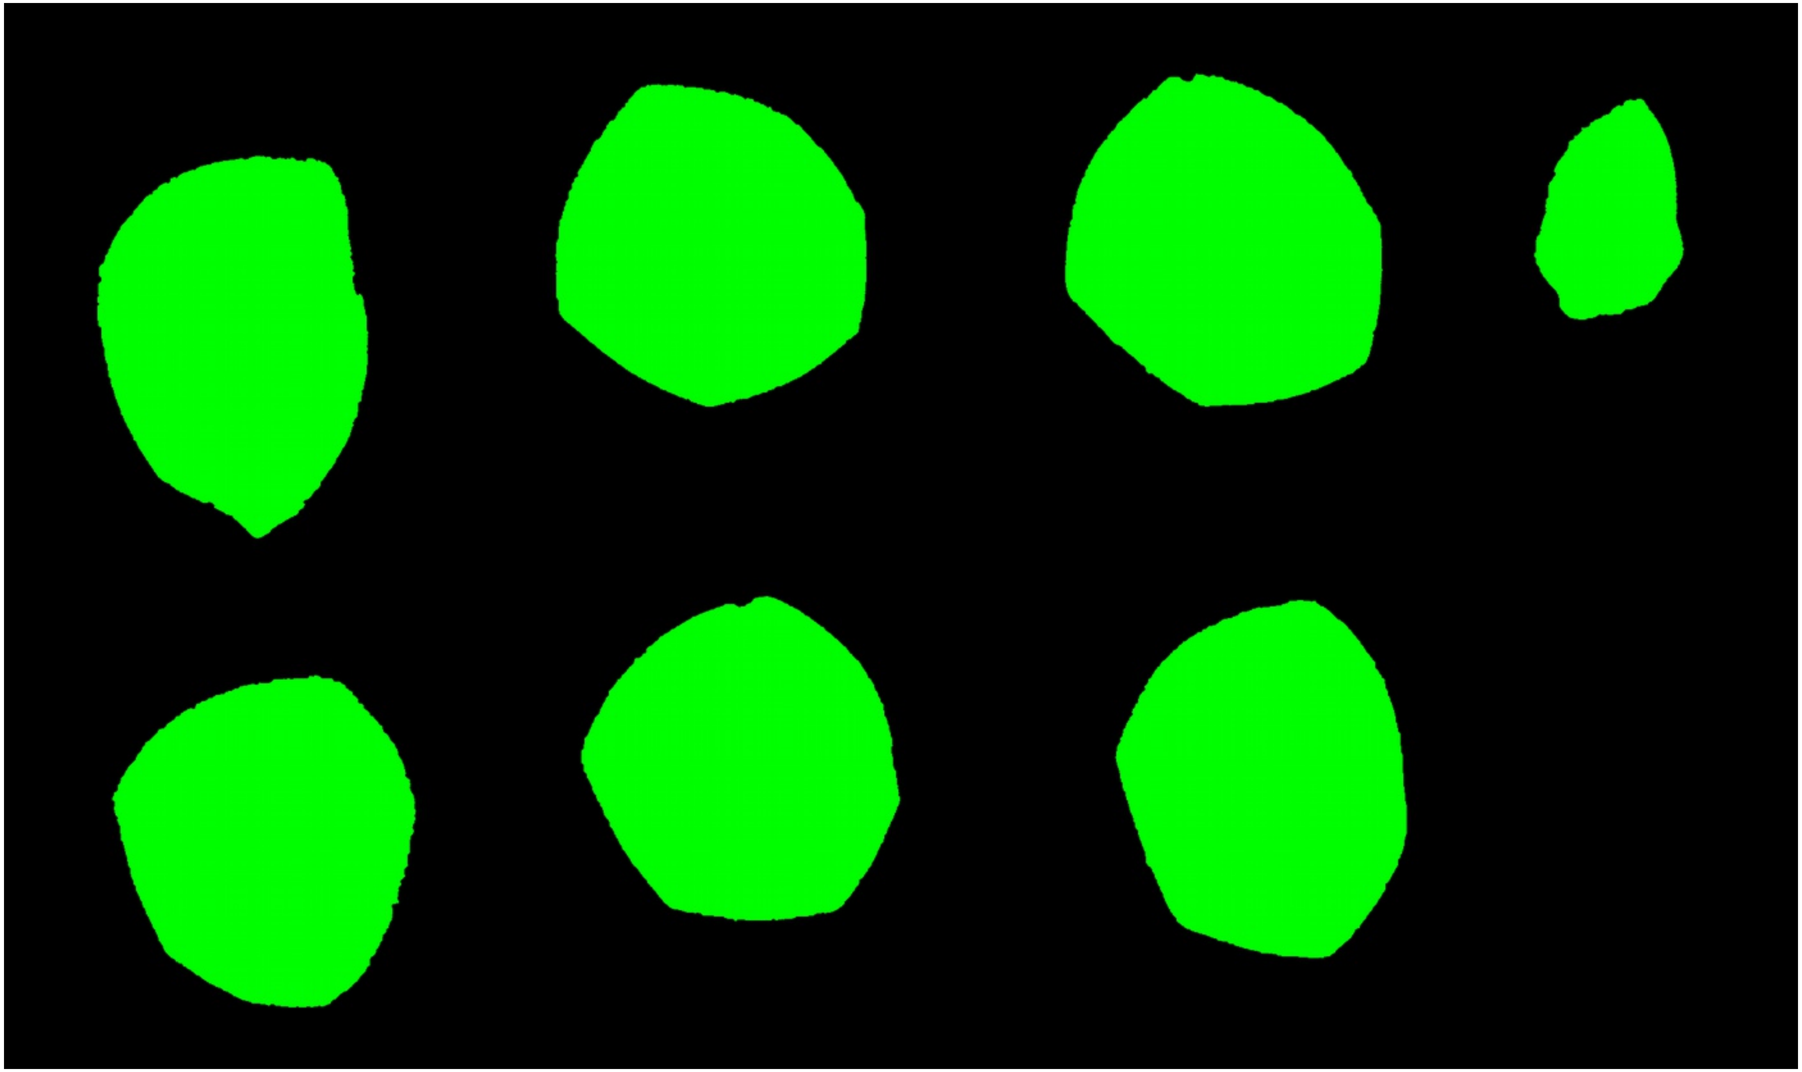

Cris Luengo
luengo@ieee.org

The Cavalieri principle

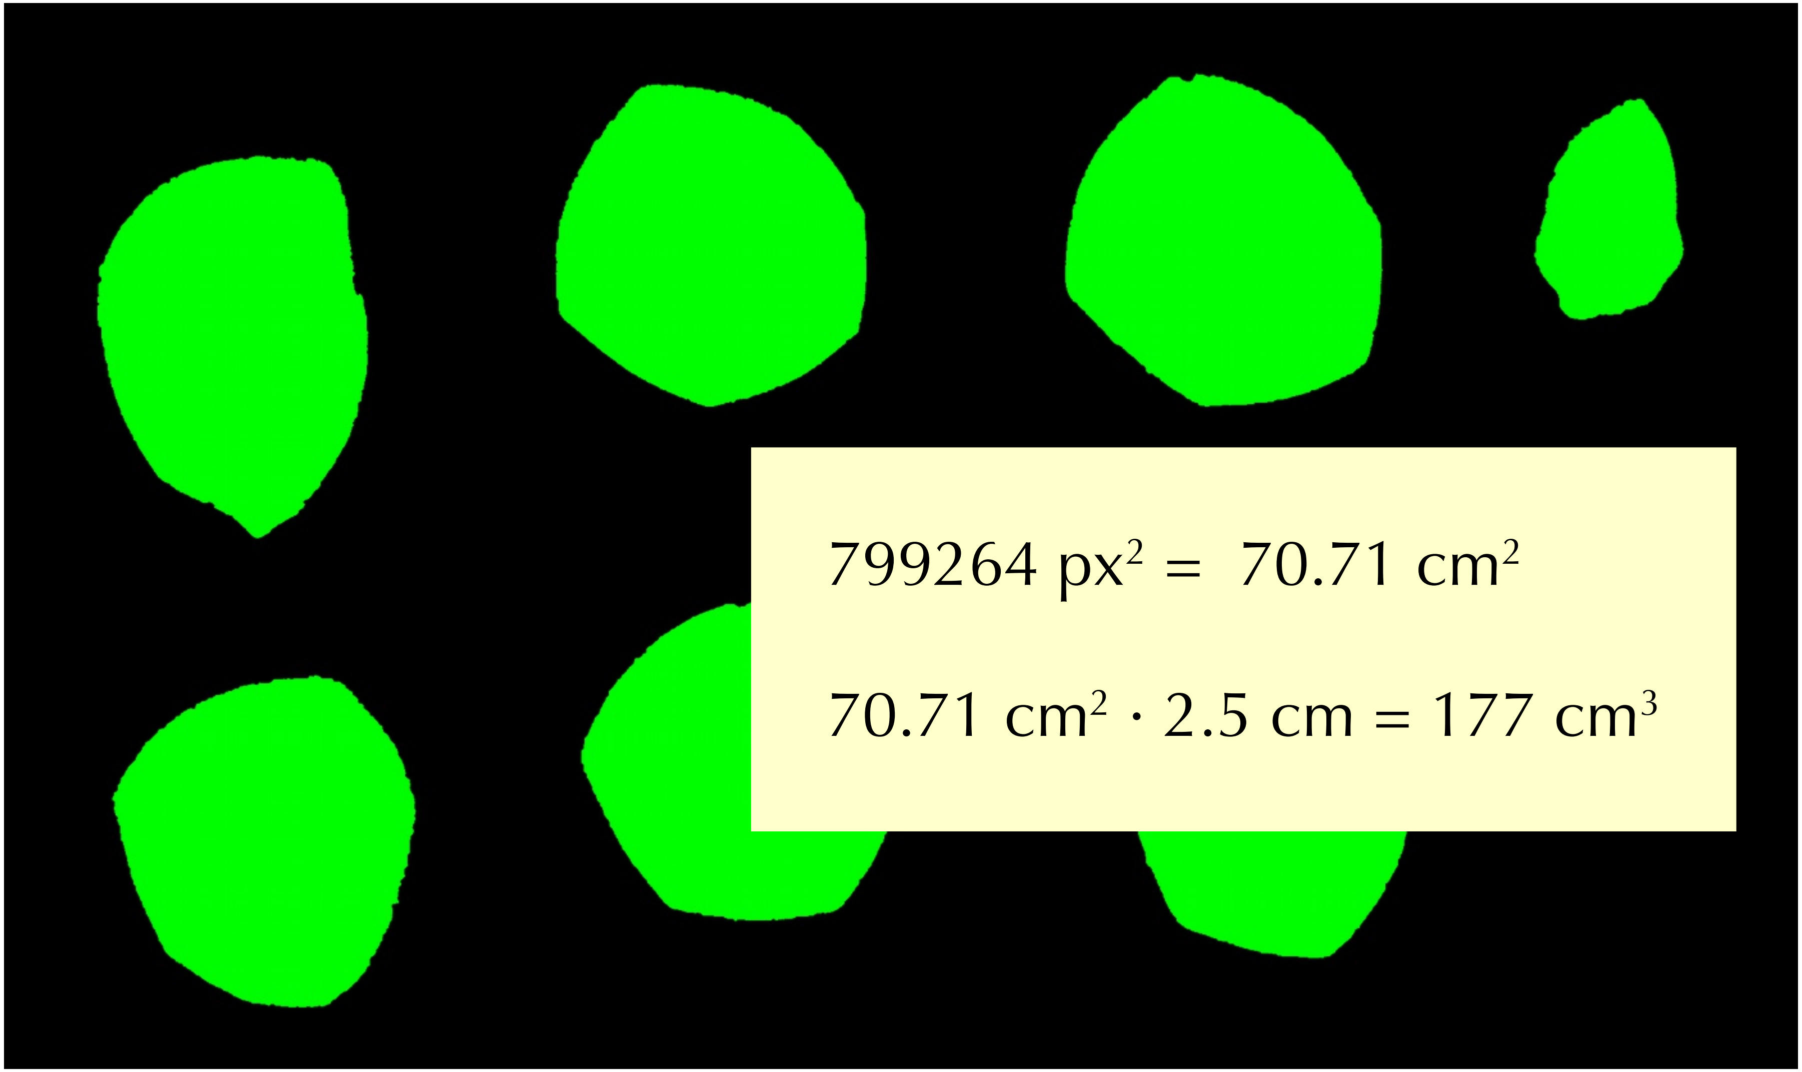

$$799264 \text{ px}^2 = 70.71 \text{ cm}^2$$

$$70.71 \text{ cm}^2 \cdot 2.5 \text{ cm} = 177 \text{ cm}^3$$

Cris Luengo
luengo@ieee.org

The Cavalieri principle

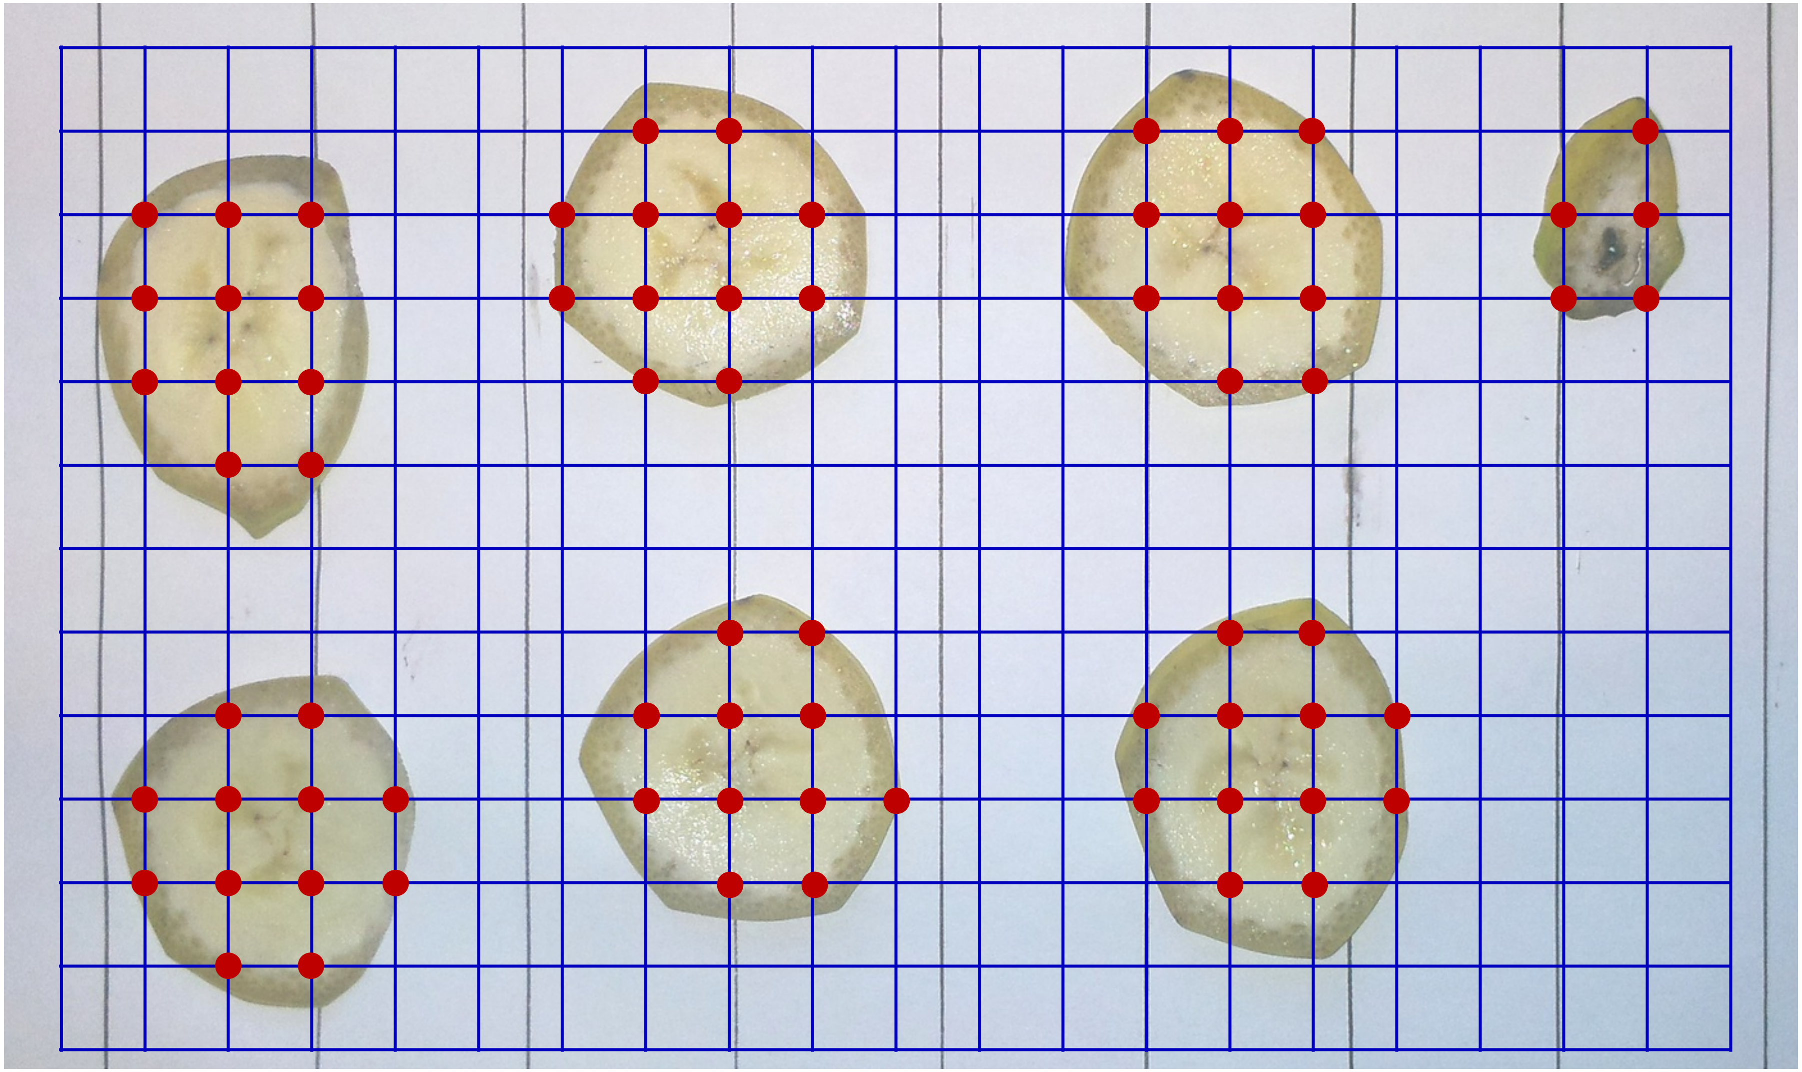

The Cavalieri principle

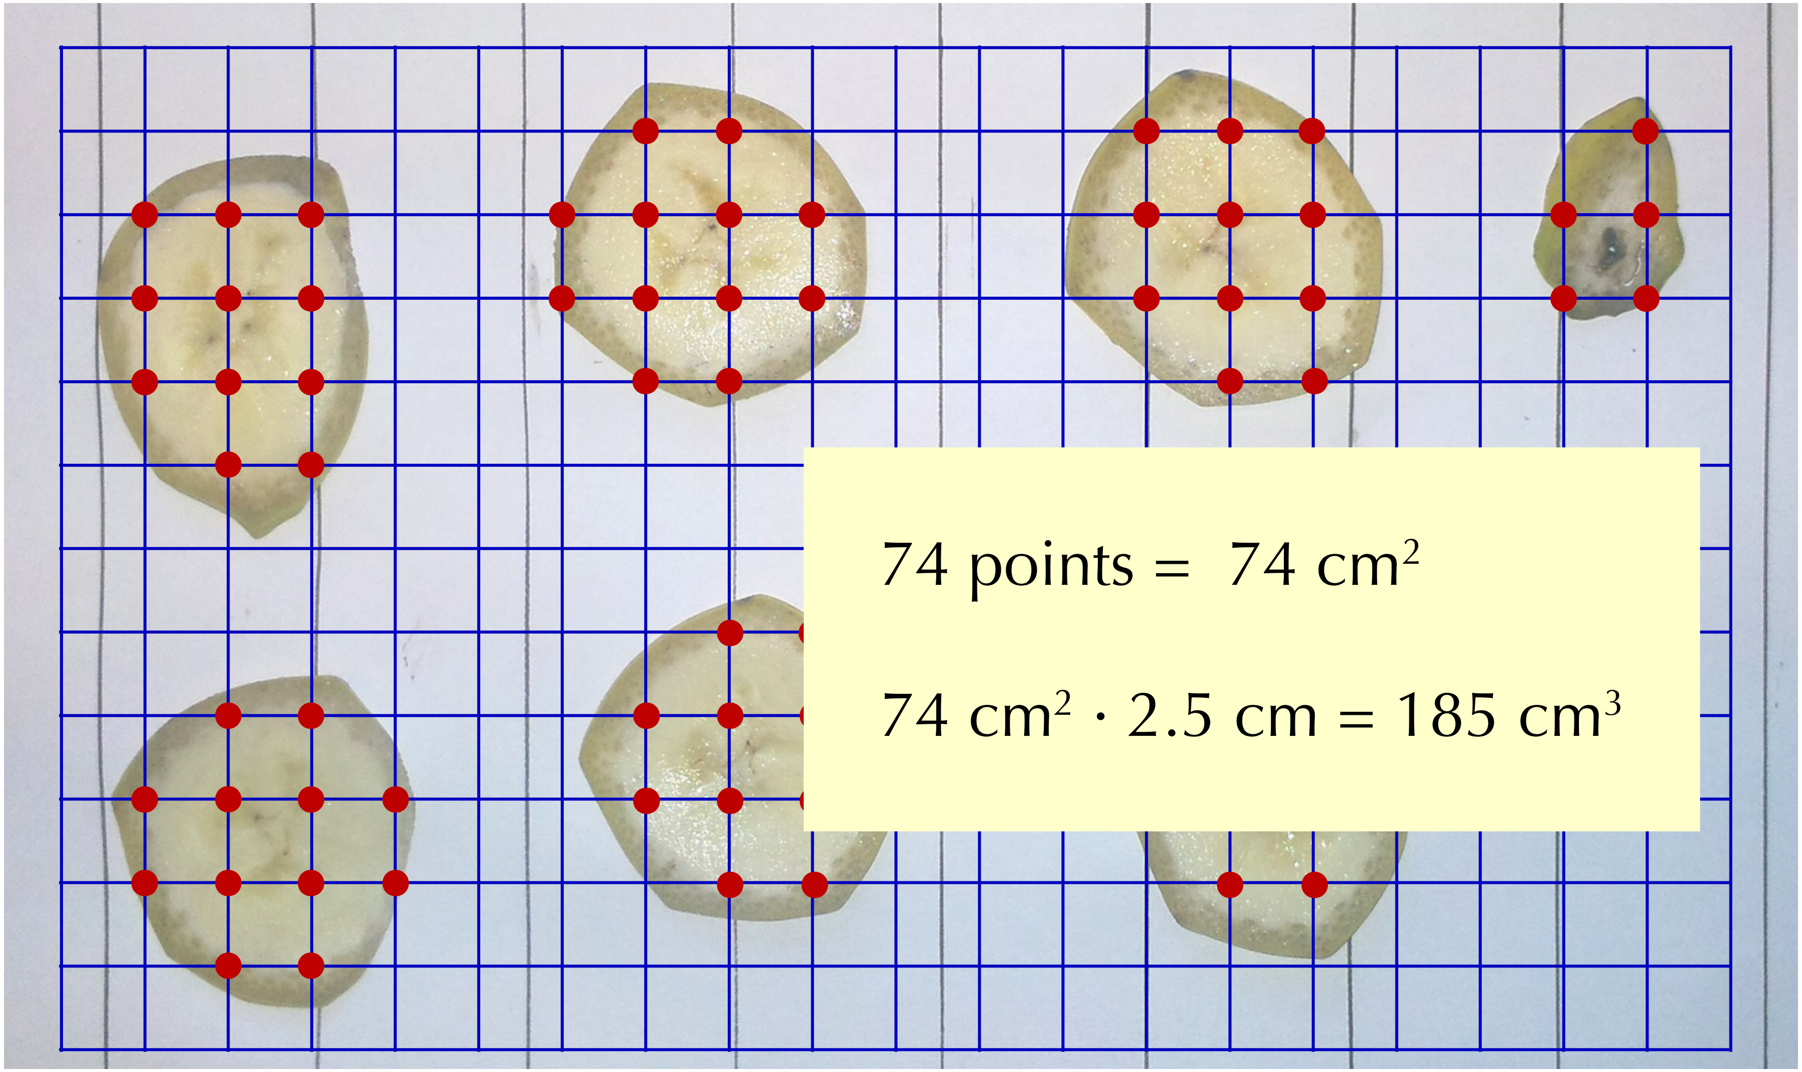

Cris Luengo
luengo@ieee.org

Estimating surface area
(or perimeter in a 2D section)

Cris Luengo
luengo@ieee.org

Boundary length

Benoit Mandelbrot showed why perimeter measurements are scale-dependent (late 1970's)

"Nature exhibits not simply a higher degree but an altogether different level of complexity. The number of distinct scales of length of natural patterns is for all purposes infinite. That is, the closer one looks, the more biological surface is present." (Mandelbrot, 1983)

"The infinite coastline of Britain"

Boundary length / surface area

- Translate the parquet lines from Buffon's problem to line density:
  - Lines of width of reference area, distance D apart
  - Reference area is W x H in size
  - H / D lines of length W give $L_{\text{total}} = W H / D$ total length
  - $L_{\text{A}} = 1 / D$ is the perimeter density of the parquet
  - A similar concept is the surface density $S_{\text{V}} = S / V_{\text{ref}}$
- We count number of intersections N with line probe of length L_{probe}
- We compute: $L_{\text{A}} = {}^{\pi}\!/_{2}\, N / L_{\text{probe}}$ $\qquad S_{\text{V}} = 2\, N / L_{\text{probe}}$

Cris Luengo
luengo@ieee.org

Boundary length / surface area

Example:

- µCT image of 1 mm across: $V_{ref} = 10^9$ µm^3
- Line probe of length 25 µm
- 1000 random probe throws yield 125 surface intersections
- $S_V = 2\ N\ /\ L_{probe} = 2 \cdot 125\ /\ (1000 \cdot 25) = 0.01$ µm^{-1}
- $S = S_V\ /\ V_{ref} = 0.01 \cdot 10^9 = 10 \cdot 10^6$ µm^2

Cris Luengo
luengo@ieee.org

Surface area

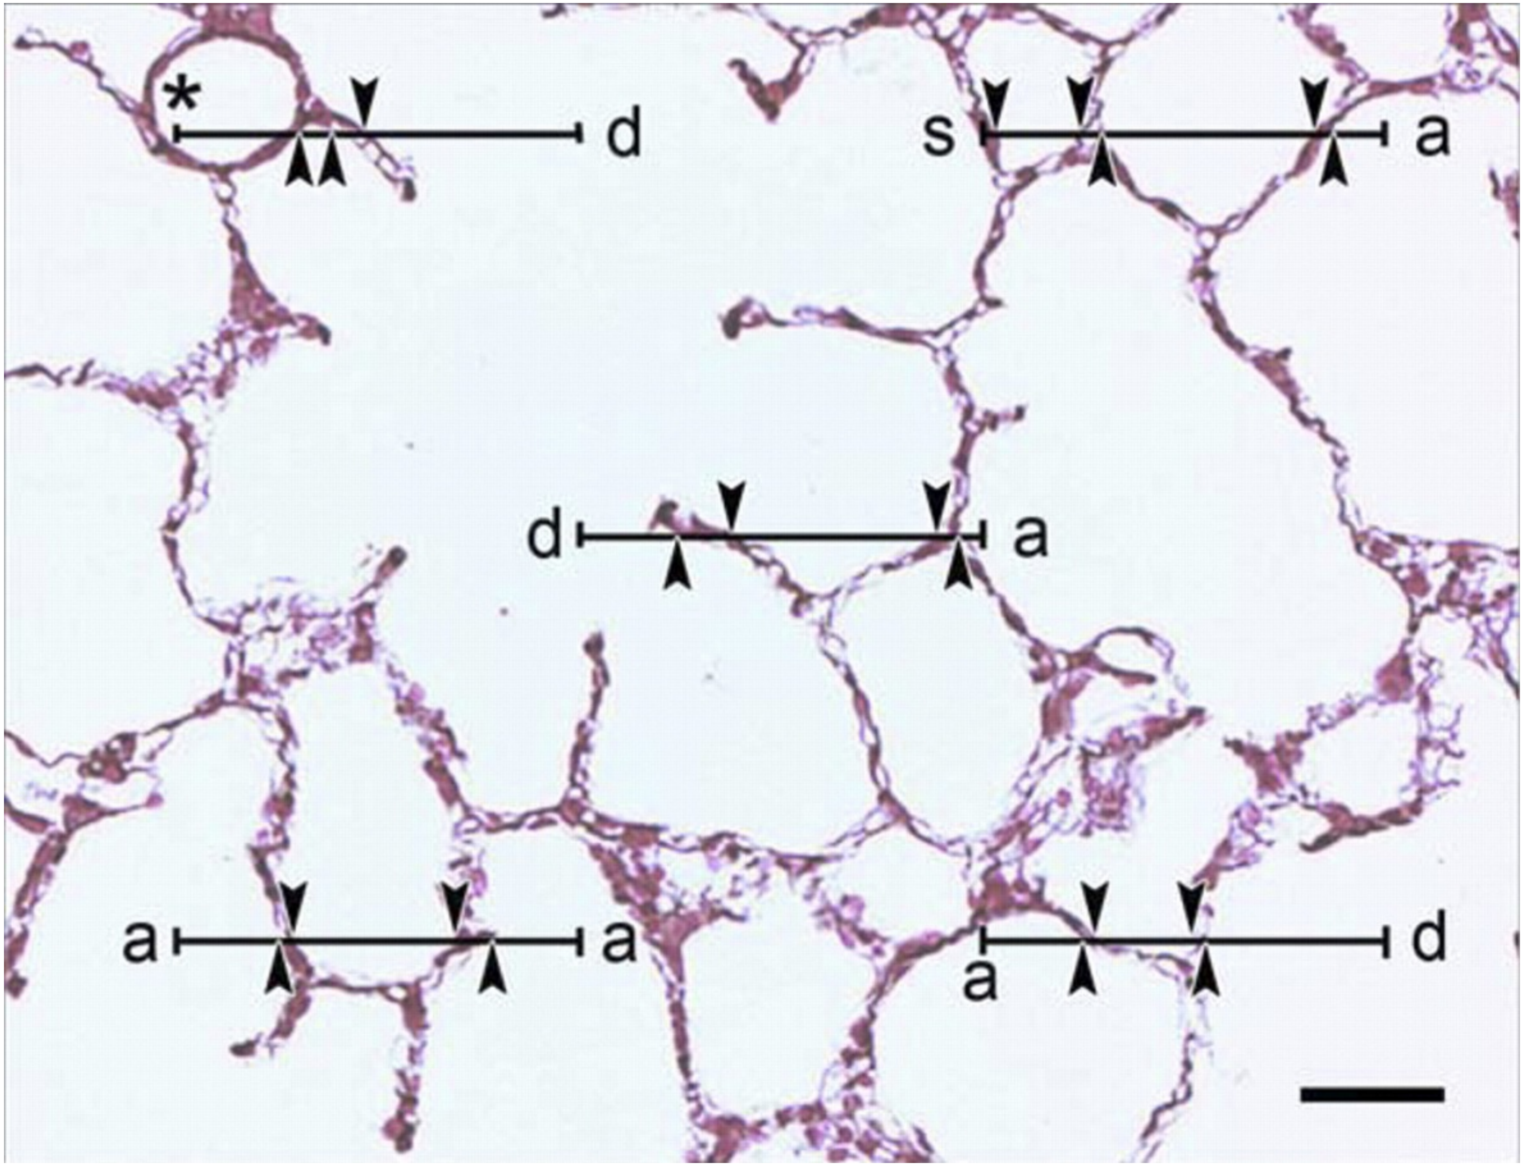

Cris Luengo
cris@cb.uu.se

Surface area

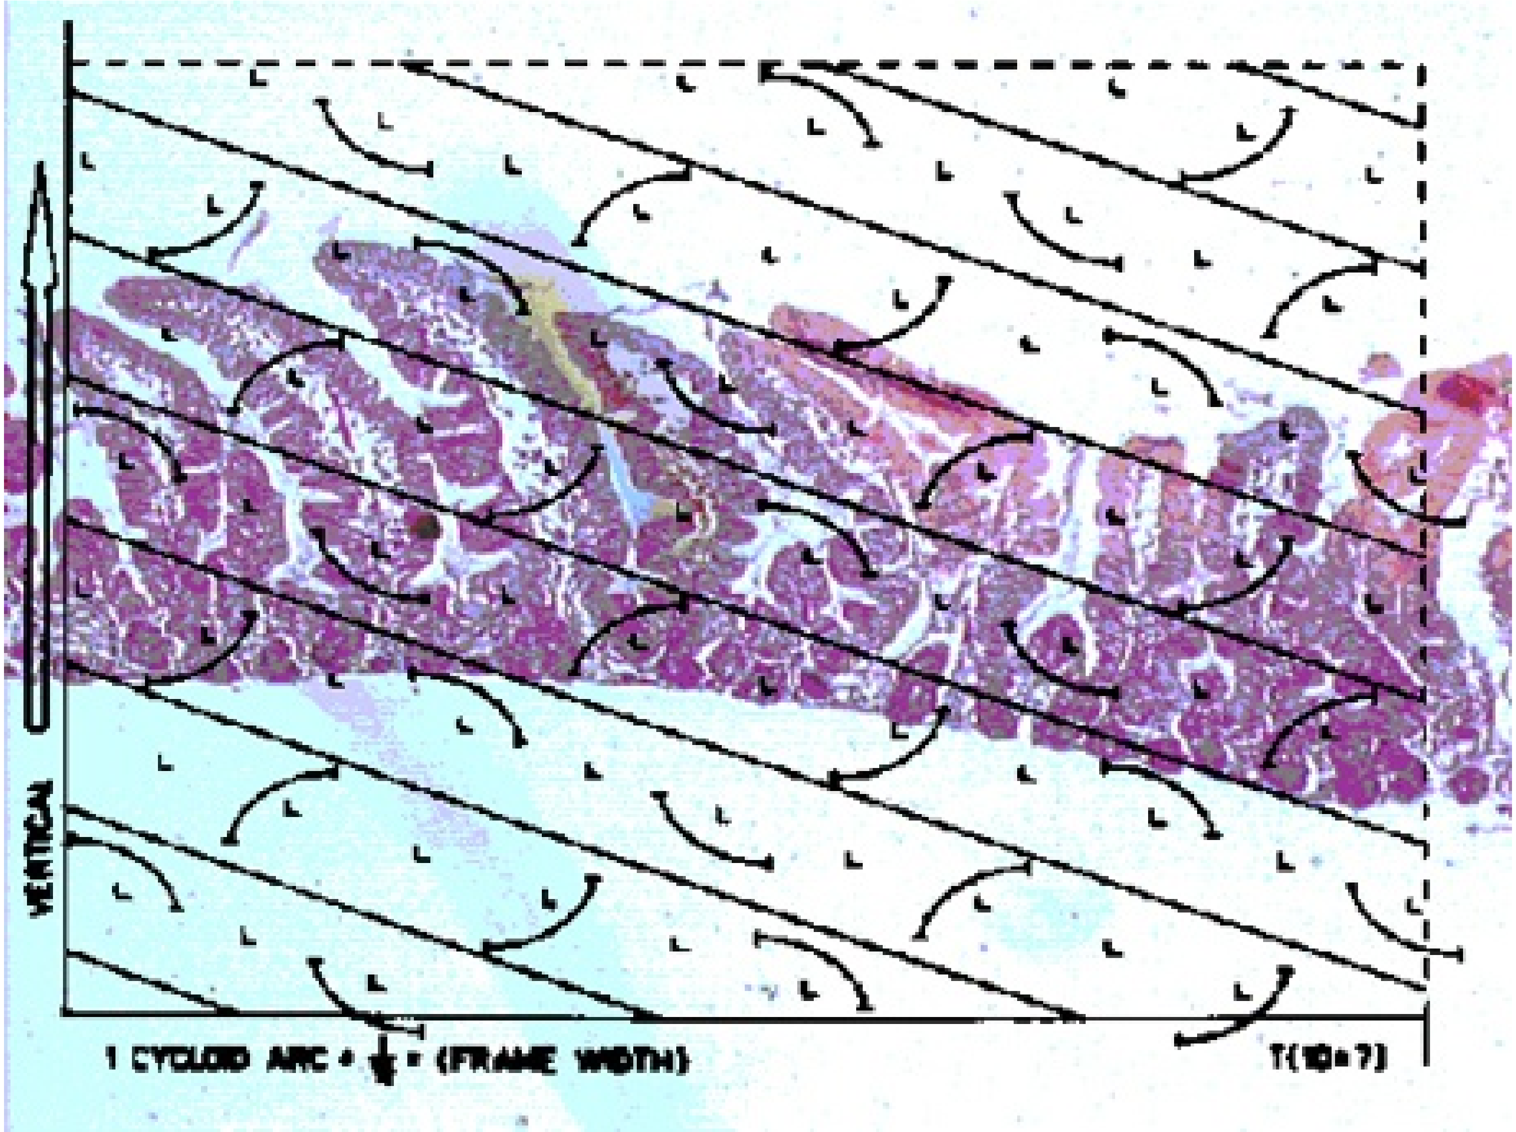

Counting

Counting objects

Any 2D shape in a 2D image, any 3D shape in a 3D image can be used for counting objects

- placement must be such that any part of the sample has equal probability of being sampled

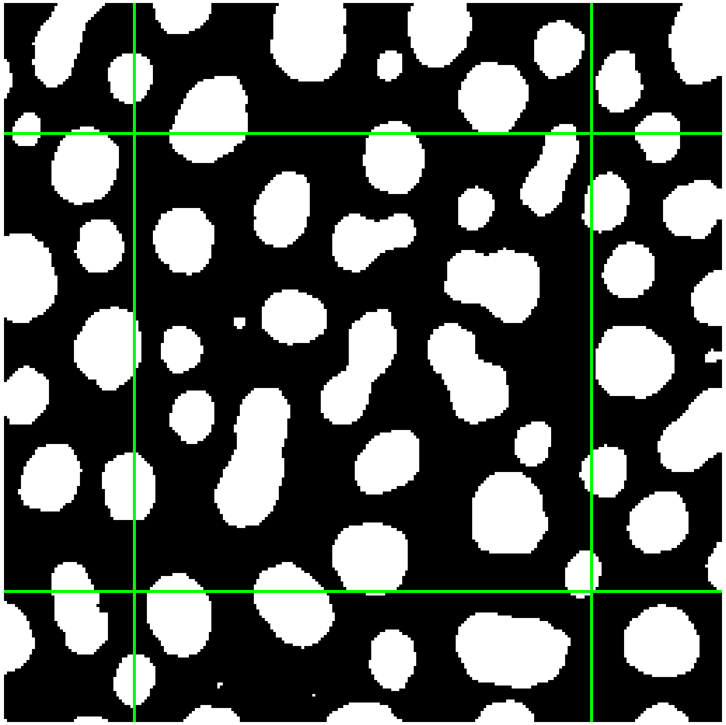

Counting objects

How to count objects partially inside the probe?

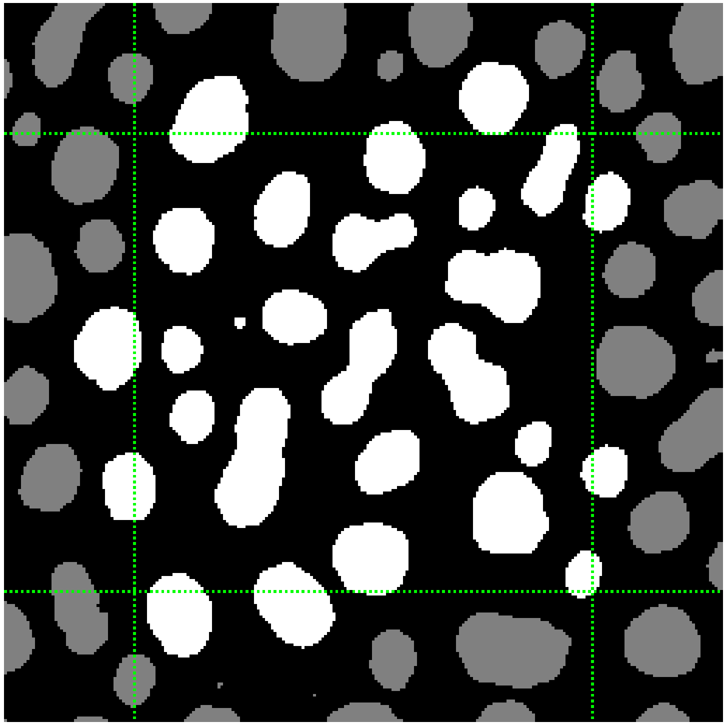

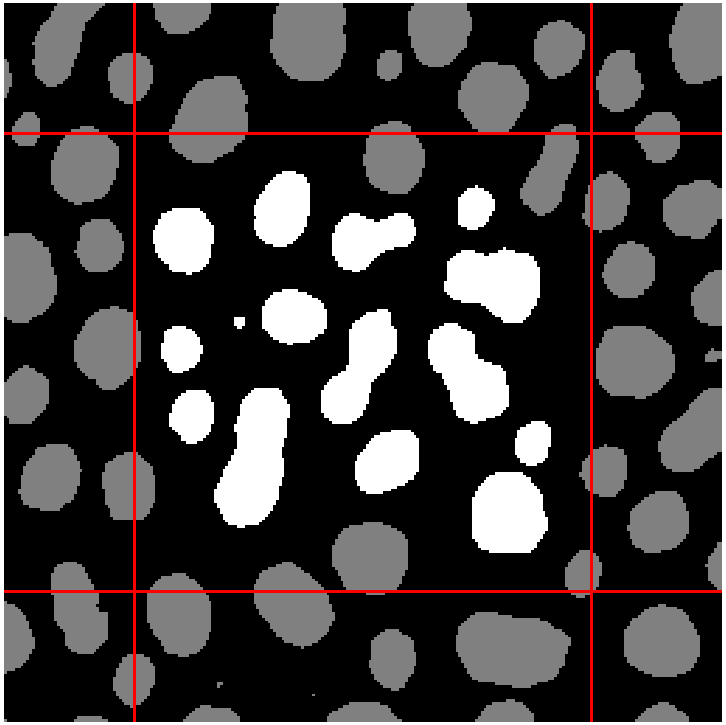

Cris Luengo
luengo@ieee.org

Counting objects

How to count objects partially inside the probe?

Use a counting frame!

(Gundersen, 1977)

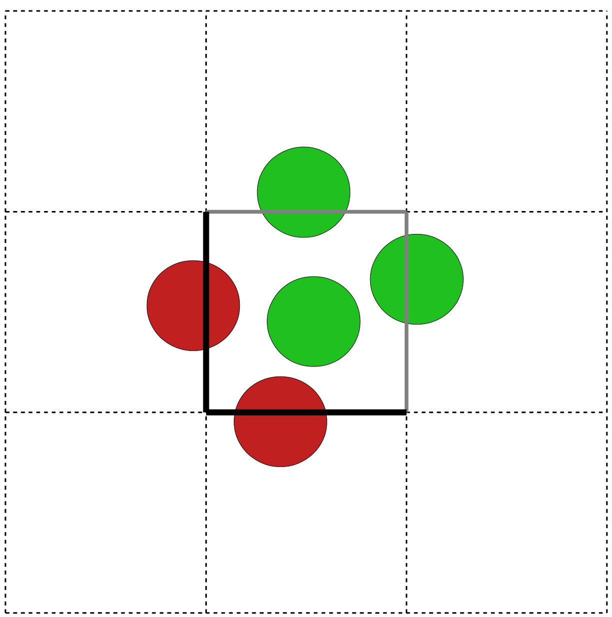

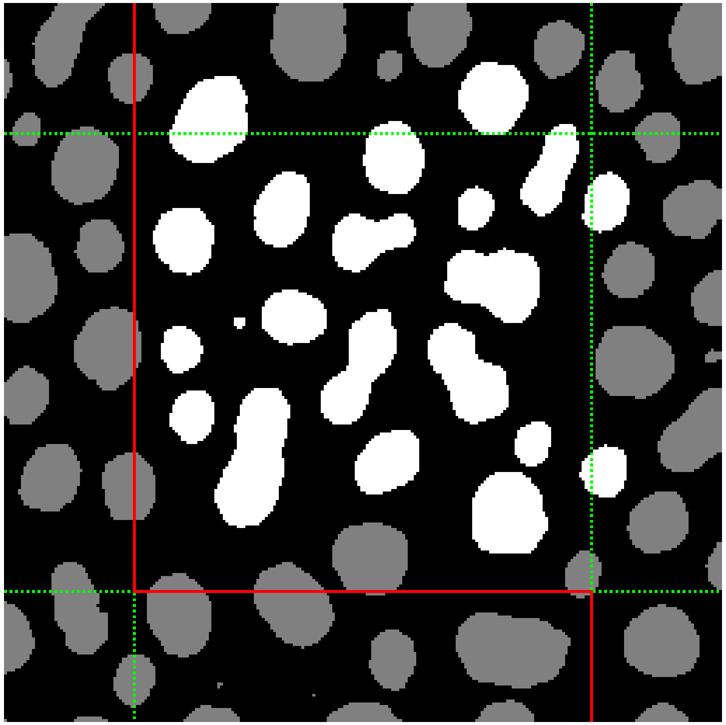

Cris Luengo
luengo@ieee.org

Counting objects

How to count objects partially inside the probe?

Use a counting frame!

(Gundersen, 1977)

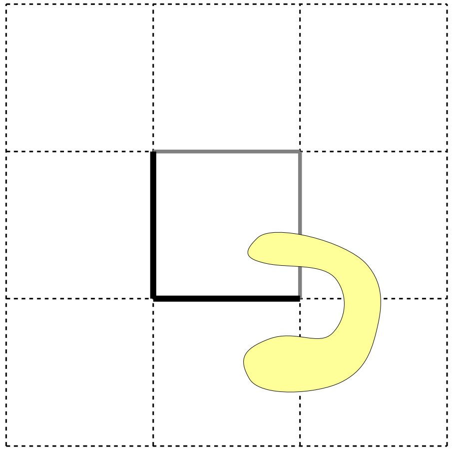

Biased count

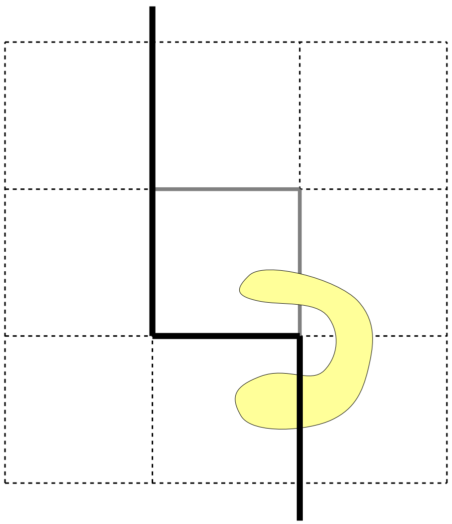

Unbiased count

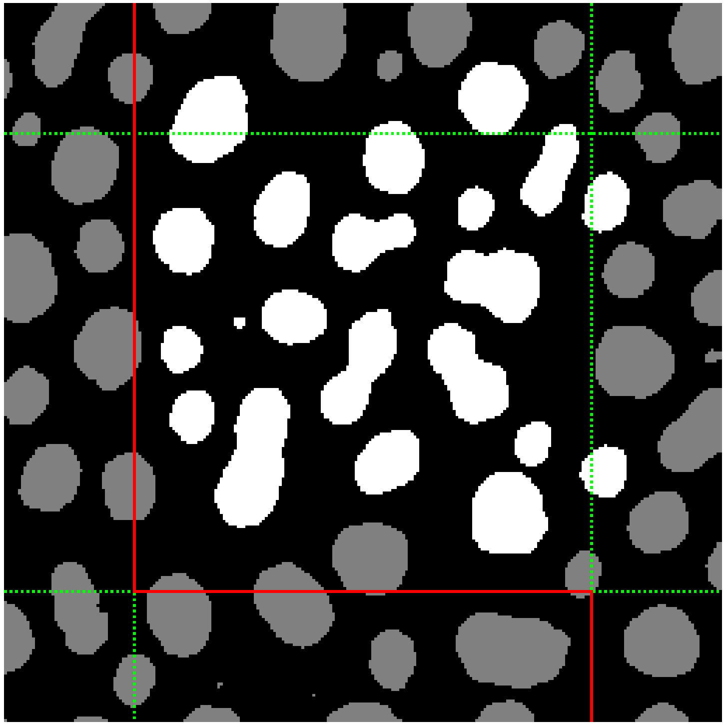

The disector

D.C. Sterio (1984) extended counting frame to 3D

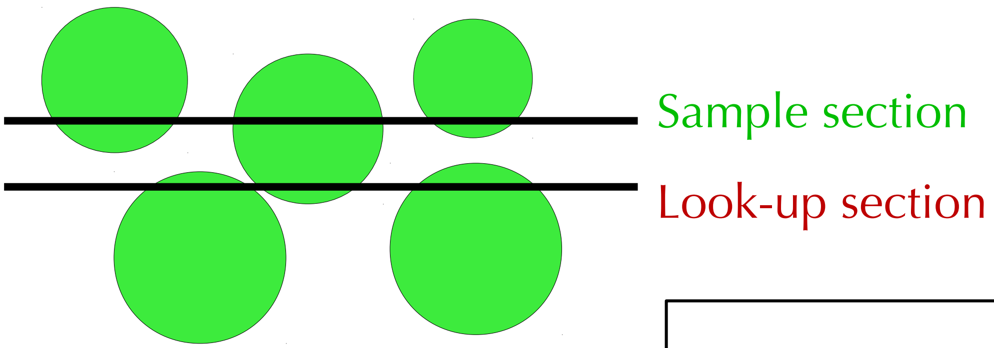

$$\frac{\text{objects}}{2 \cdot \text{volume of disector}}$$

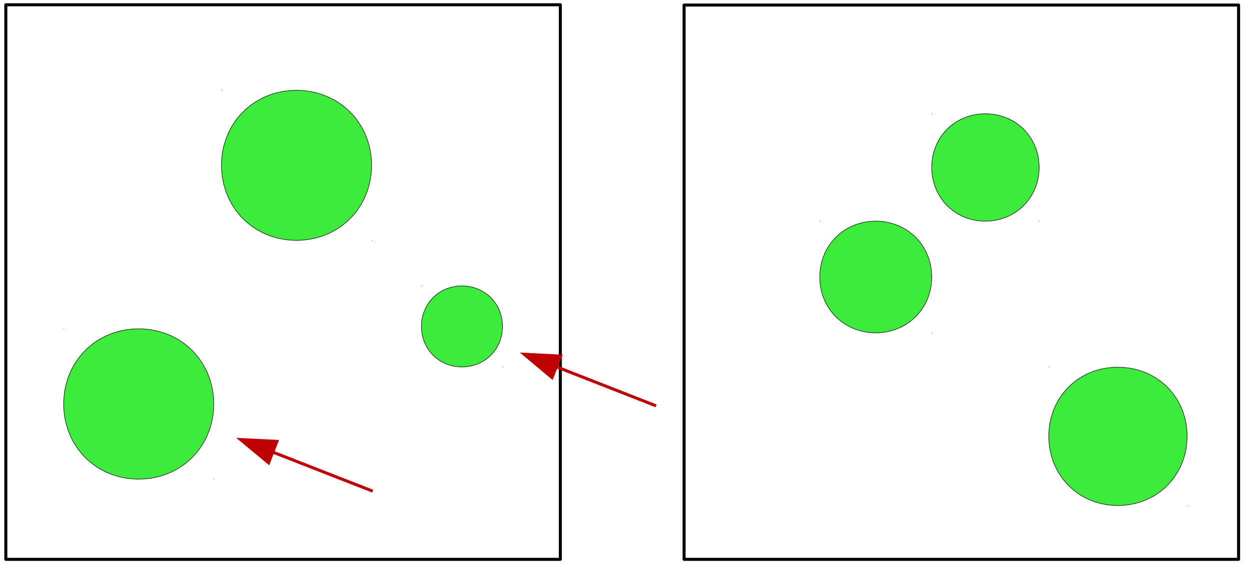

Sample section　　　　Look-up section

Counting objects in 3D

Total number of objects obtained by combining the fractionator and the disector

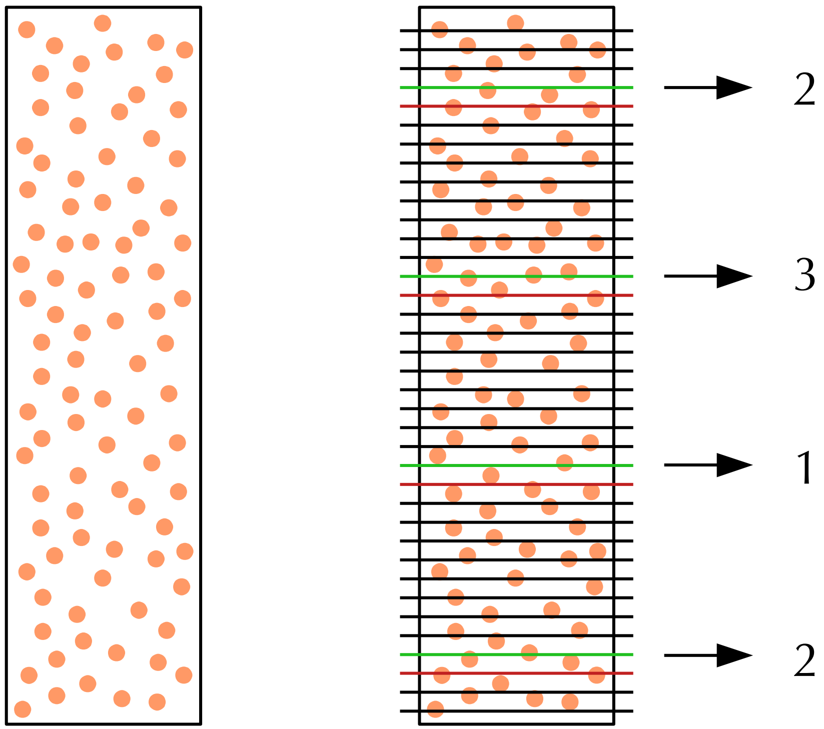

7 cells counted
used 1 of every 10 slices

$$\frac{7}{1/10} = 70$$

estimate = 70 cells

"dilute" by sampling 1 of every N slices
(systematic random uniform sampling)

Summary

- Stereology
  - a stochastic approach to measurement
  - focus is on unbiased measures
- Use systematic uniform random sampling to select:
  - individuals to take samples from
  - slices and fields of view to image
  - objects within the field of view to measure
- Use it with the fractionator principle
- Using probes for measurement:
  - probe dimensionality + feature dimensionality $\geq$ 3
  - probes randomly placed (and oriented)
  - counting frame/disector is the volume probe
  - aim to count 100-200 events per individual
- Object counting, and per-object measures
  - use the counting frame!

Tutorial on Image-Based Measurement

Part II: Measuring binarized objects

Cris Luengo
luengo@ieee.org

Overview

- Area in 2D images / volume in 3D images
- Perimeter in 2D images / surface area in 3D images
- Curvature and bending energy (2D)
- Bounding boxes, Feret diameters, moments
- Number (count)

Cris Luengo
luengo@ieee.org

Area in 2D images
Volume in 3D images

2D area / 3D volume

- The volume of an object is given by the number of pixels
- This is an unbiased estimate, assuming point-sampling

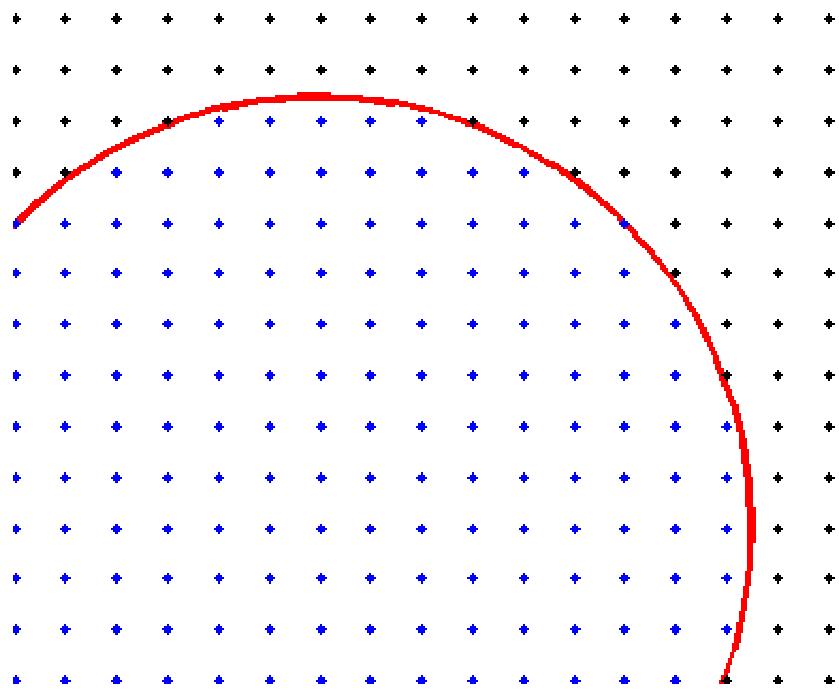

Cris Luengo
luengo@ieee.org

2D area / 3D volume

- Assuming random sampling: $\mathrm{CV} = \frac{1}{\sqrt{a}} = \frac{1}{\sqrt{pi}} r$
- CV is smaller for grid sampling: $\mathrm{CV} \propto r^{-3/2}$

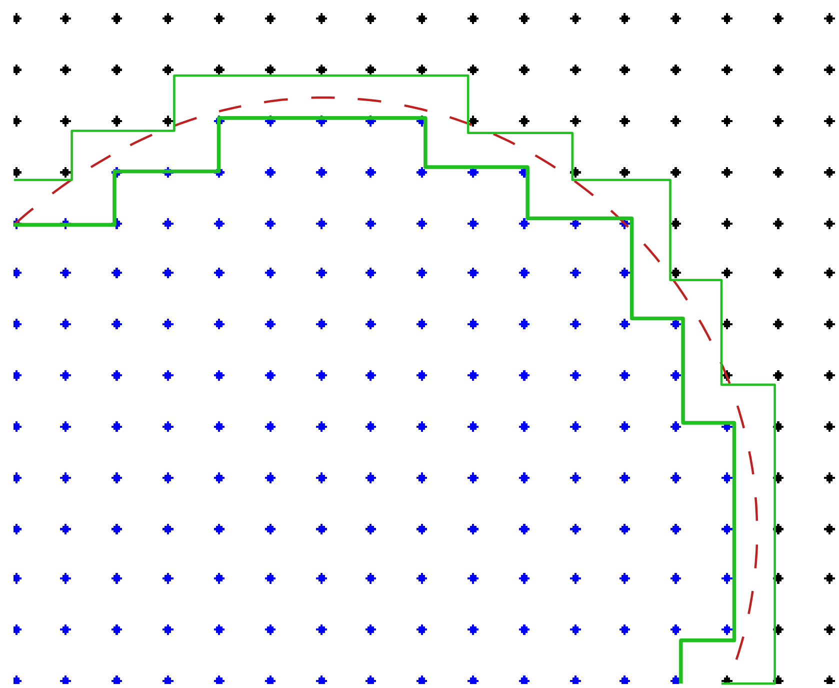

2D area / 3D volume

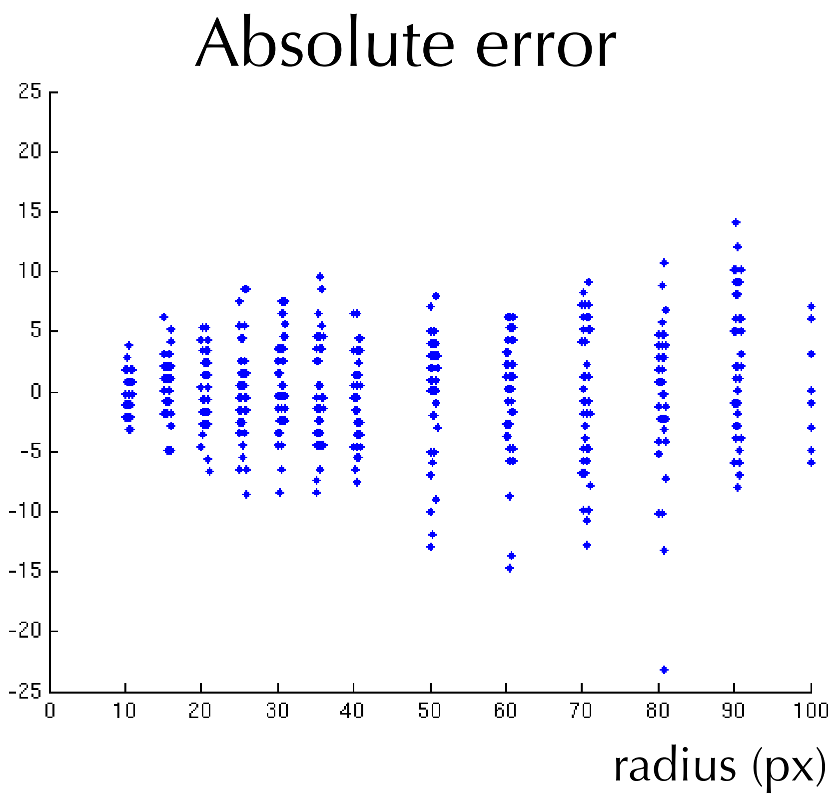

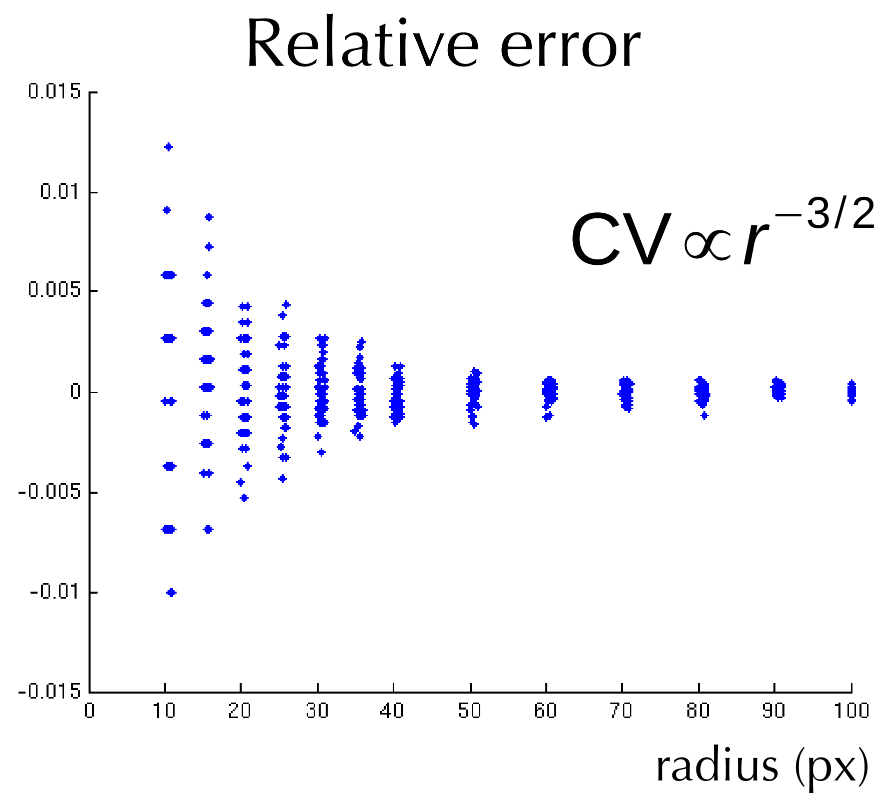

(digitally generated disks of given radius, with random shift w.r.t. sampling grid)

Cris Luengo
luengo@ieee.org

Perimeter in 2D images
Surface area in 3D images

Cris Luengo
luengo@ieee.org

2D perimeter / 3D surface area

- In 2D: measuring length – difficult
  - Trace object contour, yields a digital line
  - Length measurement on line
  - Smoothness assumption of object contour
- In 3D: measuring surface area – much more difficult!
  - Extract surface pixels
  - Estimate a 3D surface through these points
  - Estimate area of 3D surface

Cris Luengo
luengo@ieee.org

Tracing a 2D object's contour

- Start at a random point, e.g. top left pixel
- Next pixel on contour must be to the right, down-right, down, or down-left
- We now iterate:
  - take previous direction, change it counter-clockwise by 1
  - check that position for an object pixel – if not, change direction clockwise until we find an object pixel
  - add this pixel to the list, and make it "current pixel"
- Iteration finishes when we get to initial pixel **and** initial direction

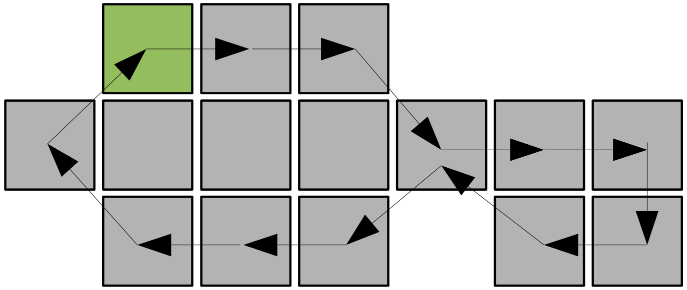

Cris Luengo
luengo@ieee.org

Chain codes

- When walking along an object's contour, we do not need to keep coordinates for each object
- The step direction from one pixel to the next is enough to store the shape information
- Together with the coordinates of the first pixel, yields all information on object
- (A.K.A. Freeman codes)

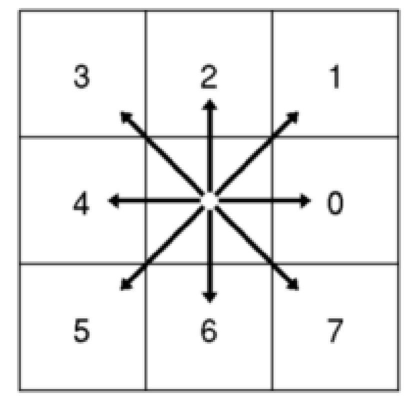

0,0,7,0,0,6,4,3,5,4,4,3,1

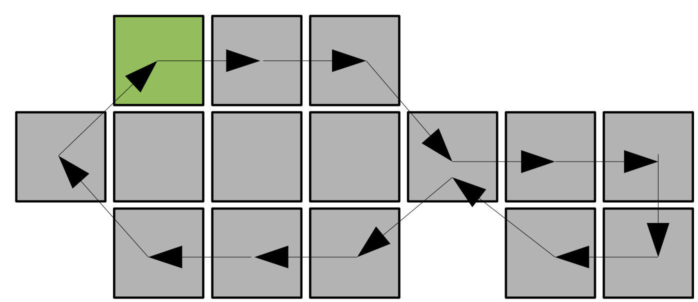

Cris Luengo
luengo@ieee.org

Chain code length

- The length of a chain code is given by the length of each step taken
  - even codes (0, 2, 4, 6) are vertical and horiz. steps
  - odd codes (1, 3, 5, 7) are diagonal steps
- Step sizes of 1 and √2 overestimate length

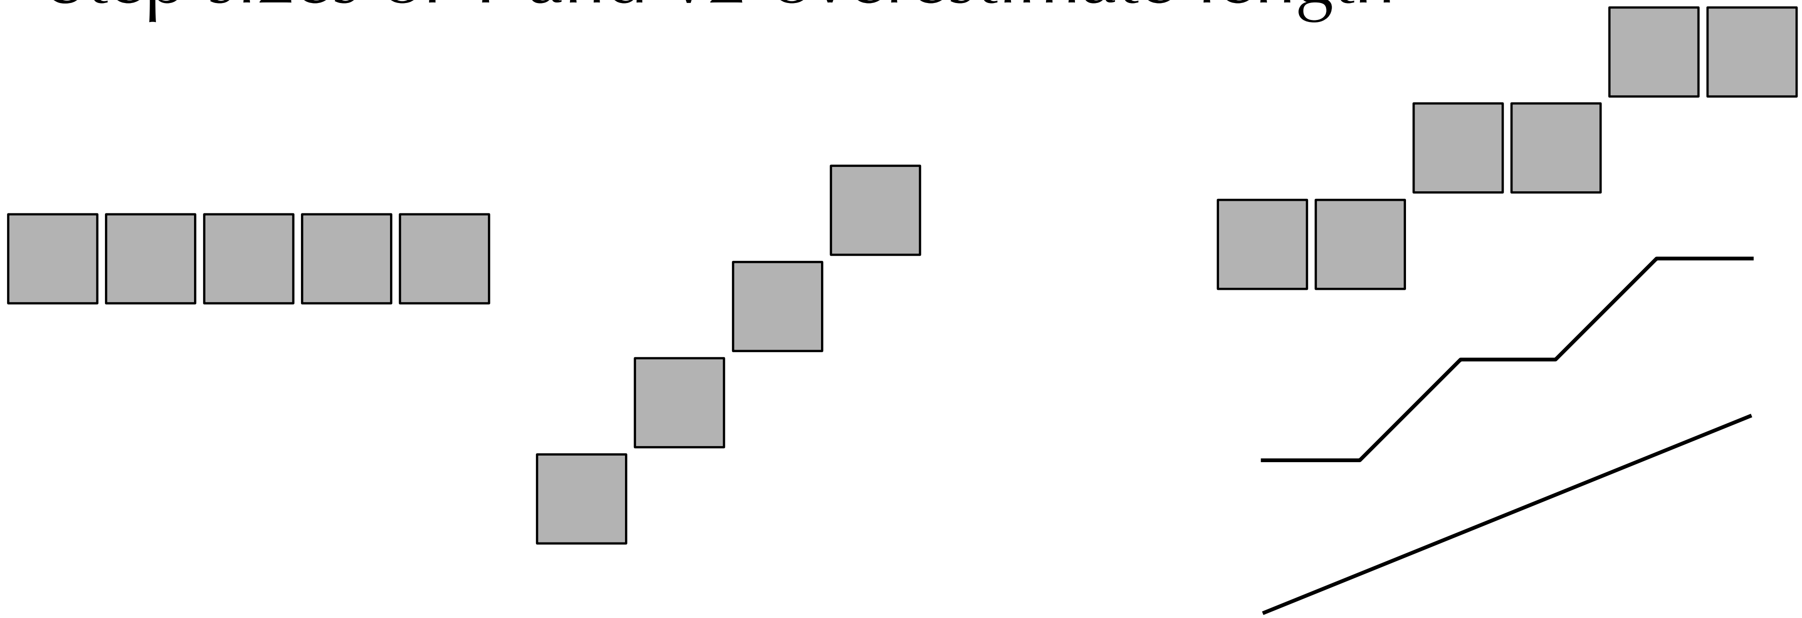

Cris Luengo
luengo@ieee.org

Chain code length

- The length of a chain code is given by the length of each step taken
  - even codes (0, 2, 4, 6) are vertical and horizontal steps
  - odd codes (1, 3, 5, 7) are diagonal steps
- Step sizes of 1 and $\sqrt{2}$ overestimate length
  (Freeman, 1970)
- Step sizes of 0.948 and 1.340 yield unbiased measure
  (Kulpa, 1977)
- Additionally, add "corner count"
  (Vossepoel, 1982)
  - Indicates changes of direction

Cris Luengo
luengo@ieee.org

Chain code length

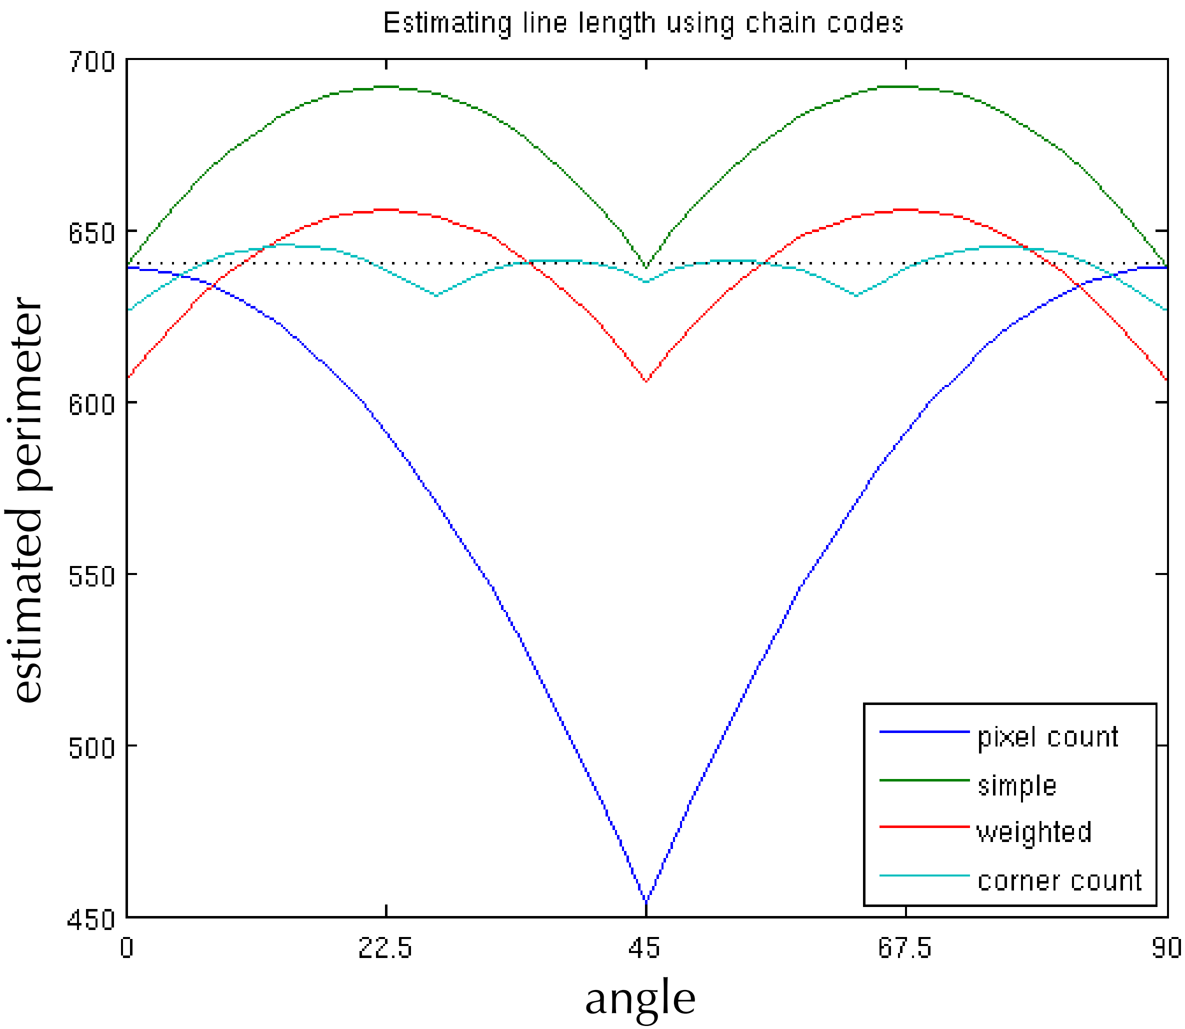

N_{pixels}

$N_{even} + \sqrt{2}\, N_{odd}$

$0.948\, N_{even} + 1.340\, N_{odd}$

$0.980\, N_{even} + 1.406\, N_{odd} - 0.091\, N_{corner}$

Chain code length

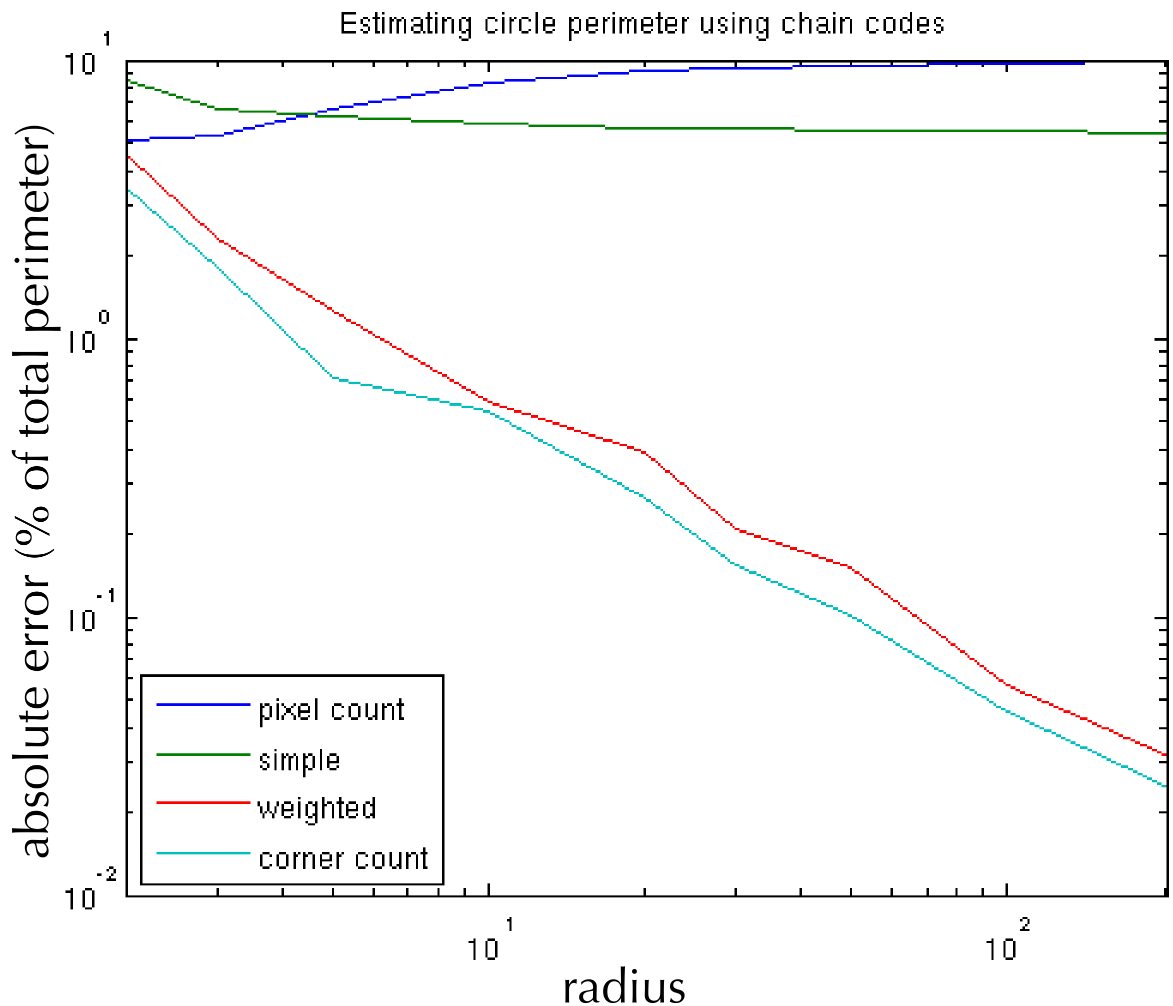

N_{pixels}

$N_{even} + \sqrt{2}\, N_{odd}$

$0.948\, N_{even} + 1.340\, N_{odd}$

$0.980\, N_{even} + 1.406\, N_{odd} - 0.091\, N_{corner}$

Cris Luengo
luengo@ieee.org

Perimeter from chain code length

- The chain code represents a polygon going through the centre of pixels on the object boundary
- The actual perimeter is half a pixel further out

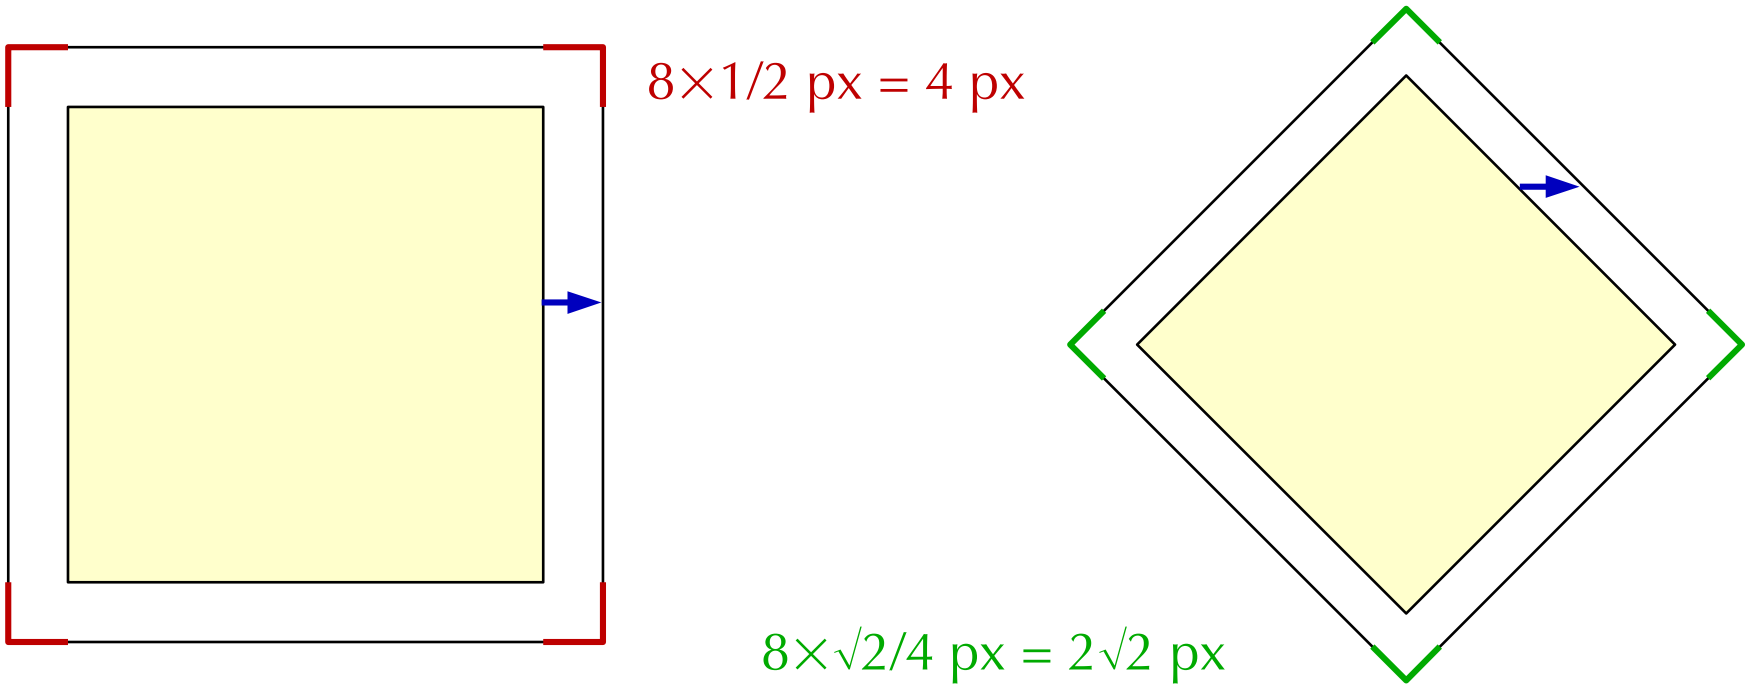

Cris Luengo
luengo@ieee.org

Perimeter from chain code length

- The chain code represents a polygon going through the centre of pixels on the object boundary
- The actual perimeter is half a pixel further out
- Theoretical value for a circle: π

true radius $\quad r' = r + 1/2$

true perimeter $\quad p' = 2\,(\,r + 1/2\,)\,\pi = p + \pi$

Cris Luengo
luengo@ieee.org

3D surface area

- Find all boundary points
  - (object points with a background neighbour)
- No chain codes possible
  - But: order of pixels wasn't important in chain codes to measure length
- One approach:
  - Classify neighbourhood type for each surface voxel
  - Determine optimal weights for each neighbourhood type
- A different approach:
  - Using marching cubes, obtain a triangulation mesh
  - Sum surface area of triangles

Curvature
Bending energy
(2D)

Cris Luengo
luengo@ieee.org

Curvature

- Curvature κ = derivative of θ along the curve

$$\kappa = \frac{d\theta}{ds} \approx \frac{\theta_{i+1} - \theta_i}{D(i, i+1)}$$

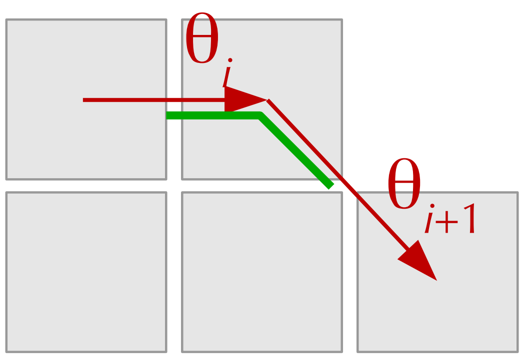

- $\theta_i = \pi/4\ [\, c_{i+1} - c_i \,]$
  difference computed in modulo arithmetic, in range [-3,3]
- $D(i,i+1) = 0.5\ (\ s(c_i) + s(c_{i+1})\ )$
- $s(c_i) = 0.948$ for even c_i and 1.340 for odd c_i

Cris Luengo
luengo@ieee.org

Curvature

- Curvature κ = derivative of θ along the curve

$$\kappa = \frac{d\theta}{ds} \approx \frac{\theta_{i+1} - \theta_i}{D(i, i+1)}$$

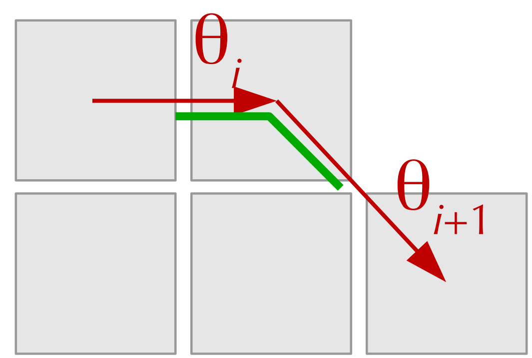

- Problem:
  Very rough estimation of point-wise curvature
  - angles discretised
  - distances discretised
- Solution: Smooth κ value over several neighbours
  - effect is similar to smoothing boundary before computation

Cris Luengo
luengo@ieee.org

Bending energy

- Often used as a shape descriptor
- Given by integral along the perimeter of the square of curvature

$$\text{B.E.} = \int_{\text{contour}} \kappa^2 \, \mathrm{d}s$$

- Integrating along perimeter requires proper step lengths

$$\text{B.E.} \approx \sum_i \left(\frac{\theta_{i+1} - \theta_i}{D(i, i+1)} \right)^2 D(i, i+1)$$

Cris Luengo
luengo@ieee.org

Bounding boxes
Feret diameters
Moments

Cris Luengo
luengo@ieee.org

Bounding box

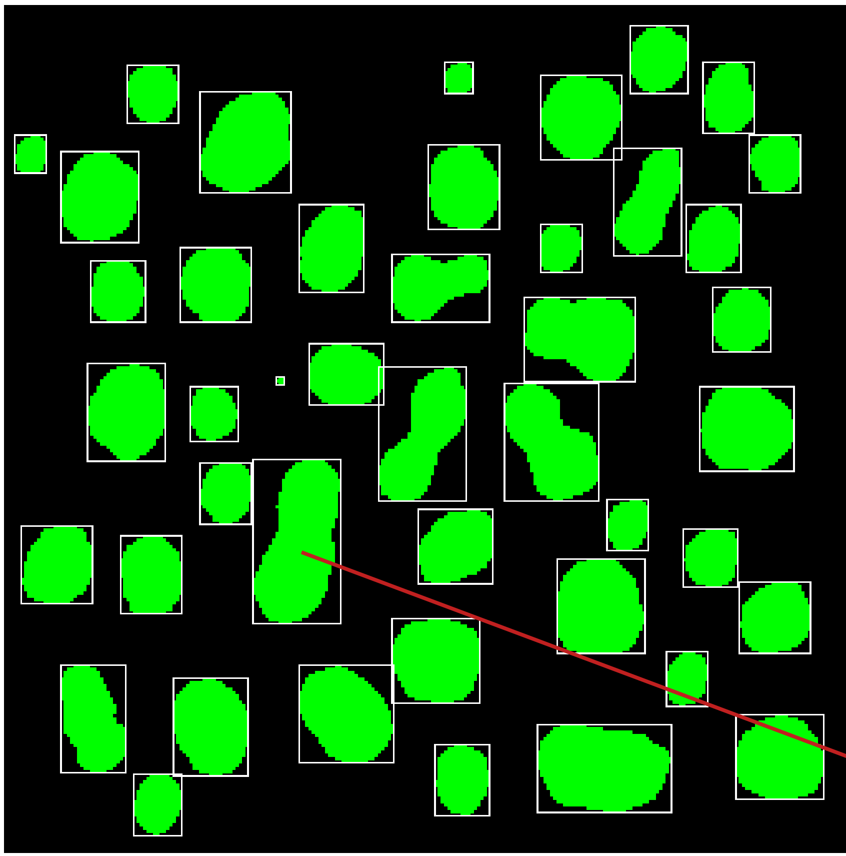

- Find minimal and maximal x and y coordinates for each connected component
- Useful for extracting individual objects from an image for further analysis (e.g. texture)

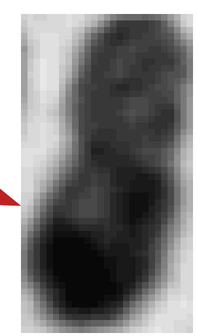

Extends easily to 3D!

Cris Luengo
luengo@ieee.org

Feret diameters

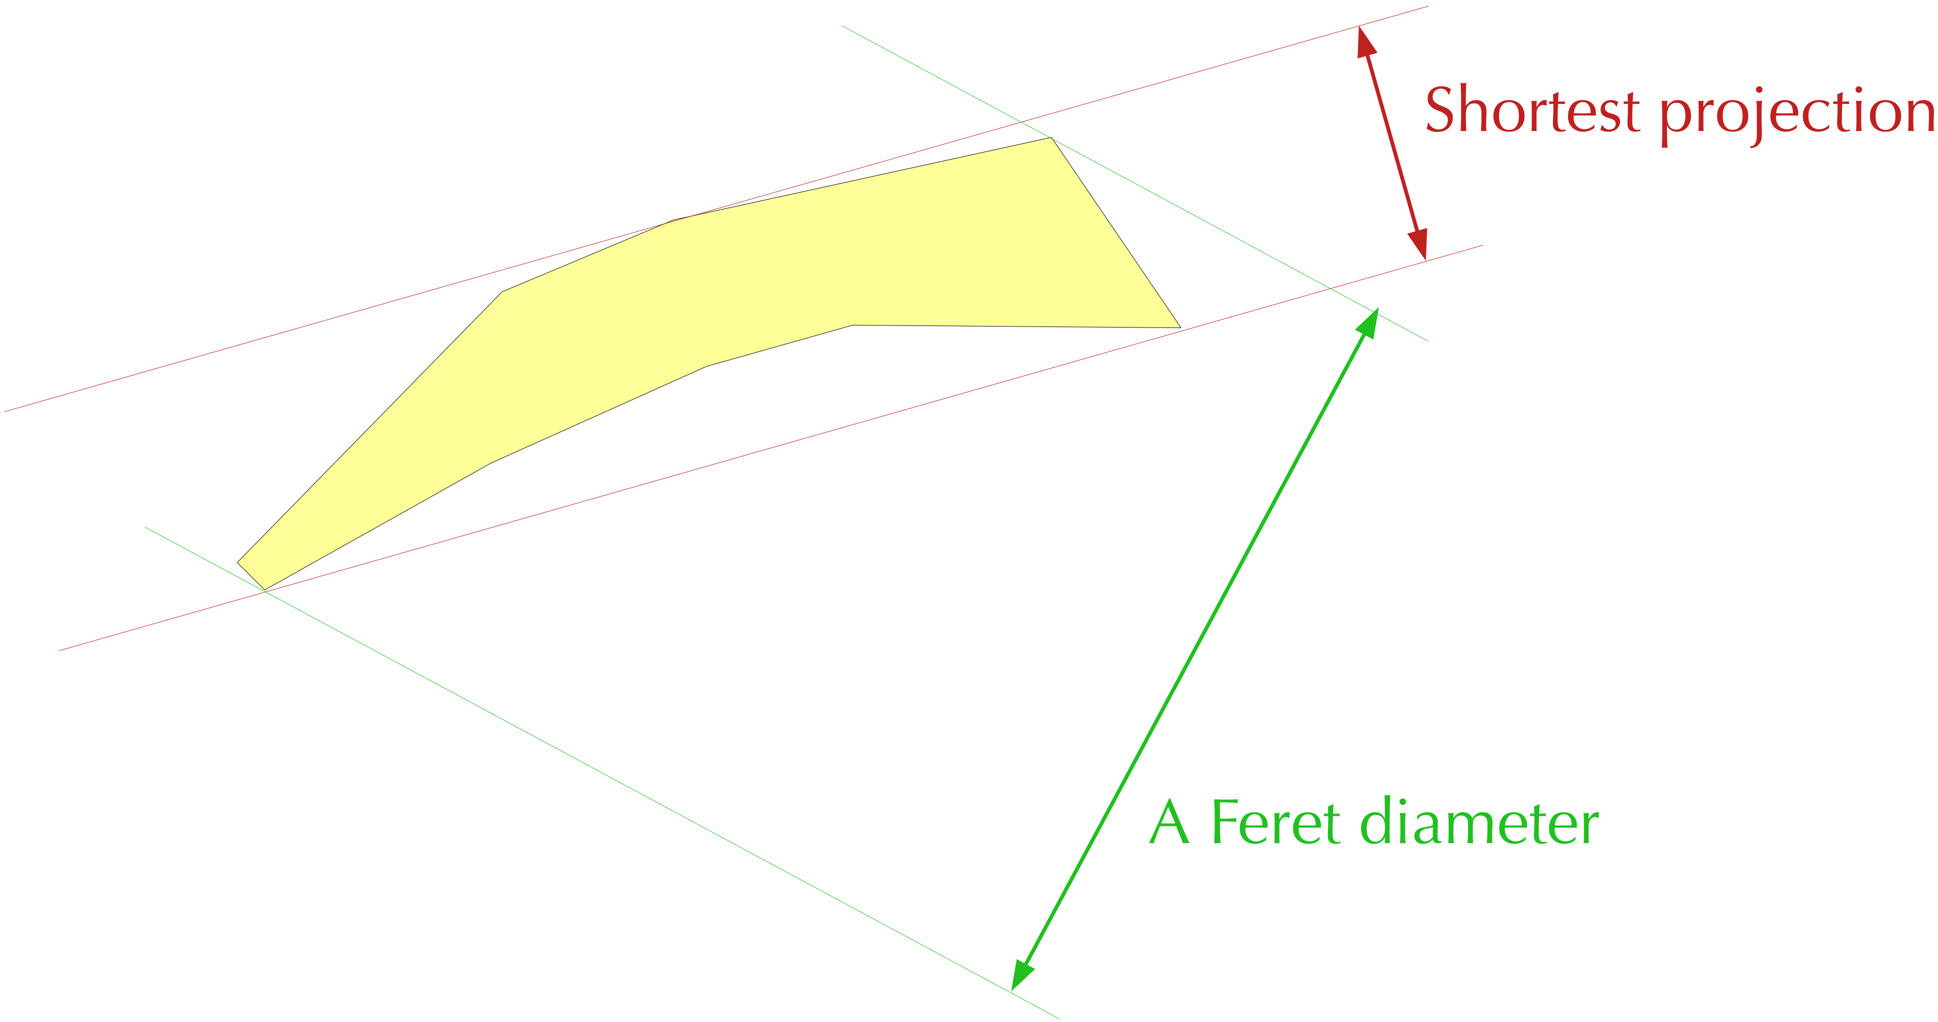

Cris Luengo
luengo@ieee.org

Feret diameters

- Easily computed using the chain code:
  - find coordinates of each pixel, in a rotated coordinate system
  - record minimum and maximum coordinate
  - rotate coordinate system in small increments
- Slightly more complex, but more precise and faster:
  - compute convex hull of object
  - use “rotating callipers” algorithm to directly determine largest and smallest projections
- As in perimeter:
  add 1 px to diameter!

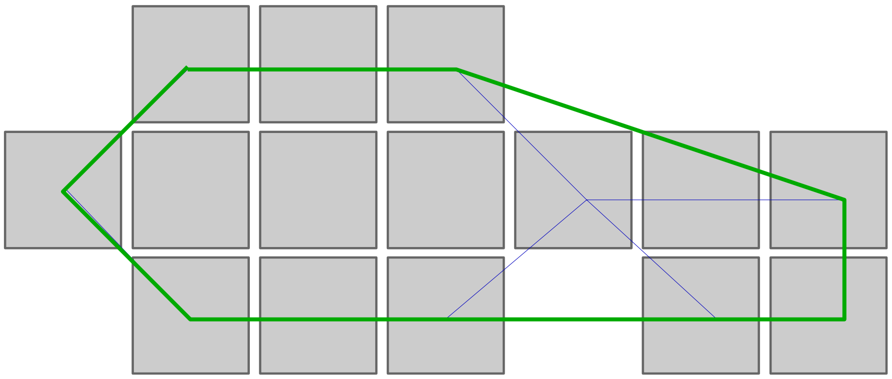

Rotating callipers algorithm

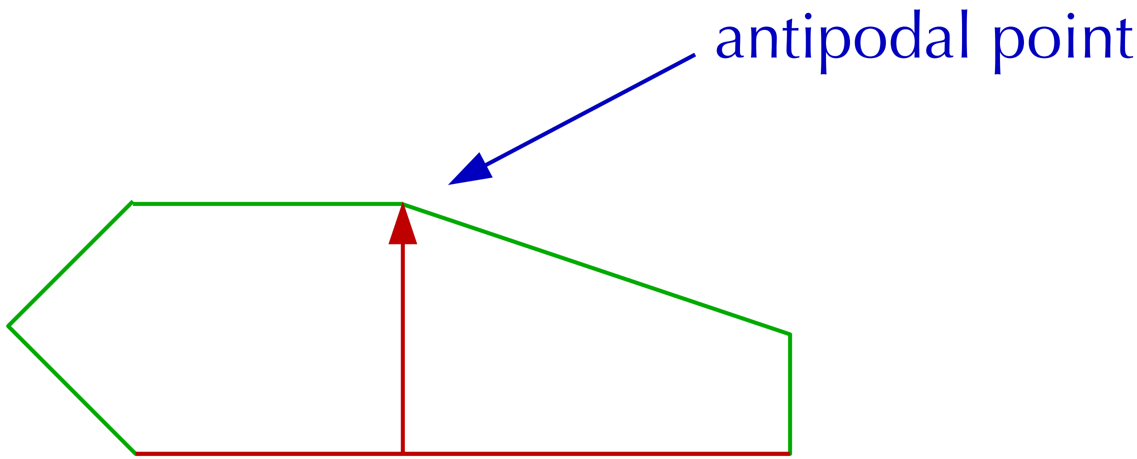

Shortest projection must be perpendicular to one of the sides of the convex hull

Cris Luengo
luengo@ieee.org

Rotating callipers algorithm

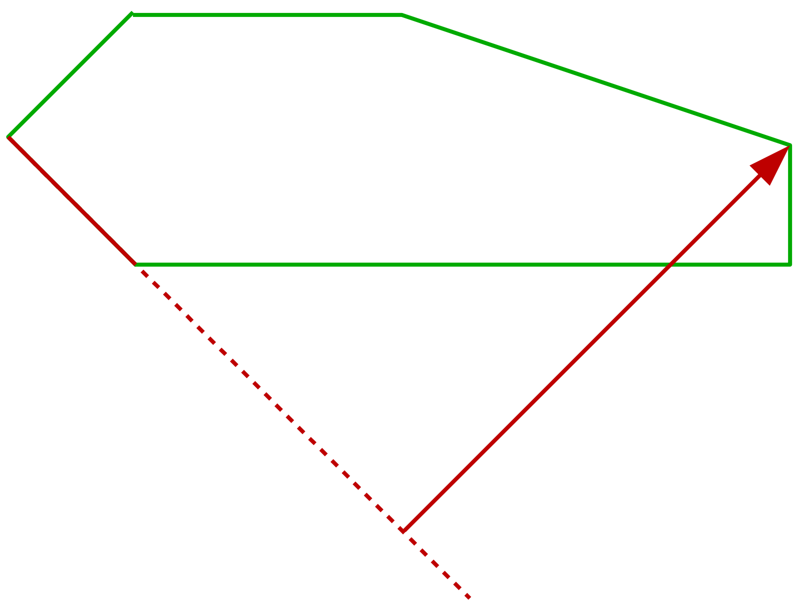

Cris Luengo
luengo@ieee.org

Rotating callipers algorithm

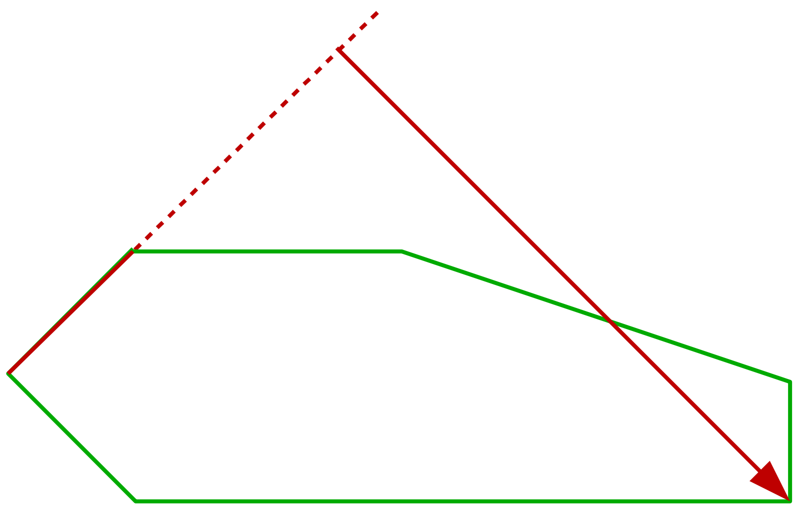

Cris Luengo
luengo@ieee.org

Rotating callipers algorithm

Cris Luengo
luengo@ieee.org

Rotating callipers algorithm

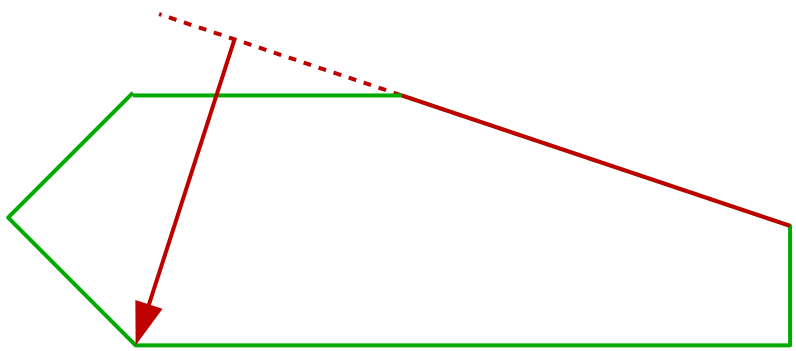

Cris Luengo
luengo@ieee.org

Rotating callipers algorithm

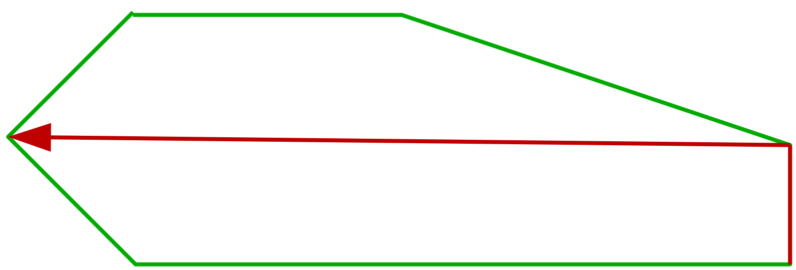

Cris Luengo
luengo@ieee.org

Feret diameters

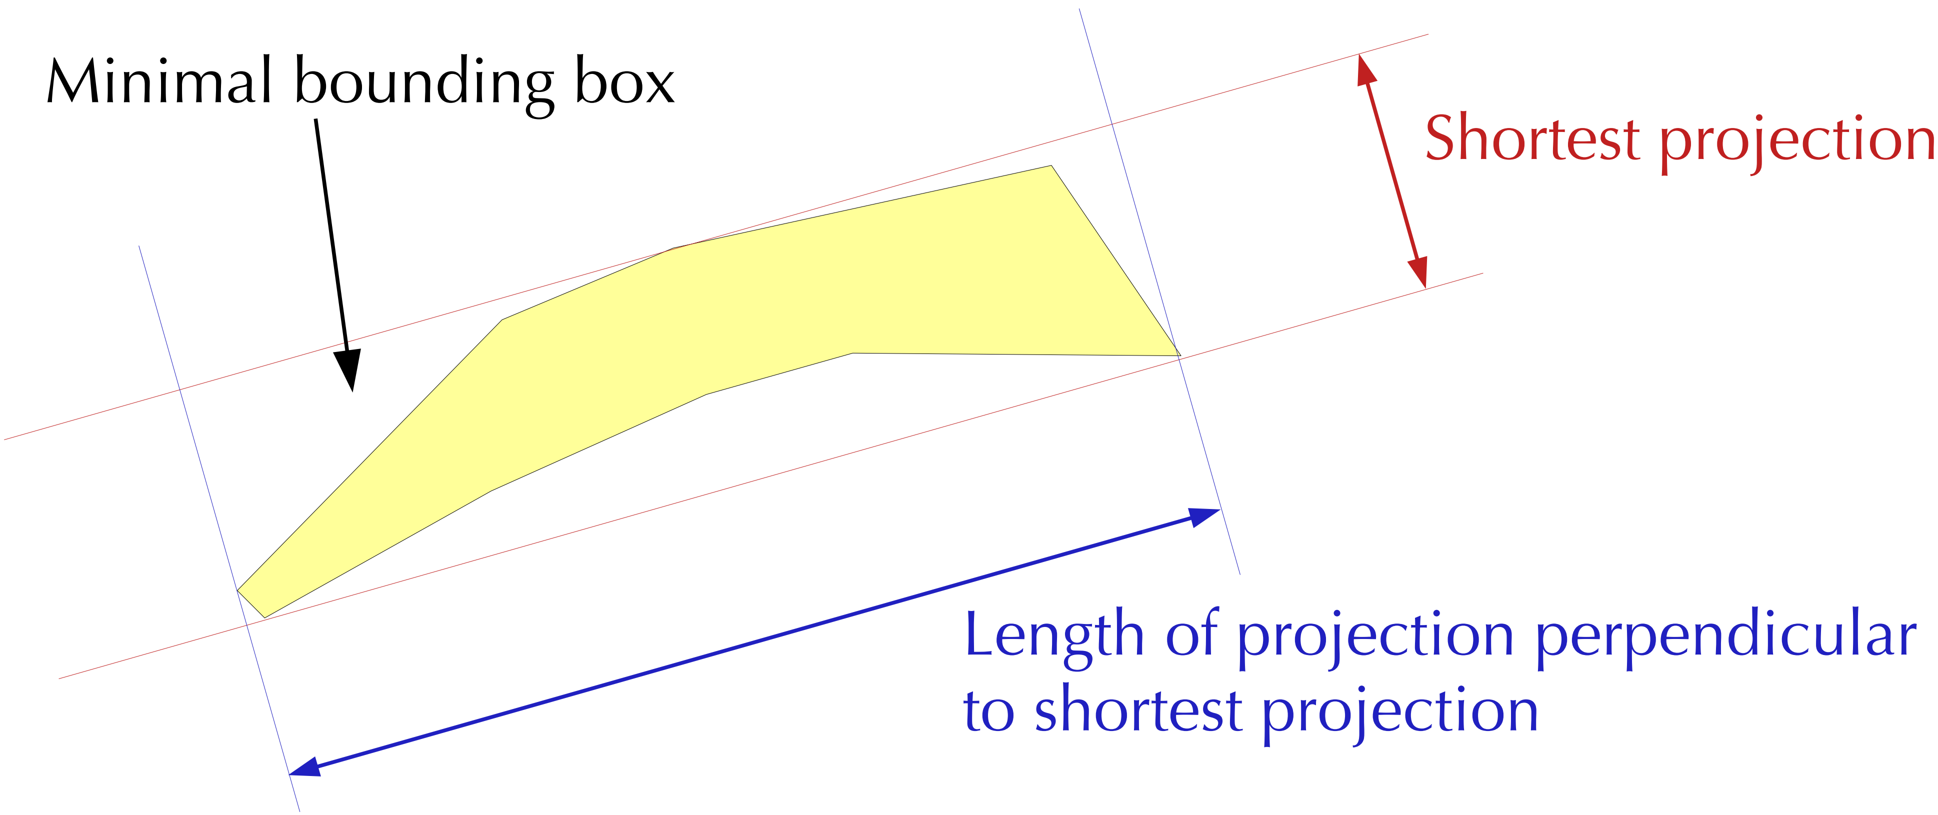

Feret diameters

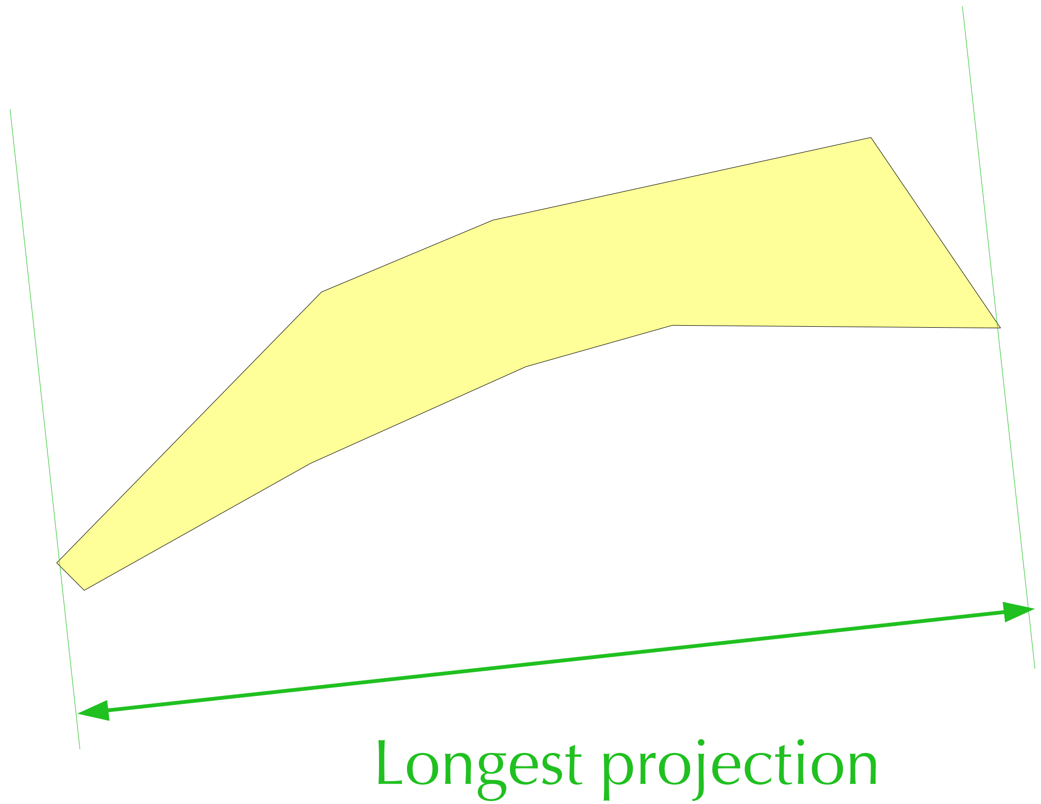

Cris Luengo
luengo@ieee.org

Minimal bounding box

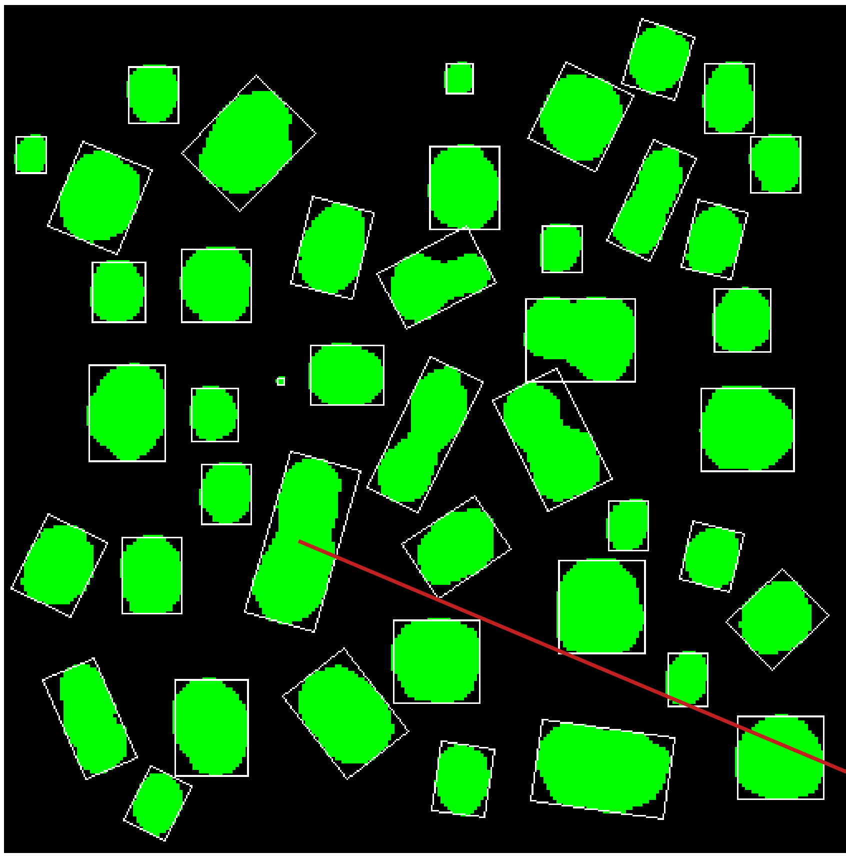

- Find angle for minimal Feret diameter
- Get minimal and maximal coordinates under given rotation
- Useful for extracting individual objects from an image for further analysis (e.g. texture)

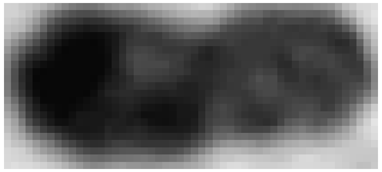

Doesn't extend to 3D

1st order moments

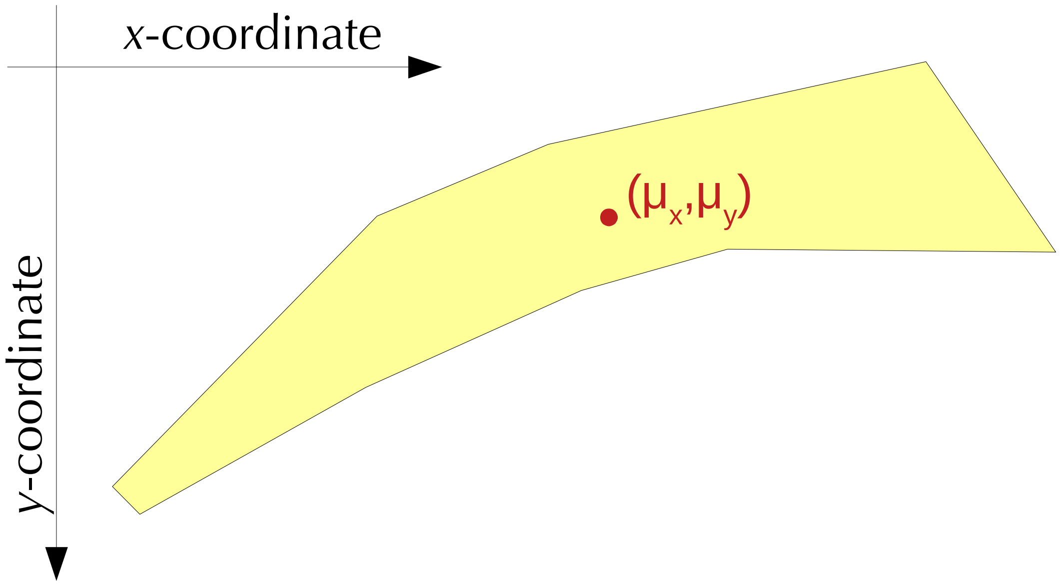

$$\mu_x = \frac{1}{N}\sum x$$

$$\mu_y = \frac{1}{N}\sum y$$

- Centre of mass
- Useful reference point for other measures
  - e.g. mean distance of boundary points to centre of mass

2nd order central moments

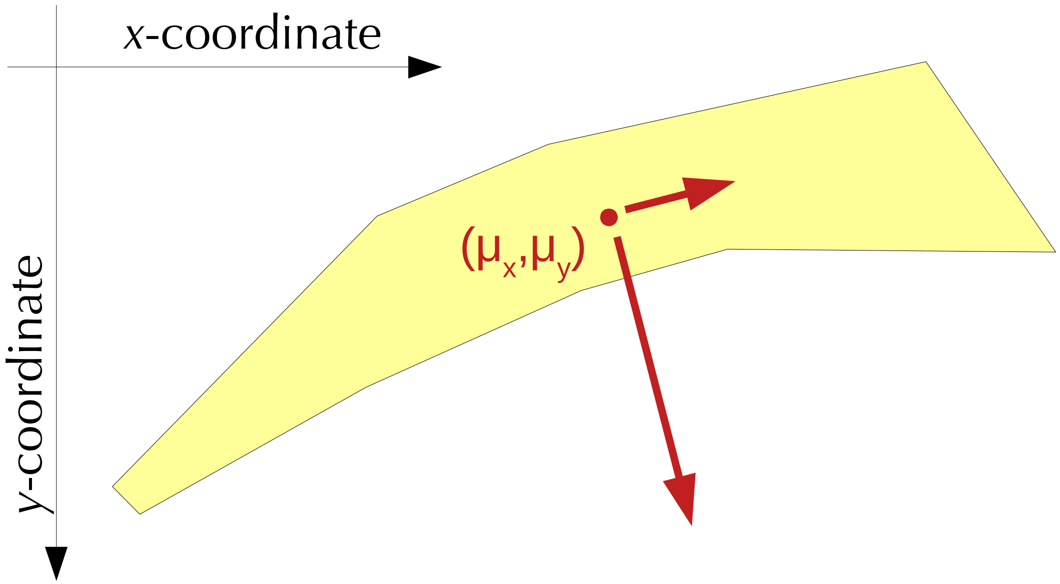

$$I_{xx} = \frac{1}{N} \sum (x - \mu_x)^2$$

$$I_{yy} = \frac{1}{N} \sum (y - \mu_y)^2$$

$$I_{xy} = \frac{1}{N} \sum (x - \mu_x)(y - \mu_y)$$

- Eigenvectors of the moment of inertia tensor give main object axes
- Eigenvalues give size of object along these axes

$$\boldsymbol{I} = \begin{bmatrix} I_{xx} & -I_{xy} \\ -I_{x,y} & I_{yy} \end{bmatrix}$$

(moment of inertia tensor)

Elliptical approximation

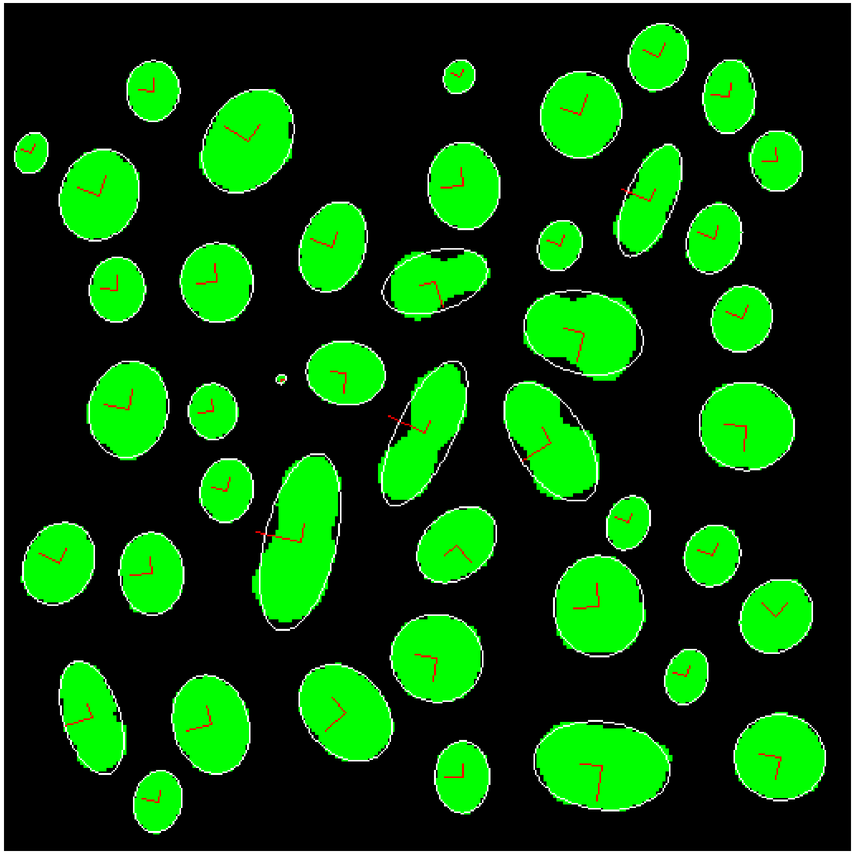

- Compute centre of gravity
- Compute inertia tensor
- Compute eigenvalue decomposition of tensor
- 2 eigenvectors are main ellipse axes
- Ellipse diameters are given by 2×eigenvalue

Extends easily to 3D!

Cris Luengo
luengo@ieee.org

Counting

Cris Luengo
luengo@ieee.org

Counting number of objects

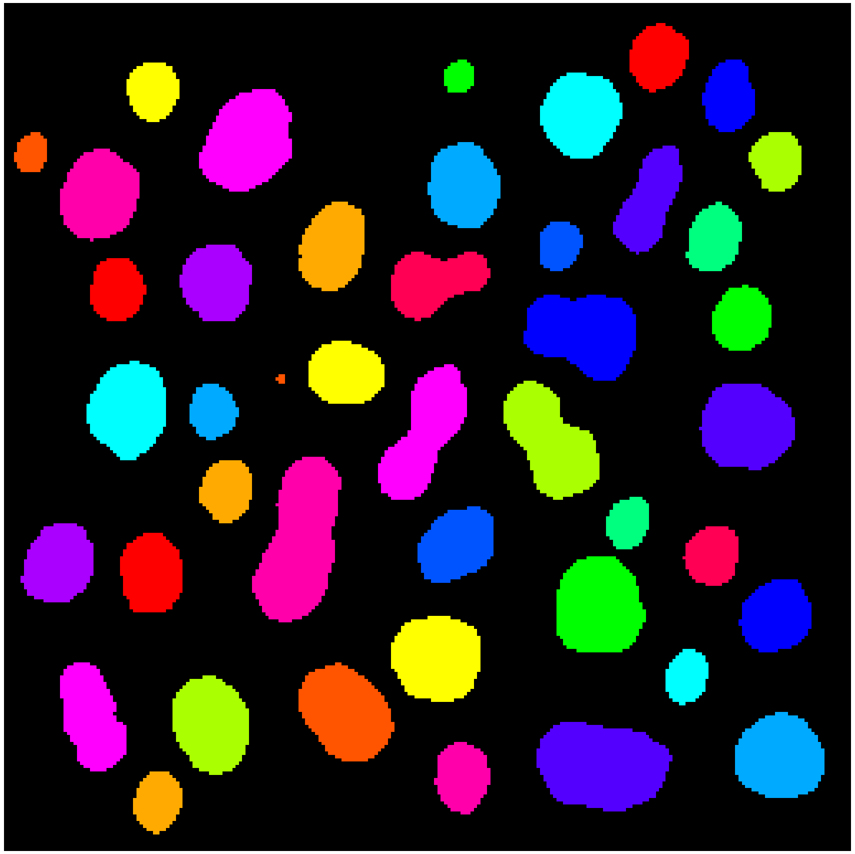

- Label image
- Count labels
- Nobj = 44

Cris Luengo
luengo@ieee.org

Euler number

- In 2D: # of objects – # of holes
- In 3D: # of objects – # of tunnels + # of cavities

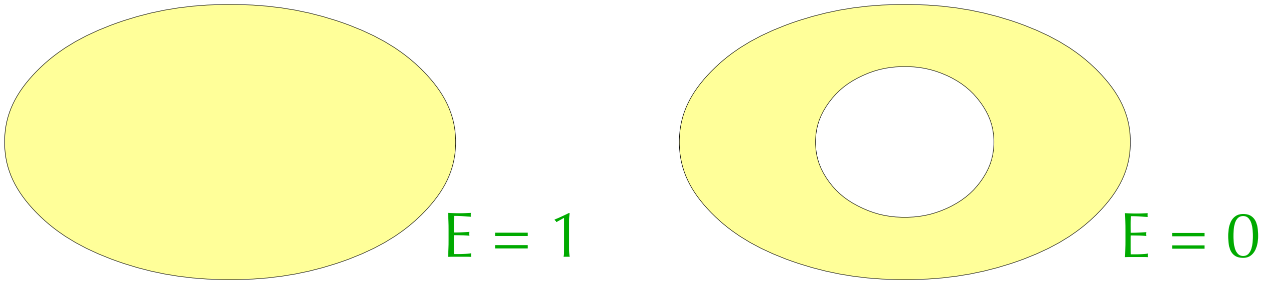

Cris Luengo
luengo@ieee.org

Gray's algorithm

- Computes Euler number based on 2x2 image regions
- $E = (C_1 - C_2 - 2C_3) / 4$
- C_1 = # of 2x2 regions with only 1 pixel set
- C_2 = # of 2x2 regions with 3 pixels set
- C_3 = # of 2x2 regions 2 pixels set in a diagonal

1	0
0	0

C_1

0	1
1	1

C_2

1	0
0	1

C_3

Cris Luengo
luengo@ieee.org

Counting objects

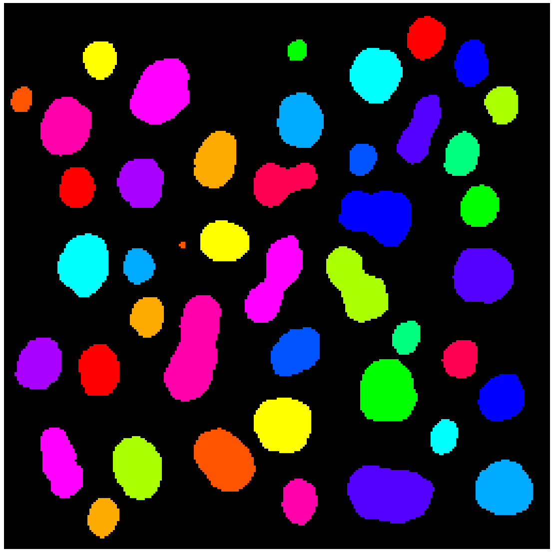

$N_{obj} = 44$

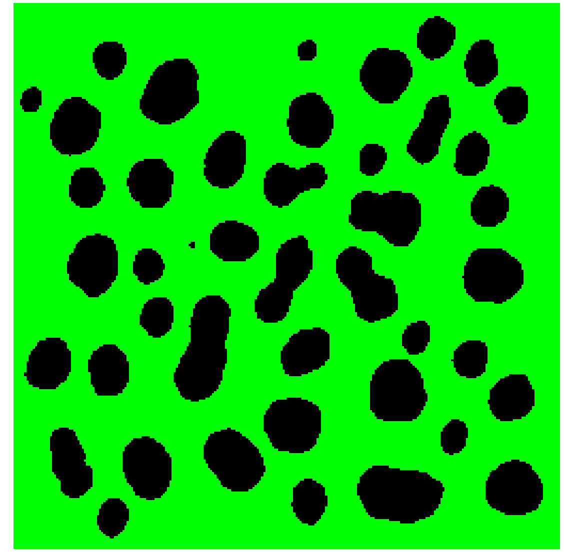

$E = -43$

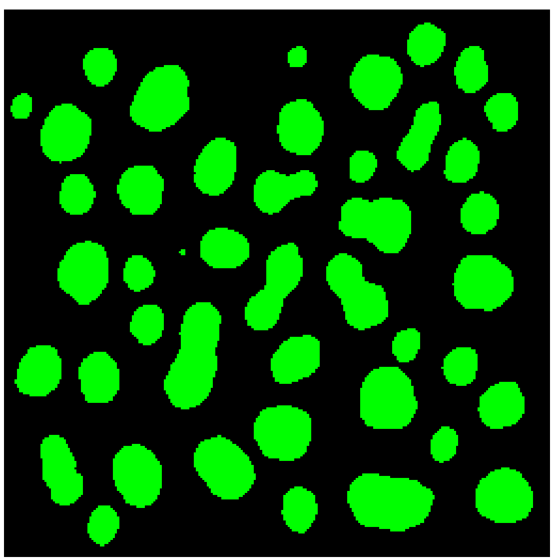

$E = 44$

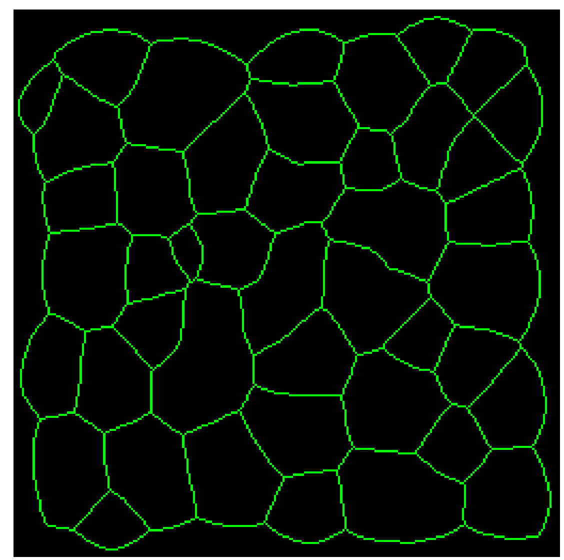

$E = -43$

Summary

- Area in 2D images / volume in 3D images
  - counting pixels
- Perimeter in 2D images / surface area in 3D images
  - take neighbourhood relations into account
- Curvature and bending energy (2D)
  - second derivative of boundary
- Bounding boxes, Feret diameters, moments
- Number (count)
  - Gray's algorithm as an alternative

Tutorial on Image-Based Measurement

Part III: Taking grey-values into account

Centre for Image Analysis
Uppsala University

Cris Luengo
luengo@ieee.org

UPPSALA
UNIVERSITET

Overview

- The point-sampling model:
  - image formation, band limit, sampling, Fourier analysis
- Soft clipping
- Measurement:
  - area
  - perimeter
  - curvature
  - bending energy
  - Euler number (object count)

Cris Luengo
luengo@ieee.org

Possible measures

- Area (2D) / volume (3D)
  - integral over image (sum of grey values)
  - effectively dimensionality-independent
- Perimeter (2D) / surface area (3D)
  - we convert the problem to a volume problem
  - effectively dimensionality-independent
- (Isophote) curvature (2D/3D)
  - based on 2^{nd} derivative along the contour
- Bending energy (2D/3D)
  - integrating squared curvature along contour
- Euler number (object count, 2D)
  - integral of curvature along contour is constant

Cris Luengo
luengo@ieee.org

The point sampling model

Cris Luengo
luengo@ieee.org

The point-sampling model

- Point sampling is what is assumed in signal theory
- Point sampling is only useful if the image is band limited
  - otherwise we get aliasing
  - *sampling frequency* > 2 · *band limit* (Nyquist)
- CCDs do not point-sample
  - but: same as a uniform filter followed by point sampling

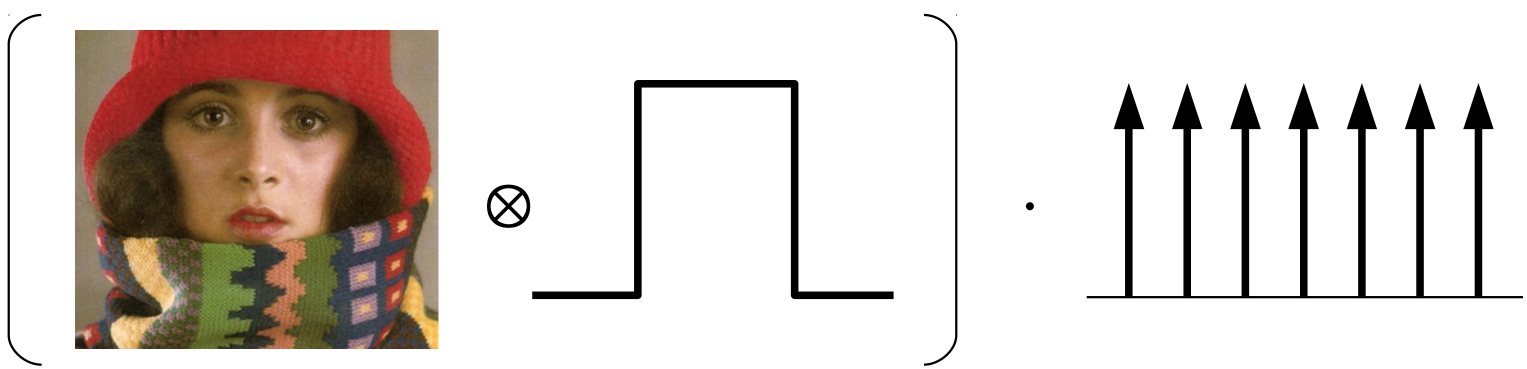

Cris Luengo
luengo@ieee.org

Band-limited images

- Any optical image formation system imposes a band limit
- A sampled band-limited image *exactly* represents the continuous band-limited image
  - if sampled properly
- The continuous band-limited image is a version of real world that lacks very high frequencies
- This smooth image preserves large-scale geometric properties of the imaged objects, but not small scale ones
  - as in “the infinite coastline of Britain”

The sampling property

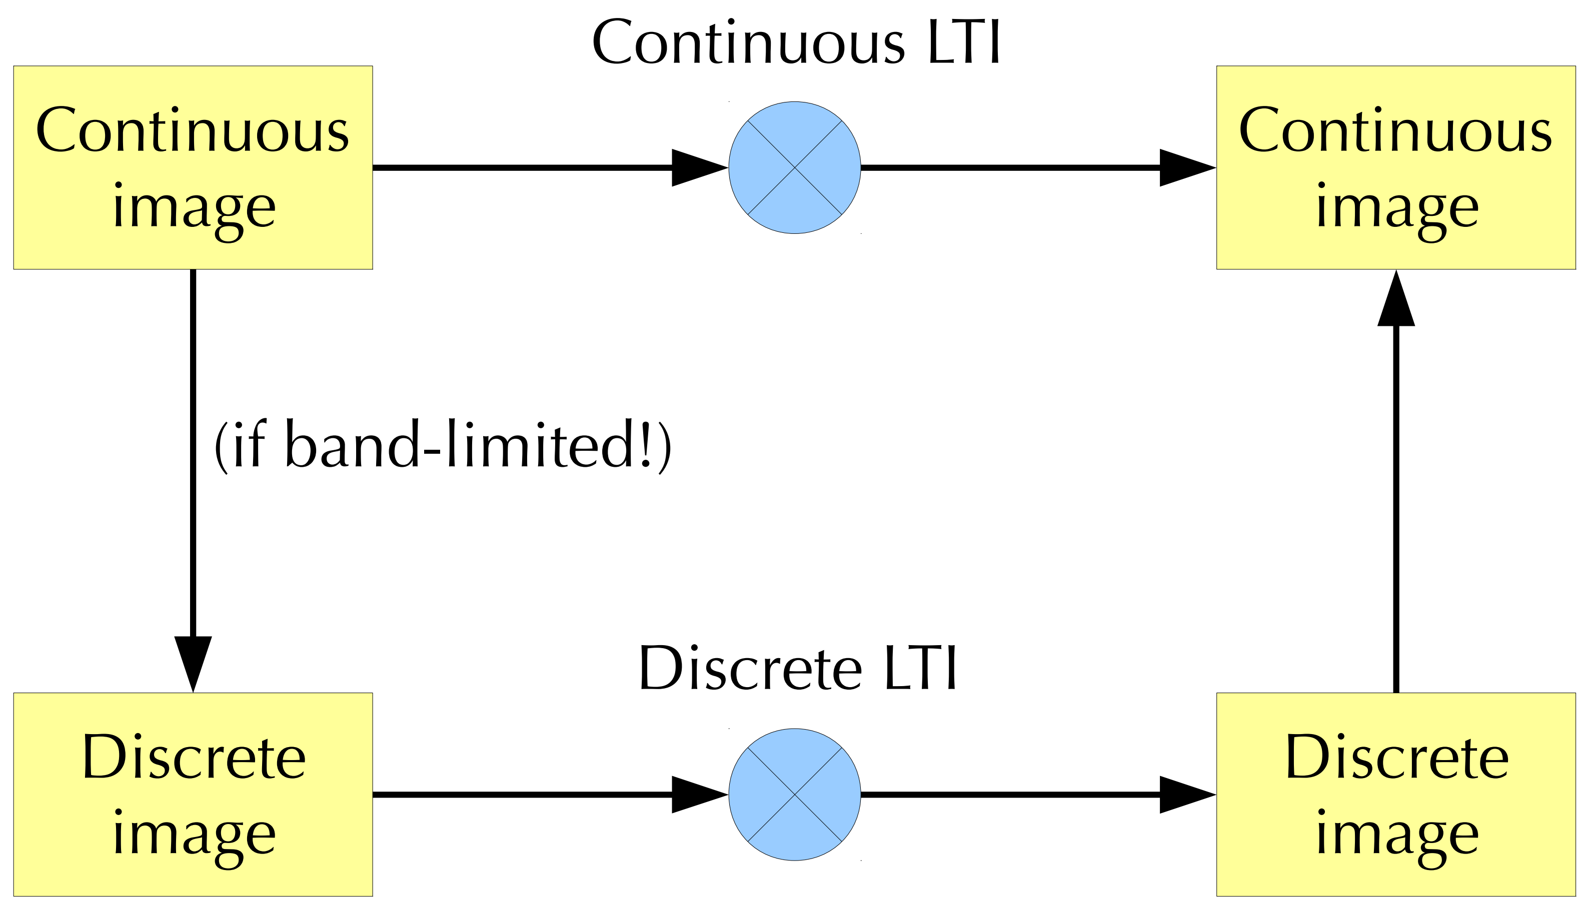

Cris Luengo
luengo@ieee.org

Optical image formation

- Image formation system (e.g. lenses) creates a band-limited image — imposes resolution

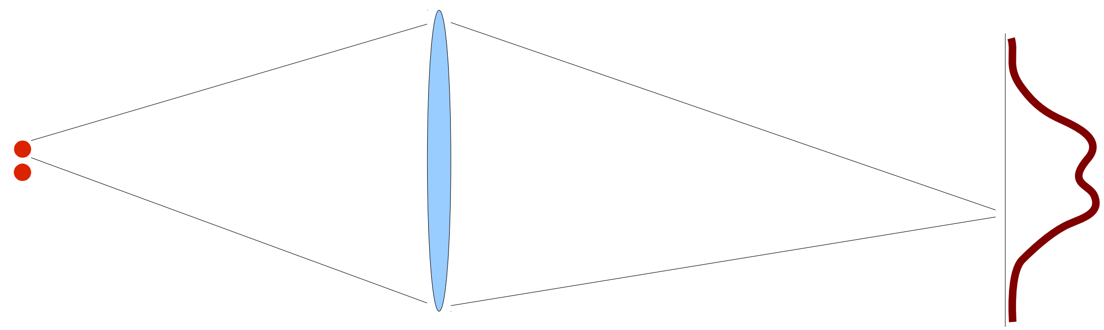

- Standard optics' point-spread function (PSF) can be approximated by a Gaussian
  - ideal lens has Airy function for PSF
  - but lens imperfections are unavoidable

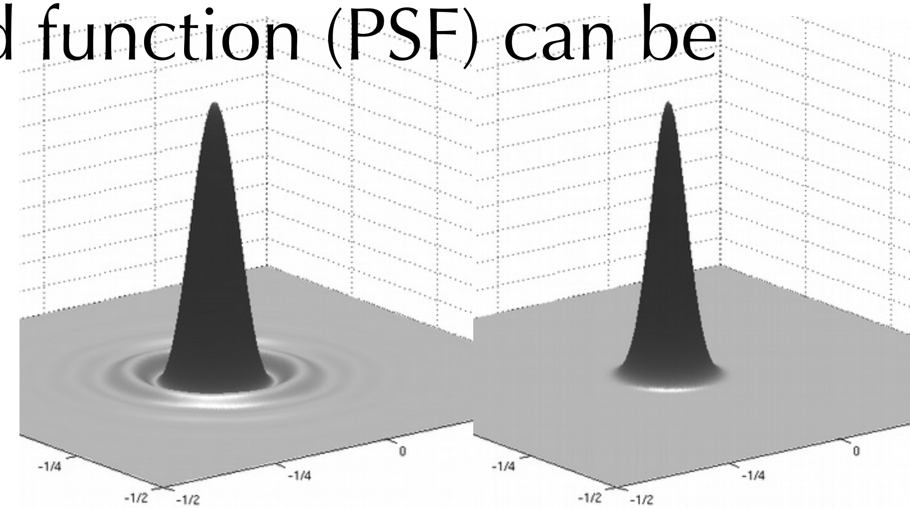

Cris Luengo
luengo@ieee.org

Optical image formation

- The image is smoothed by a PSF (convolution!) before sampling

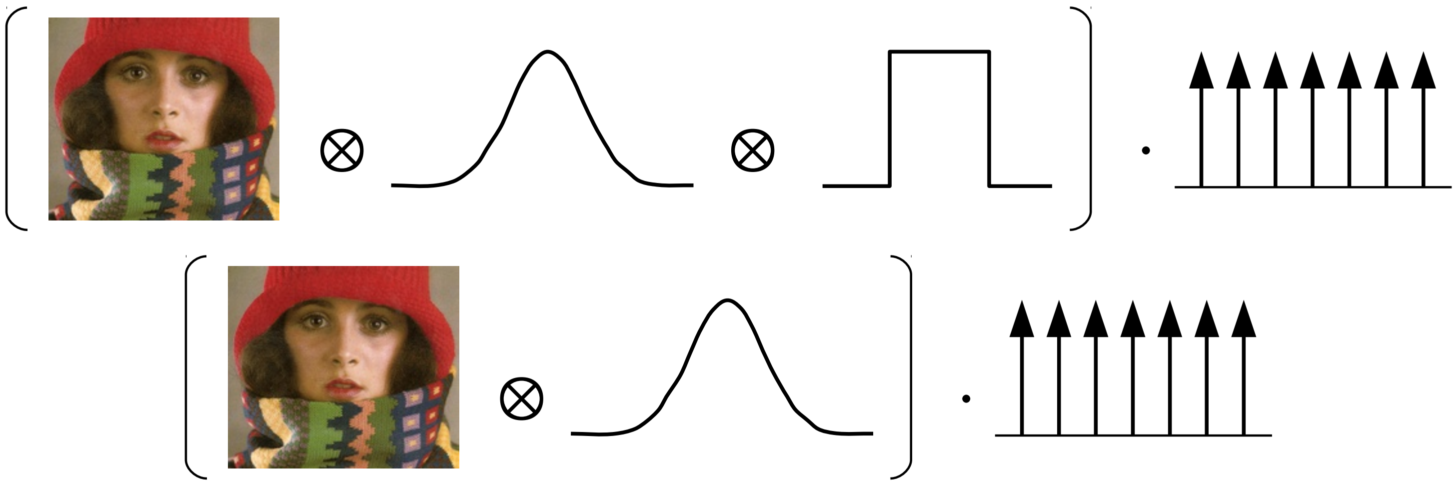

- Neither the smoothing nor the sampling change the total amount of light in the image

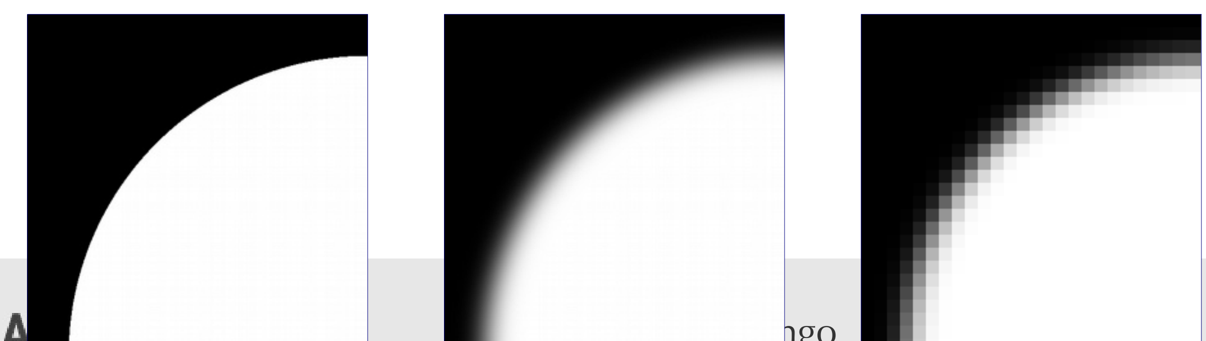

What happens in the Fourier domain

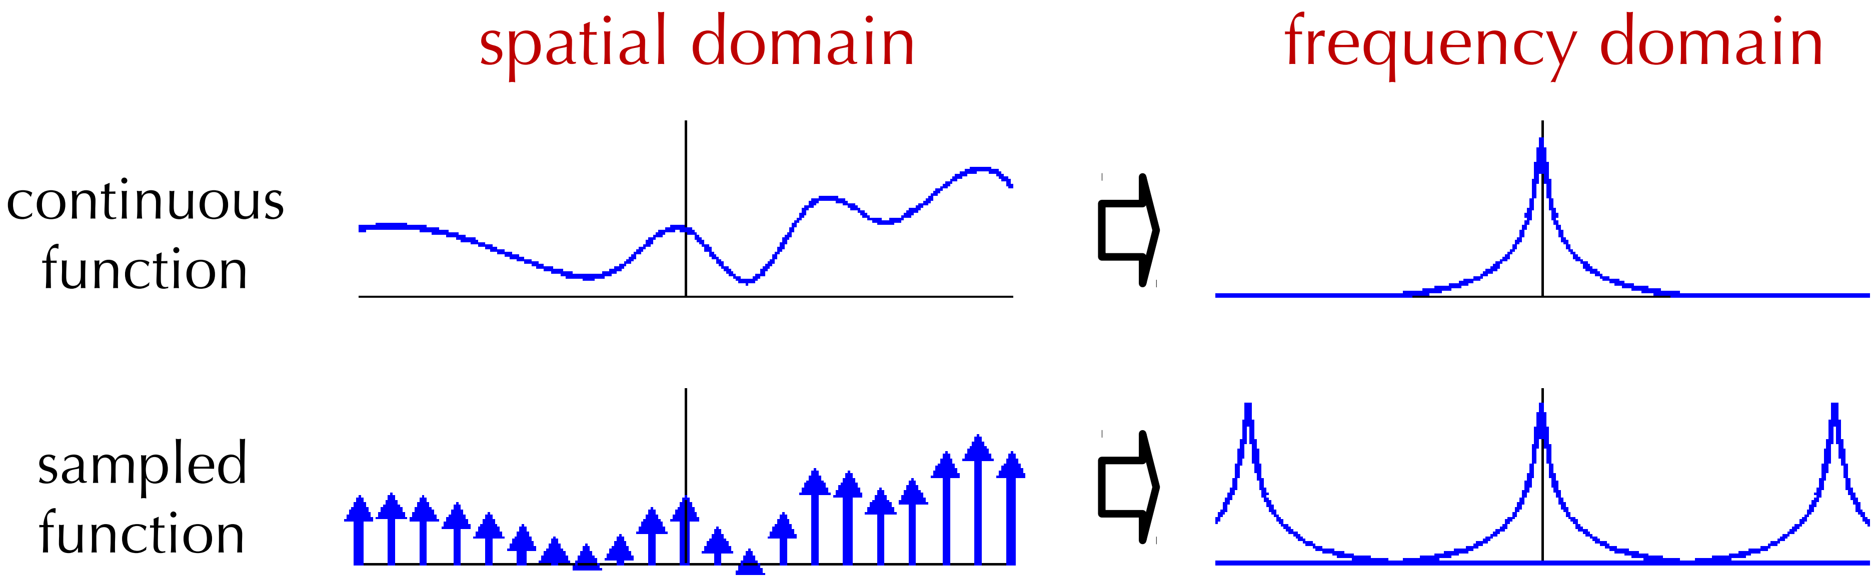

- The 0th frequency is proportional to the total amount of light
- 0th frequency is unaltered by sampling
- Sum of samples is equal (proportional) to integral over continuous function

Cris Luengo
luengo@ieee.org

What happens in the Fourier domain

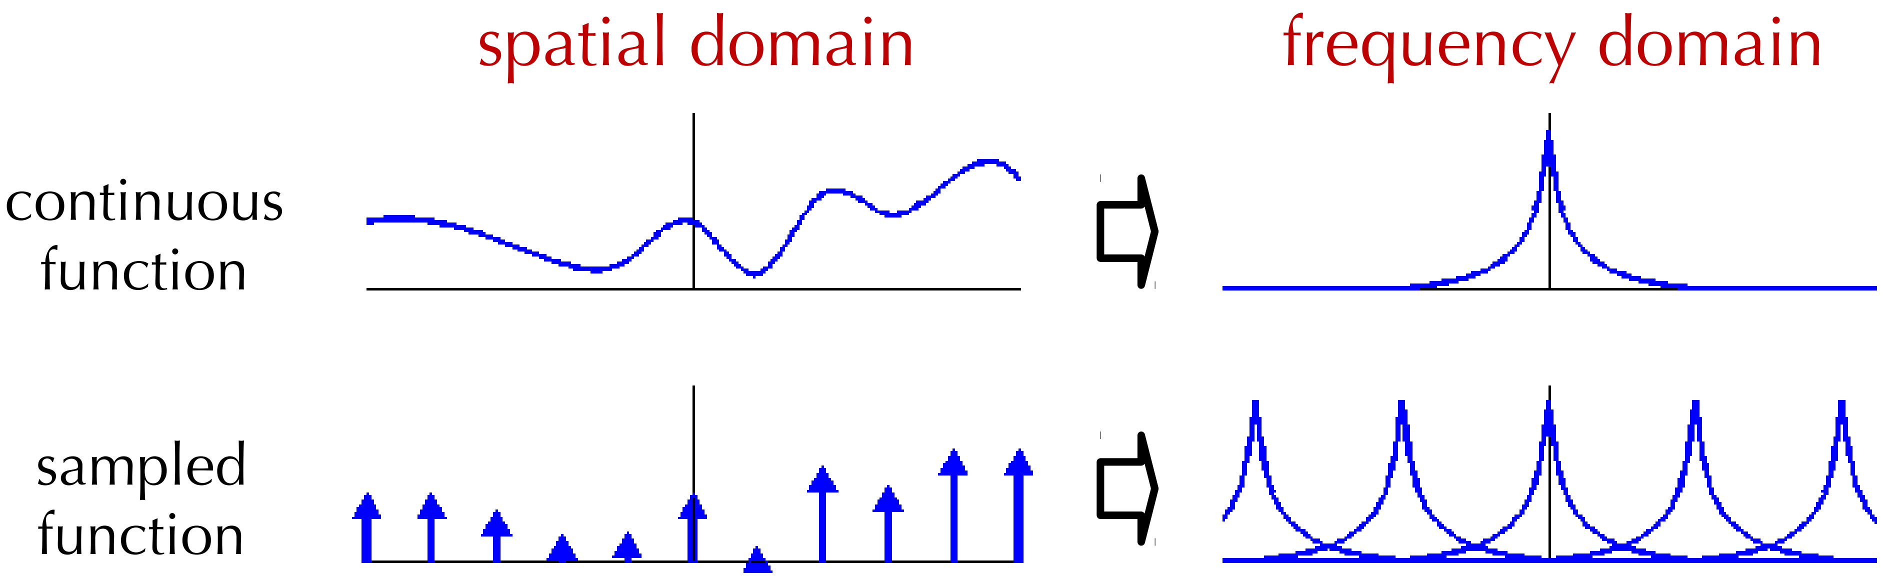

- But: aliasing can affect the 0th frequency!
- Sum of samples is equal (proportional) to integral over continuous, band-limited function <u>if sampled correctly</u>

Soft clipping

Threshold vs. soft clipping

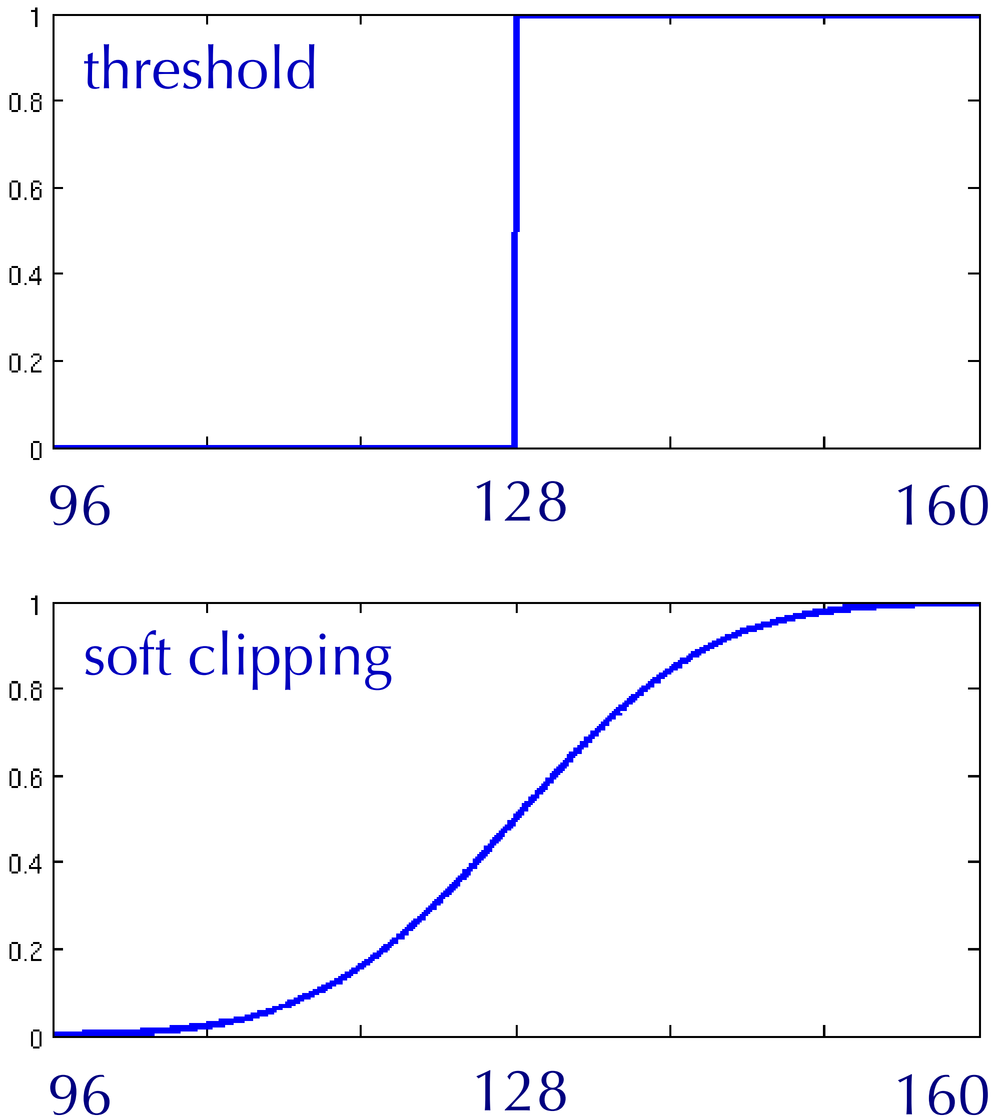

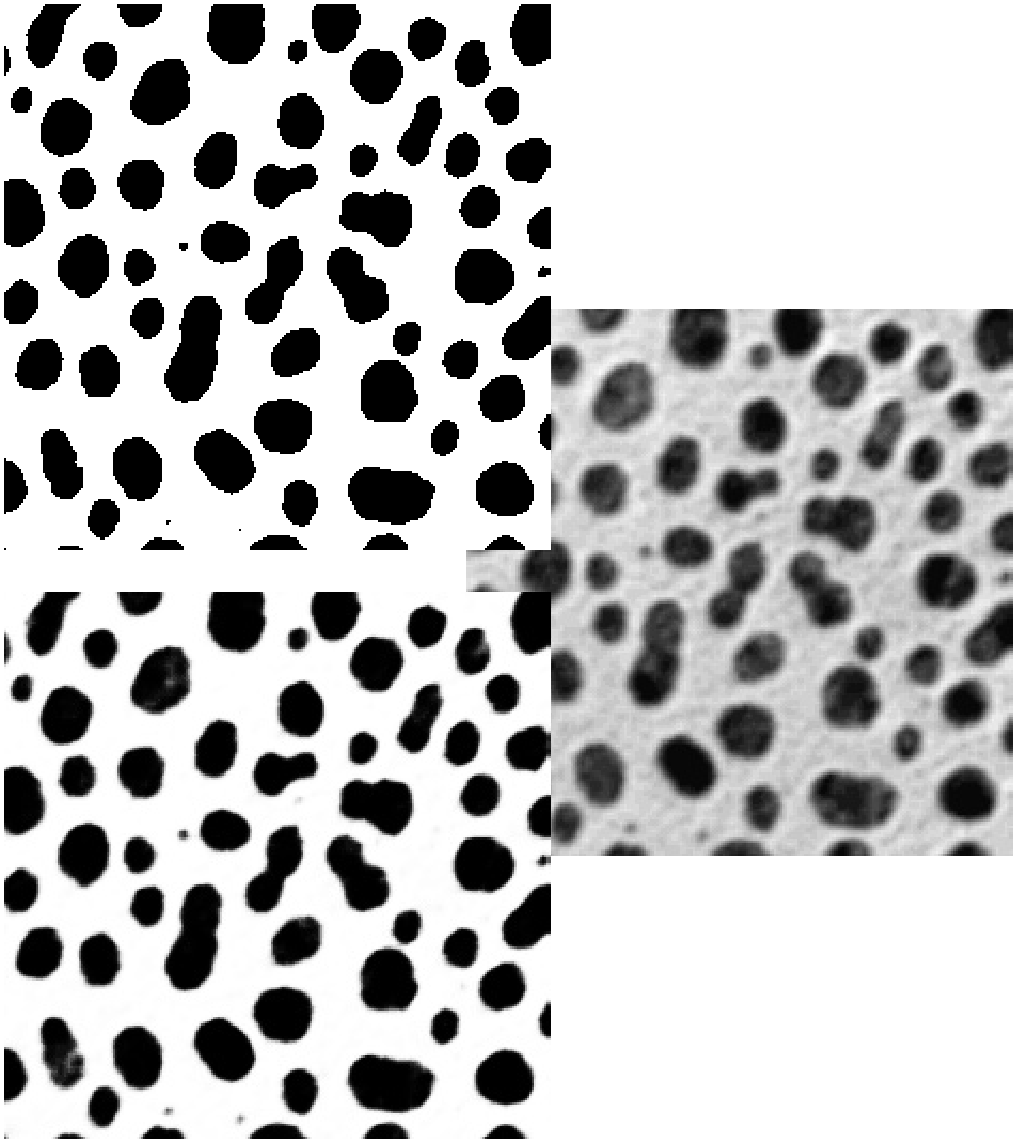

Threshold vs. soft clipping

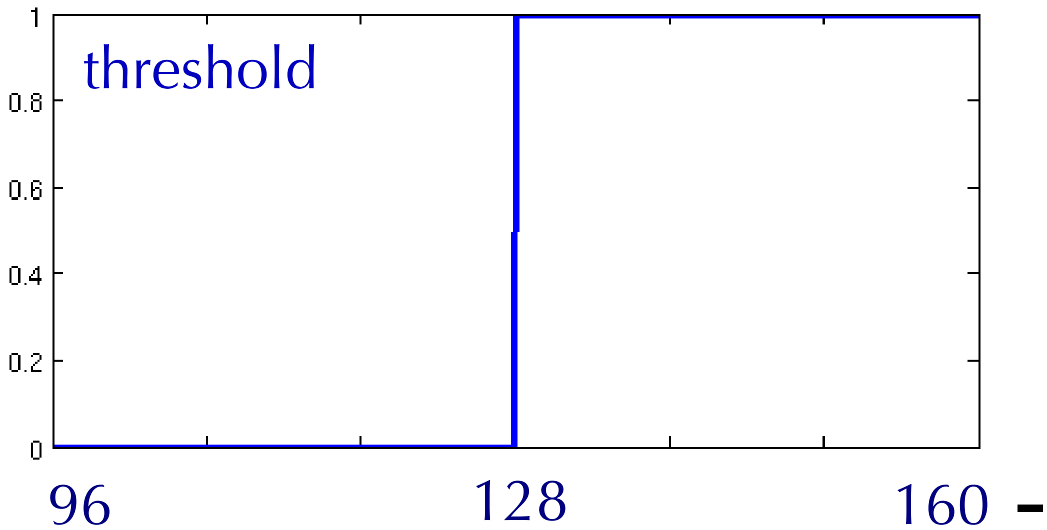

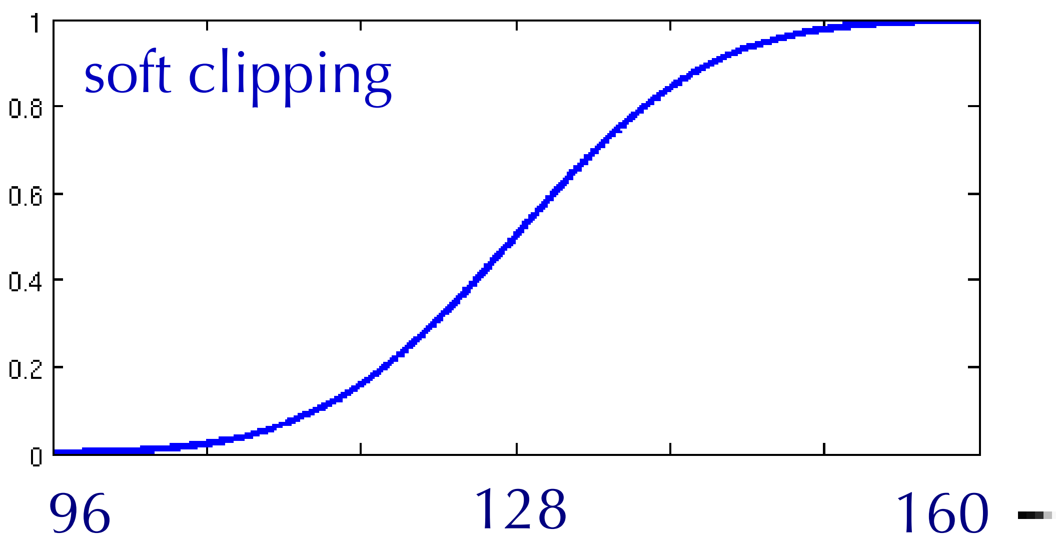

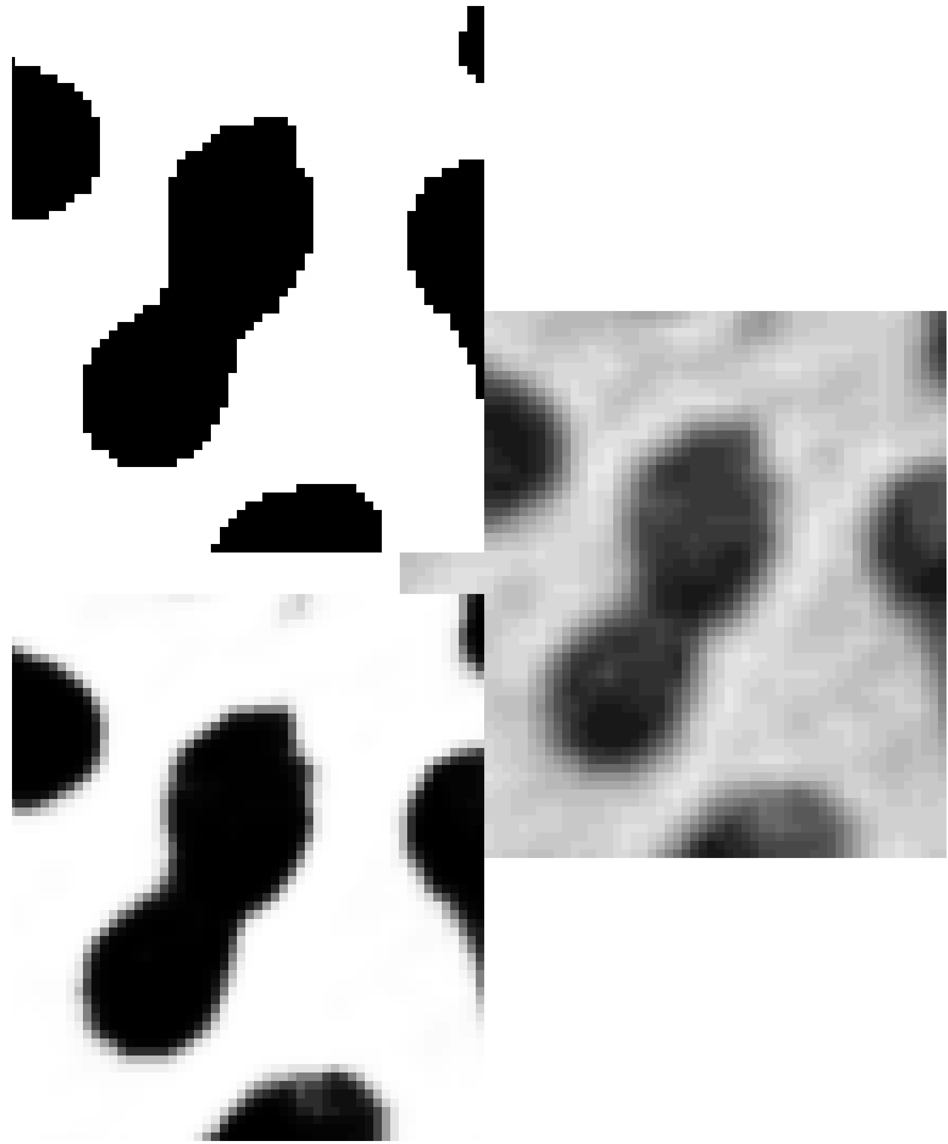

Cris Luengo
luengo@ieee.org

Soft clipping

Selecting a proper range is important:

- too small: introduction of aliasing
- too large: background and foreground not uniform

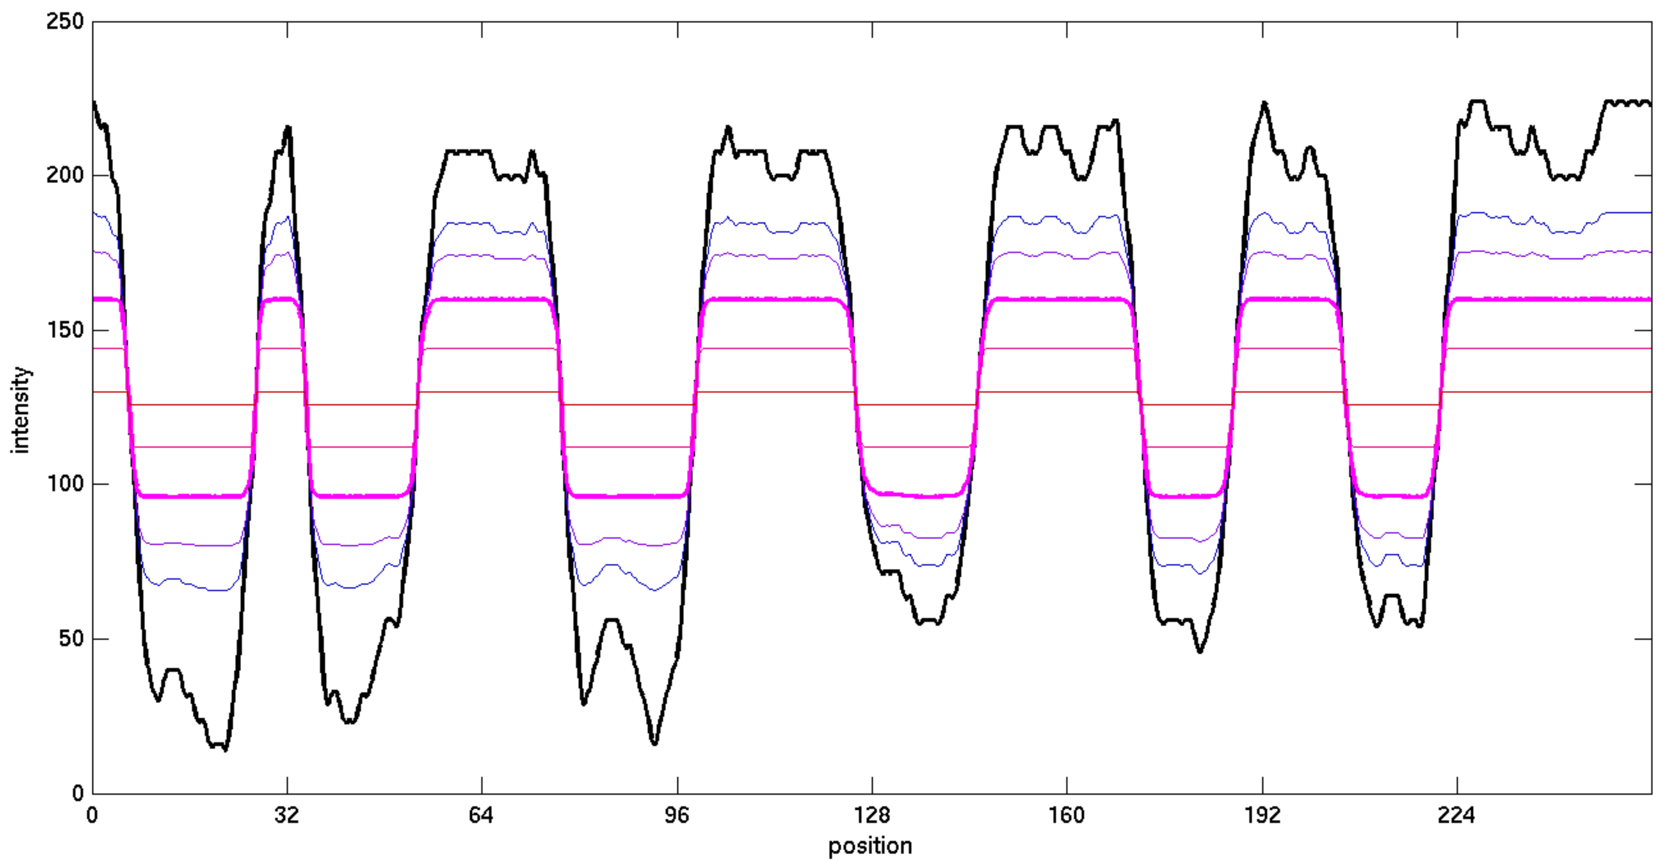

Soft clipping

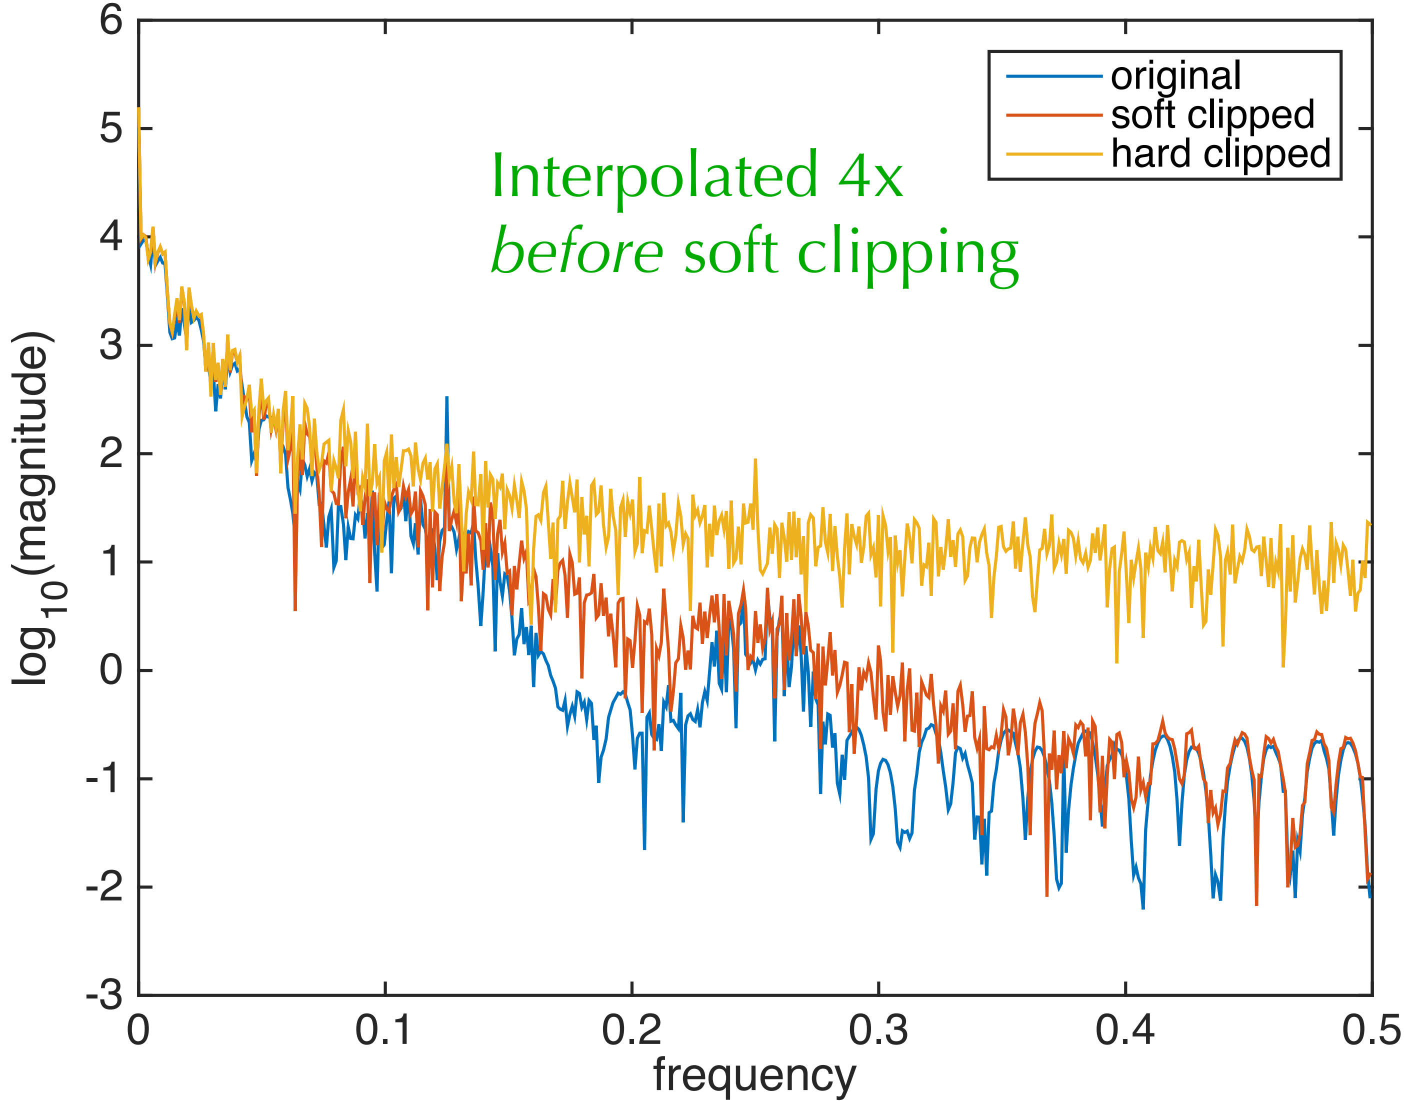

Gaussian filtering

Cris Luengo
luengo@ieee.org

Low-pass filters

The typical "box filter"

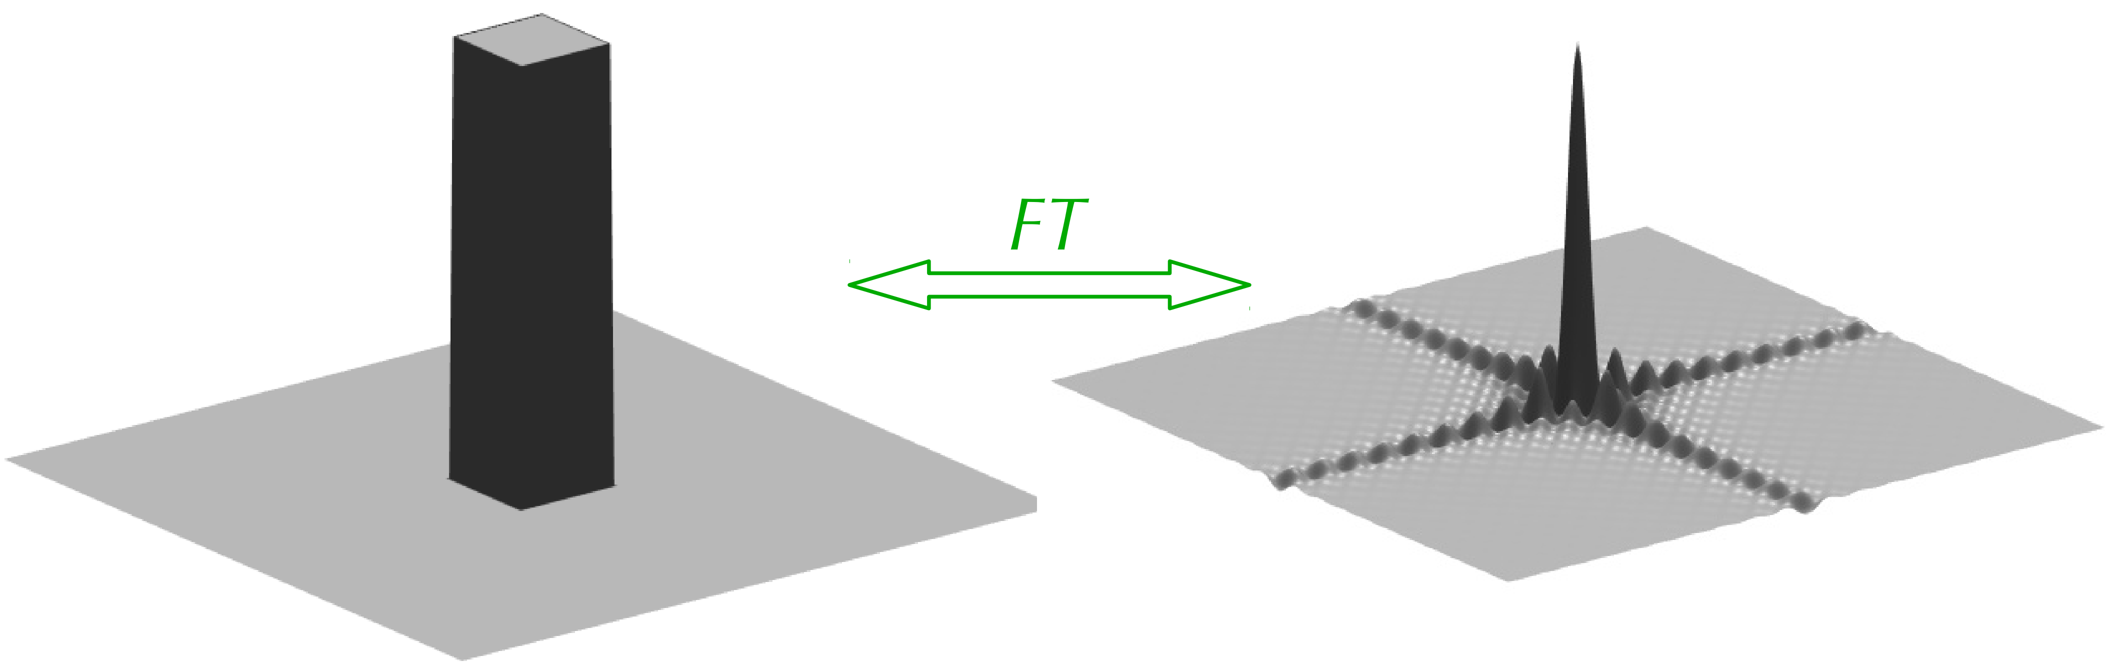

Spatial domain

Frequency domain

Centre for Image Analysis
Uppsala University

Cris Luengo
luengo@ieee.org

UPPSALA
UNIVERSITET

Low-pass filters

The “ideal” low-pass filter

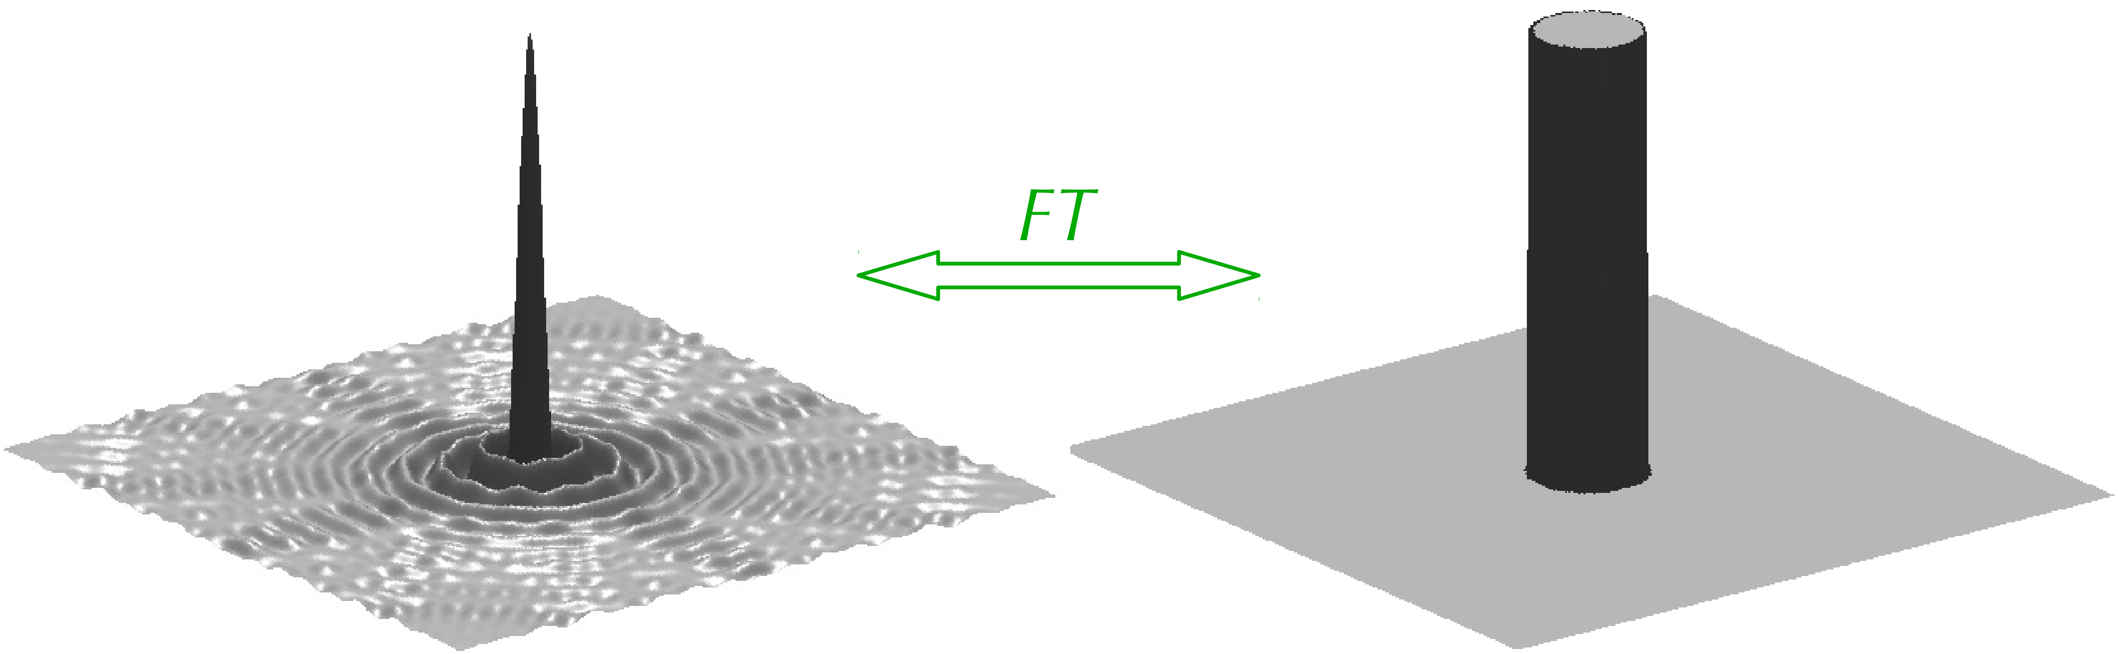

Spatial domain

Frequency domain

Cris Luengo
luengo@ieee.org

Low-pass filters

The Gaussian filter

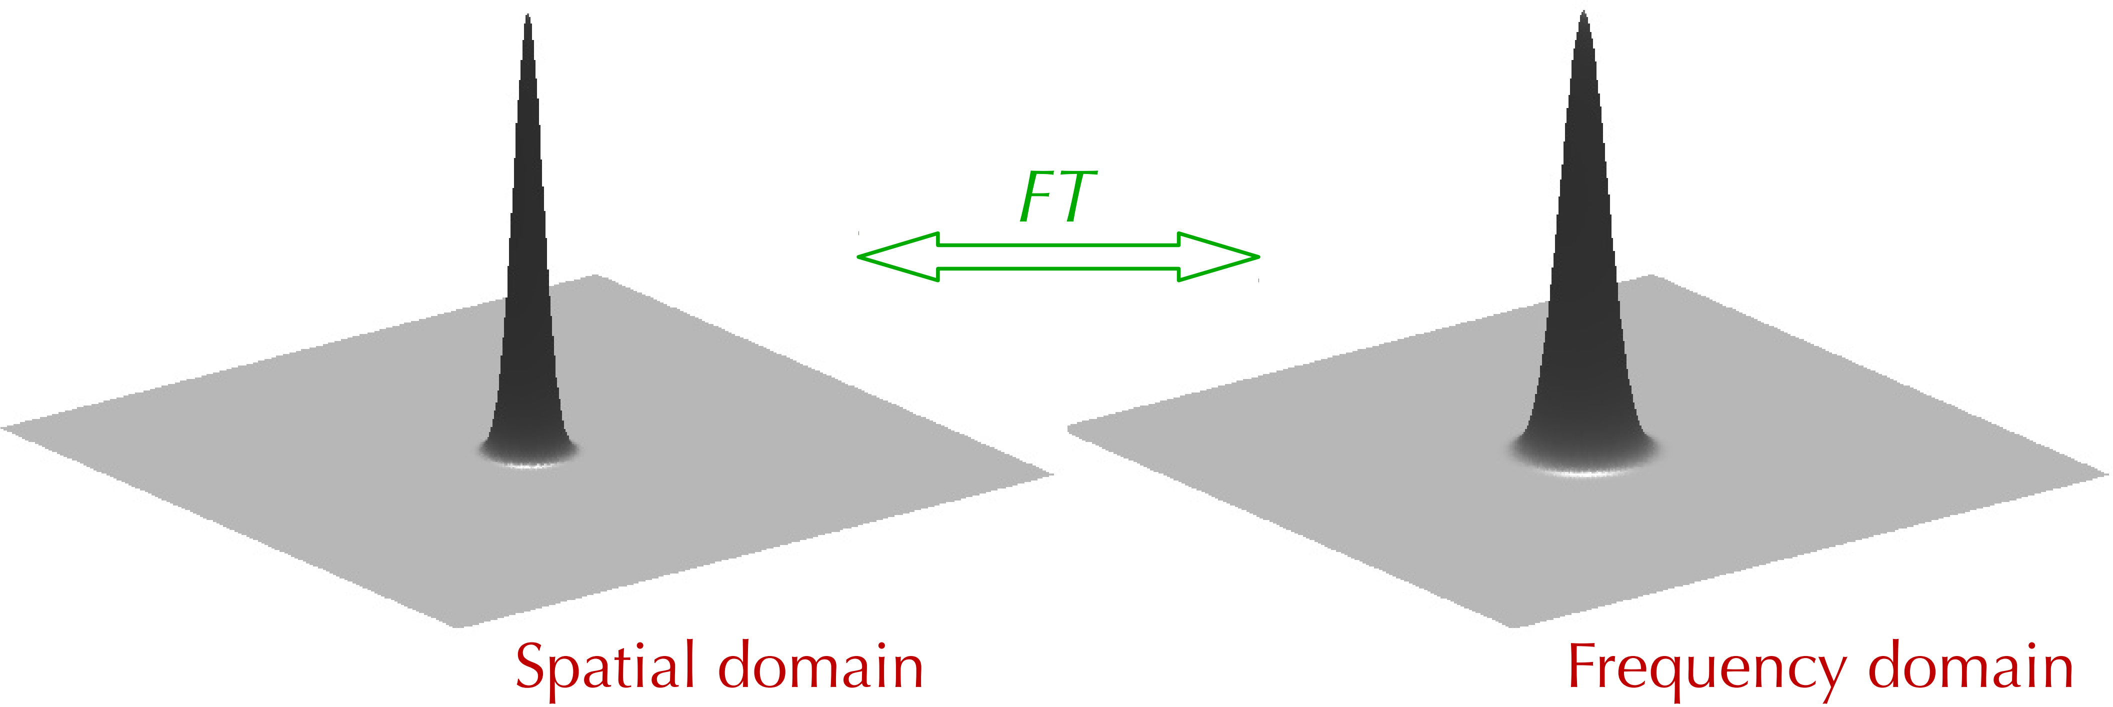

Cris Luengo
luengo@ieee.org

Gaussian filters

- Provide optimal compromise between compactness in spatial and frequency domain
- Isotropic but separable
- Derivatives computed with derivative of Gaussian:

$$\frac{\partial}{\partial x}\left(f \otimes G\right) = f \otimes \frac{\partial}{\partial x} G$$

Cris Luengo
luengo@ieee.org

Area in 2D images
Volume in 3D images

Cris Luengo
luengo@ieee.org

2D area / 3D volume

- The equivalent of counting pixels in the binary case: integrating intensity
- Unbiased estimate for size if:
  - pixels inside the object have value 1
  - pixels outside the object have value 0
  - pixels on the boundary have a grey-value given by the point-spread function of the optical system
- Low-pass filtering the image does not modify the result
  - (depending on boundary condition)

Cris Luengo
luengo@ieee.org

2D area (ideal case)

Area of 100 disks (r = 21 px) with sub-pixel shifts

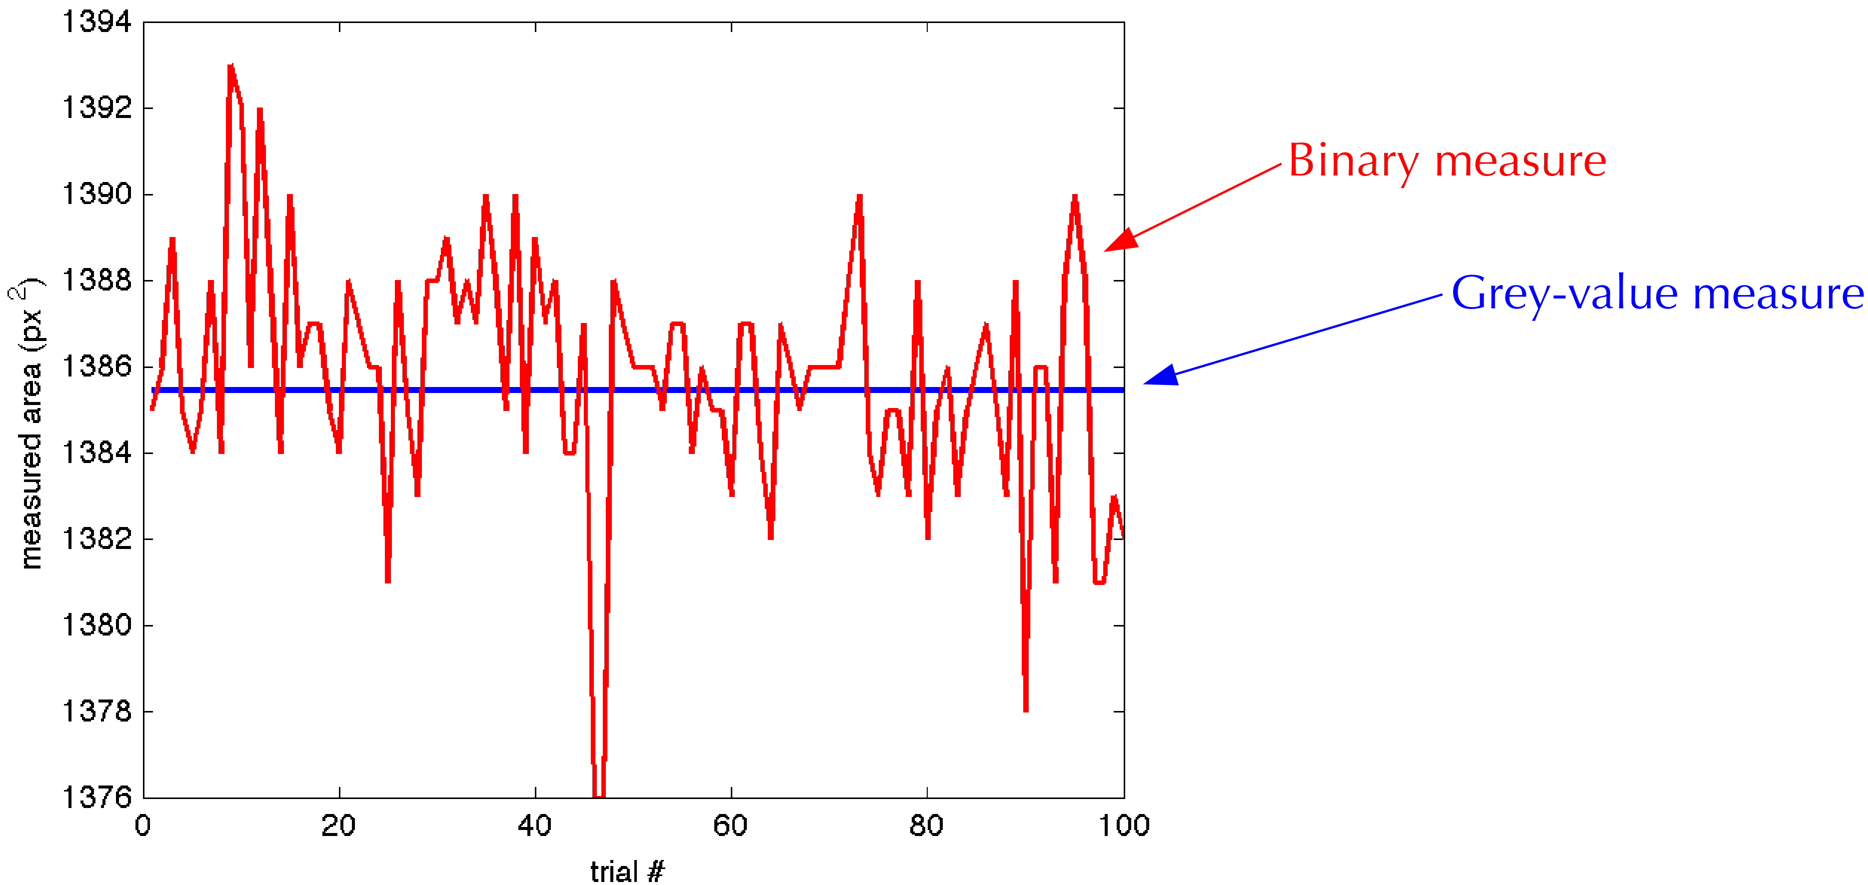

Expected measure: 1385.442360 px^2

Grey-value measure: 1385.442352 ± 0.000001 px^2 (std = 0.000006)

Binary measure: 1385.8 ± 0.6 px^2 (std = 2.9)

2D area (with soft clipping)

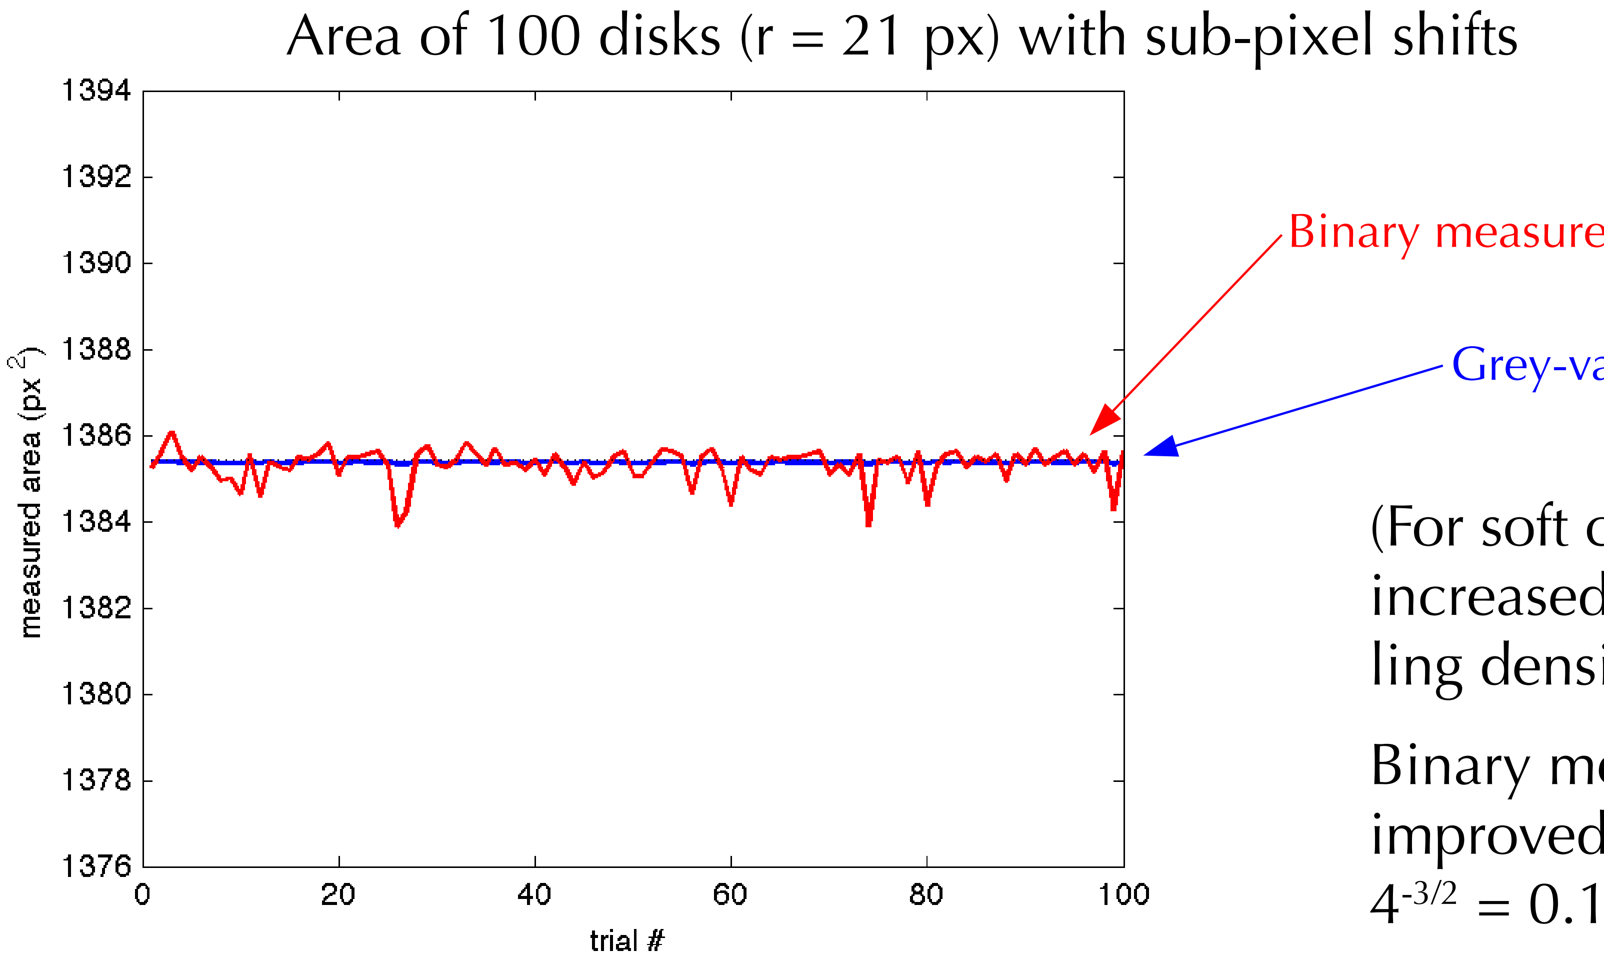

(For soft clipping I increased the sampling density 4 times)

Binary measure improved by factor $4^{-3/2} = 0.125$

Expected measure: 1385.442360 px^2

Grey-value measure: 1385.370 ± 0.003 px^2 (std = 0.013)

Binary measure: 1385.31 ± 0.08 px^2 (std = 0.39)

Perimeter in 2D images
Surface area in 3D images

Cris Luengo
luengo@ieee.org

Perimeter

- Given a grey-value object with a constant intensity *H*
- If we extend the object by a fixed distance *D*, the volume of the extension is given by: *P D H*
  (*P* = perimeter)

$$D\,H = \int_x \left(f(x) - f(x+D) \right) dx$$

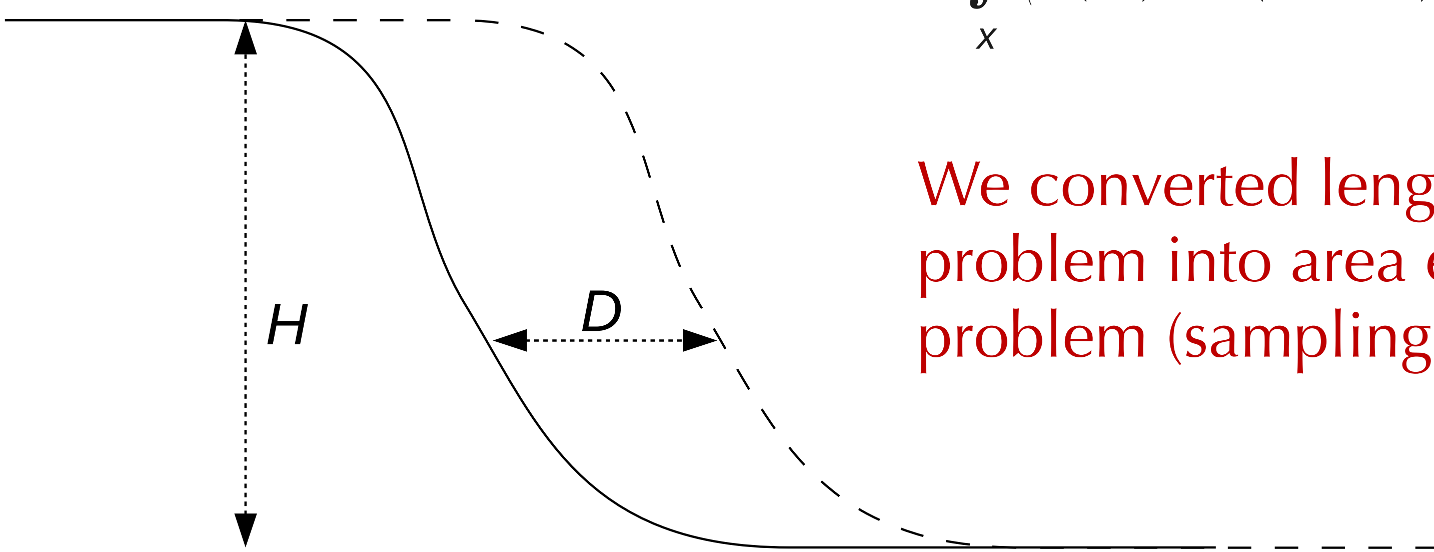

We converted length estimation problem into area estimation problem (sampling-invariant!)

Cris Luengo
luengo@ieee.org

Perimeter

- Given a grey-value object with a constant intensity *H*
- If we extend the object by a fixed distance **ε**, the volume of the extension is given by: *P* **ε** *H*
  (*P* = perimeter)

$$P = \frac{1}{H} \int_x \lim_{\varepsilon \to 0} \frac{f(x) - f(x + \varepsilon)}{\varepsilon} dx$$

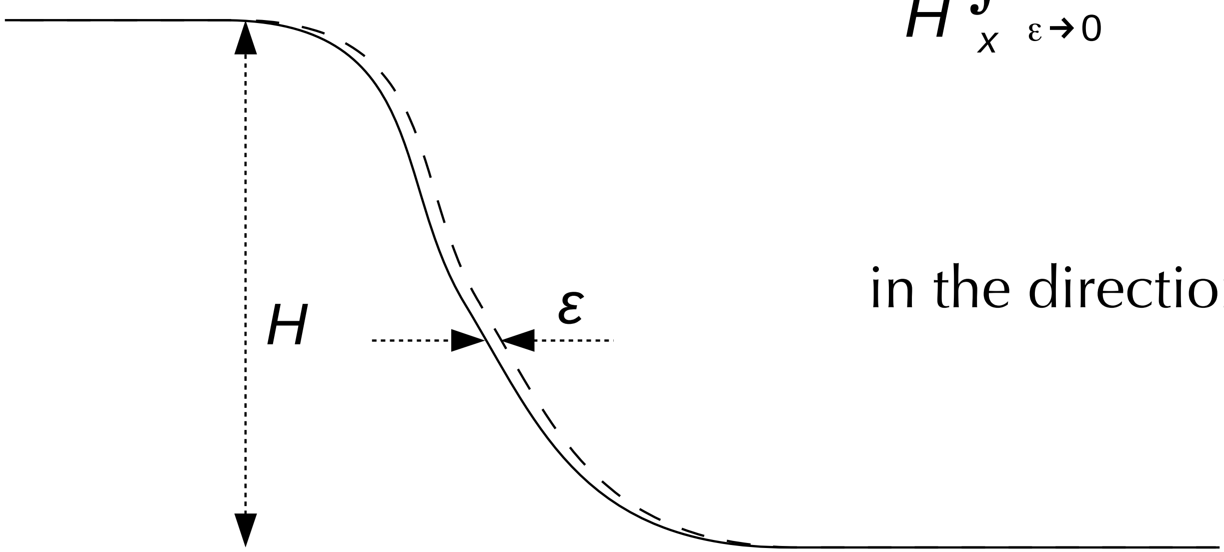

in the direction of the gradient....

Cris Luengo
luengo@ieee.org

Perimeter

- Soft clipping
- Gradient magnitude $\sqrt{(\frac{\partial f}{\partial x})^2+(\frac{\partial f}{\partial y})^2}$
- Integration (sum)

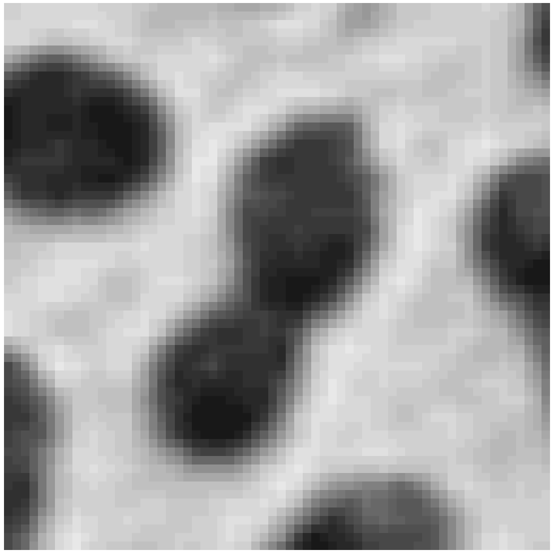

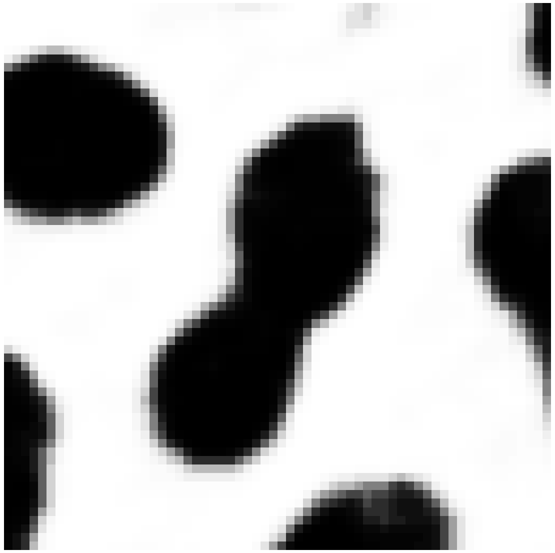

Perimeter

Perimeter of 100 disks (r = 21 px) with sub-pixel shifts

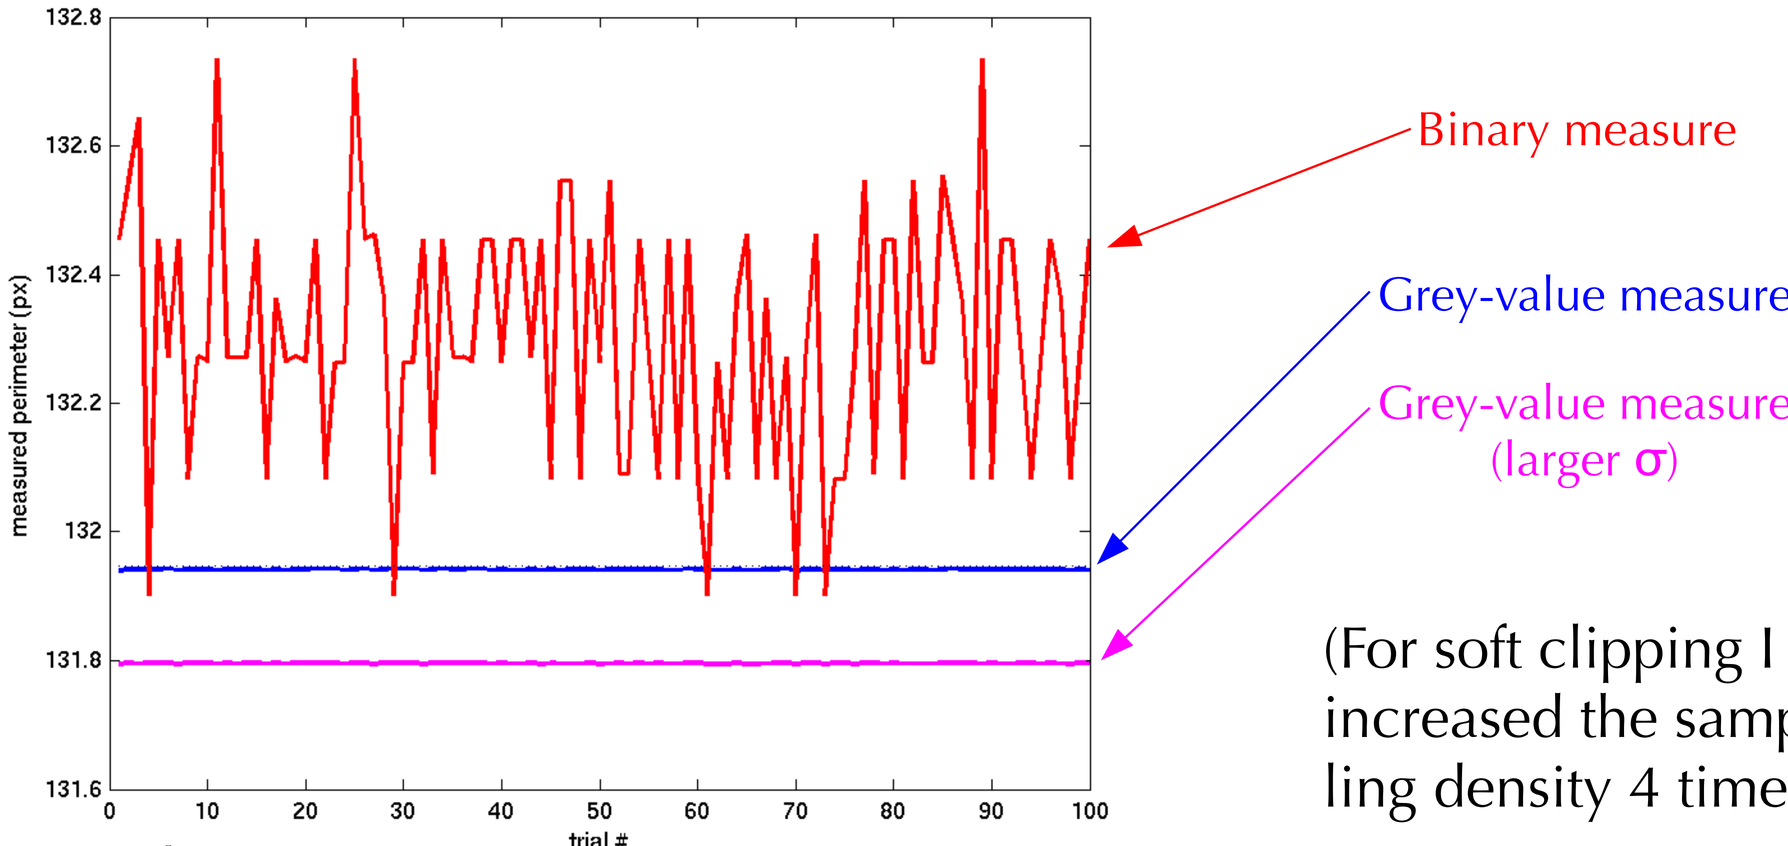

(For soft clipping I increased the sampling density 4 times)

Expected measure:	131.946891 px	
Grey-value measure:	131.9415 ± 0.0001 px	(std = 0.0006)
with larger σ:	131.7958 ± 0.0001 px	(std = 0.0006)
Binary measure:	132.04 ± 0.01 px	(std = 0.07)

Curvature
Bending energy
(2D)

Cris Luengo
luengo@ieee.org

Isophote curvature

- Contour direction: $\vec{c}=(-f_y, f_x)$

$$\theta=\arccos\left(\frac{-f_y}{|g|}\right)=\arcsin\left(\frac{f_x}{|g|}\right)=\arctan\left(\frac{-f_y}{f_x}\right)$$

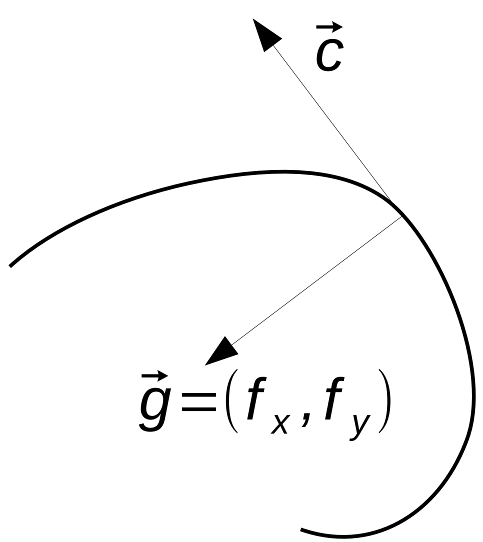

- To differentiate along the curve:

$$\frac{\mathrm{d}}{\mathrm{d}s}=\cos\theta\frac{\partial}{\partial x}+\sin\theta\frac{\partial}{\partial y}=\frac{-f_y}{|g|}\frac{\partial}{\partial x}+\frac{f_x}{|g|}\frac{\partial}{\partial y}$$

- Curvature κ = derivative of θ along the curve

$$\kappa=\frac{\mathrm{d}\theta}{\mathrm{d}s}=-\frac{f_{xx}f_y^2-2f_{xy}f_xf_y+f_{yy}f_x^2}{\left(f_x^2+f_y^2\right)^{3/2}}=\frac{-f_{cc}}{|g|}$$

Bending energy

- Given by integral along the perimeter of the square of curvature
- Integrate along perimeter by multiplying by |g| and integrating over the image

$$\text{B.E.} = \int_{\text{contour}} \kappa^2 \, \mathrm{d}s = \iint_{\text{image}} \kappa^2 |g| \, \mathrm{d}x \, \mathrm{d}y = \iint_{\text{image}} \frac{f_{cc}^2}{|g|} \, \mathrm{d}x \, \mathrm{d}y$$

Cris Luengo
luengo@ieee.org

Bending energy

- Given by integral along the perimeter of the square of curvature
- Integrate along perimeter by multiplying by |g| and integrating over the image

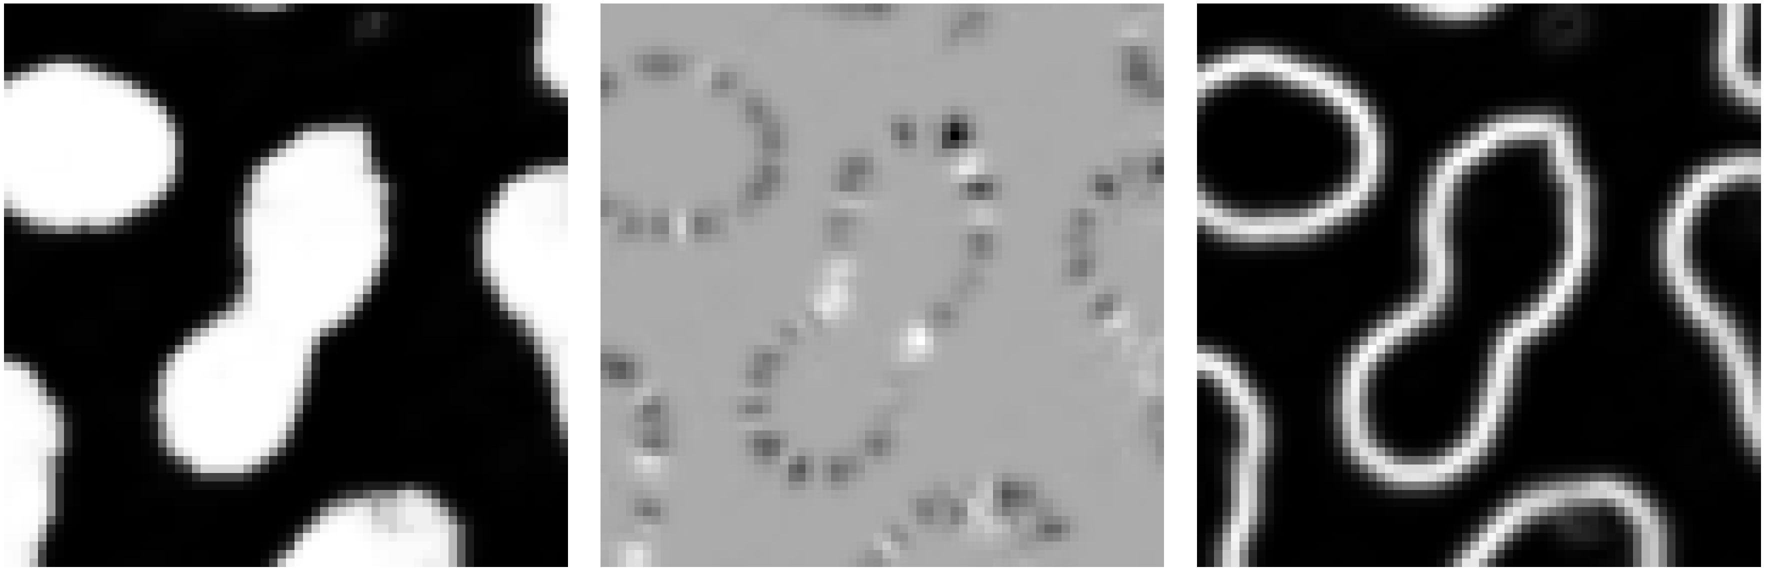

soft clip f_{cc} |*g*|

Cris Luengo
luengo@ieee.org

Counting

Cris Luengo
luengo@ieee.org

Euler number

- In 2D: # of objects – # of holes
- In 3D: # of objects – # of tunnels + # of cavities

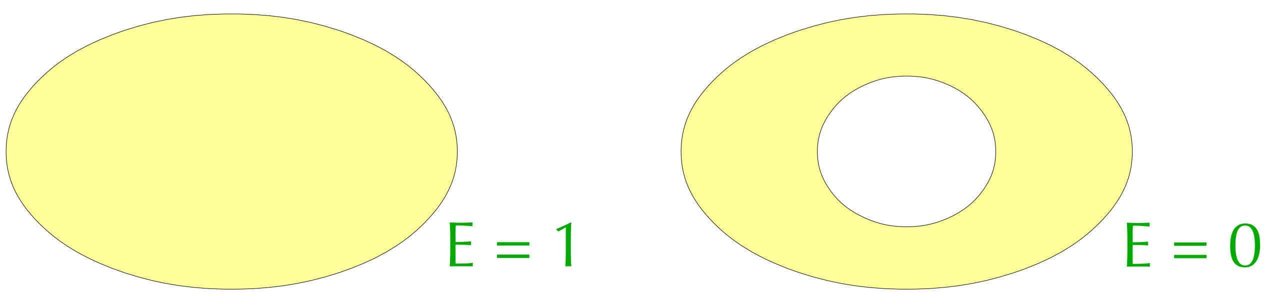

Cris Luengo
luengo@ieee.org

Euler number

- Integral of curvature along a closed contour is always 2π, a hole in an object contributes with -2π

$$\text{Euler number} = \frac{1}{2\pi}\int_{\text{contour}} \kappa \, \mathrm{d}s = \frac{1}{2\pi}\iint_{\text{image}} \kappa |g| \, \mathrm{d}x \, \mathrm{d}y = \frac{1}{2\pi}\iint_{\text{image}} -f_{cc} \, \mathrm{d}x \, \mathrm{d}y$$

- Integral of second derivative in gradient direction also yields a constant 2π for a closed contour

$$\text{Euler number} = \frac{1}{2\pi}\iint_{\mathit{image}} f_{gg} \, \mathrm{d}x \, \mathrm{d}y$$

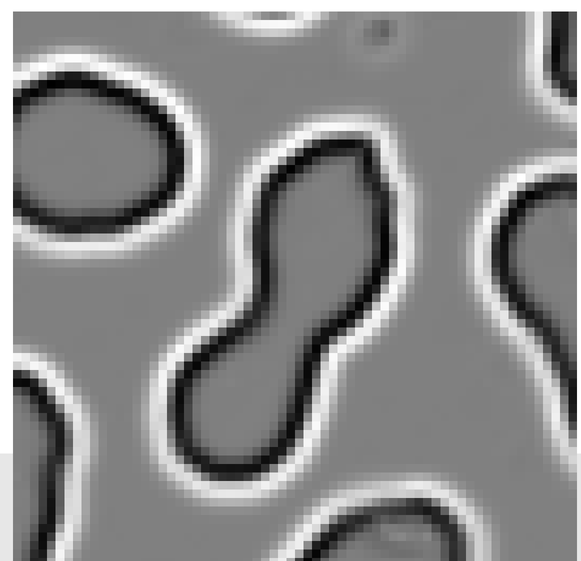

Summary

- Area/volume = integral over image
- Perimeter/surface area
  - obtained by converting to area measurement problem
- Curvature
  - computed through 2nd derivative along contour
  - bending energy & Euler number
- Prepare image by soft clipping
  - (equivalent to thresholding, but without loss of band limitation)

Tutorial on Image-Based Measurement

Part IV: Granulometries

Cris Luengo
luengo@ieee.org

The segmentation problem

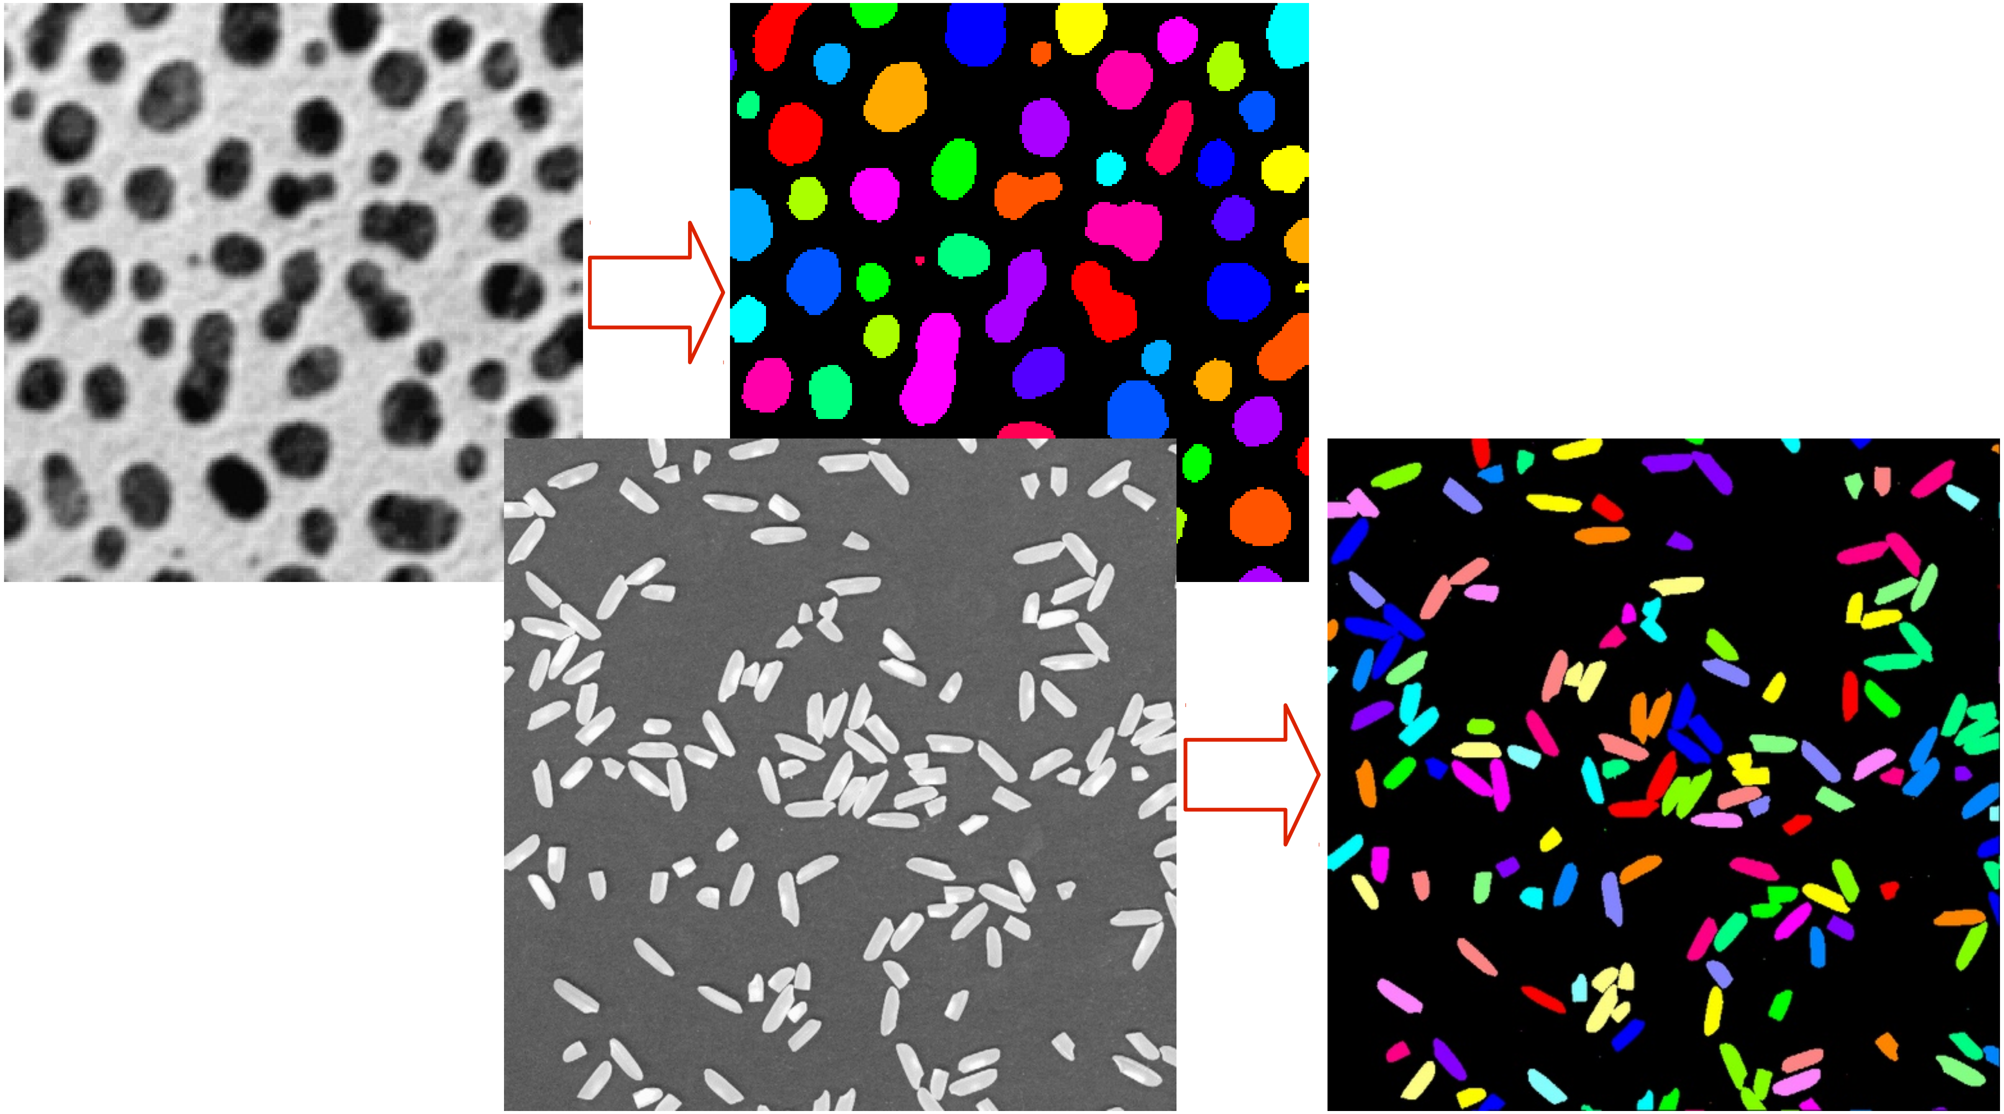

Cris Luengo
luengo@ieee.org

Overview

- The closing and opening
  - remove objects from image by size
- The granulometry
  - estimates size distributions without segmentation
- The path closing and opening
  - removes objects from image by length
  - use with granulometry to estimate length distributions
  - I will also discuss recent and current improvements

Cris Luengo
luengo@ieee.org

Using an image as a landscape

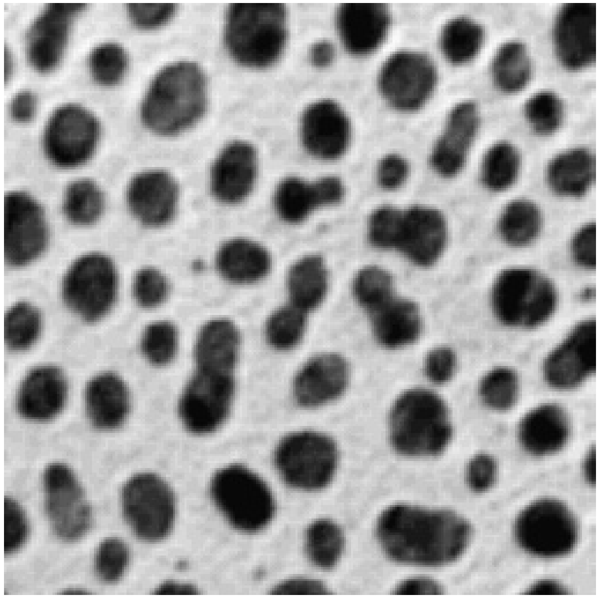

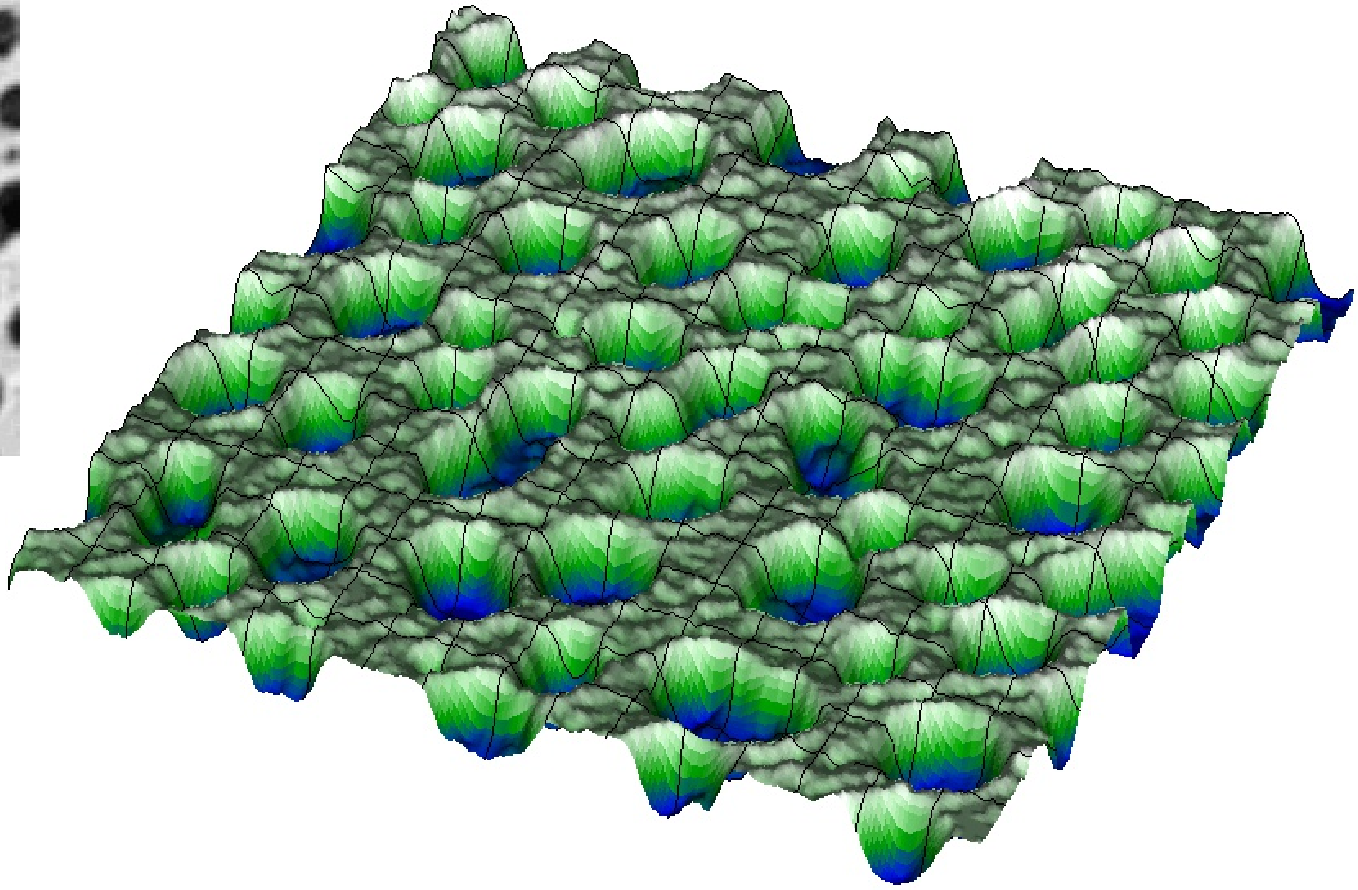

Cris Luengo
luengo@ieee.org

Basic operations math. morphology

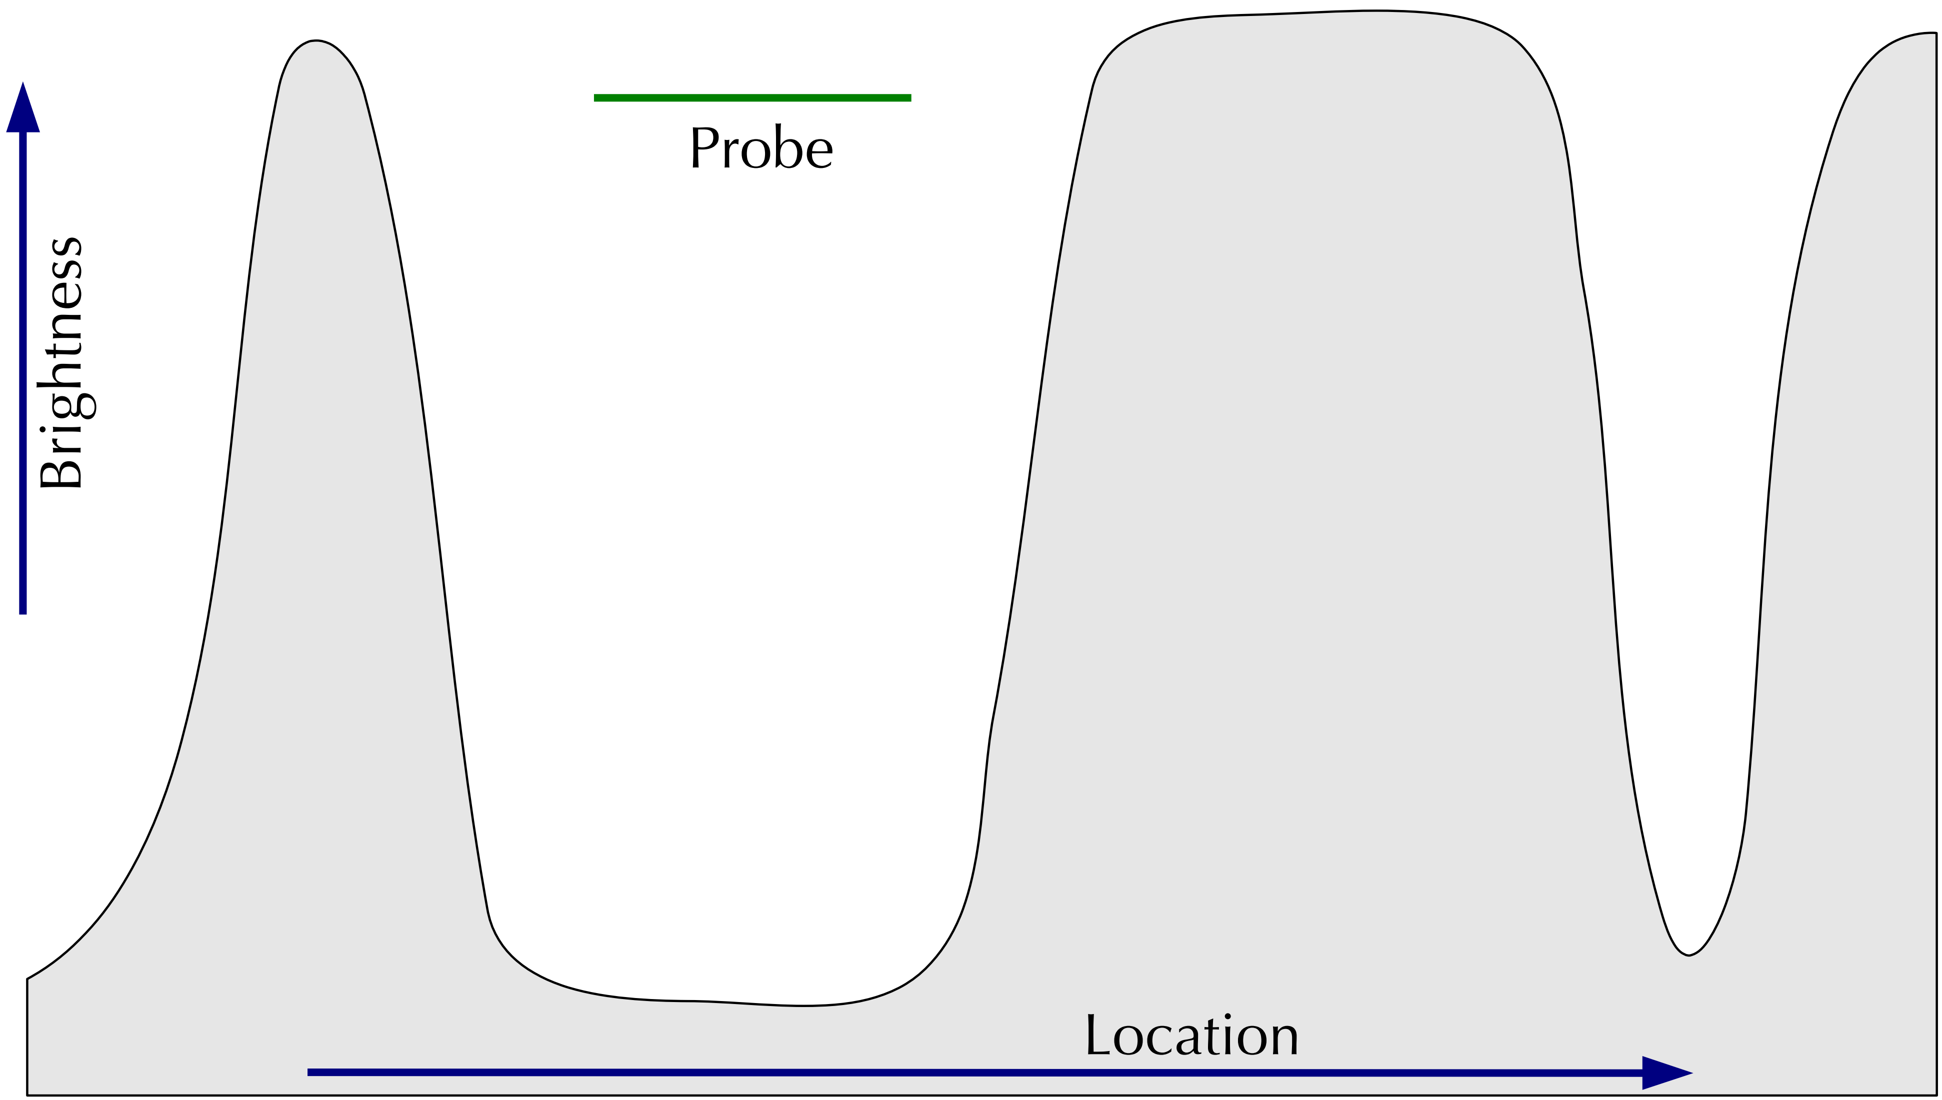

Cris Luengo
luengo@ieee.org

Erosion = local minimum

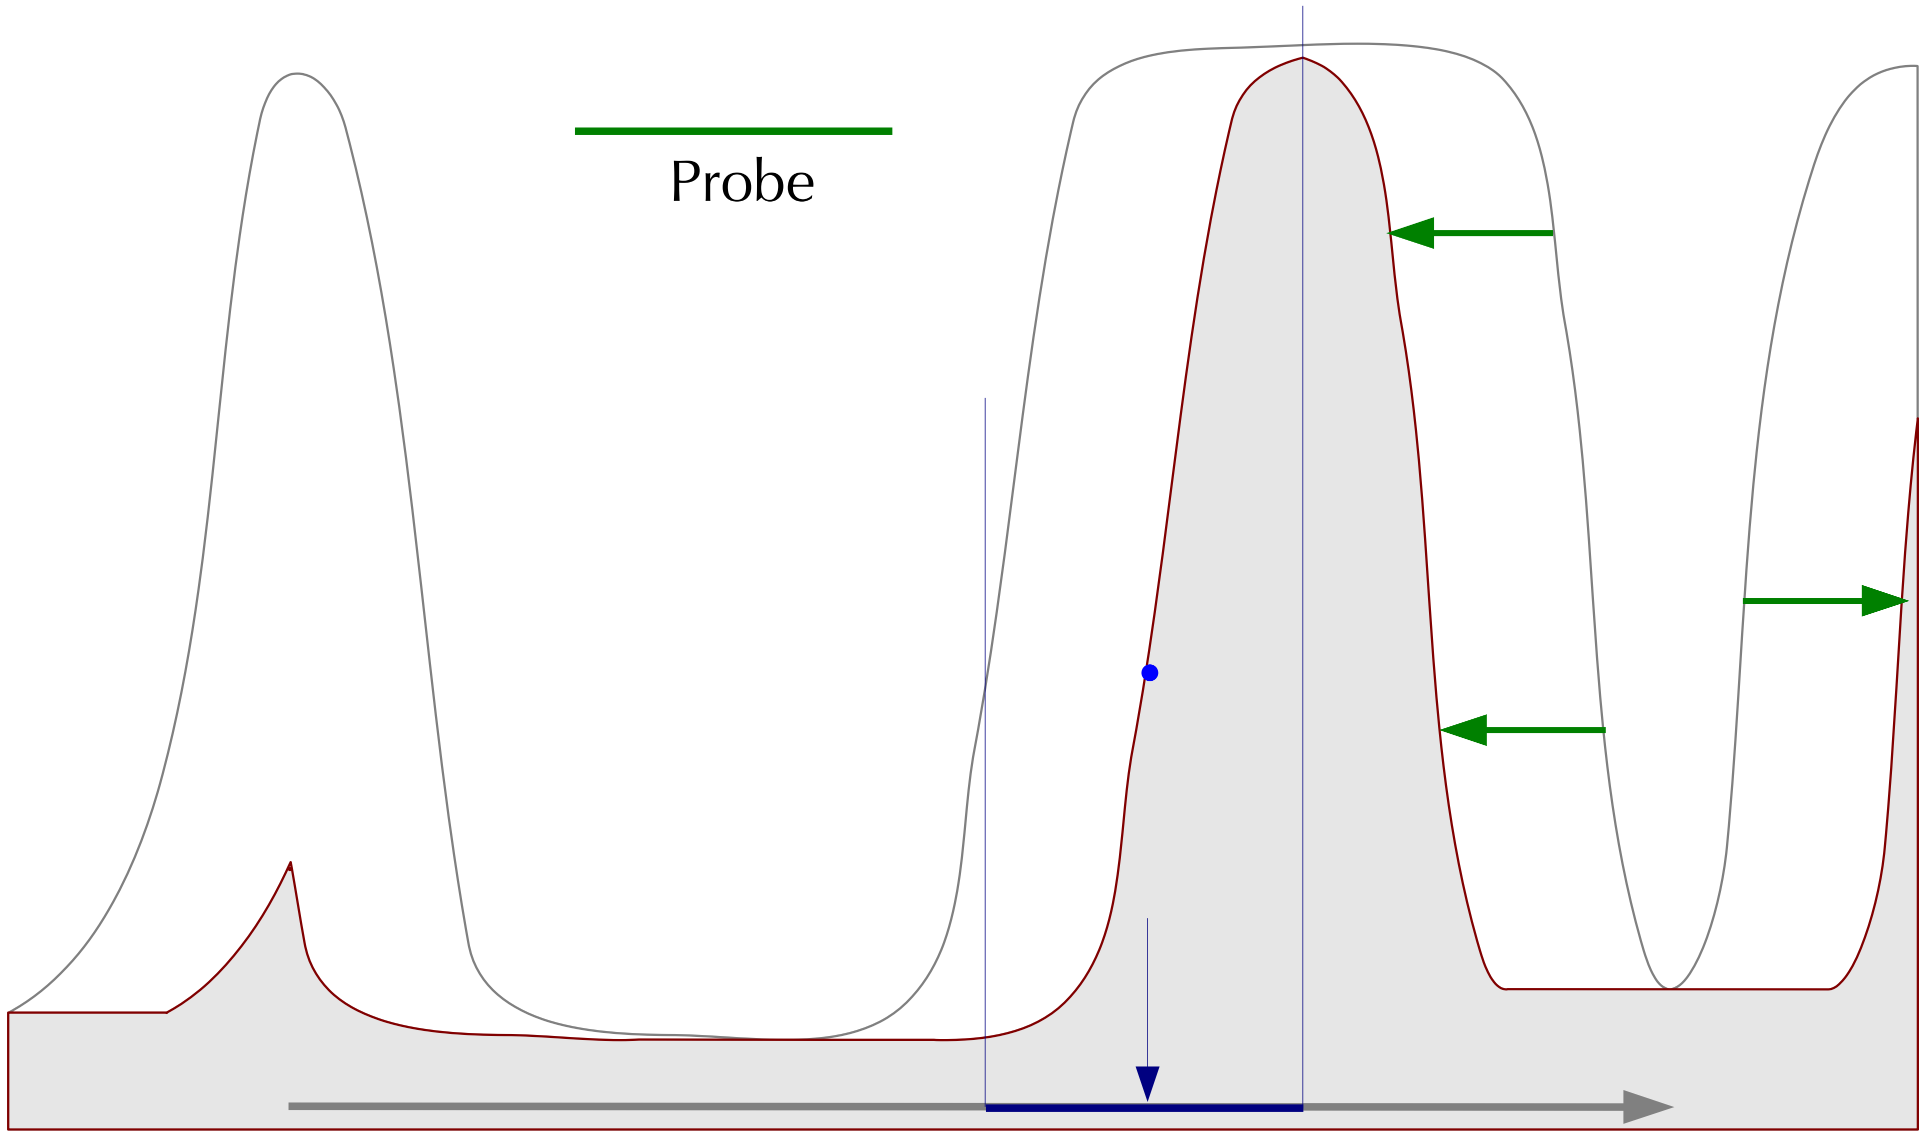

Erosion = local minimum

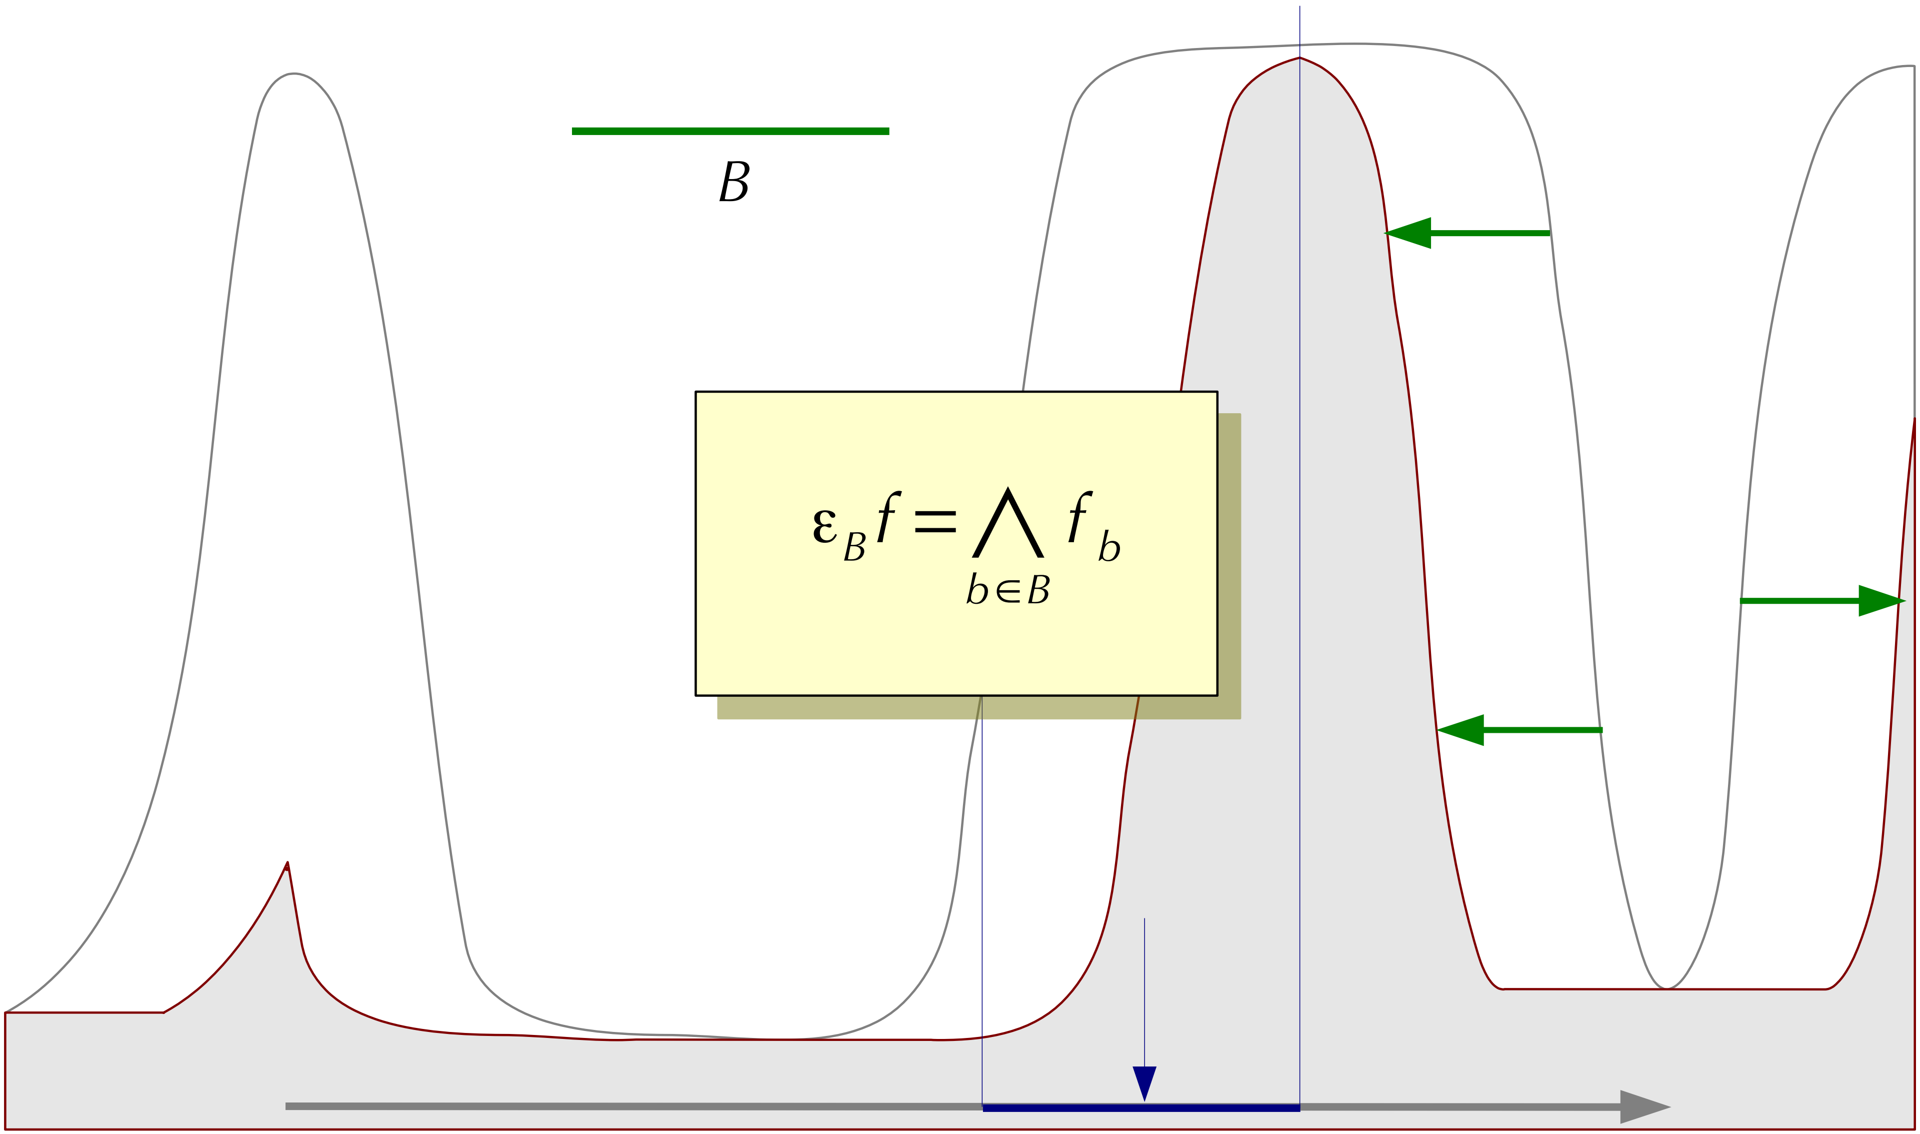

$$\varepsilon_B f = \bigwedge_{b \in B} f_b$$

Cris Luengo
luengo@ieee.org

Dilation = local maximum

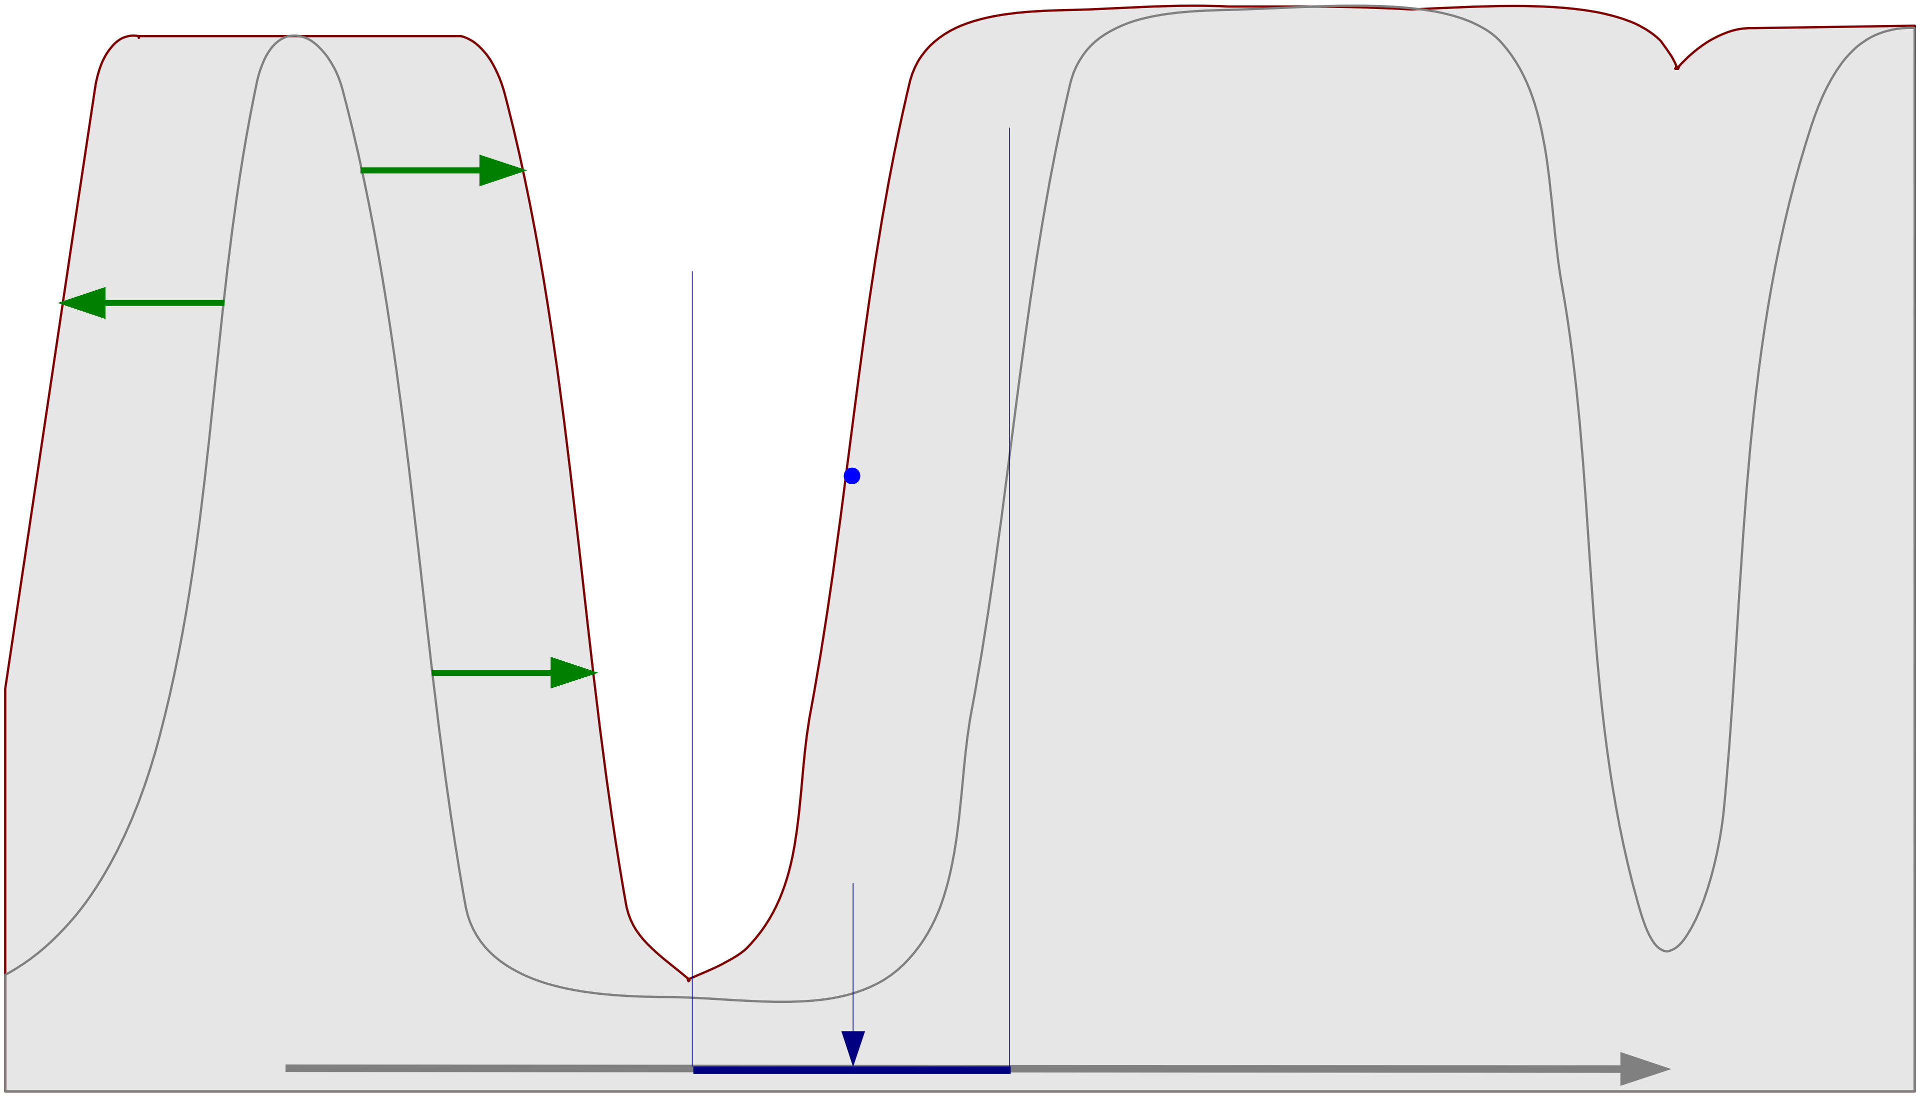

Dilation = local maximum

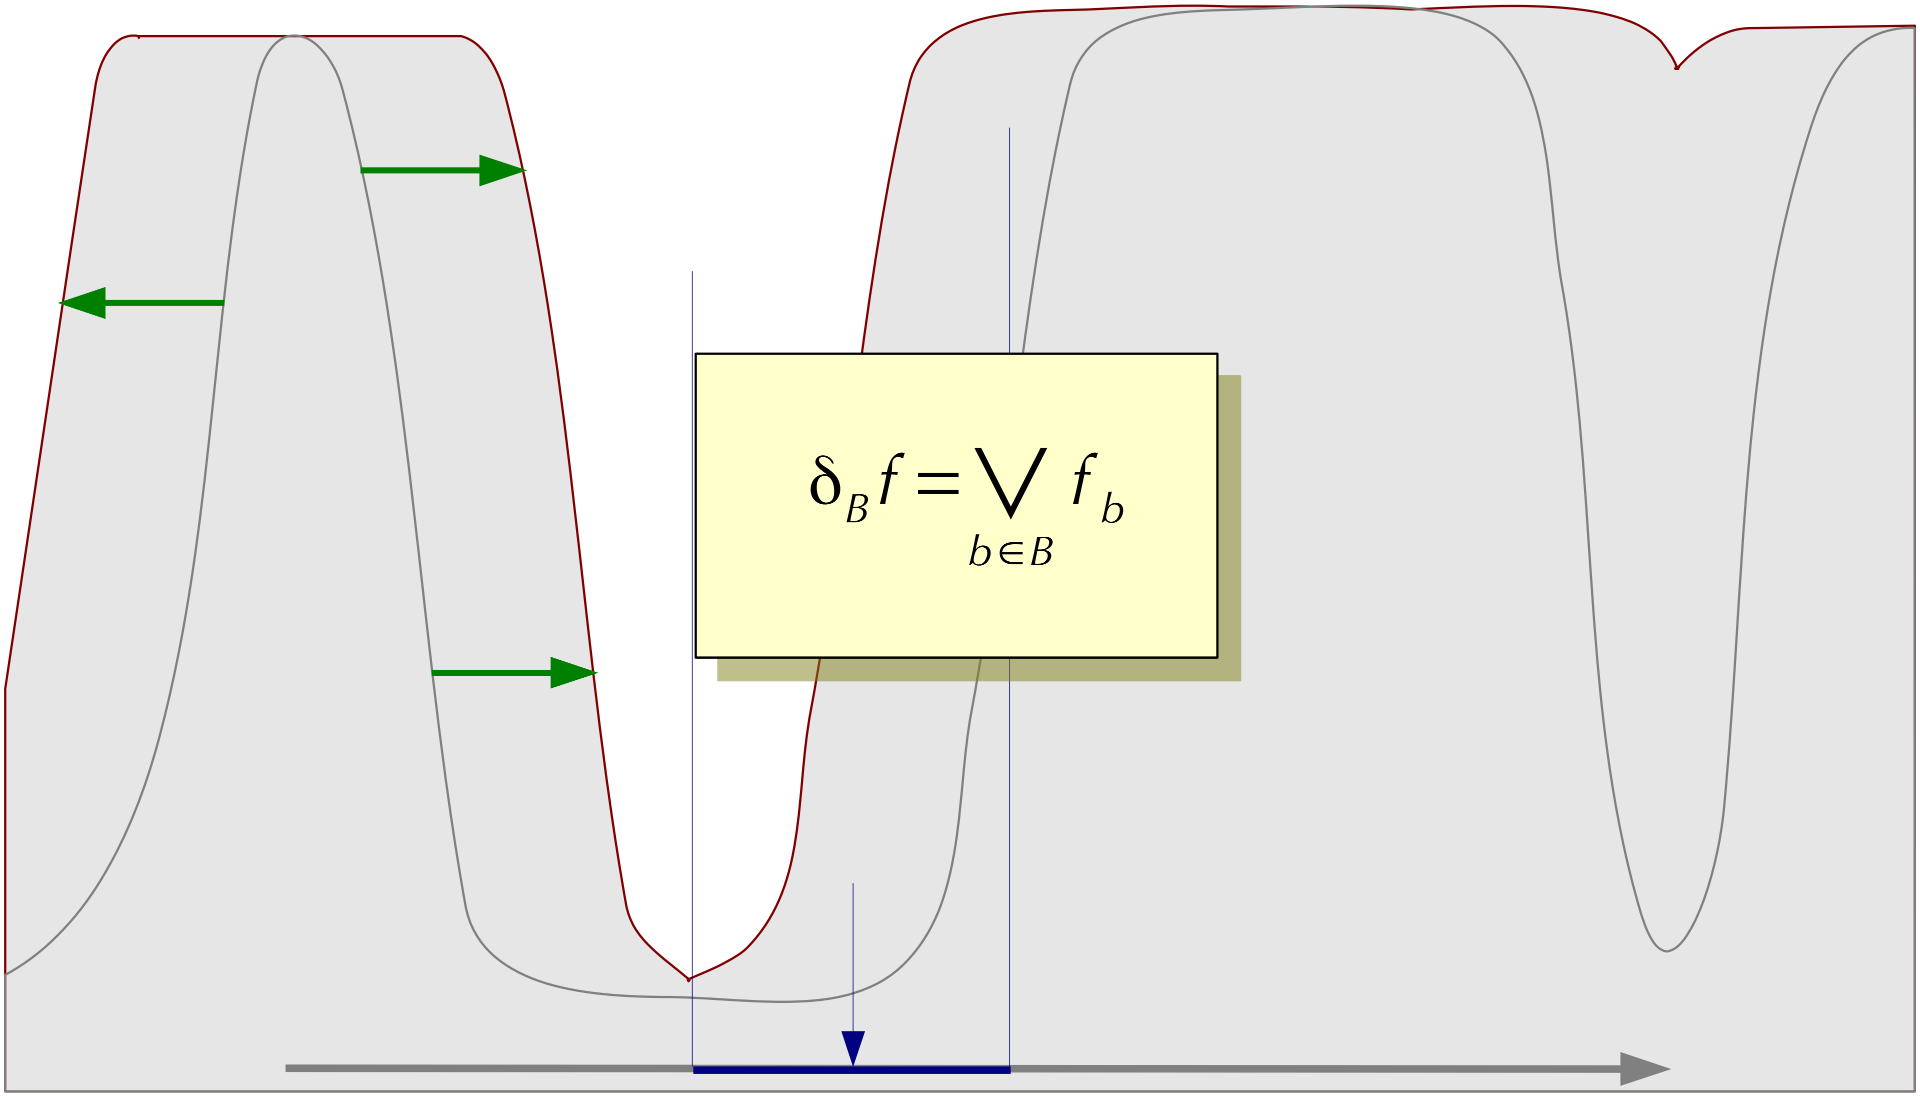

$$\delta_B f = \bigvee_{b \in B} f_b$$

Dilation + erosion = closing

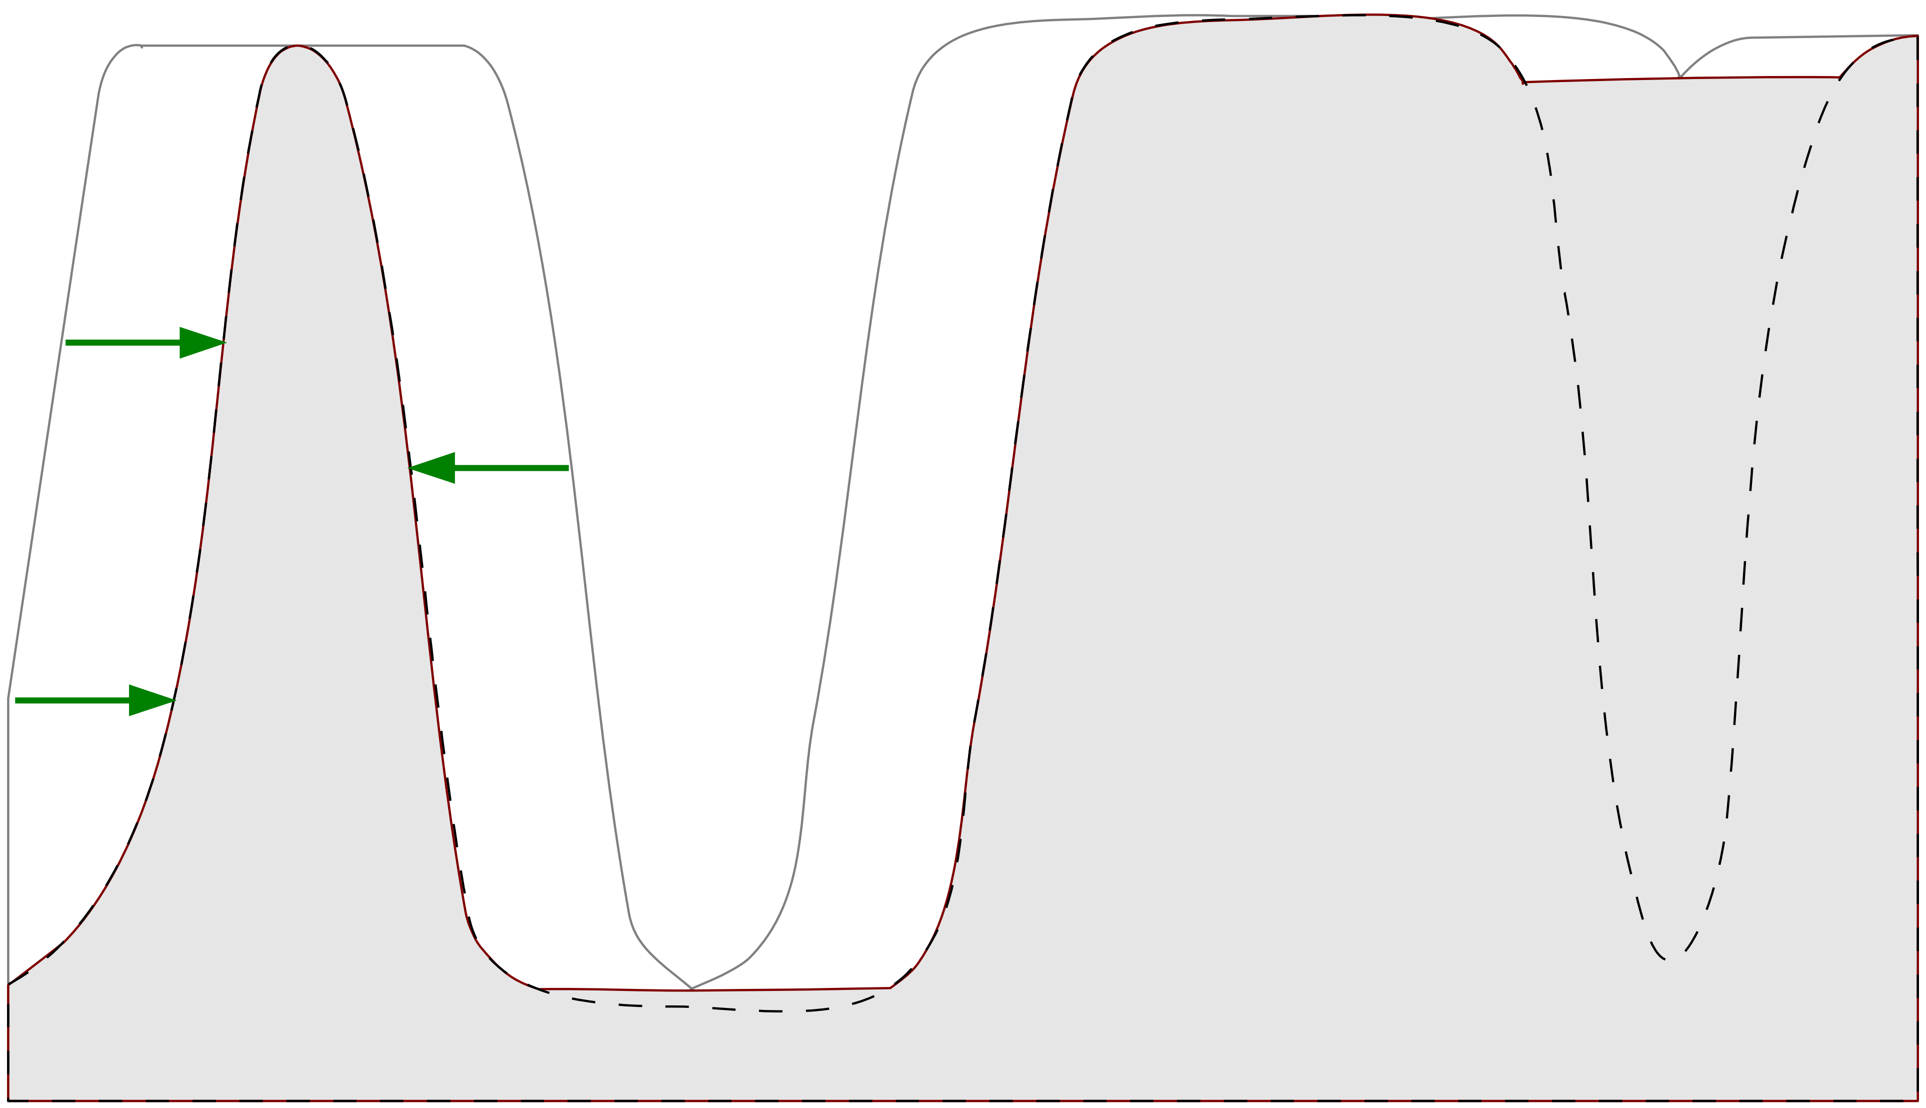

Dilation + erosion = closing

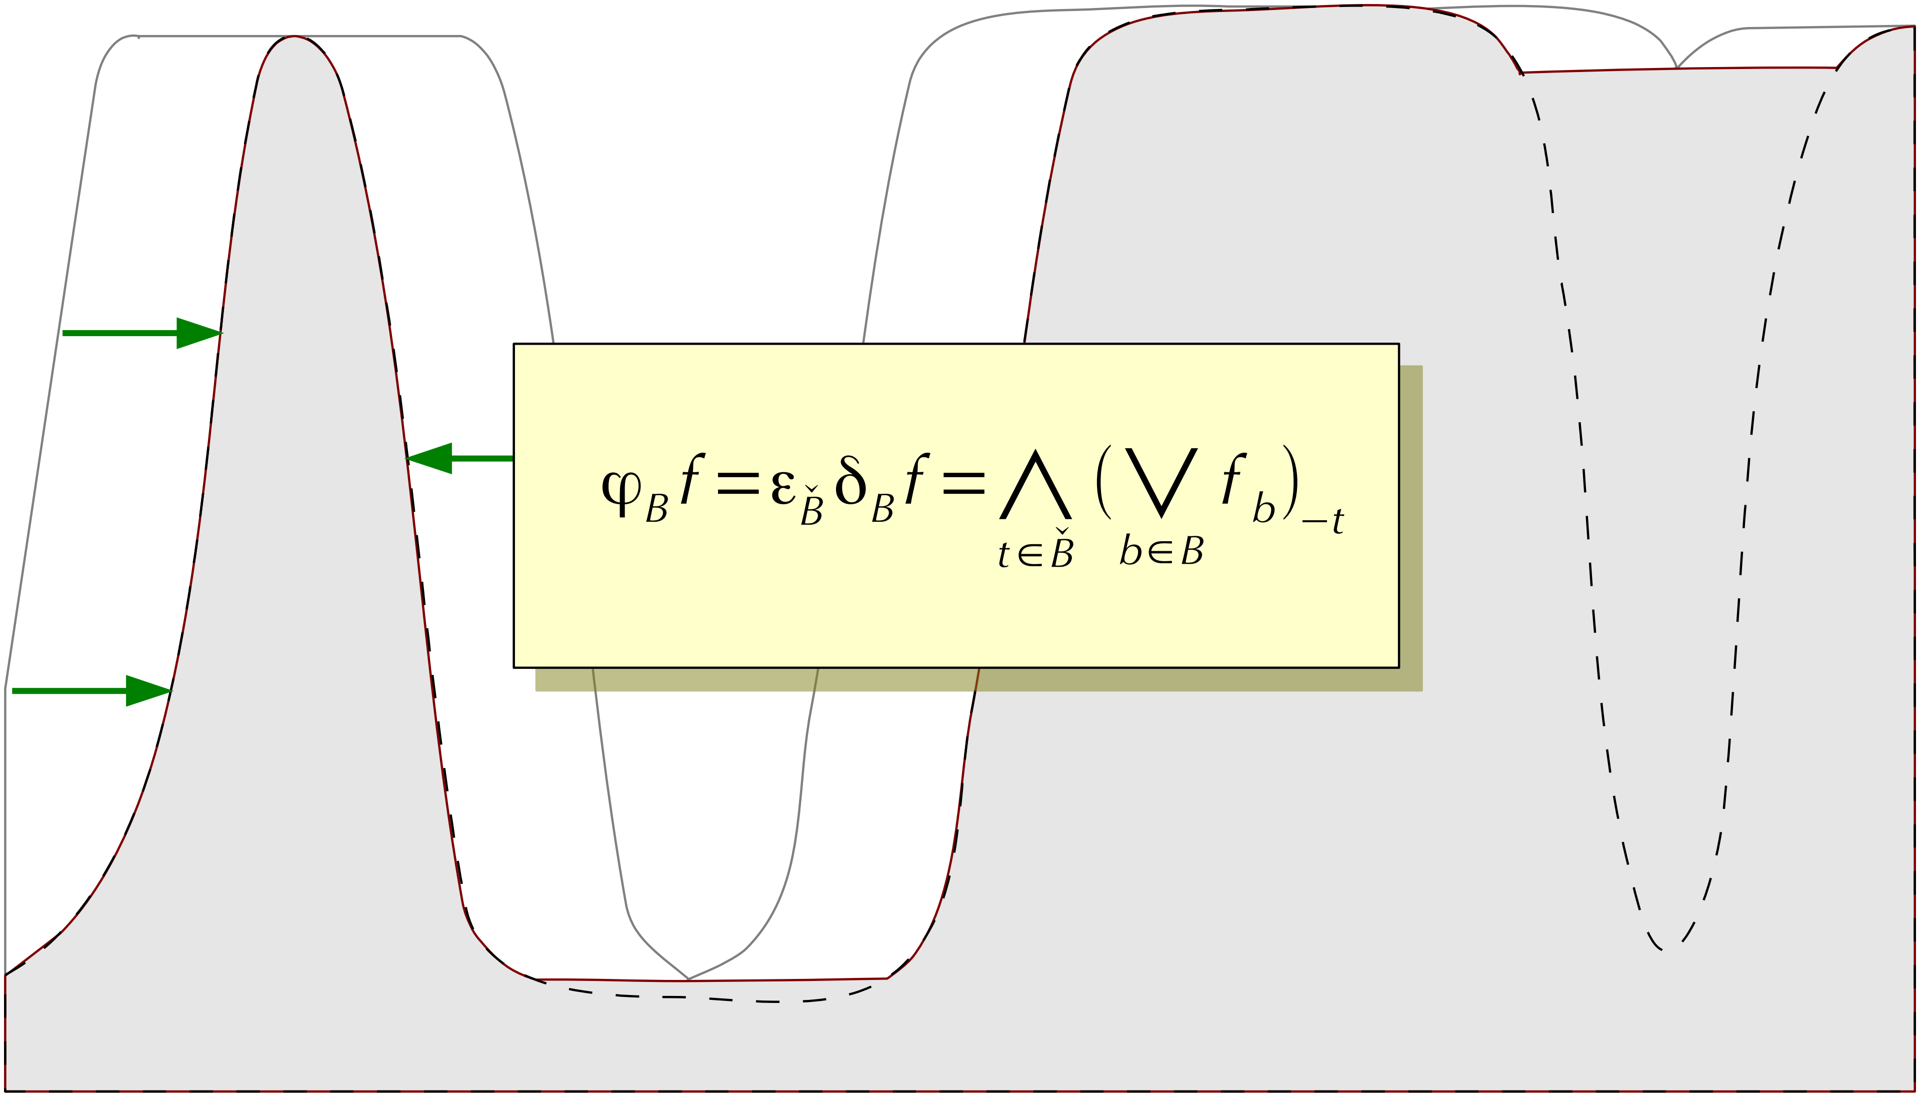

$$\varphi_B f = \varepsilon_{\check{B}} \delta_B f = \bigwedge_{t \in \check{B}} \Big(\bigvee_{b \in B} f_b \Big)_{-t}$$

Dilation + erosion = closing

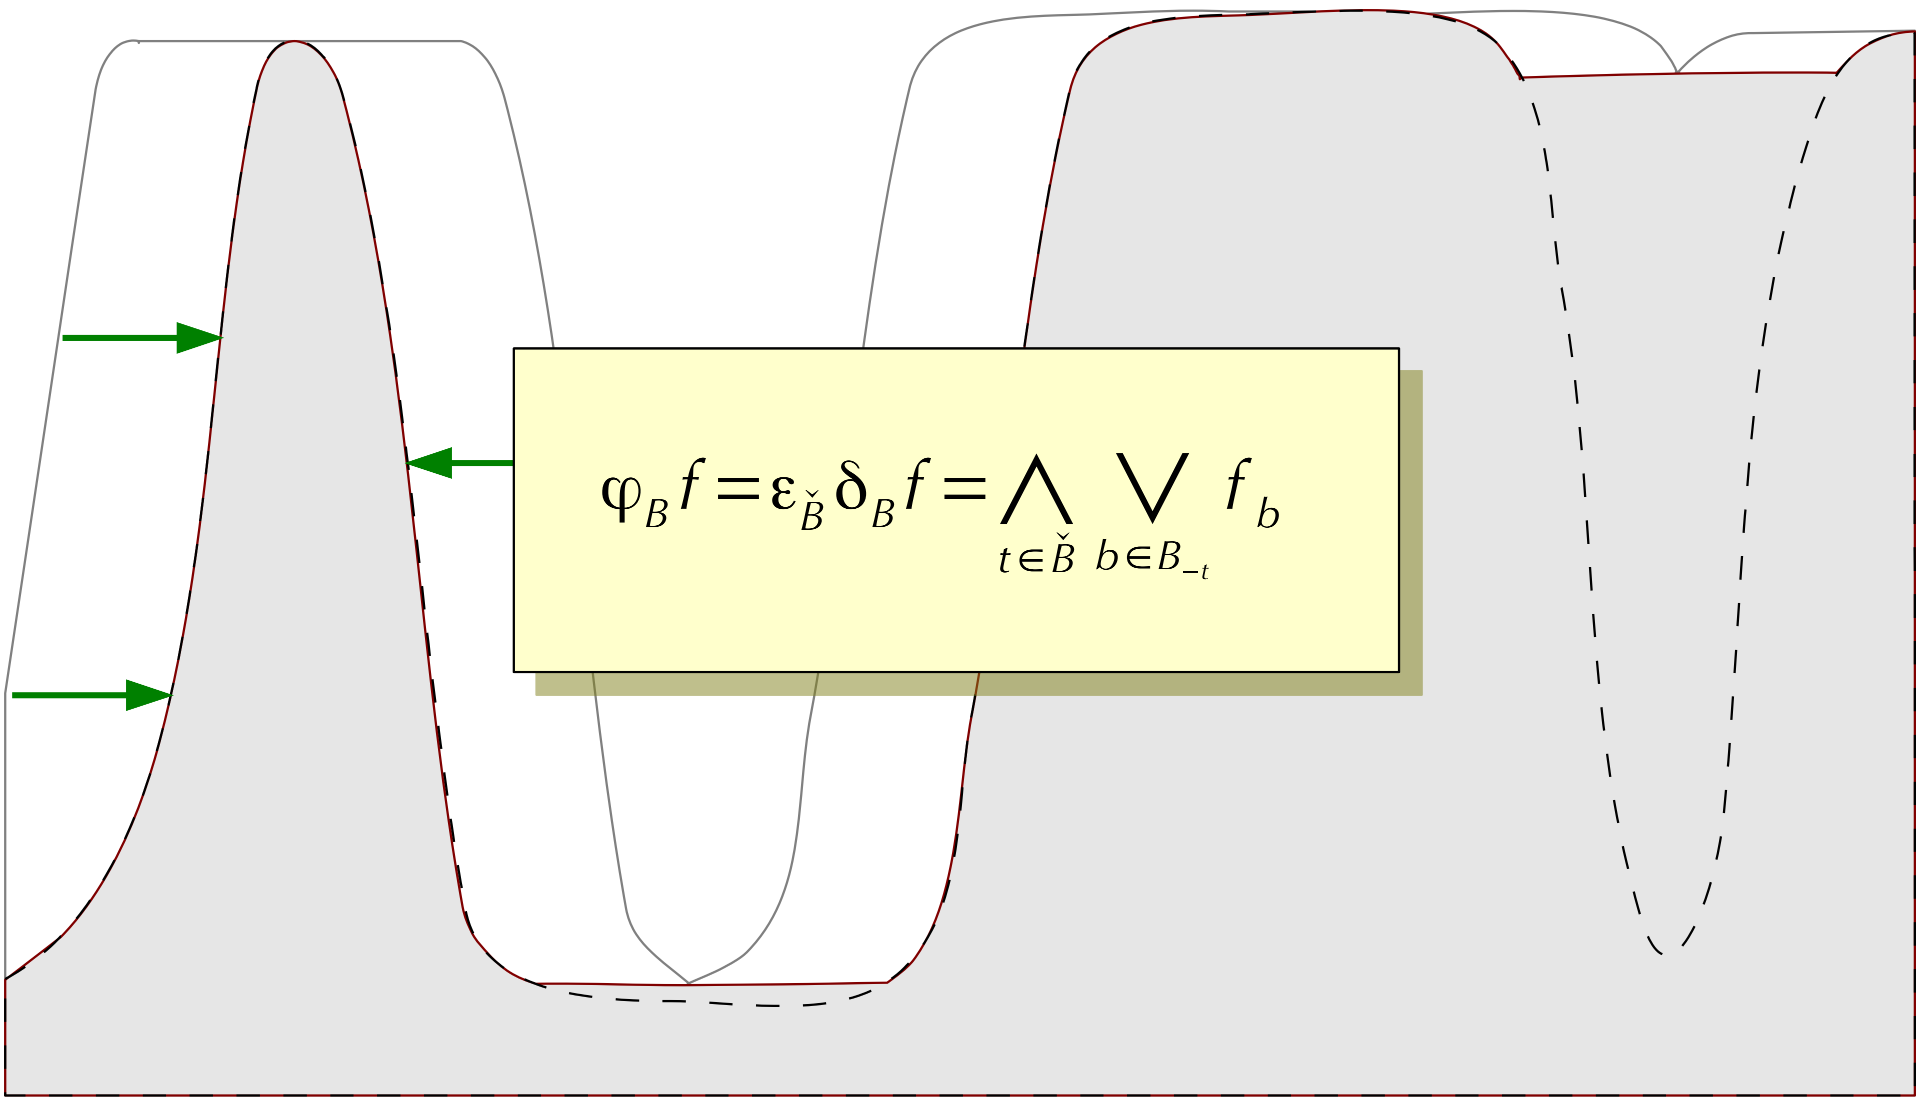

$$\varphi_B f = \varepsilon_{\check{B}} \delta_B f = \bigwedge_{t \in \check{B}} \bigvee_{b \in B_{-t}} f_b$$

Dilation + erosion = closing

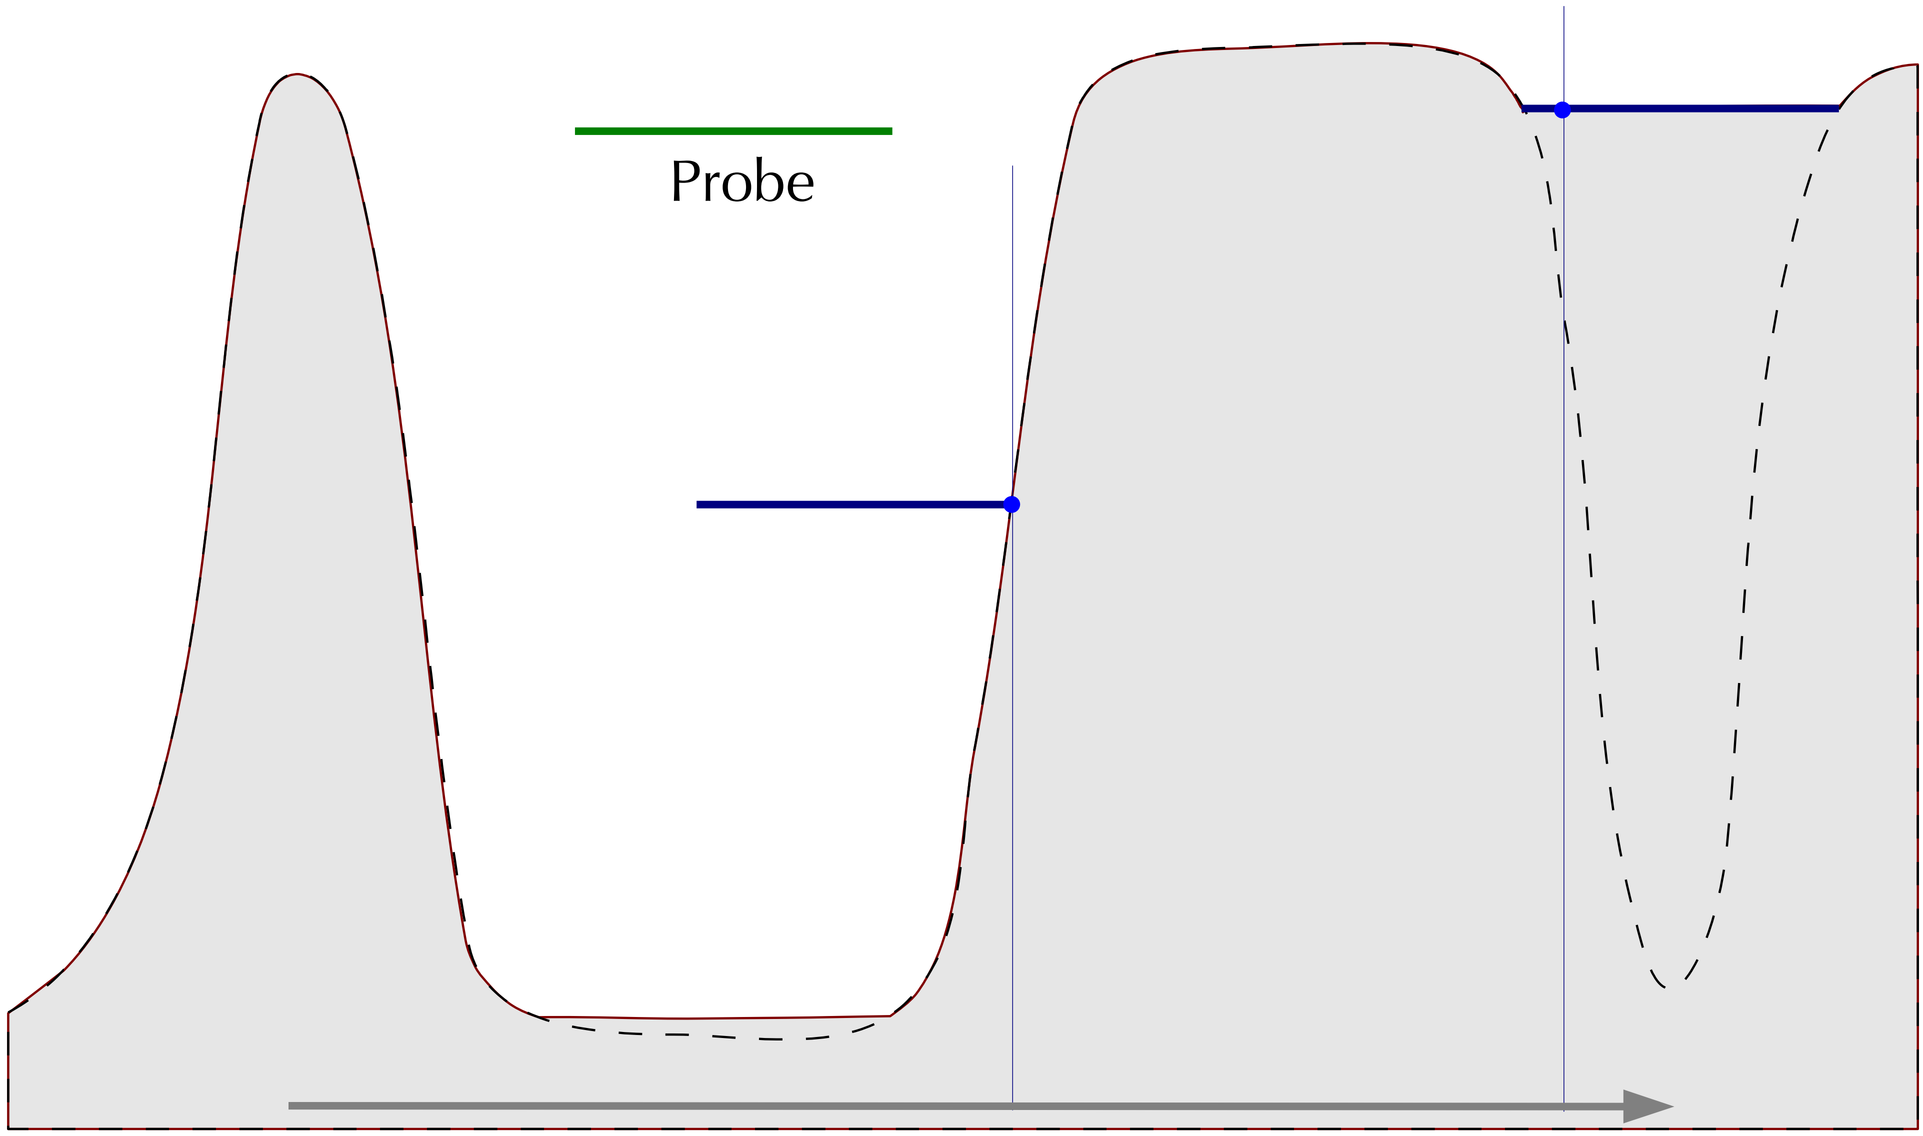

The closing on an image

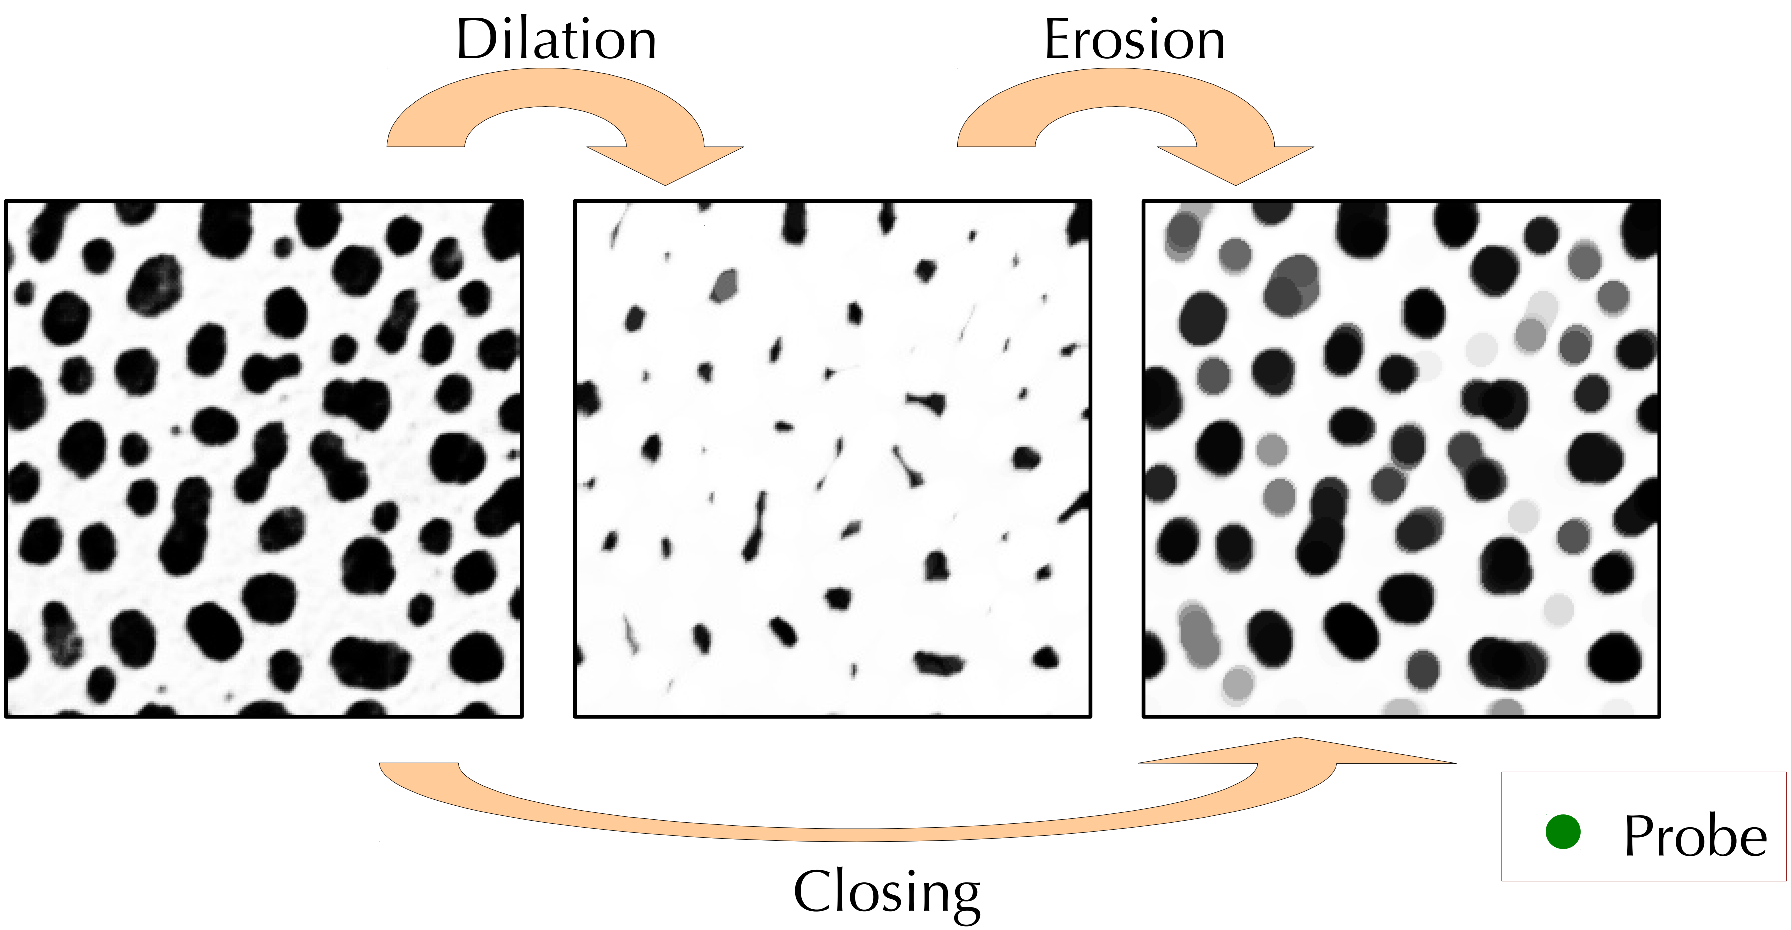

The opening

Opposite of the closing: erosion first, then dilation

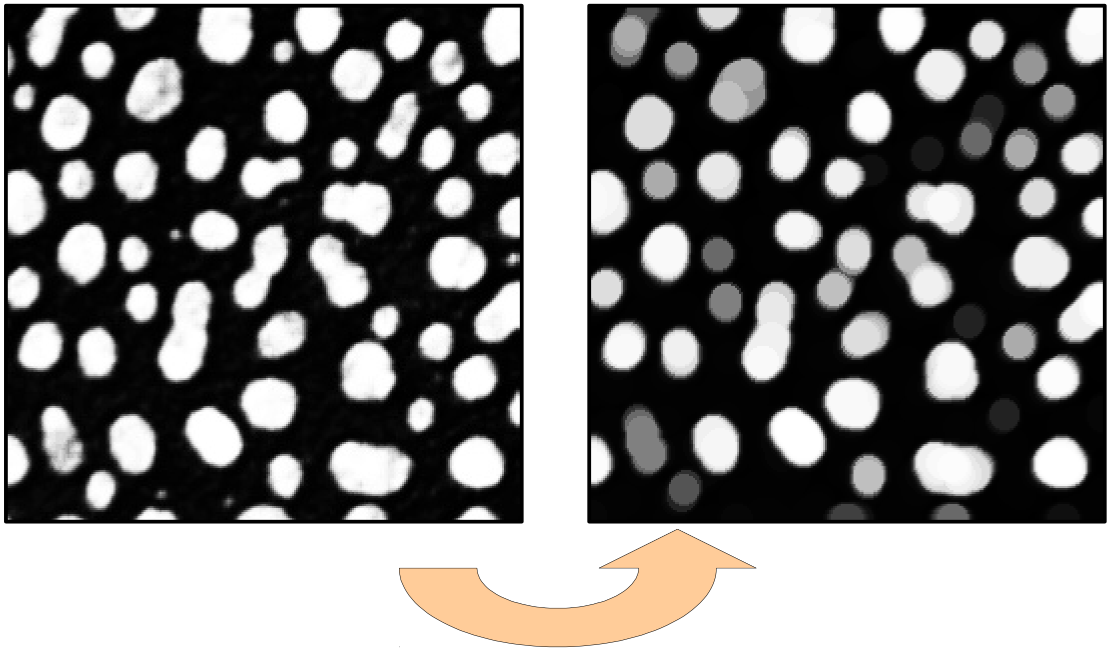

Opening

Closing at different scales

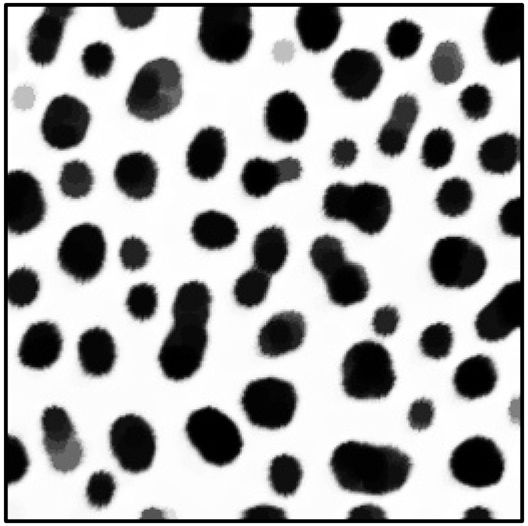

size = 10

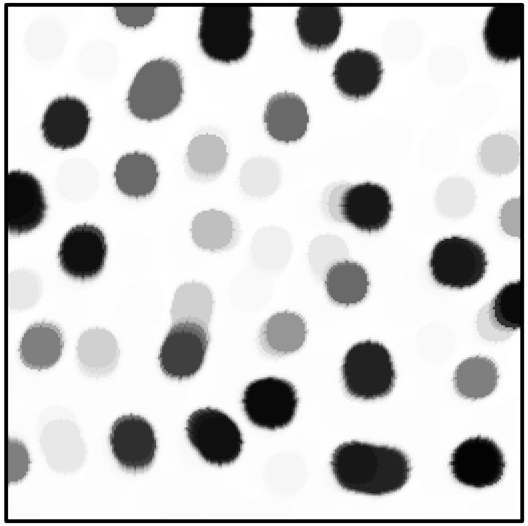

size = 20

size = 30

The granulometry

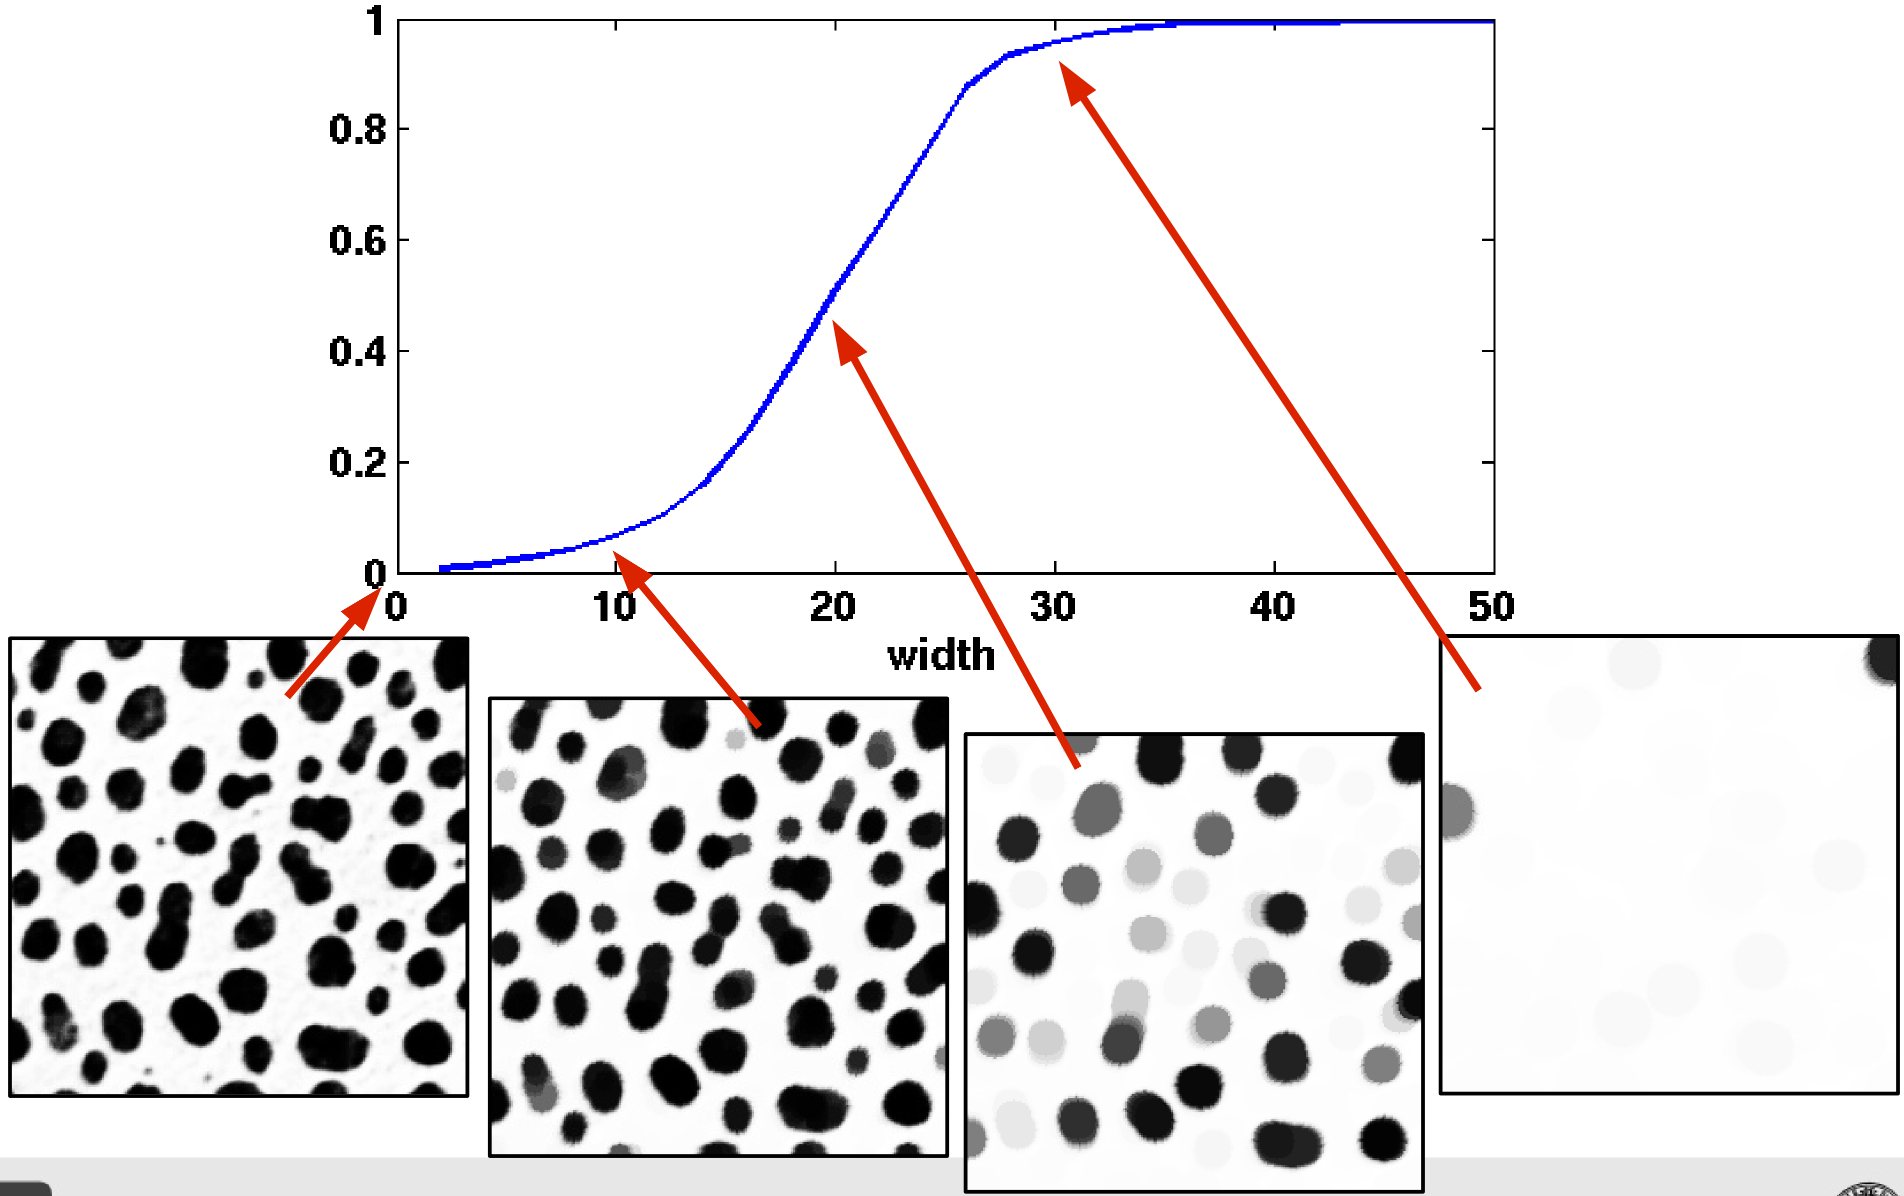

Cris Luengo
luengo@ieee.org

The granulometry

- A cumulative size distribution
- Volume-weighted:
  - Objects are weighted by the number of pixels and the contrast
  - Objects with higher contrast contribute more
- We are not counting objects!
- To compare with real-world measurements one must:
  - Prepare the image correctly
    - all objects have the same contrast
    - objects and background have uniform intensity
  - Normalize the granulometry appropriately
  - Apply a counting frame (?)

Cris Luengo
luengo@ieee.org

The size distribution

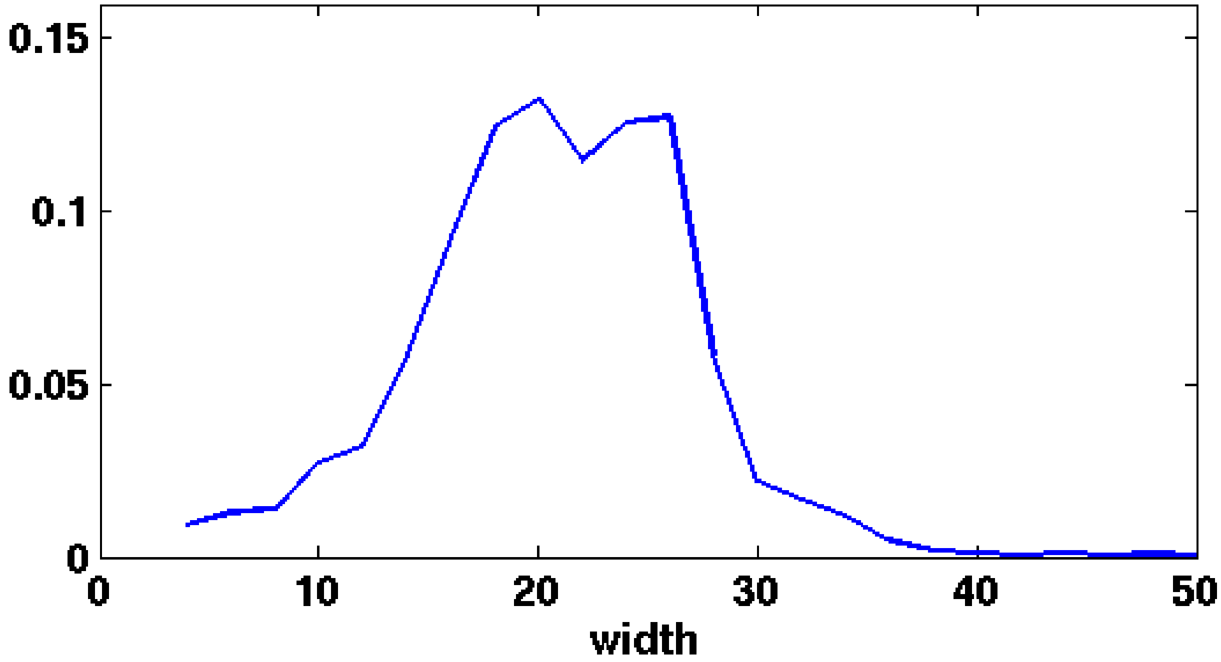

If the granulometry is a cumulative size distribution, then its derivative is a size distribution

Pore size distribution

- A: untreated milk gel
- B: + substrate
- C: + substrate & enzyme
- D: + substrate & enzyme

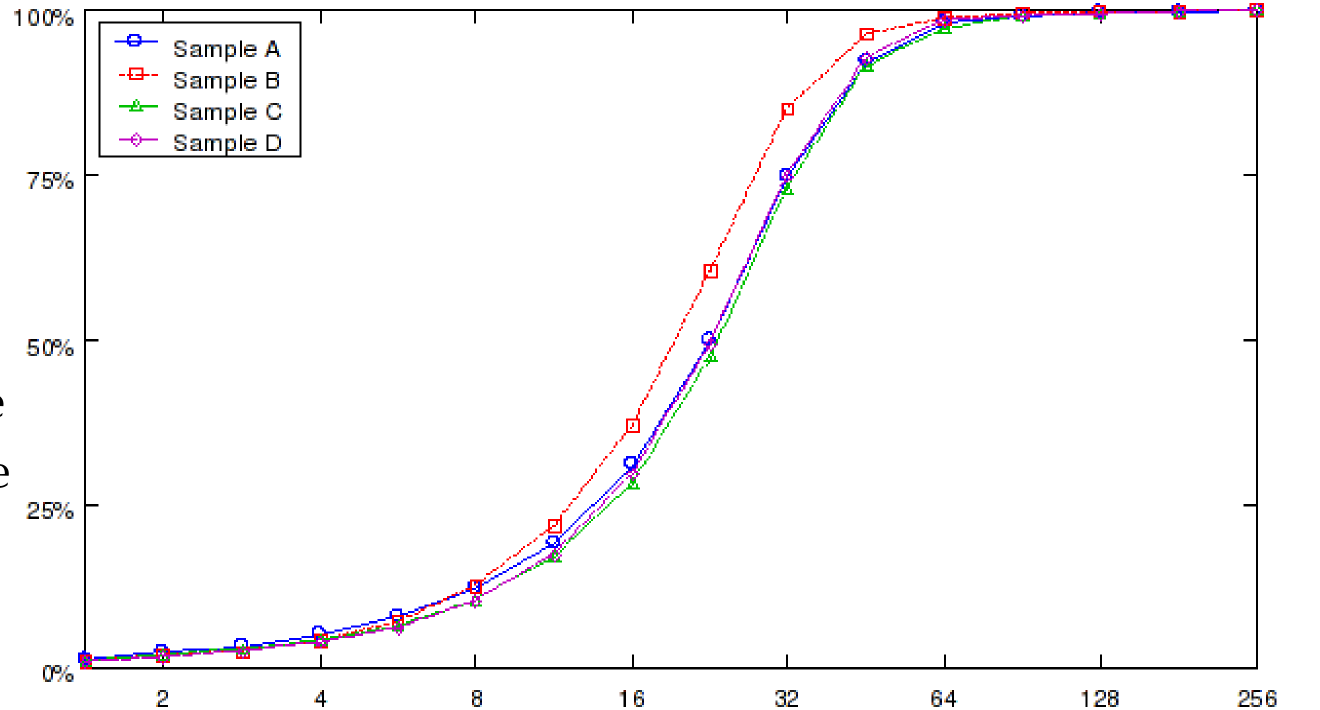

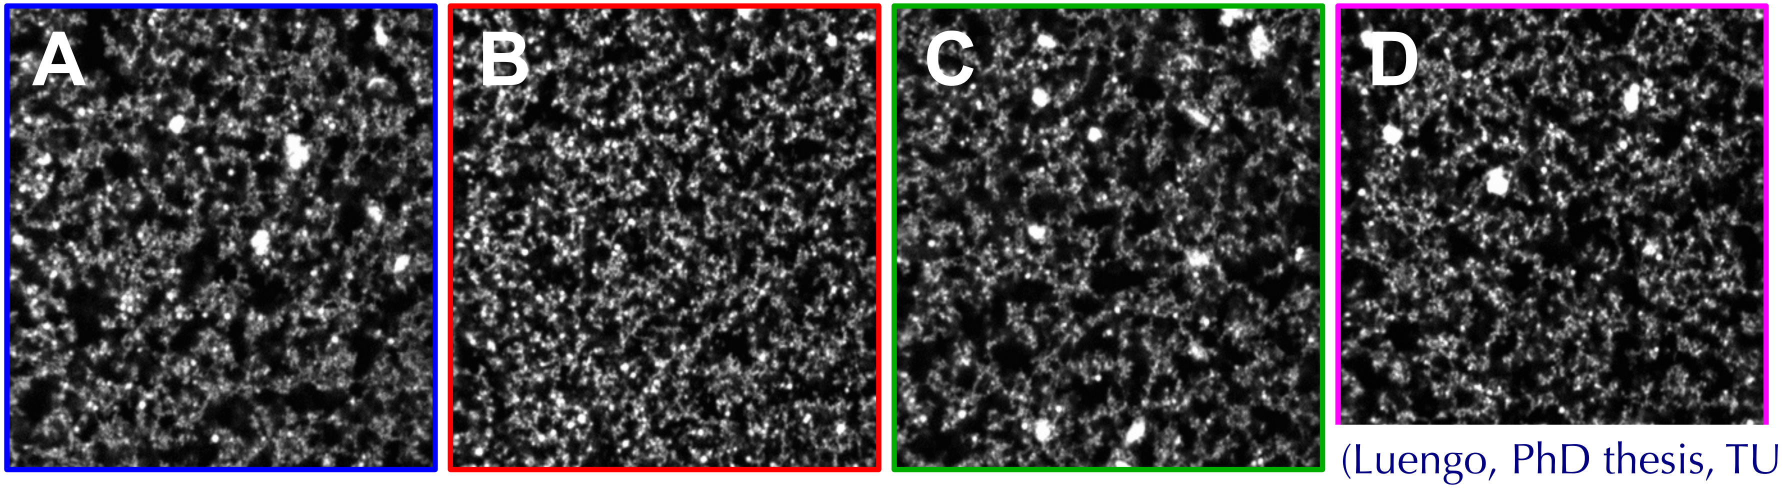

(Luengo, PhD thesis, TU Delft, 2003)

Pore size distribution

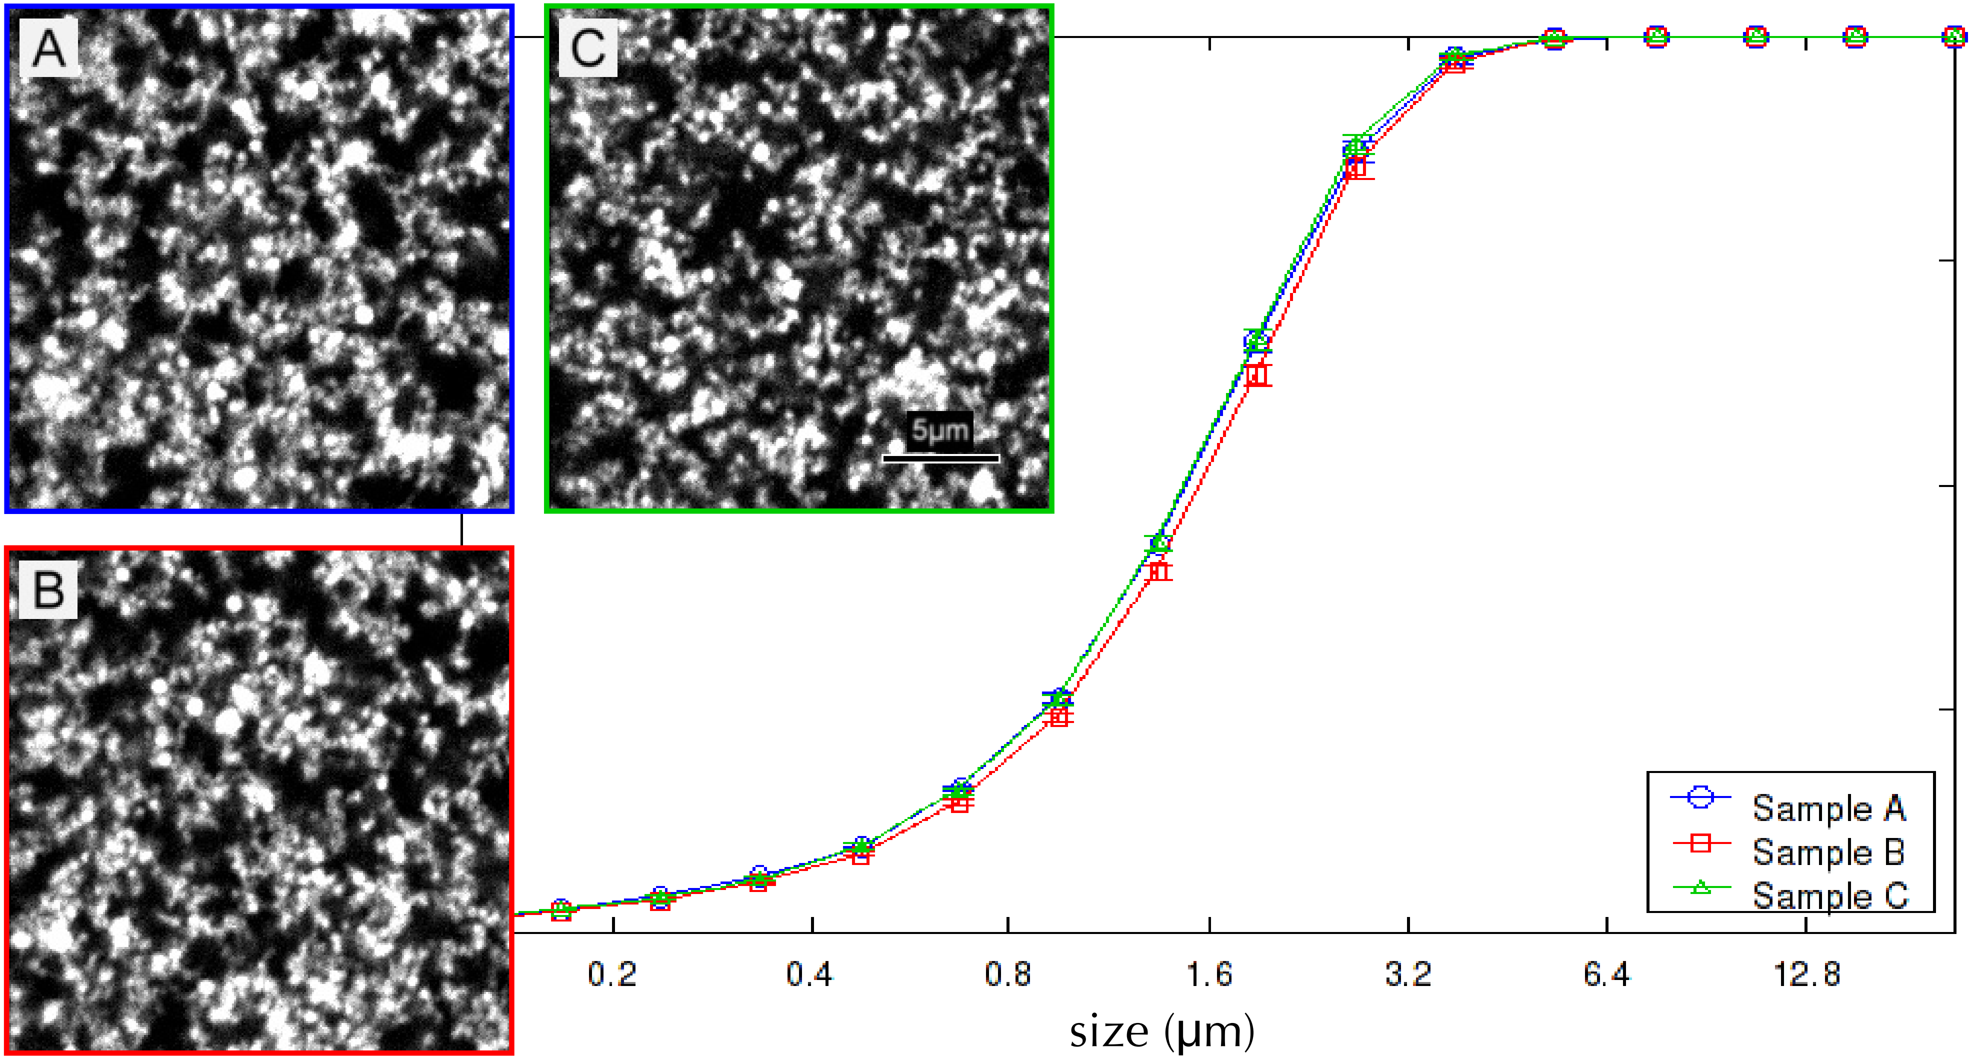

(Luengo, PhD thesis, TU Delft, 2003)

Closing with different shapes

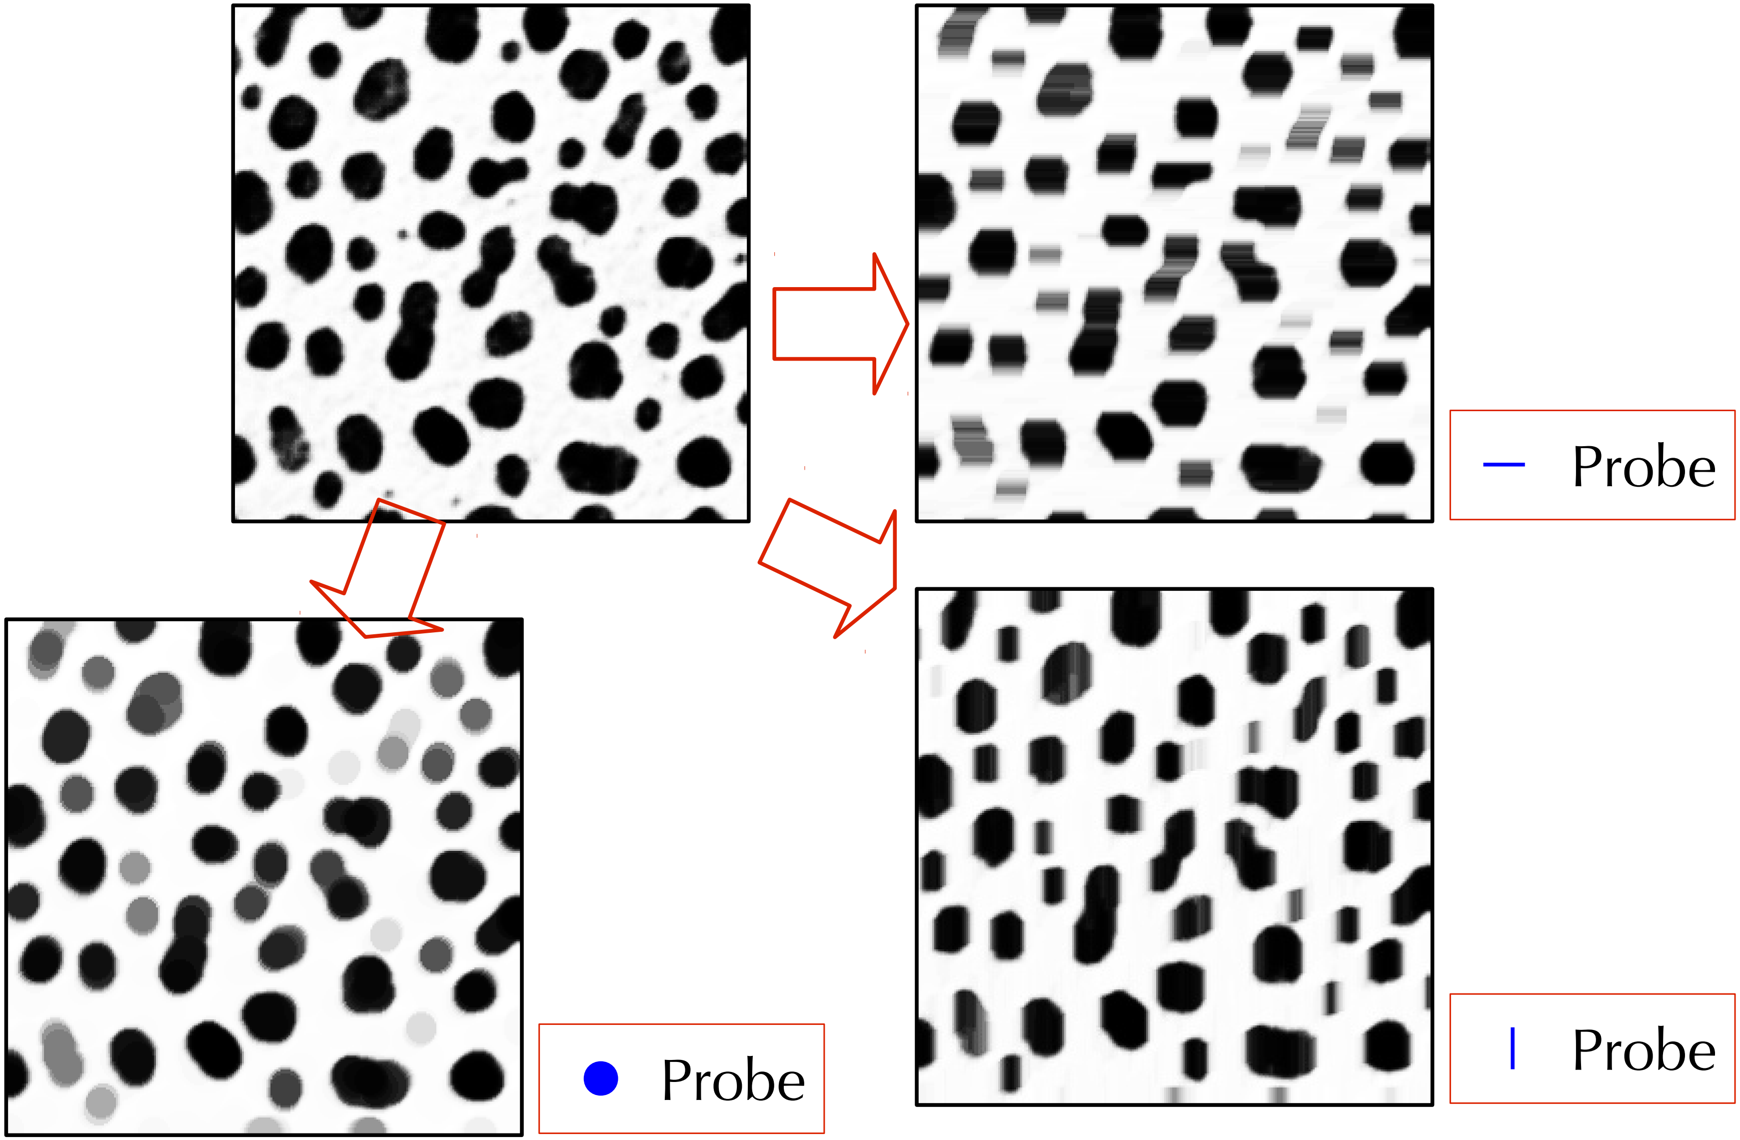

Rotation invariance

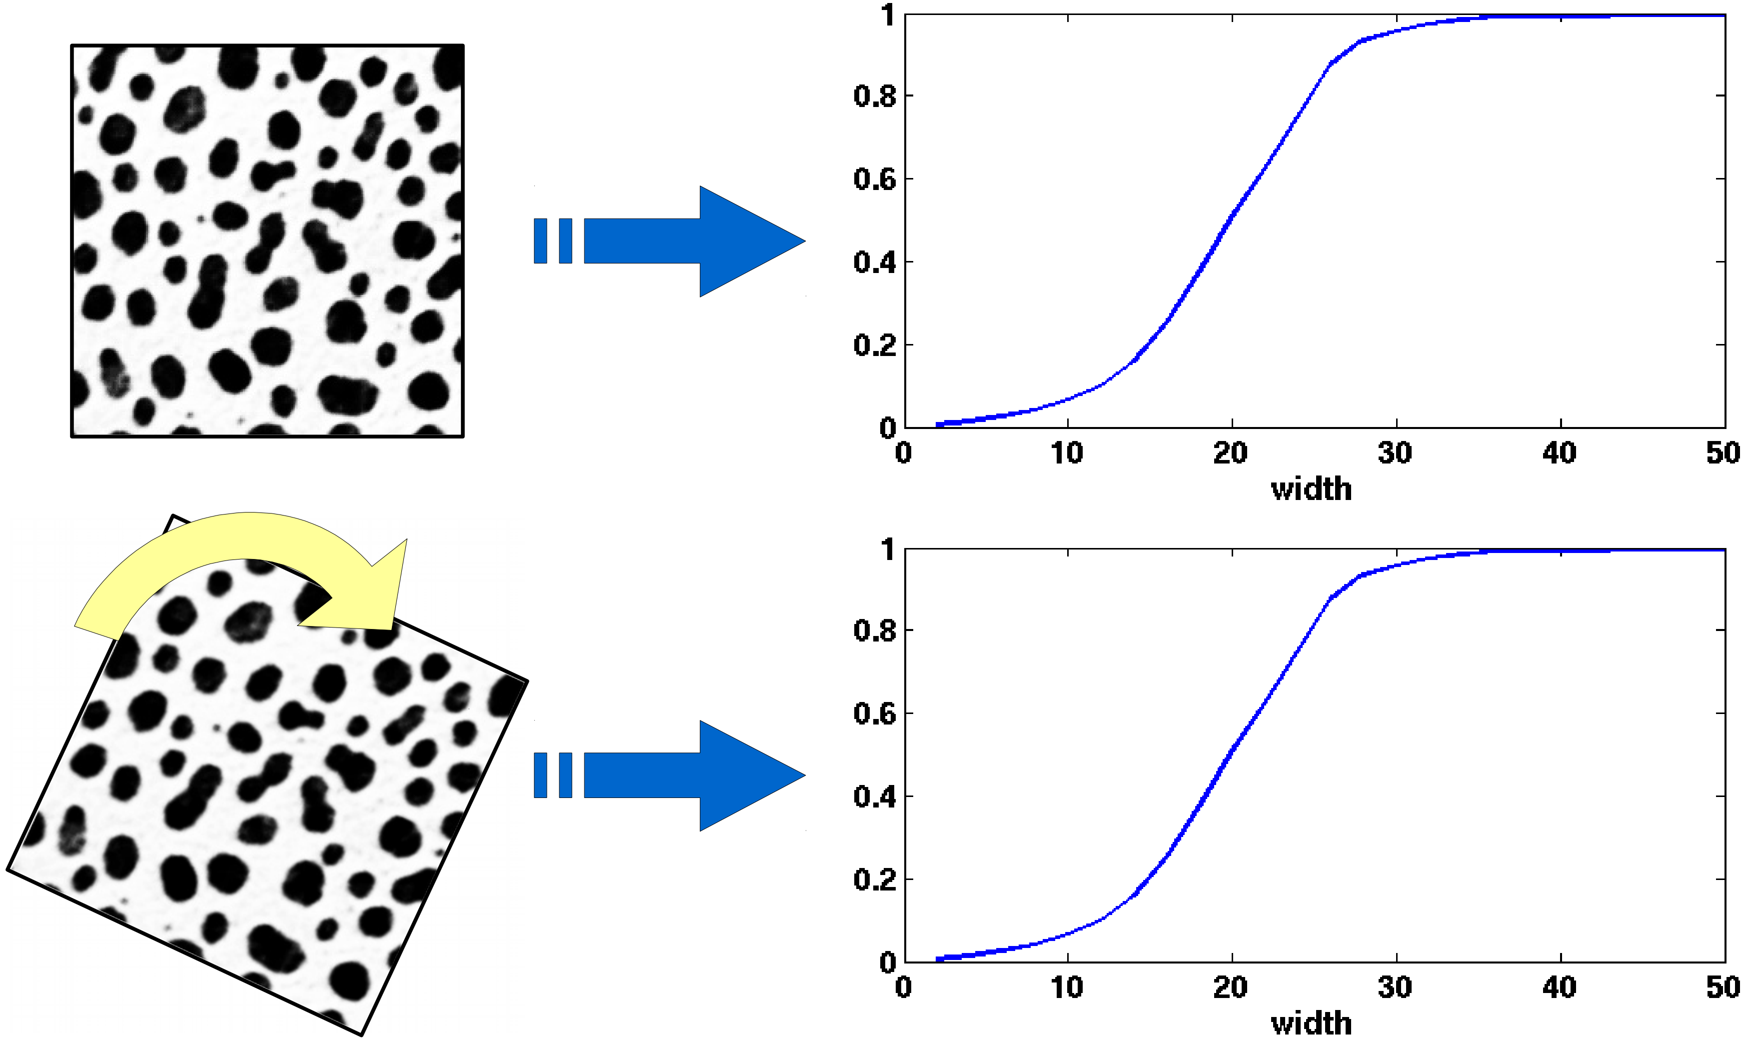

Closing with lines at many angles

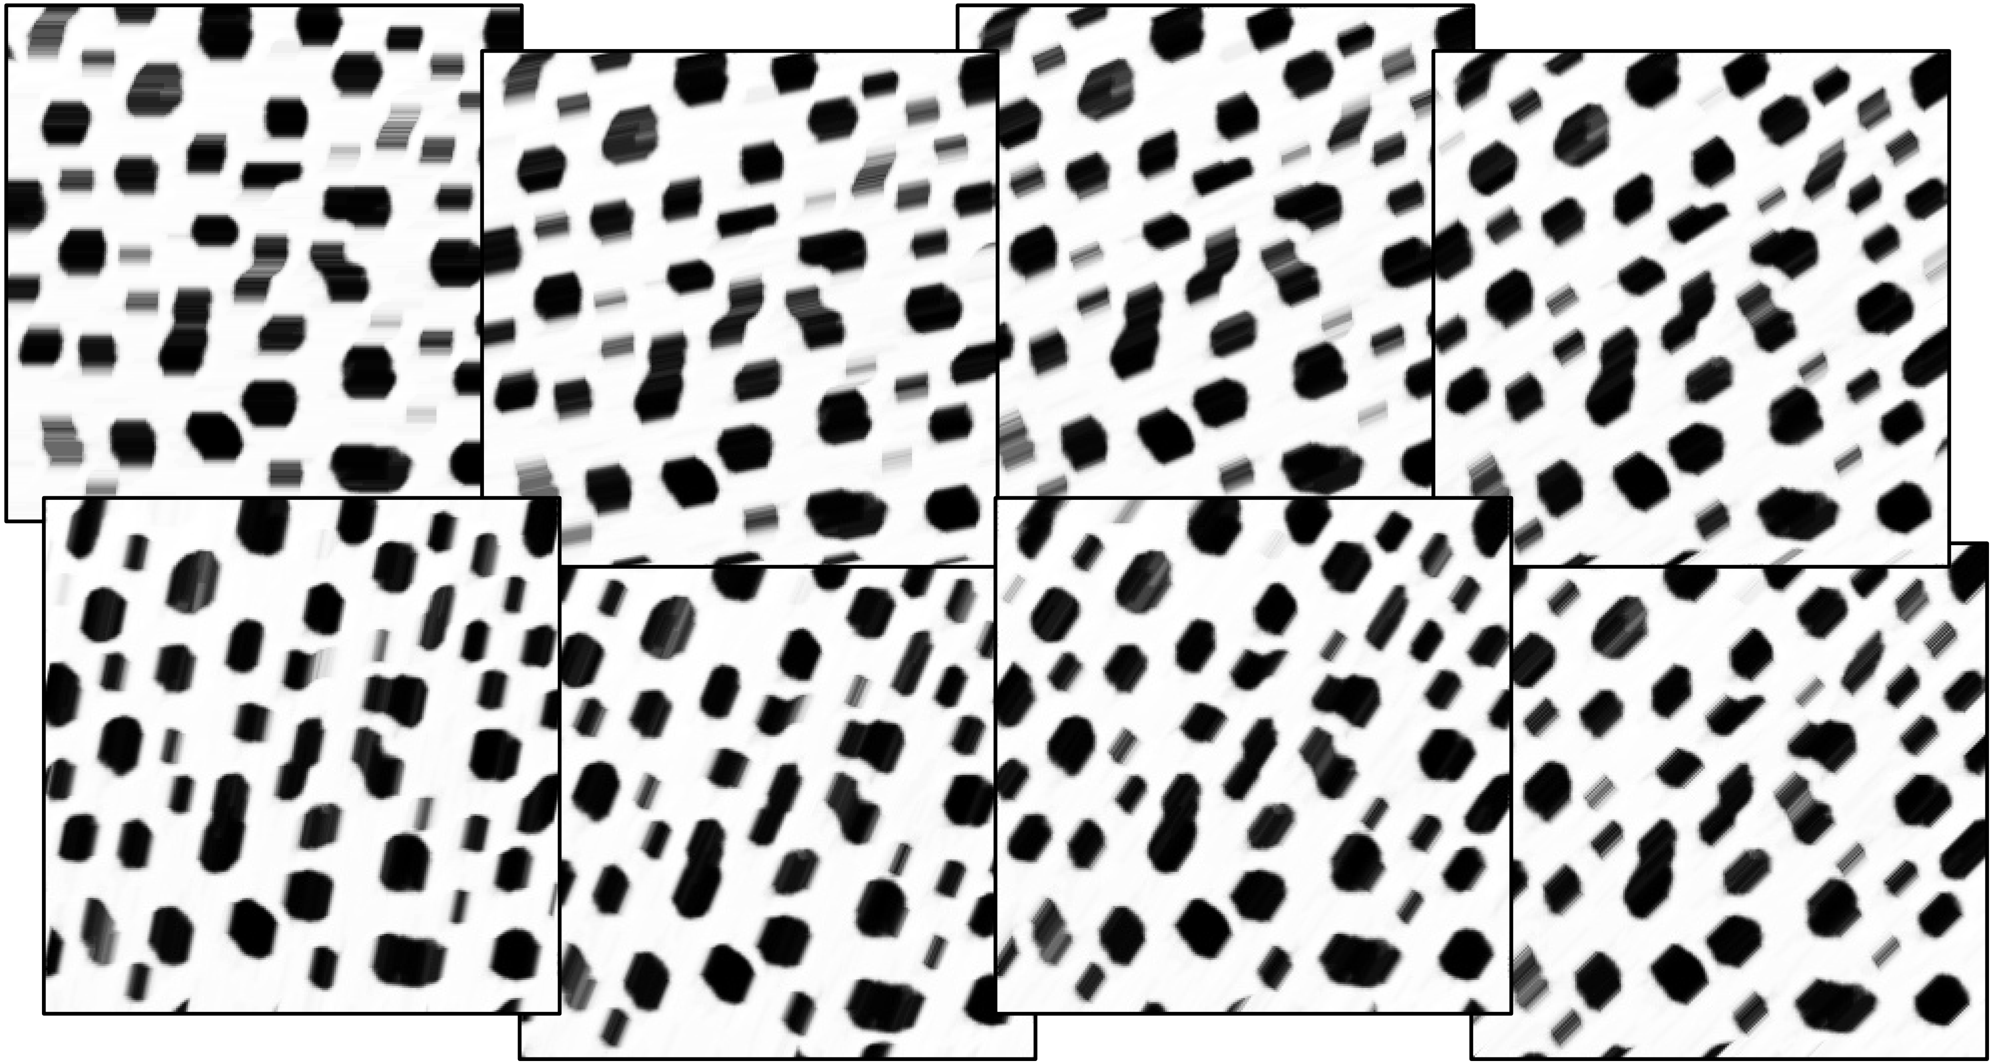

Closing with lines at many angles

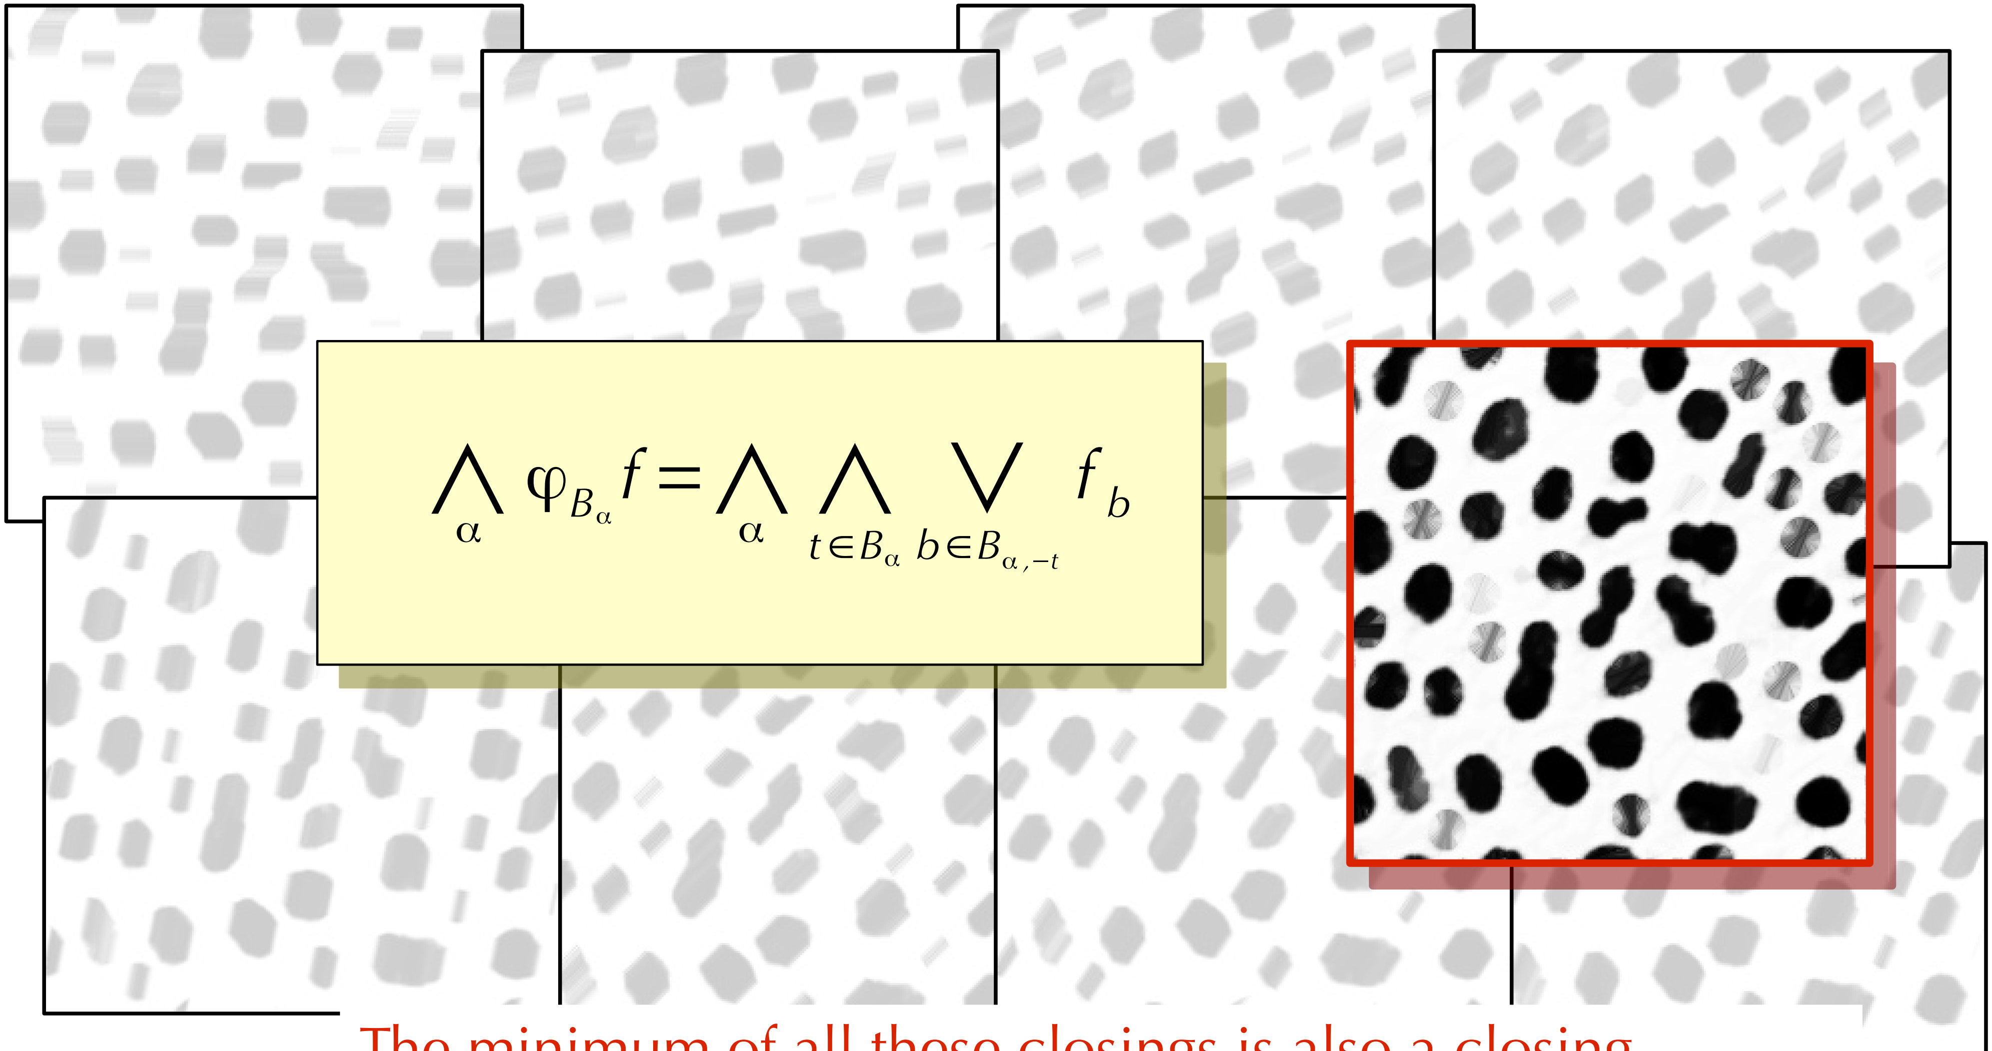

$$\bigwedge_{\alpha} \varphi_{B_\alpha} f = \bigwedge_{\alpha} \bigwedge_{t \in B_\alpha} \bigvee_{b \in B_{\alpha,-t}} f_b$$

The minimum of all these closings is also a closing

Isotropic and non-isotropic closing

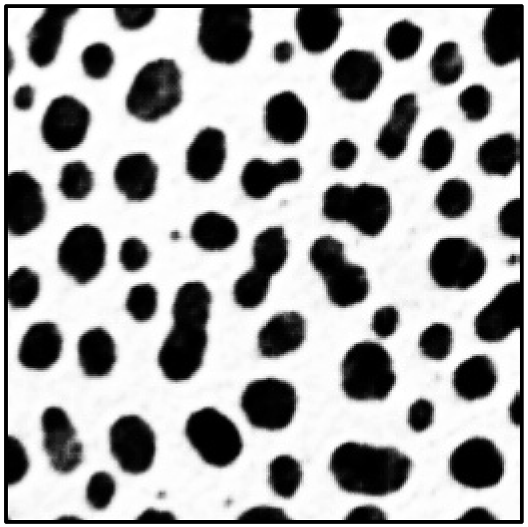

Isotropic

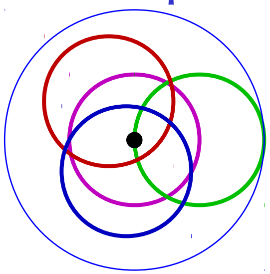

Selects on width

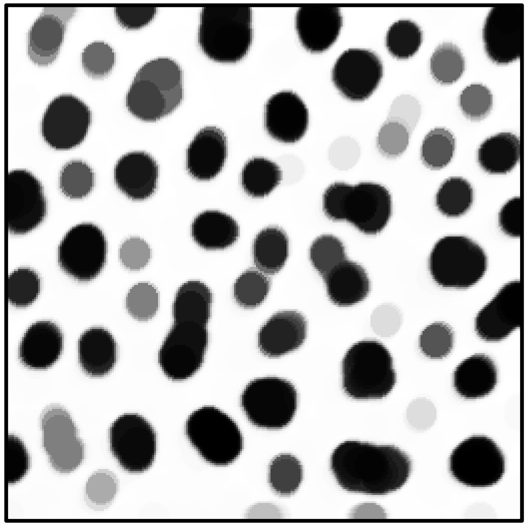

Anisotropic

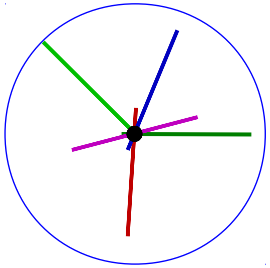

Selects on length

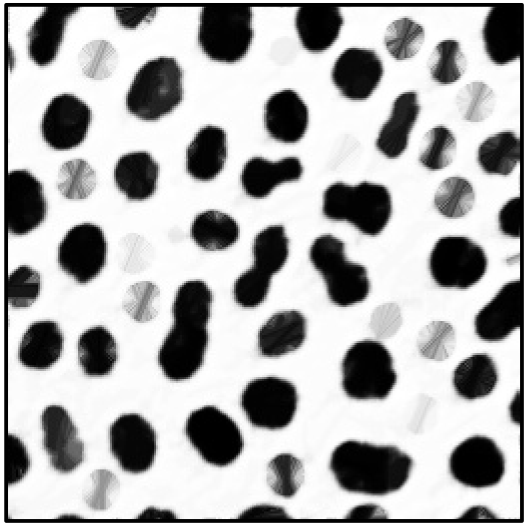

Cris Luengo
luengo@ieee.org

Length distribution

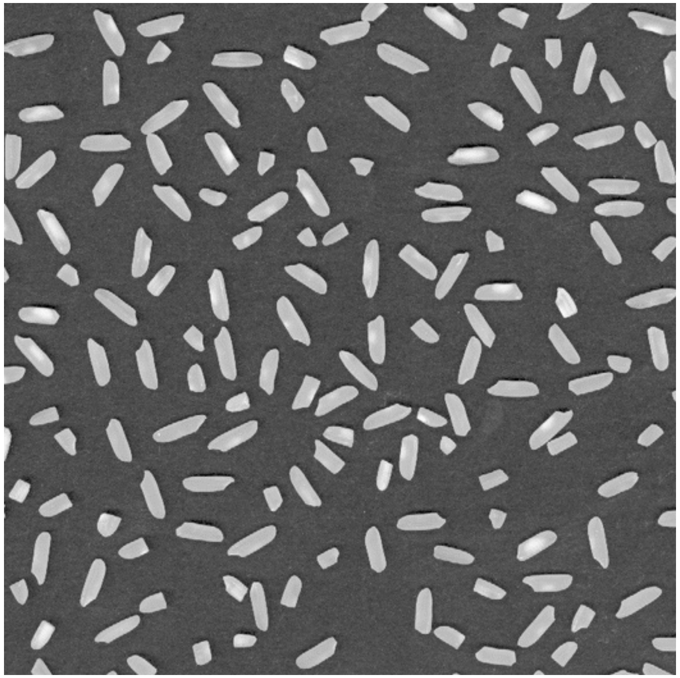

Easy to measure length
for each grain

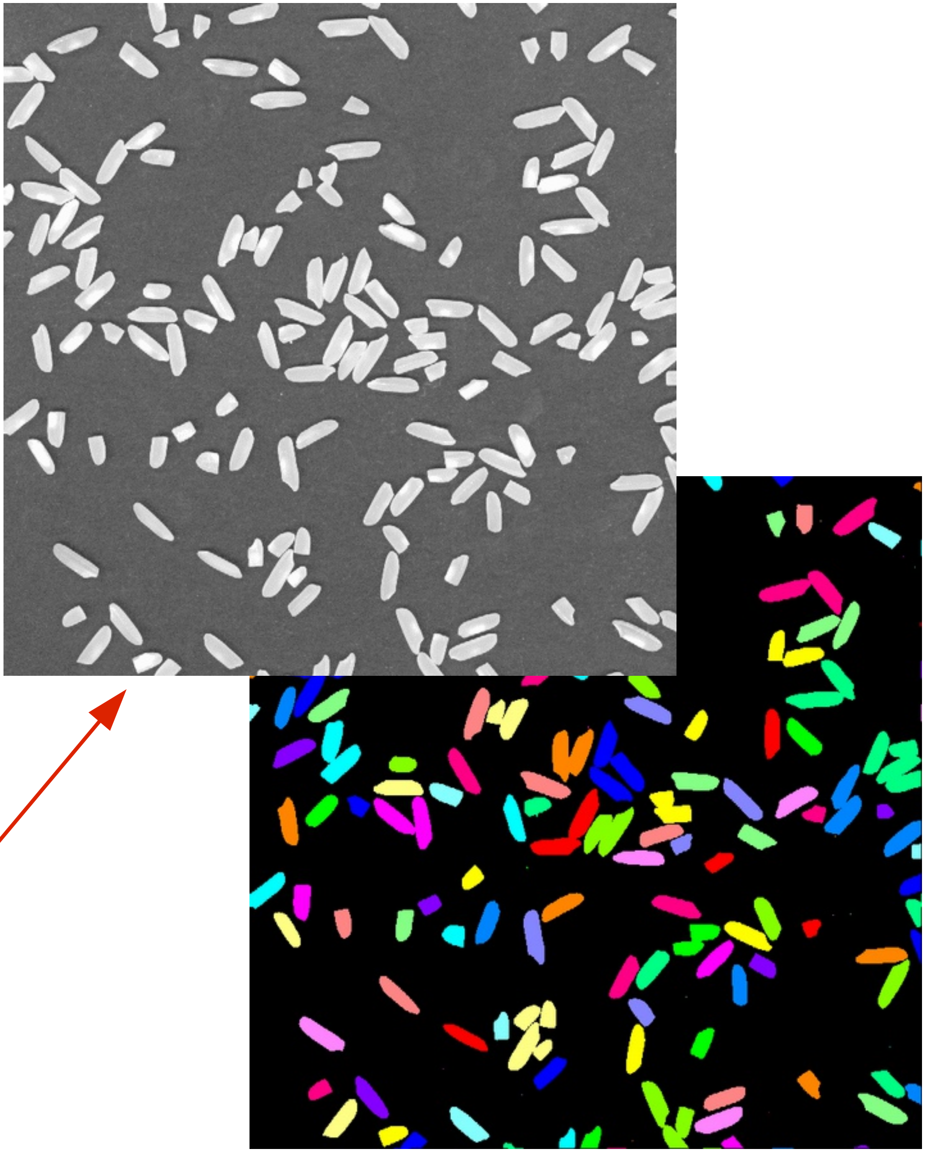

Not so easy…

Length distribution

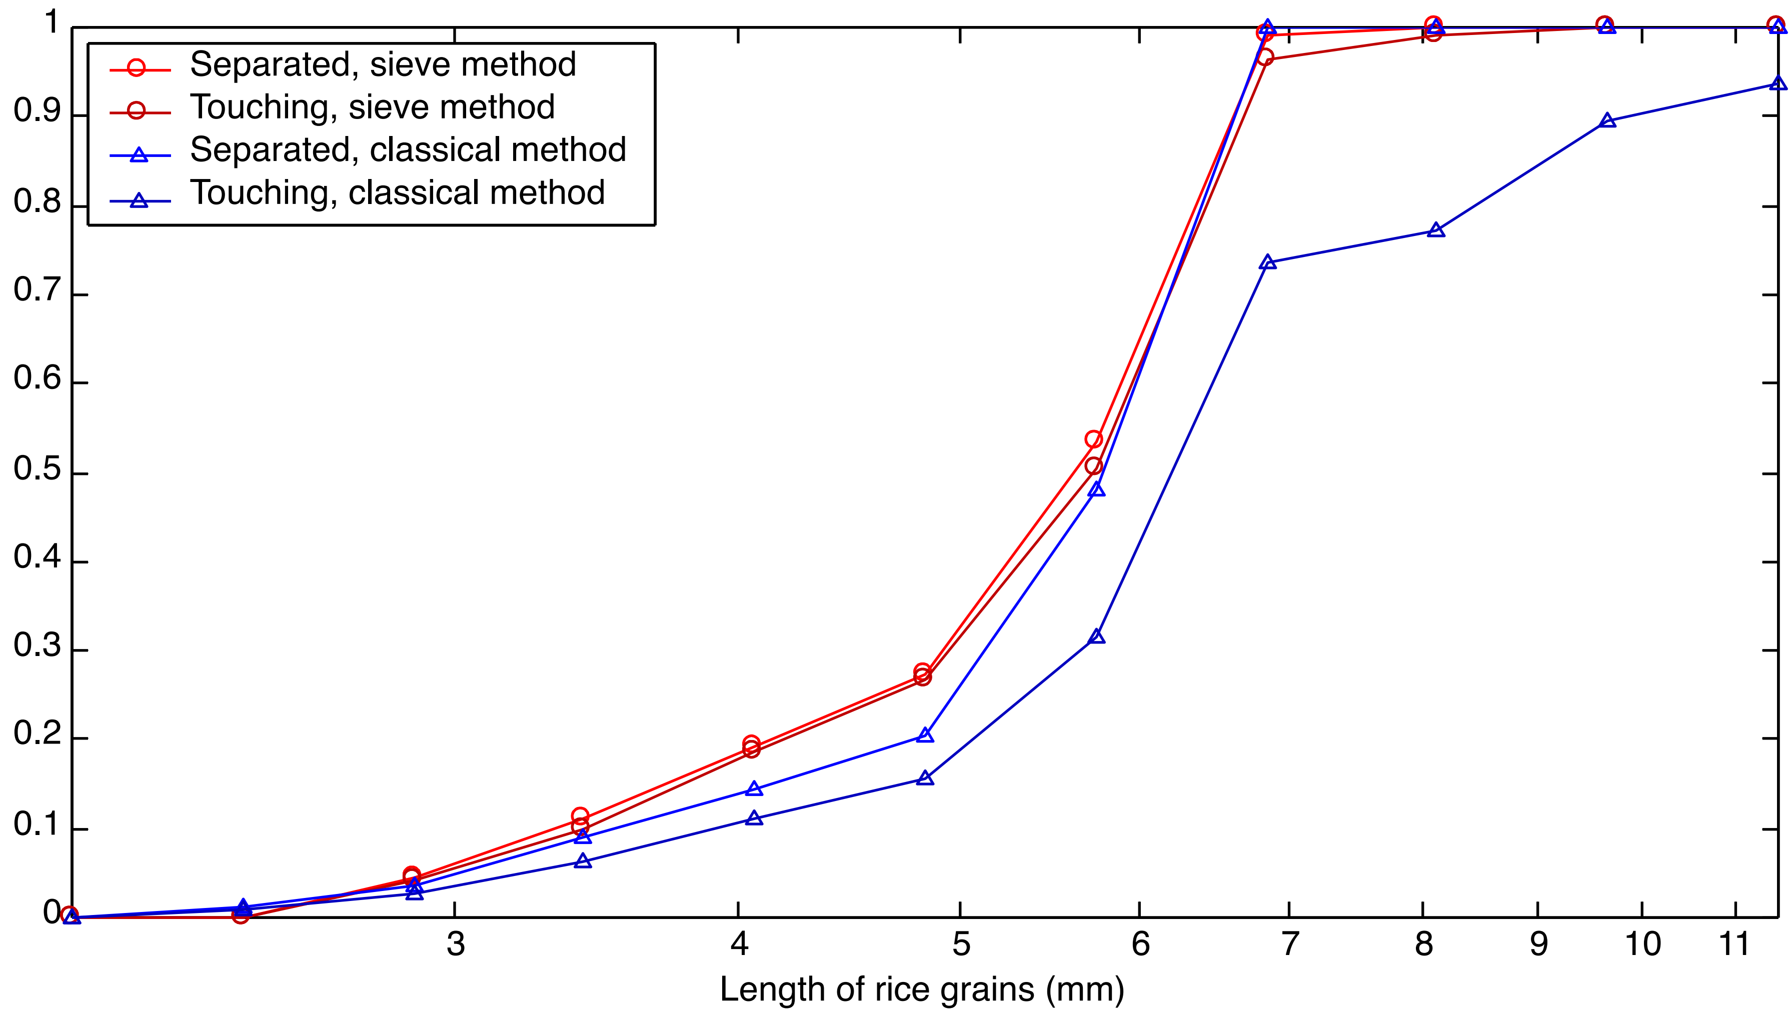

(Luengo & van Vliet, 2001)

Length distribution

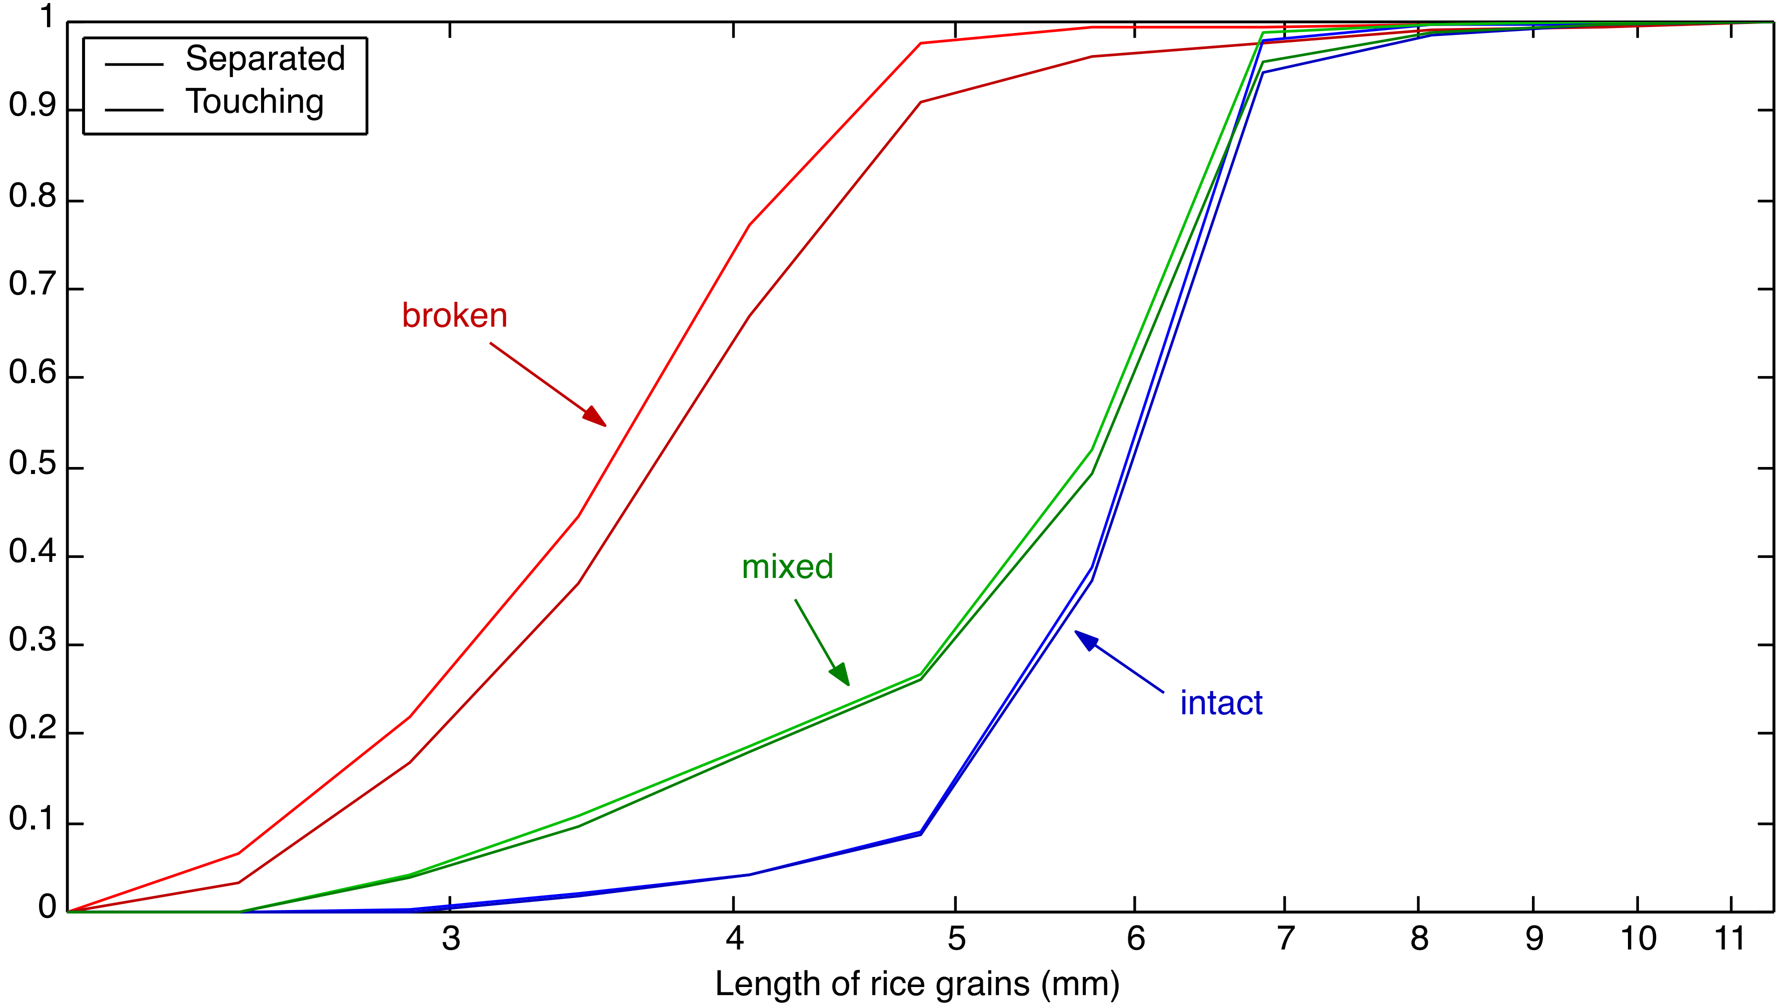

(Luengo & van Vliet, 2001)

Rotating line segments

- How many angles do we need to probe to get an accurate measurement?
  - Longer line = more angles! ($\sim l$)
  - Obviously quite expensive
- How about 3D images?
  - $\sim l^2$
  - Prohibitive (or need lots of patience!)
- Is there a better way?

Path closing and opening

- Minimum over closings with all possible lines that are approximately horizontal and composed of p pixels
- Thus: instead of taking closings with many slightly rotated lines, we take closings with many variations on the line shape
- Number of combinations grows exponentially with p

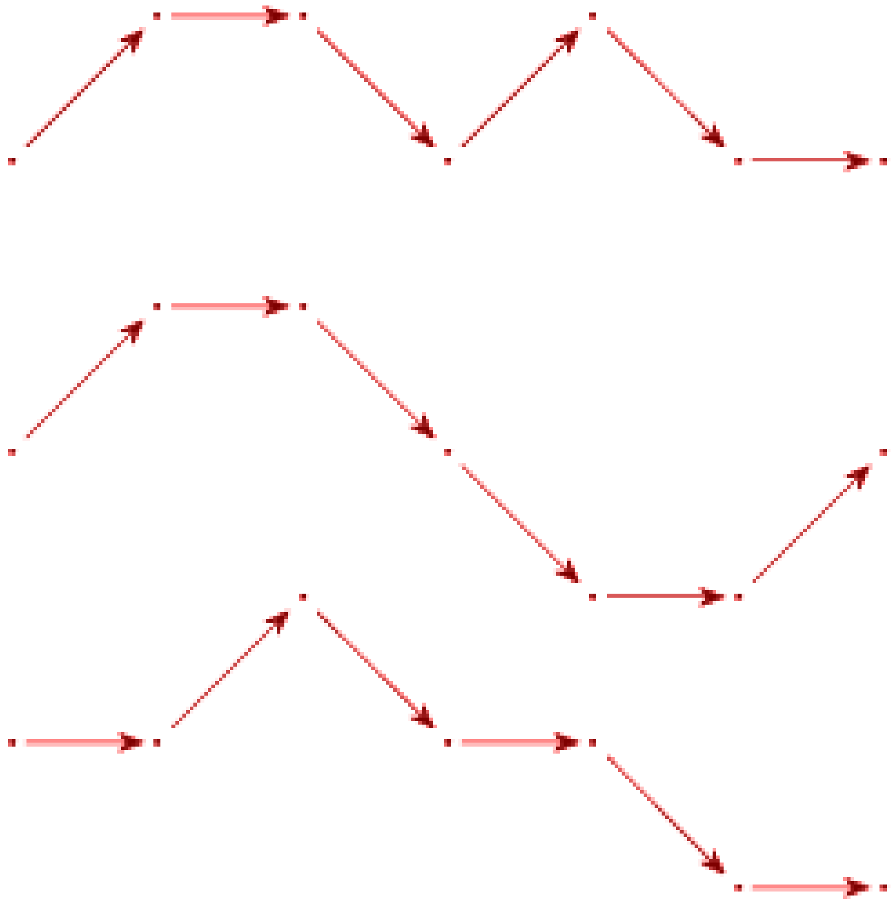

(Buckley & Talbot, 2000)

Adjacency relations

Horizontal paths:
E, NE, SE

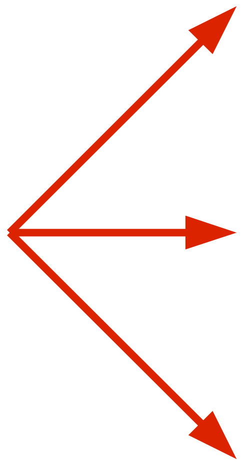

Diagonal paths:
NE, N, E

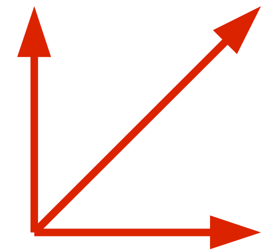

SE, E, S

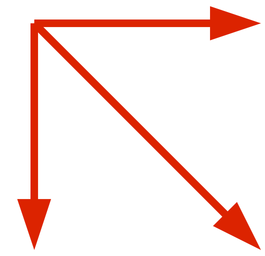

Vertical paths:
N, NW, NE

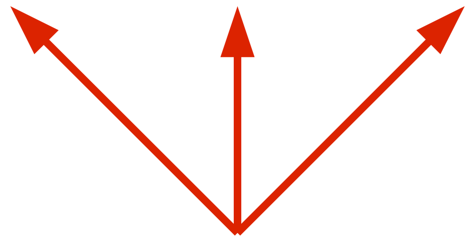

Cris Luengo
luengo@ieee.org

Path closing and opening

Horizontal paths:
E, NE, SE

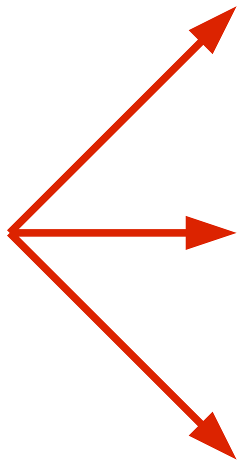

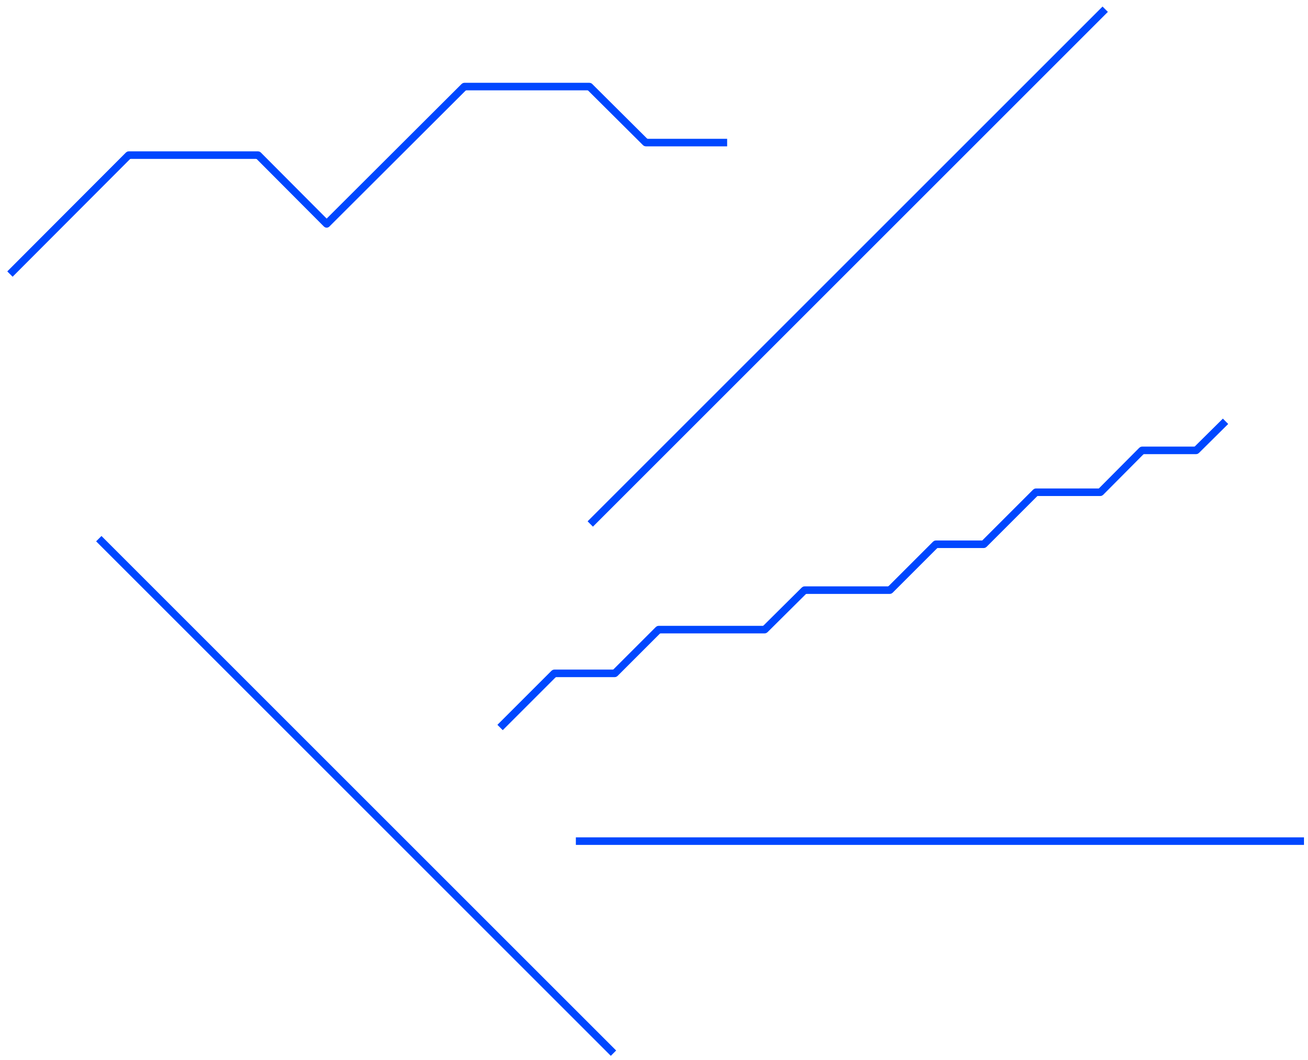

Path closing and opening

- Clever heuristic makes this algorithm $O(p \log(p))$
  - (Talbot & Appleton, Image and Vision Computing 25(4), 2007)
  - This is true for any number of dimensions
- Minimum of output of the 4 path closings gives a rotation-invariant operation (in 2D)
  - In 3D there are 13 different path closings
- $O(p \log(p))$ is much better than the $O(l^{n-1})$ complexity of rotated straight line closing (in an n-dimensional image)
- We need to add a constraint on consecutive steps to make the assumption $p \approx l$
  - (Luengo, IEEE Trans. Image Processing 19(6), 2010)

Constrained path closing and opening

Allow only one consecutive step in a direction that is not the main direction

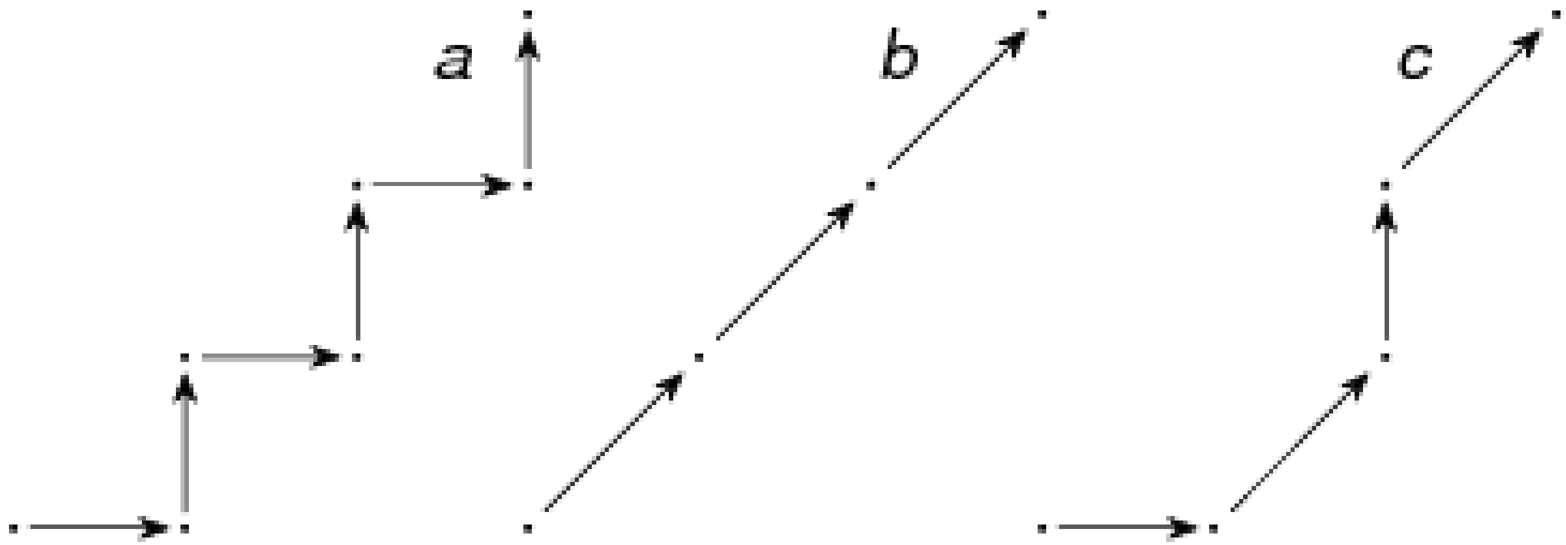

Constrained path opening

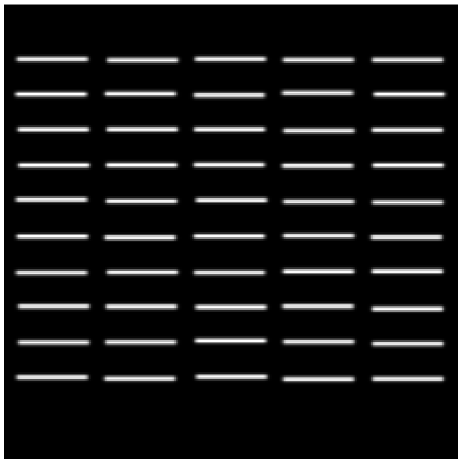

0°

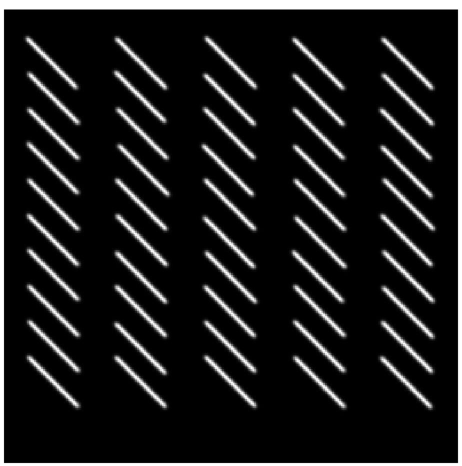

45°

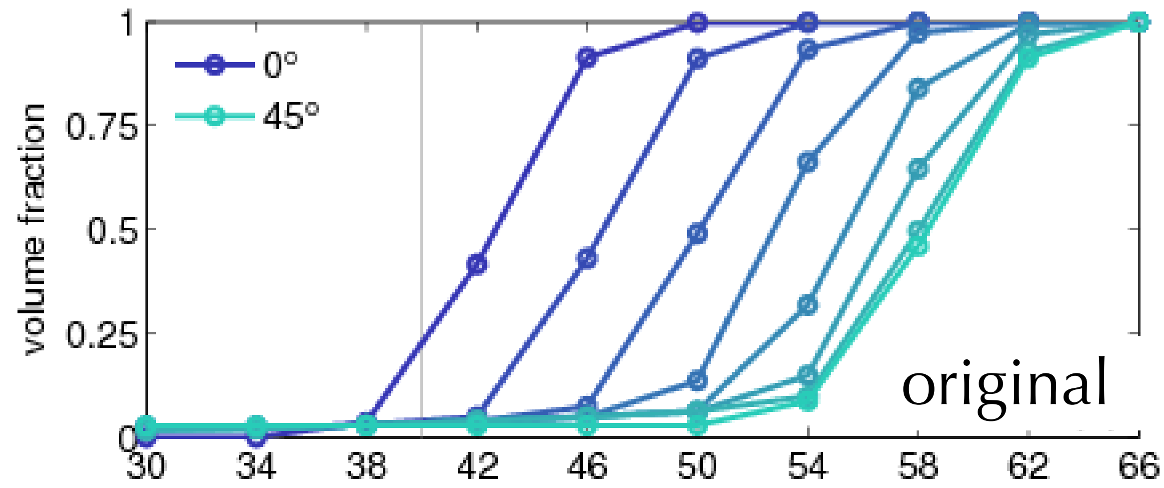

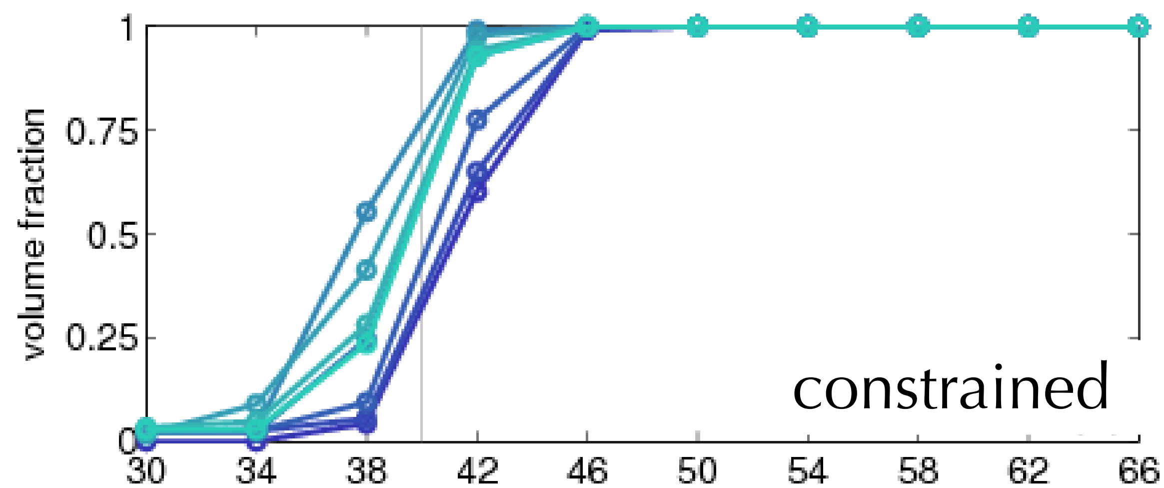

Fibre composite

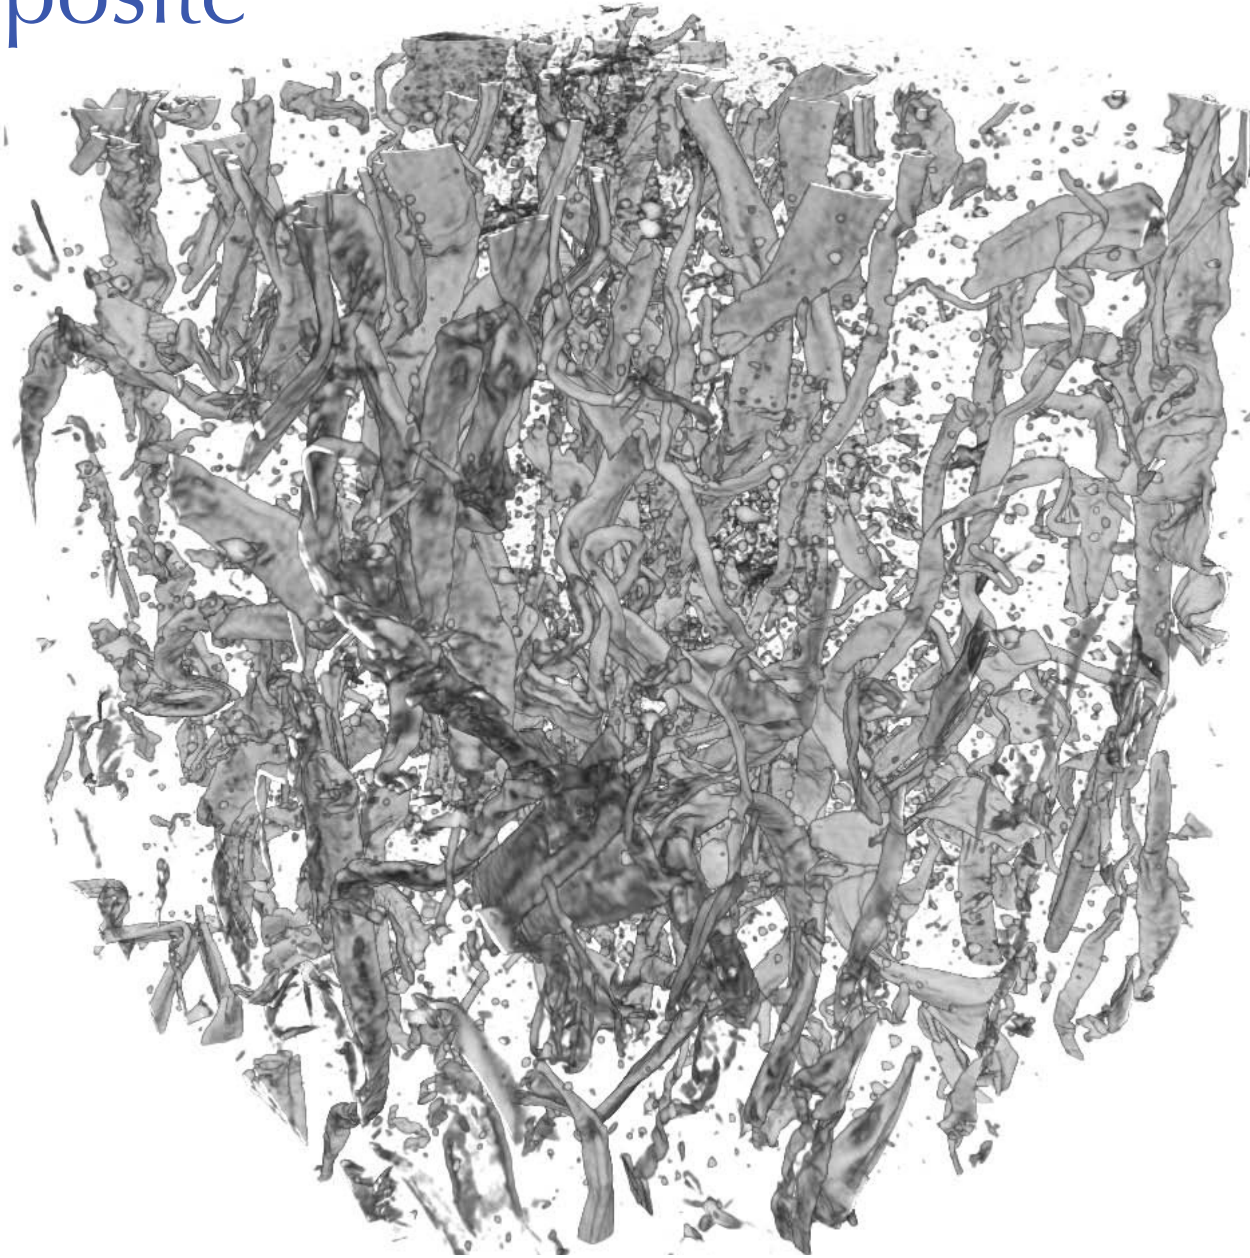

(rendering: Arttu Miettinen)

Fibre length distribution

- Volume-weighted distribution ≠ count
- Fibres partially in the image?

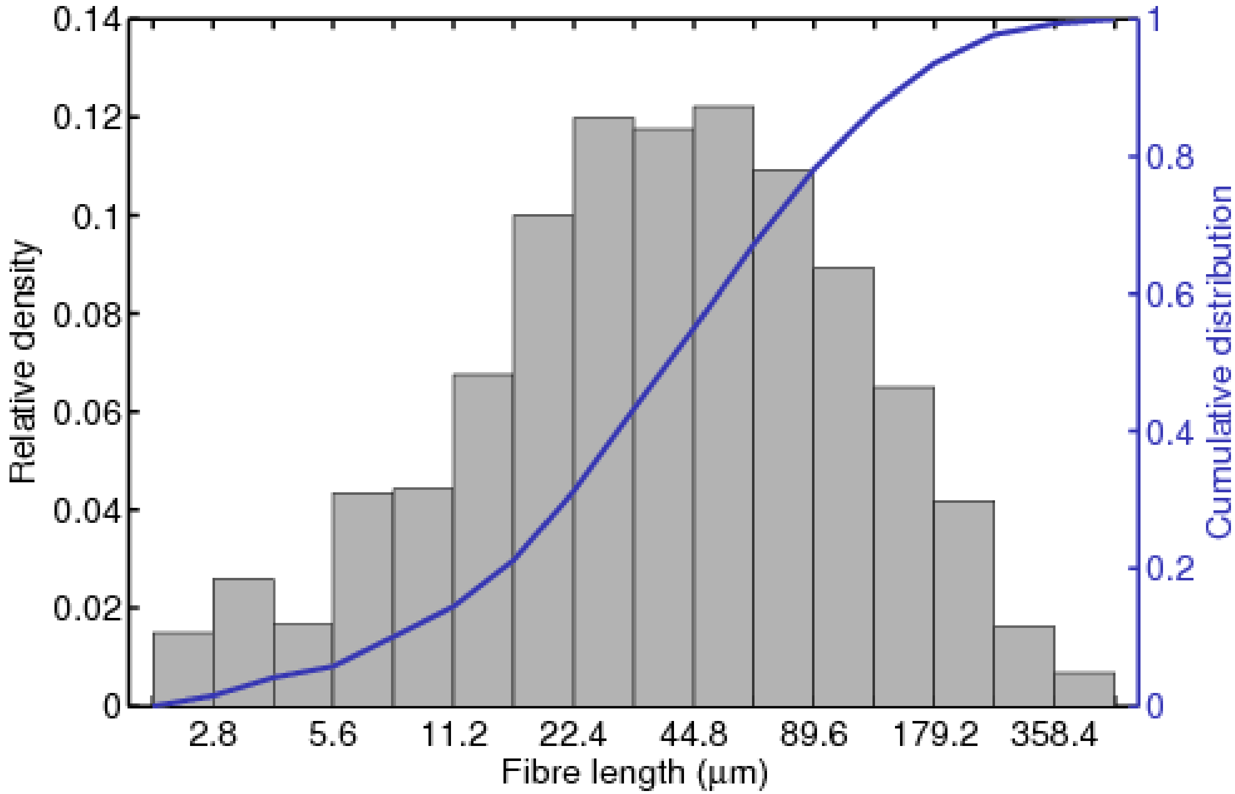

(Luengo, 2010)

Fibre length distribution

- Fibres cut at edge of image
  - Long fibres will seem shorter
  - We will underestimate the amount of long fibres, and overestimate the amount of short fibres
- Long fibres are more likely to be cut by the image edge
  - If we don't measure cut fibres, we will underestimate the amount of long fibres
- Fibres longer than image will always be cut
  - Choose a proper image size

Cris Luengo
luengo@ieee.org

Correcting the distribution

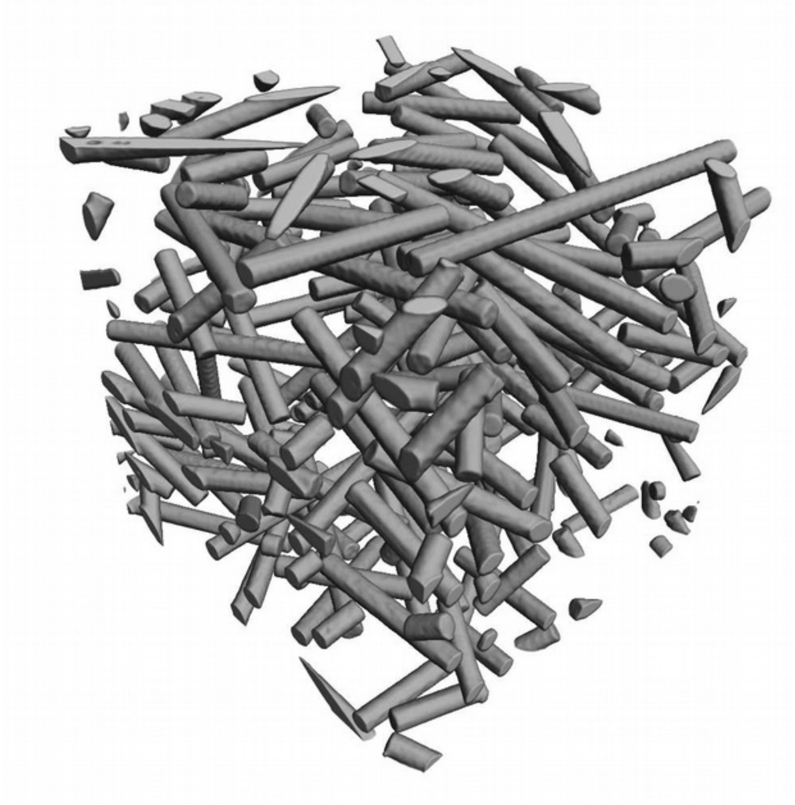

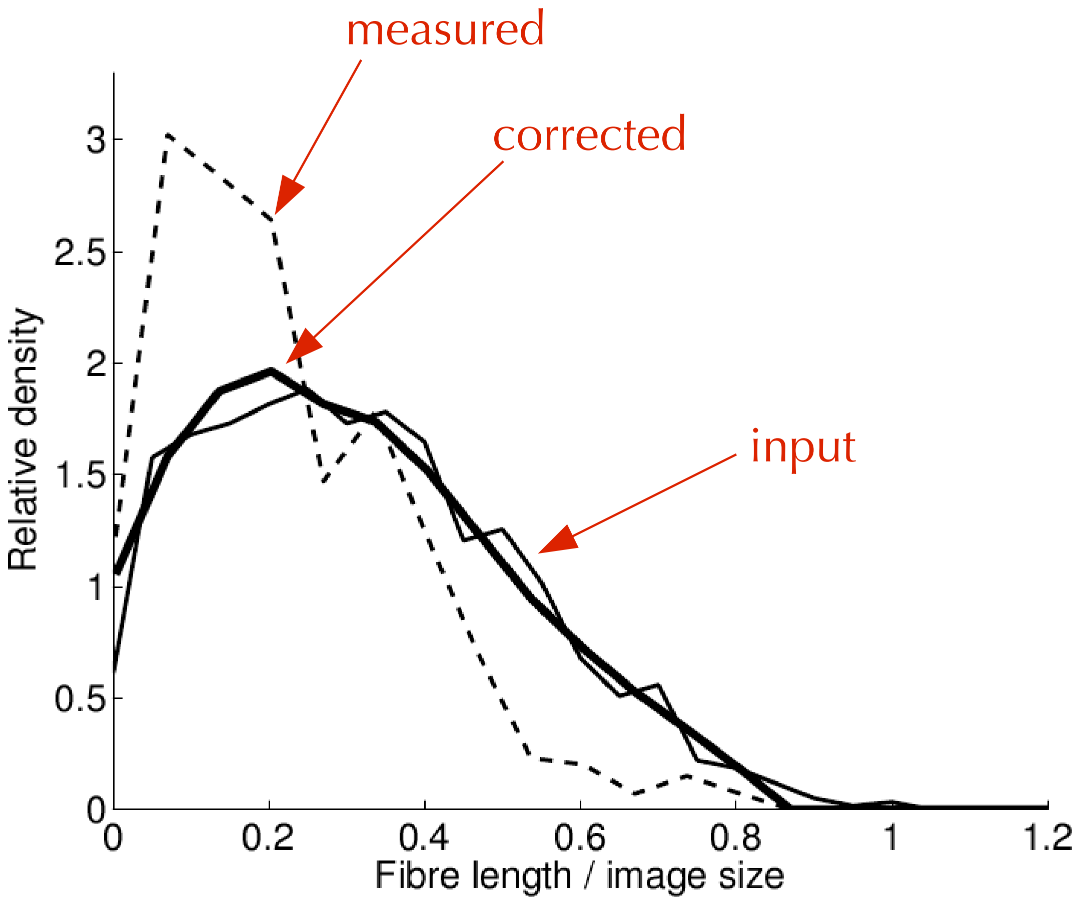

(Miettinen & al., 2012)

Correcting the distribution

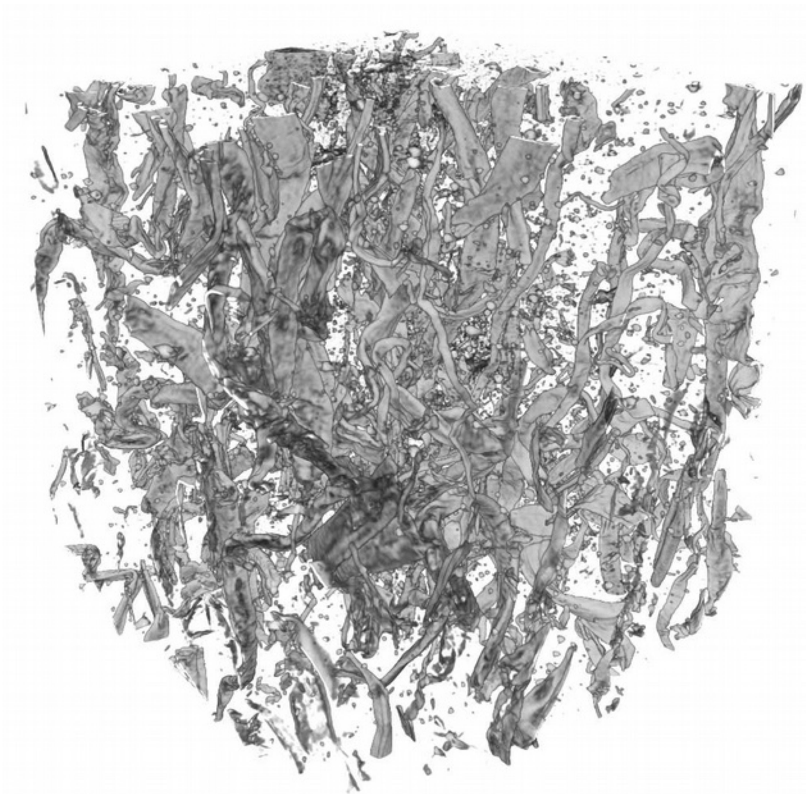

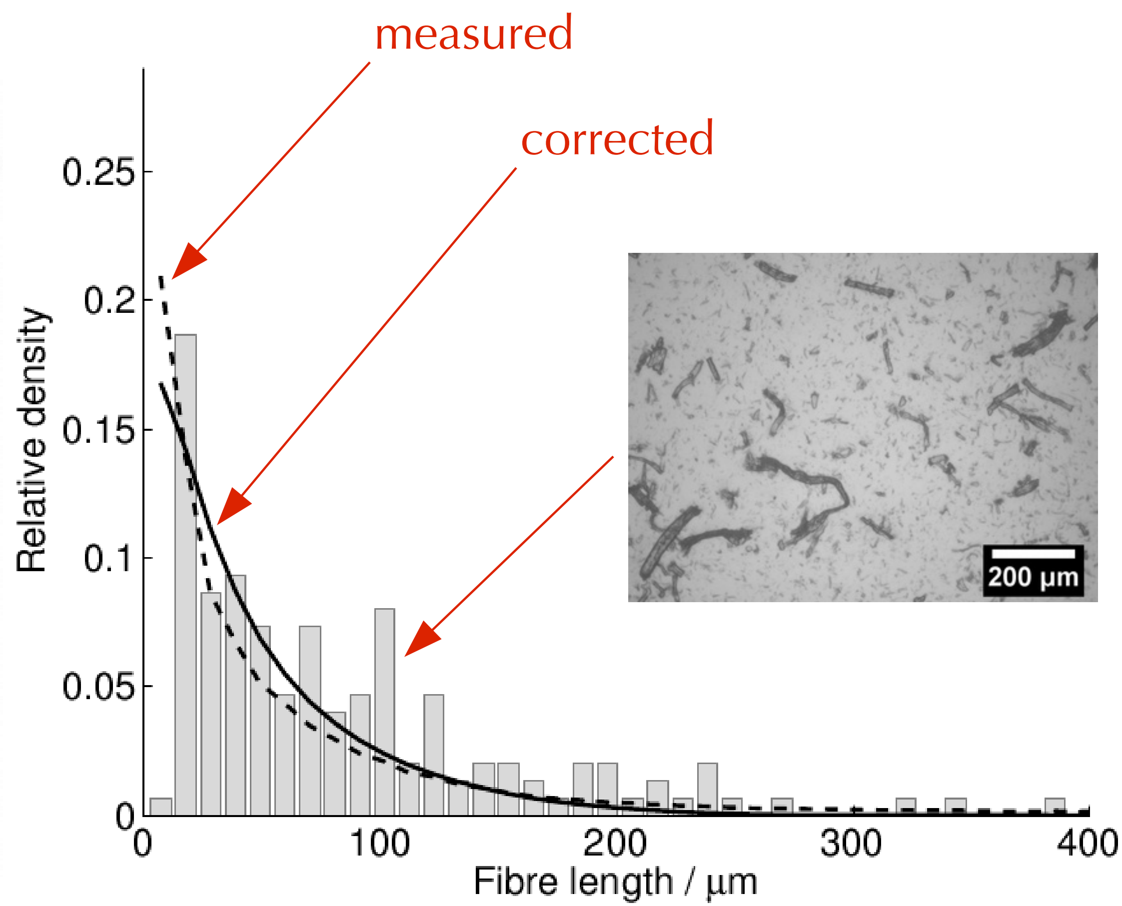

(Miettinen & al., 2012)

Incomplete and robust path opening

- Incomplete path opening:
  - rank-max opening
  - a certain number of the *l* pixels along the path are ignored
  - (by original authors of path opening)
- Robust path opening:
  - skipping at most *n* pixels in between each pixel on the path
  - based on my version of the algorithm
  - (Cokelaer, Talbot & Chanussot, 2012)

Cris Luengo
luengo@ieee.org

Parsimonious path opening

- Speeding up the path opening by:
  - selecting maximal paths from top to bottom of image (or left to right, or diagonally)
  - applying 1D openings along these paths
  - reconstruction by dilation with input image as mask
- Algorithm is linear with number of pixels in image and independent of path length
- Selecting paths in smaller strips of the image improves results

(Morard, Dokládal & Decencière, 2014)

Cris Luengo
luengo@ieee.org

Parsimonious path opening

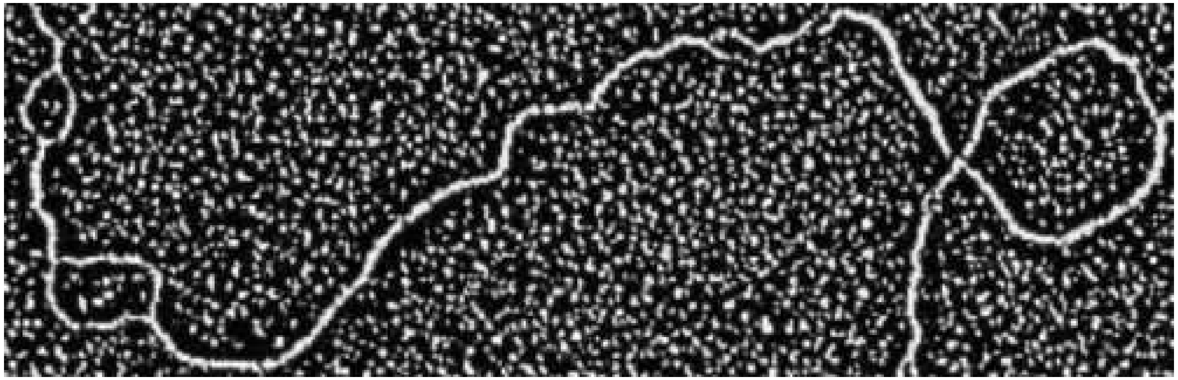

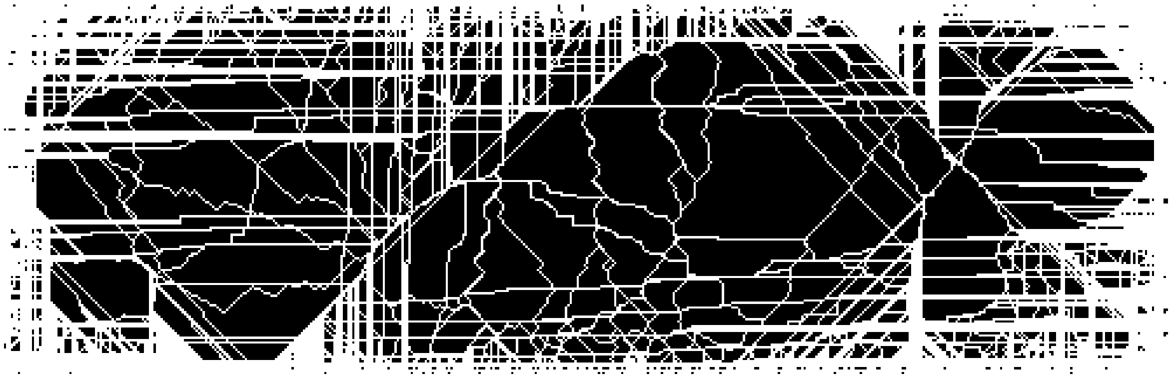

(Morard, Dokládal & Decencière, 2014)

Graph-based path opening

- Preselect paths using the grey-value skeleton
  - independent of path direction
  - avoids large gaps
- Preselected paths form a graph
- A graph version of the algorithm addresses many fewer pixels
  - O(p log p) operation becomes much closer to O(p)
  - ...but in practice larger p makes the algorithm faster
- H-minima transform simplifies the skeleton
  - for further speed improvements

Cris Luengo
luengo@ieee.org

Graph-based path opening

no H-minima transform

small H

larger H

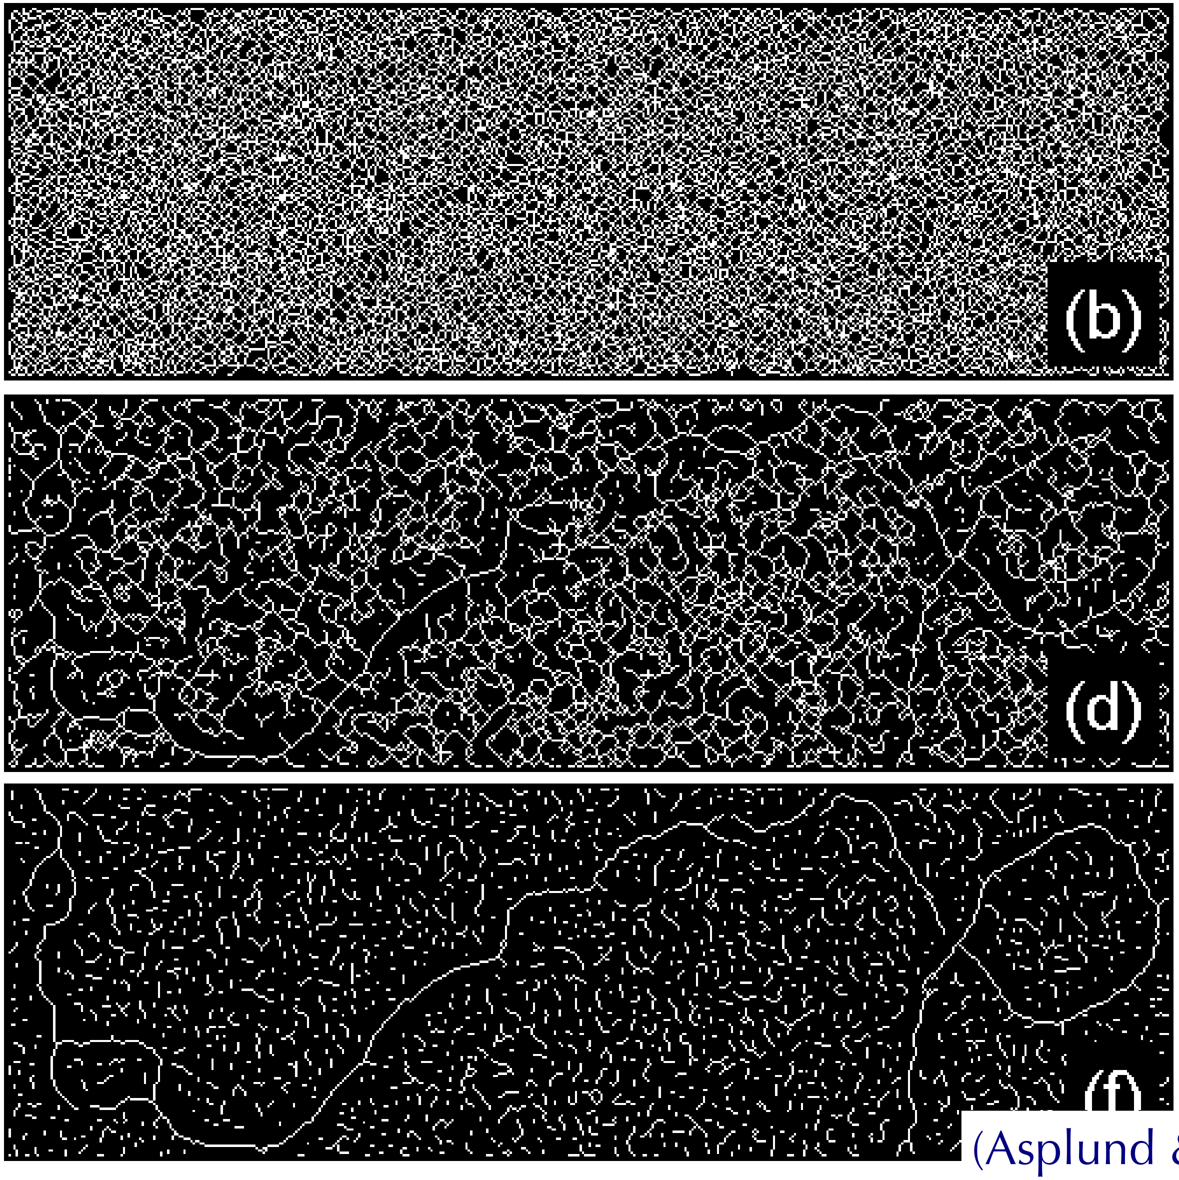

(Asplund & Luengo, 2015)

Graph-based path opening

input

standard path opening

graph-based path opening

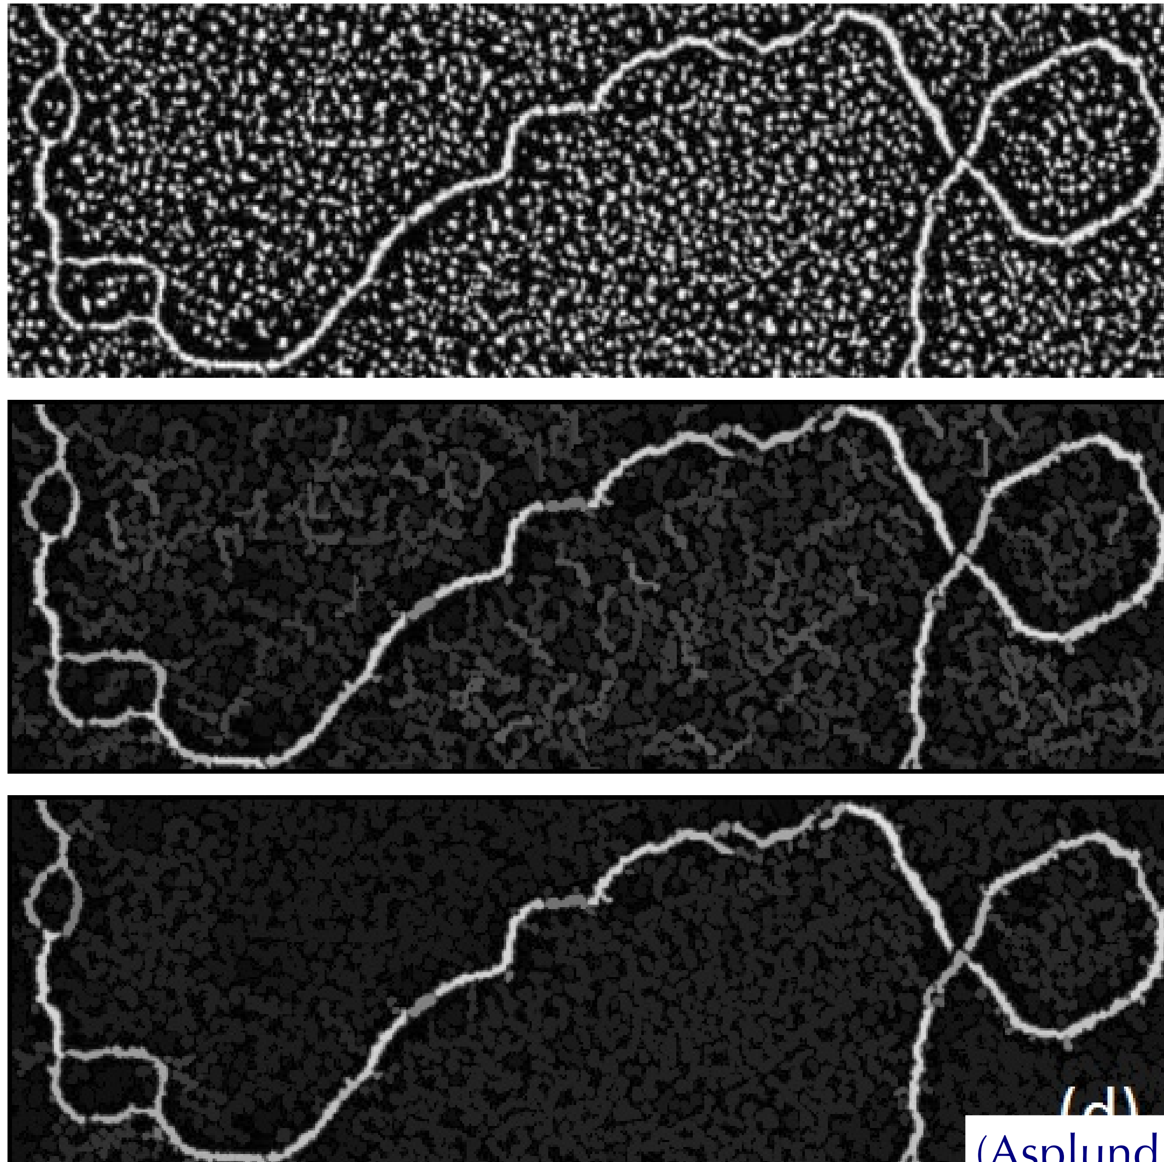

(Asplund & Luengo, 2015)

Centre for Image Analysis
Uppsala University

Cris Luengo
luengo@ieee.org

UPPSALA UNIVERSITET

Summary

- Granulometry
  - multi-scale operation
  - normalization to obtain size distribution estimate
  - applied with isotropic structuring elements:
    size is width
  - applied with rotating line structuring elements:
    size is length
- Path opening
  - fast alternative to line structuring elements at many angles
  - yields length distribution in a granulometry

Cris Luengo
luengo@ieee.org