Image Features from Phase Congruency

Peter Kovesi

Centre for Exploration Targeting
School of Earth and Environment
The University of Western Australia

Natural images are formed by the reflectance of light off objects. This, in conjunction with the shapes of objects, gives rise to images having some statistical properties

- Pixel values are correlated with their neighbours.
- Objects have a bias to being oriented in horizontal or vertical directions.
- Amplitude spectrum decays at approximately 1/frequency.

The human eye has evolved to process images having these properties

Images have a 1/f amplitude spectrum or $1/f^2$ power spectrum (approximately)

a) 

b) 

Field 1987, 1999.

These images from Hyvarinen, Hurri and Hoyer 2009.

a) 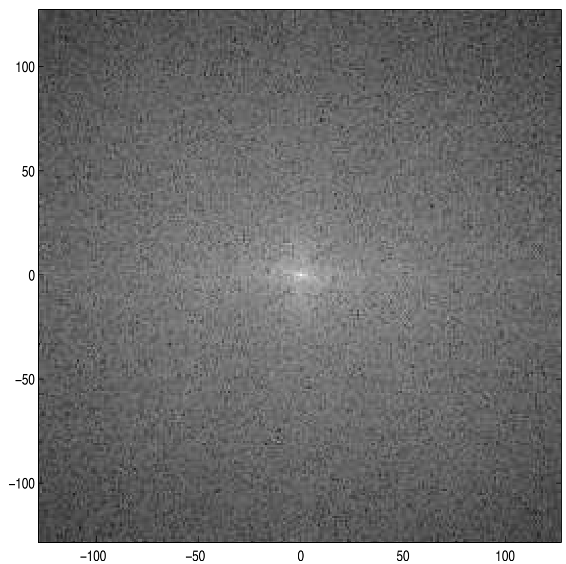

b) 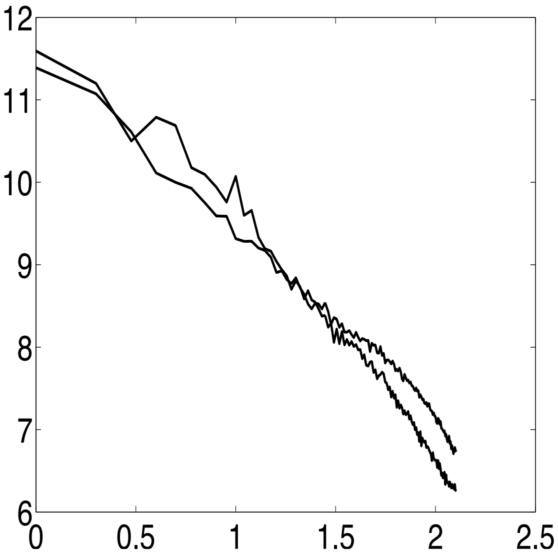

a) log of power spectrum

b) log-log plot of average cross section of power spectrum for both images. Slope is approximately $1/f^2$

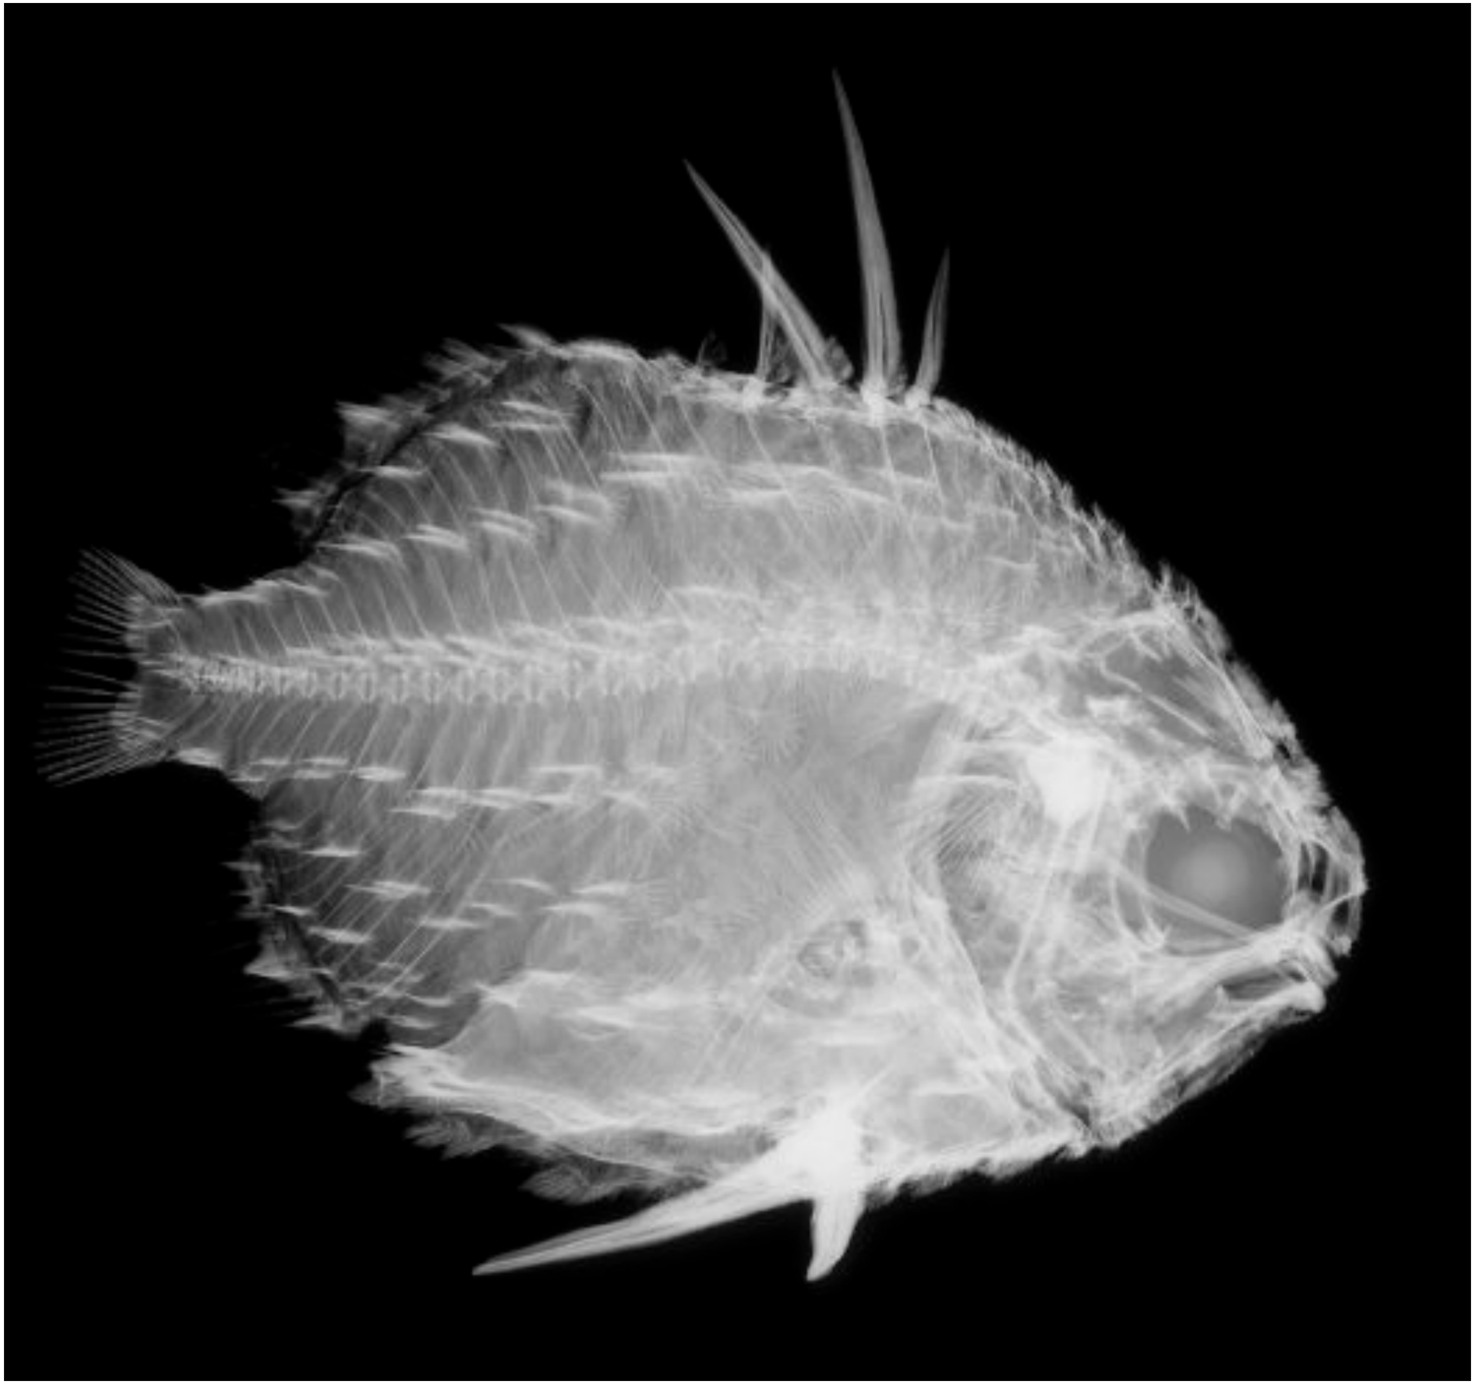

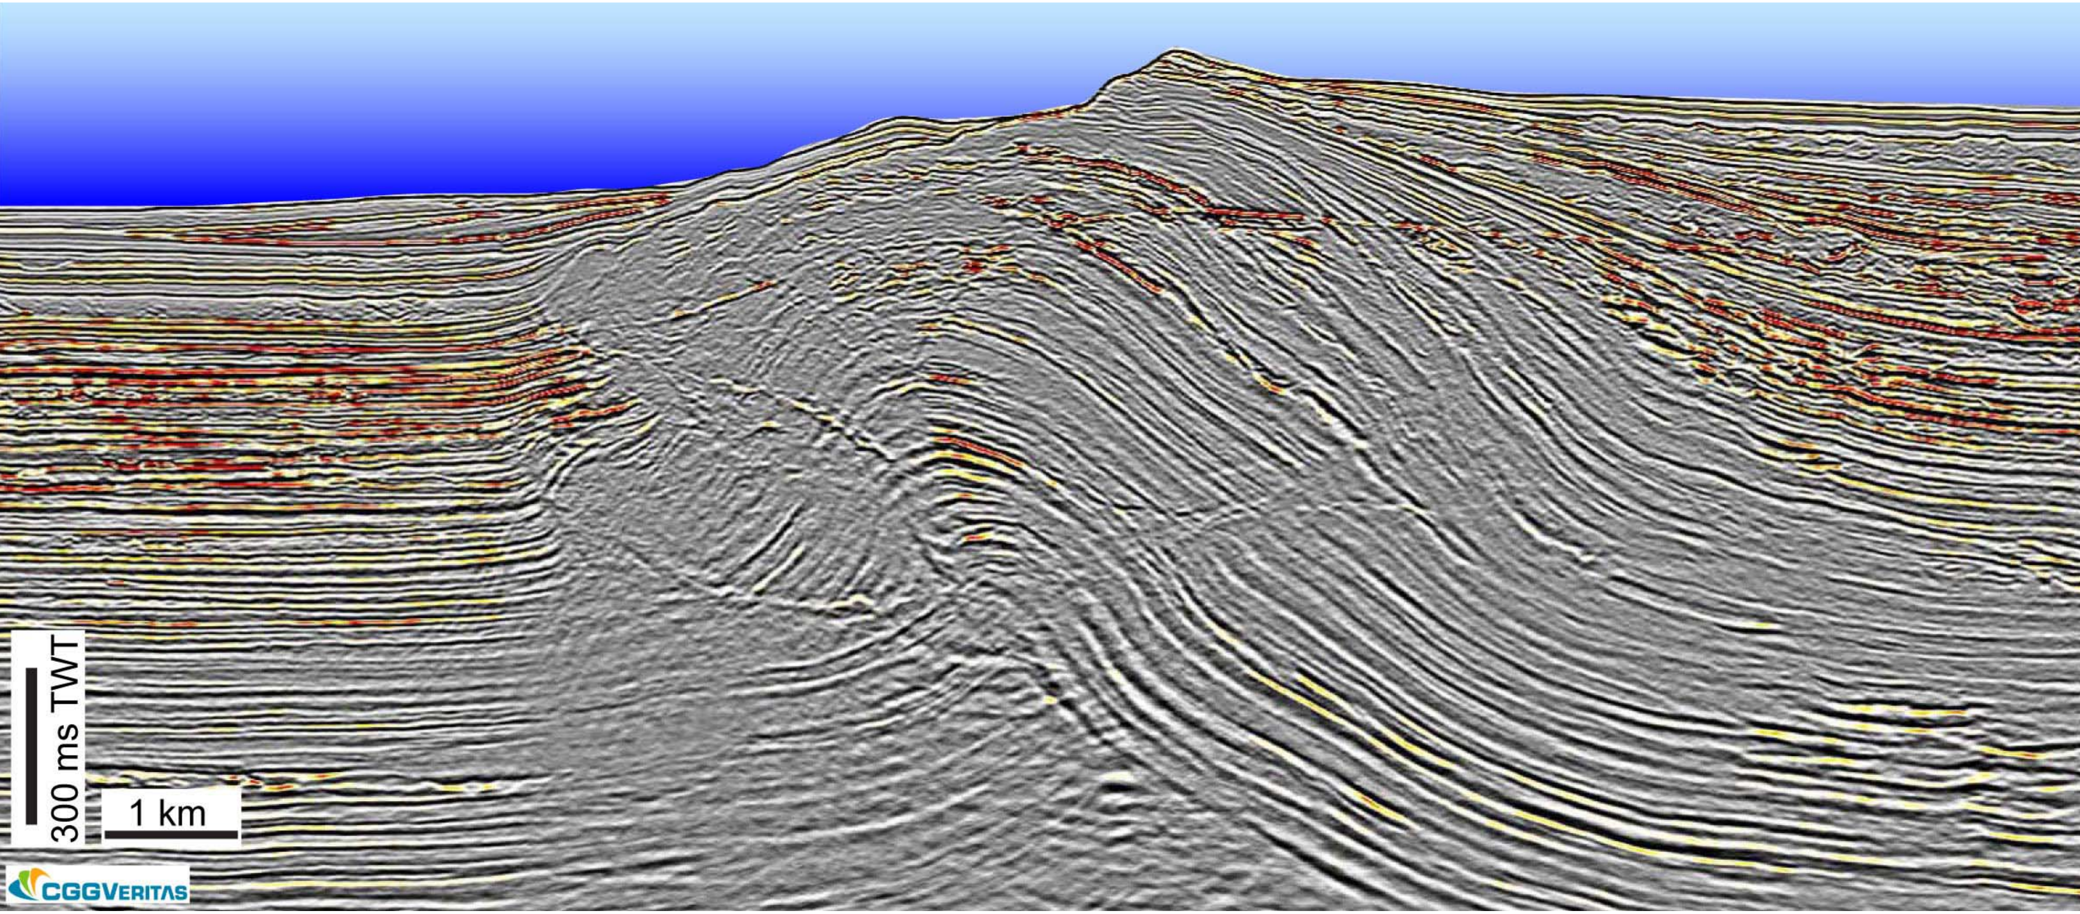
300 ms TWT
1 km
CGGVERITAS

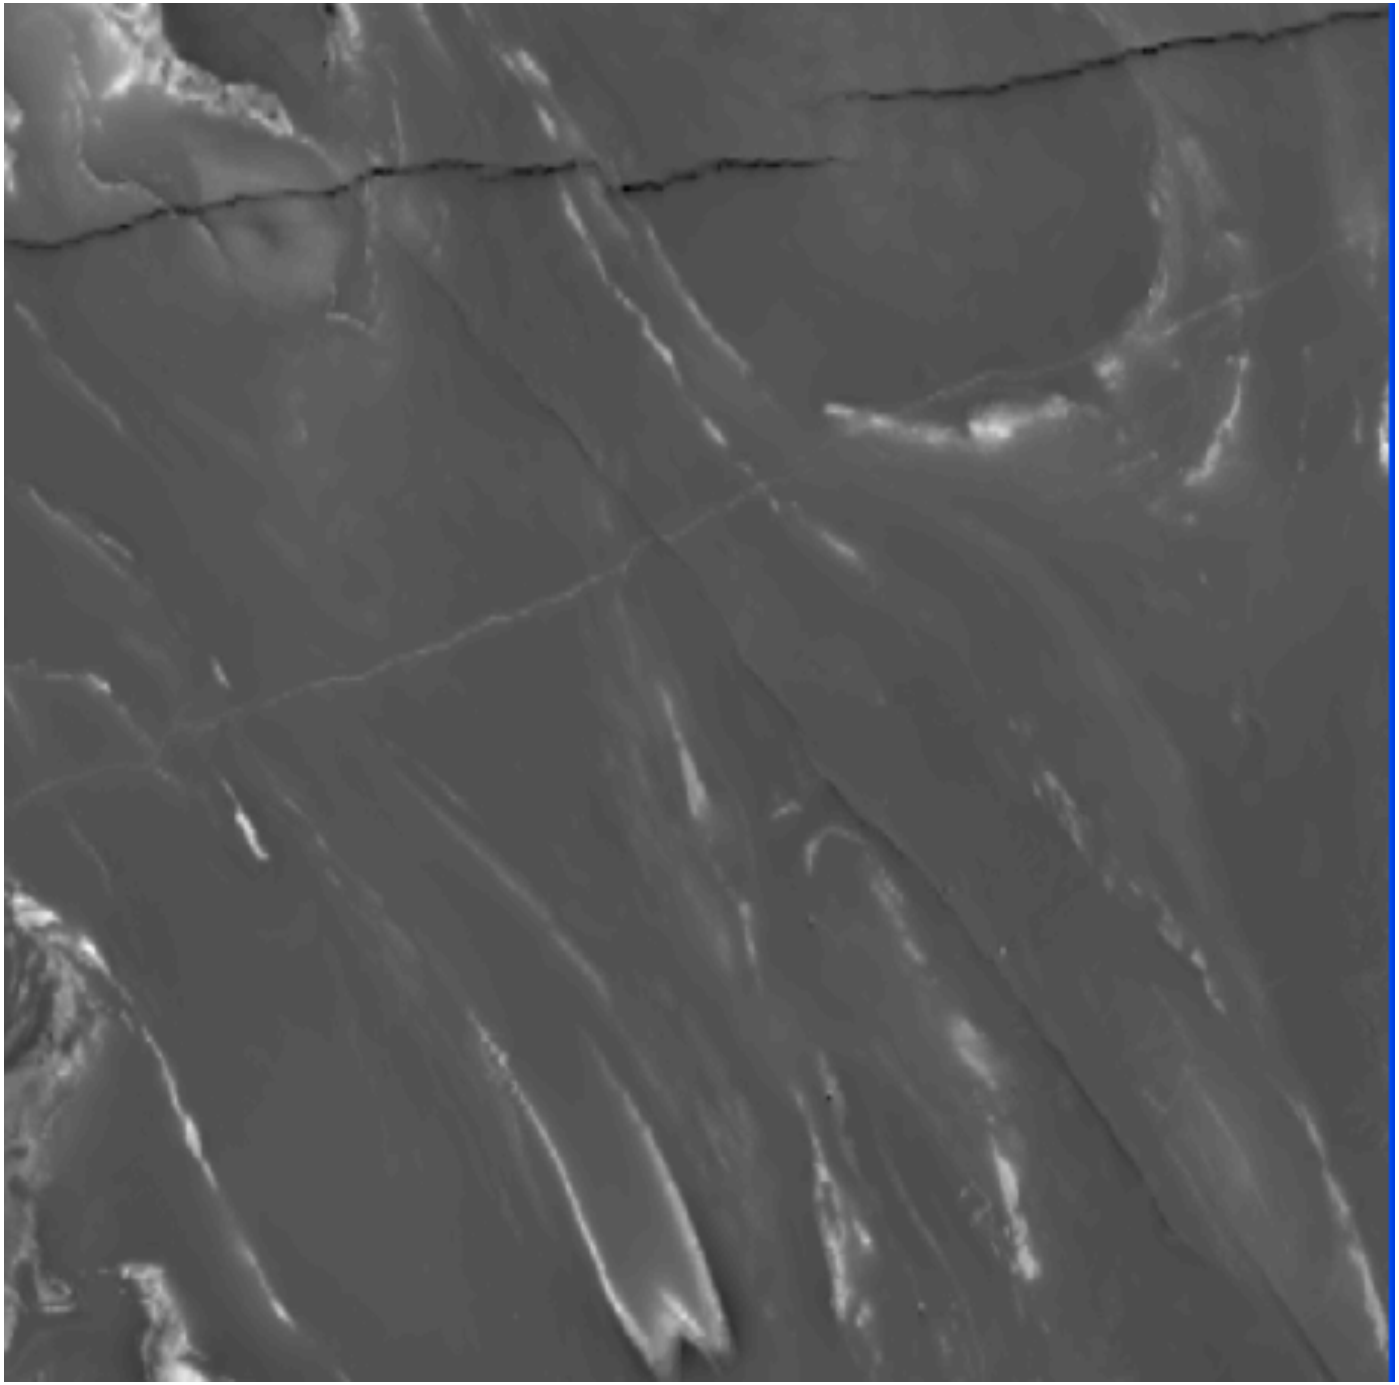

(Data belongs to
Fugro Airborne
Surveys Pty Ltd.)

The aesthetics of noise…

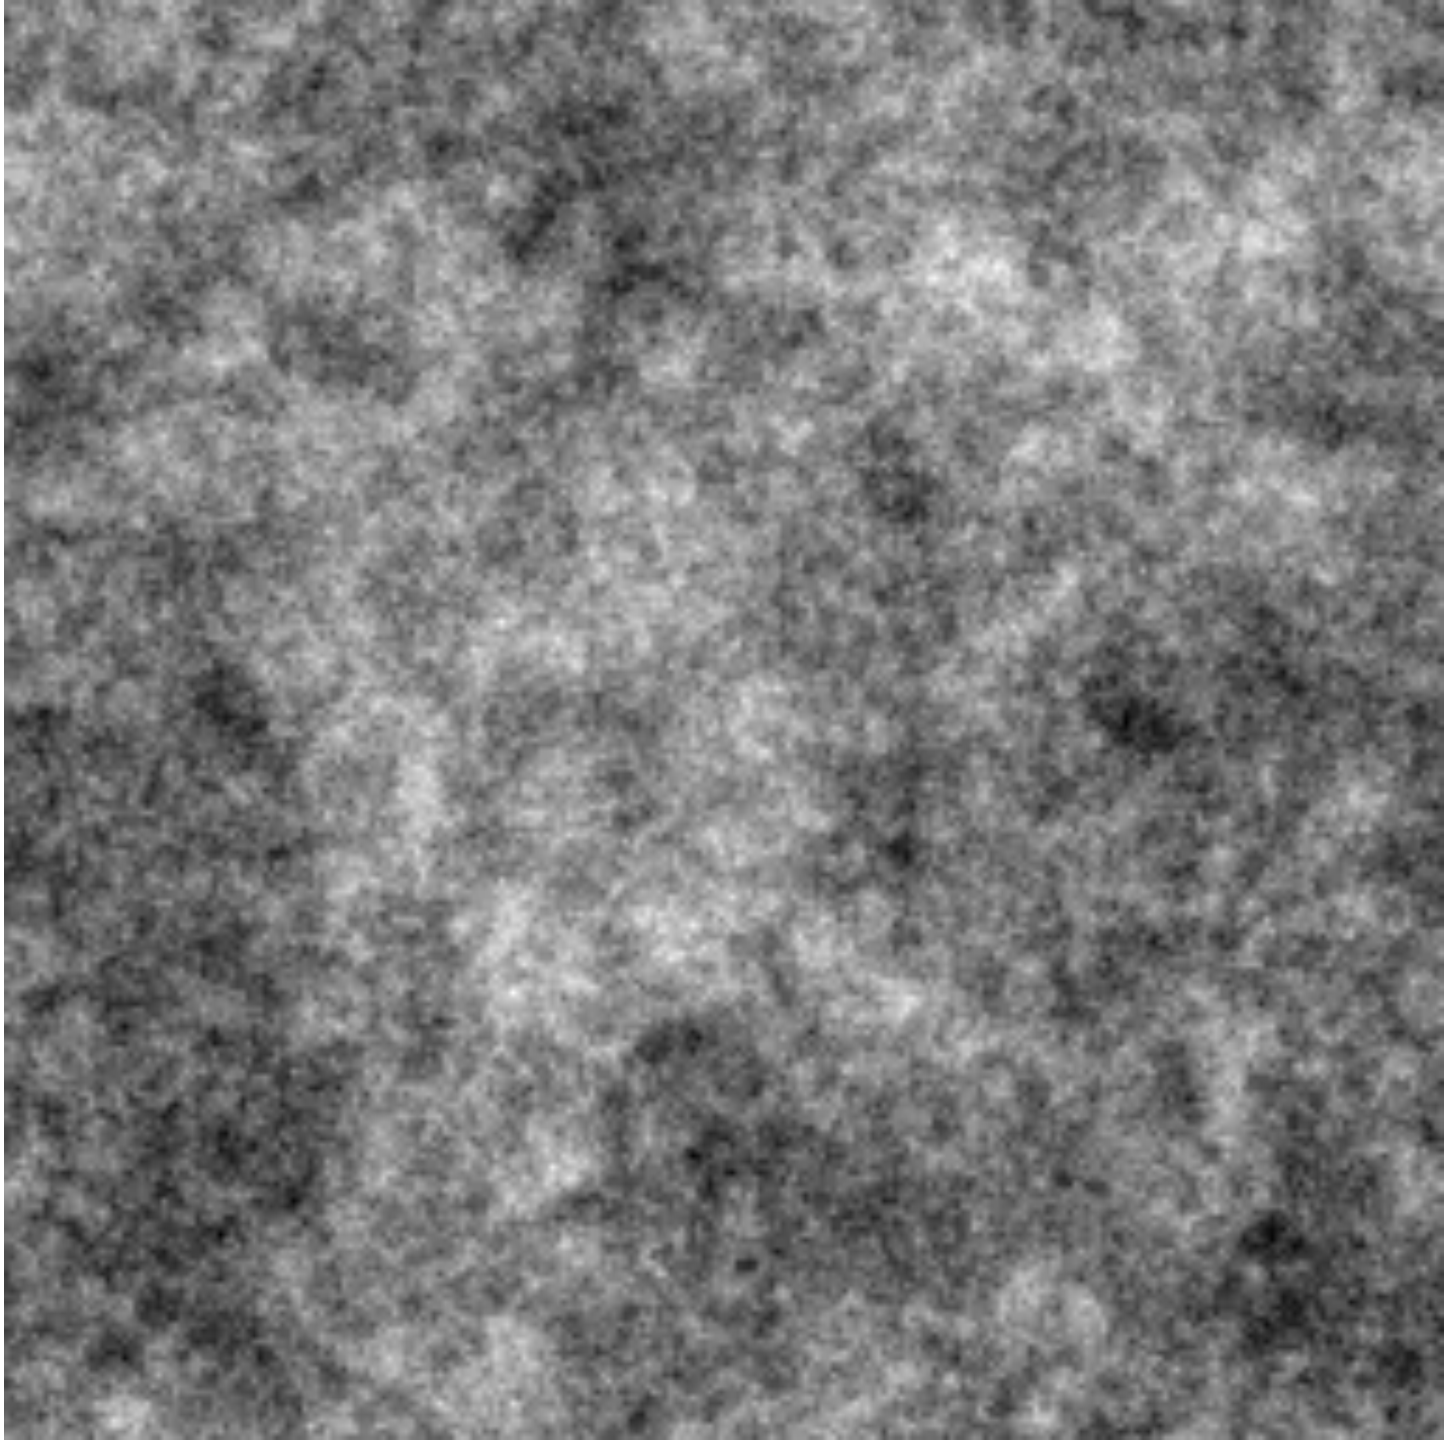

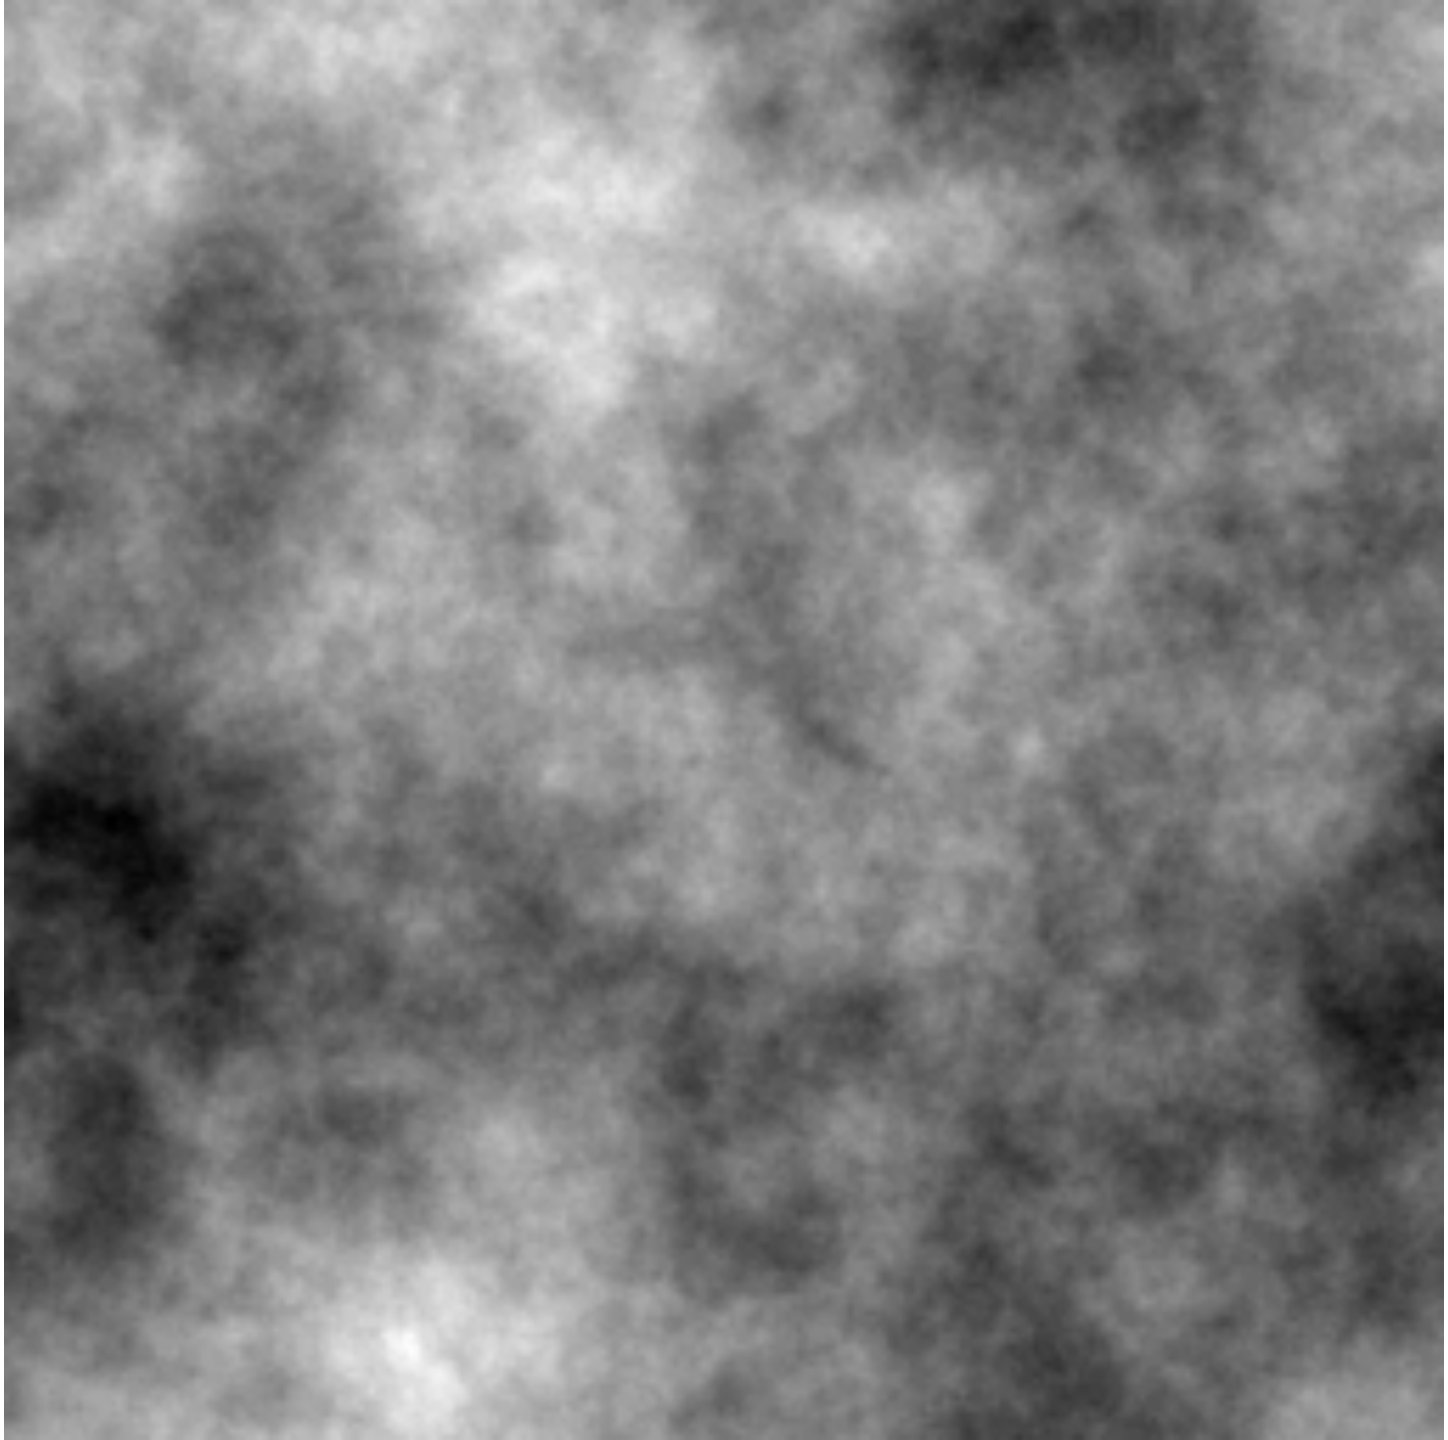

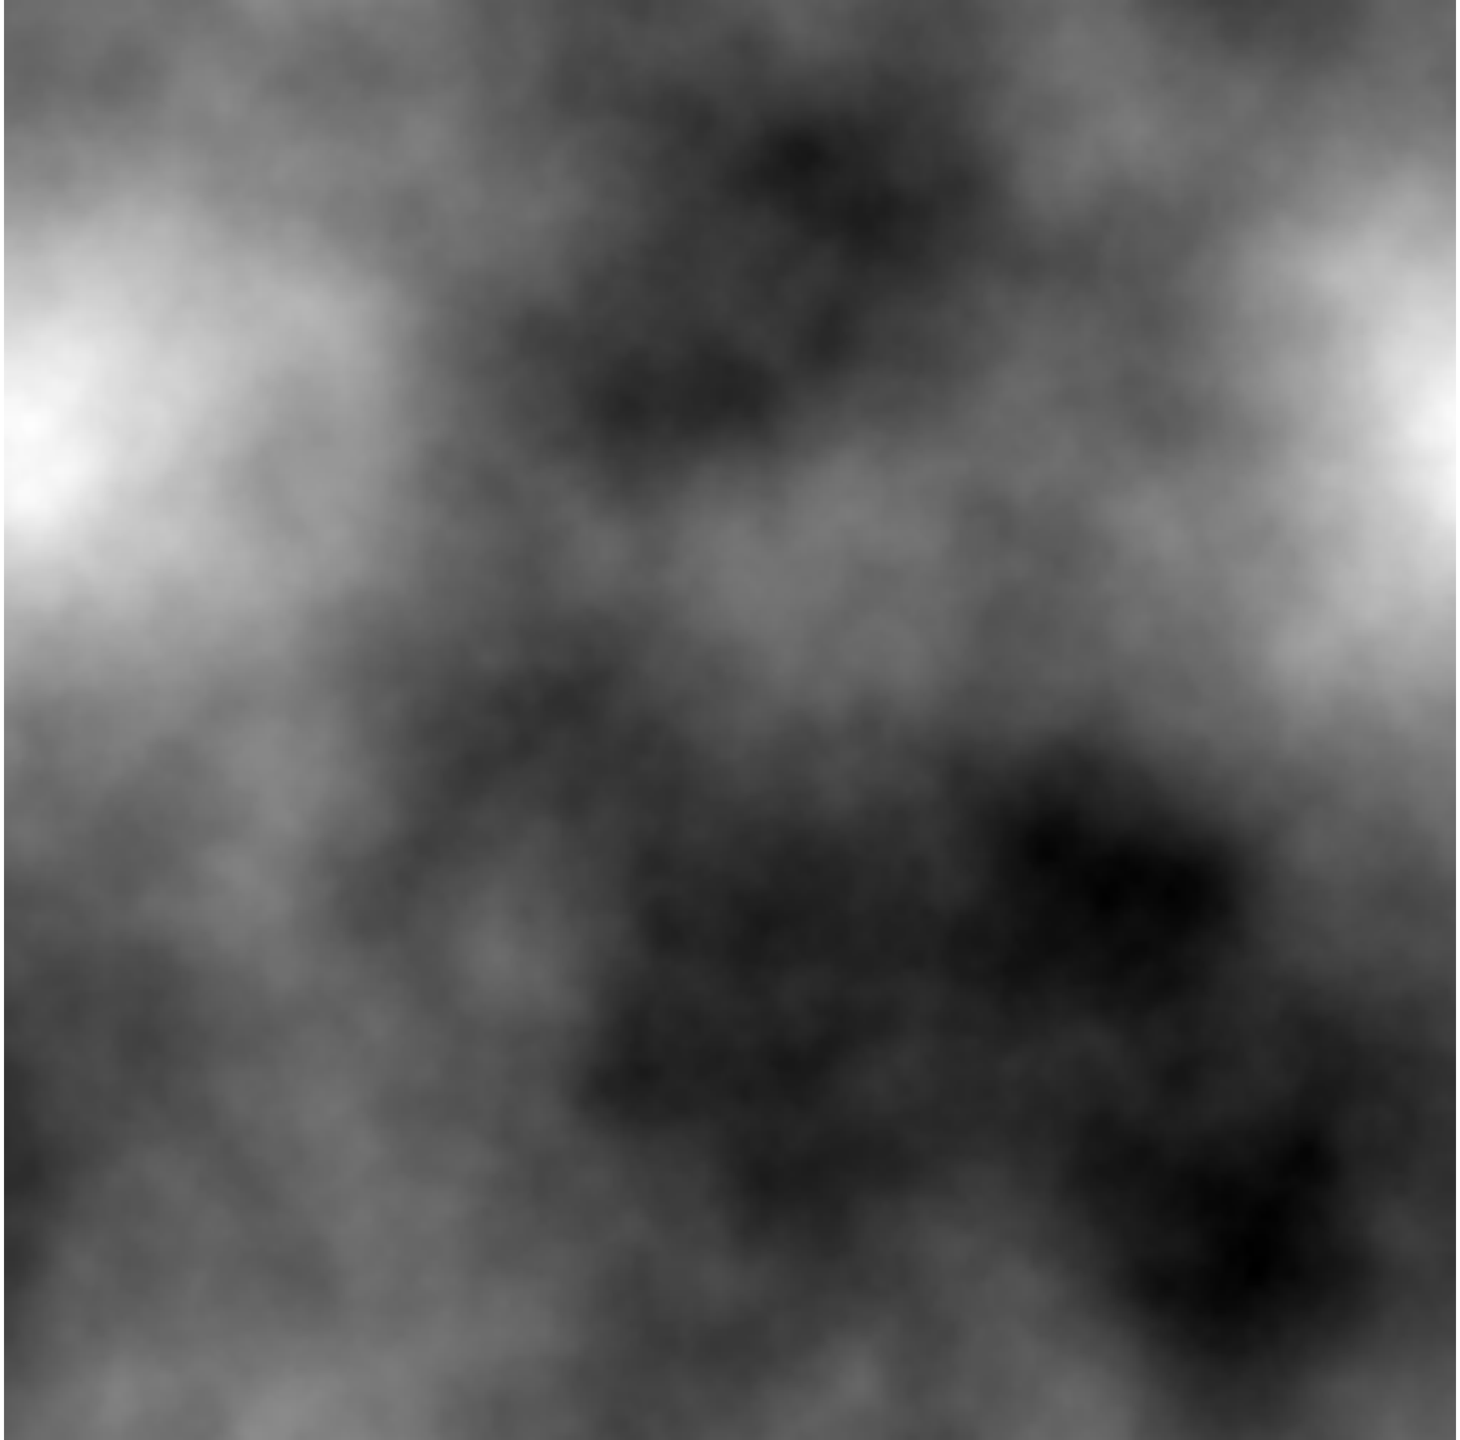

What are Image Features?

The classic model is to assume features are step discontinuities in the intensity.
Extrema in the smoothed gradient are marked as features.

Noisy step

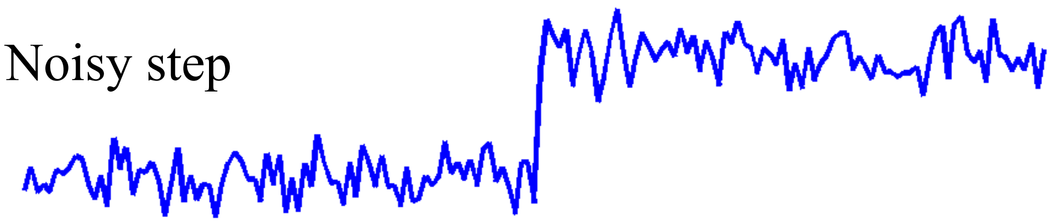

Gradient

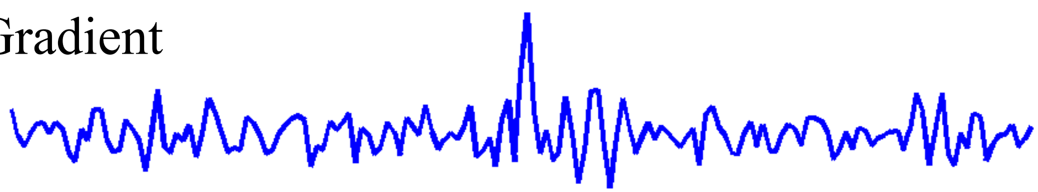

Smoothed gradient

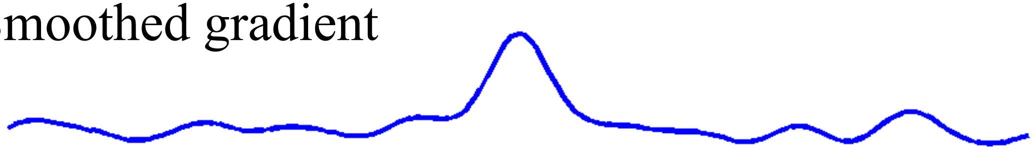

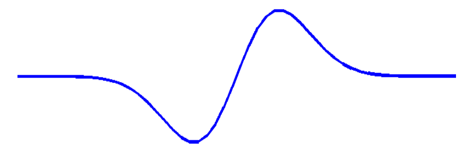

First derivative
Gaussian filter
Scale σ

This simplistic model can be very inadequate…


©2005 - TheColoringSpot.com. All rights reserved.

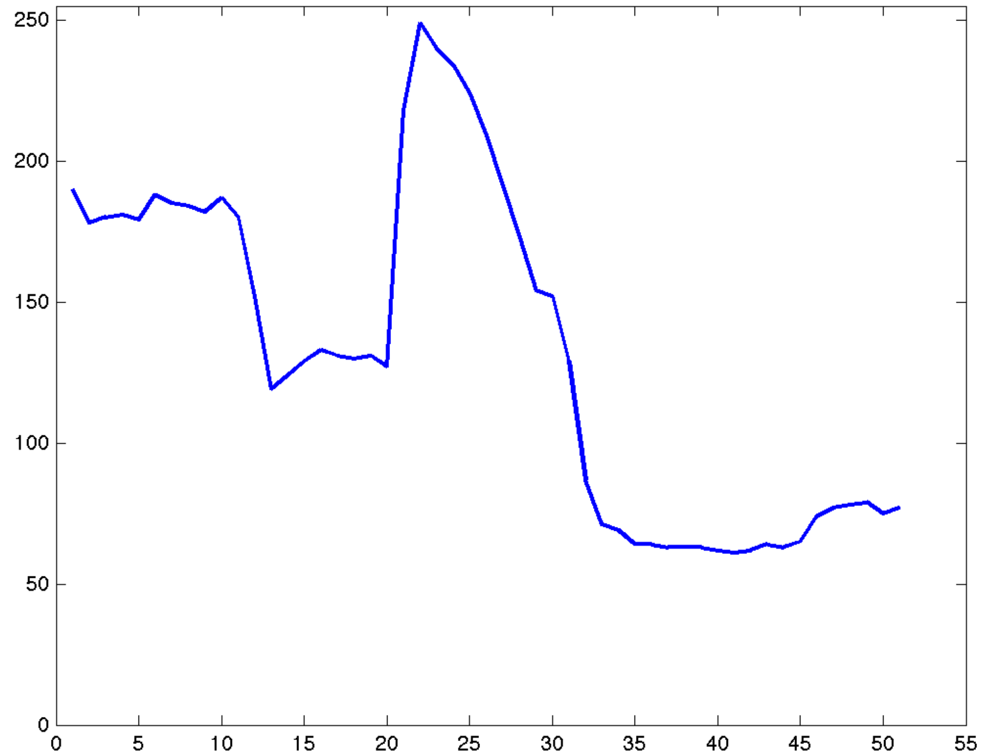

Image grey levels along image slice.

Where would you sketch a line in the image?

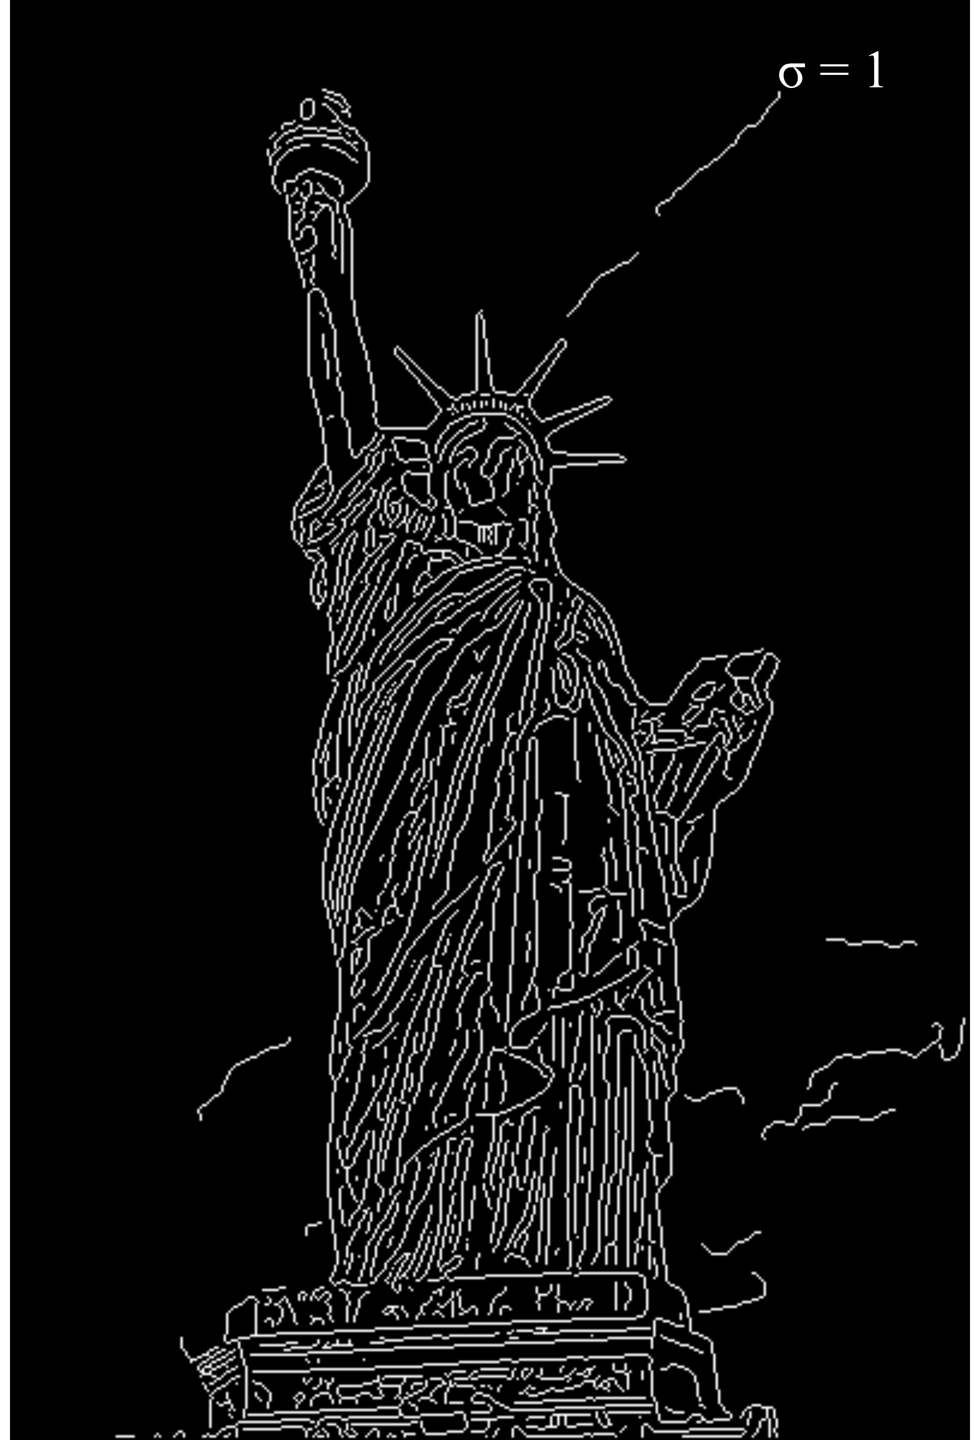
σ = 1

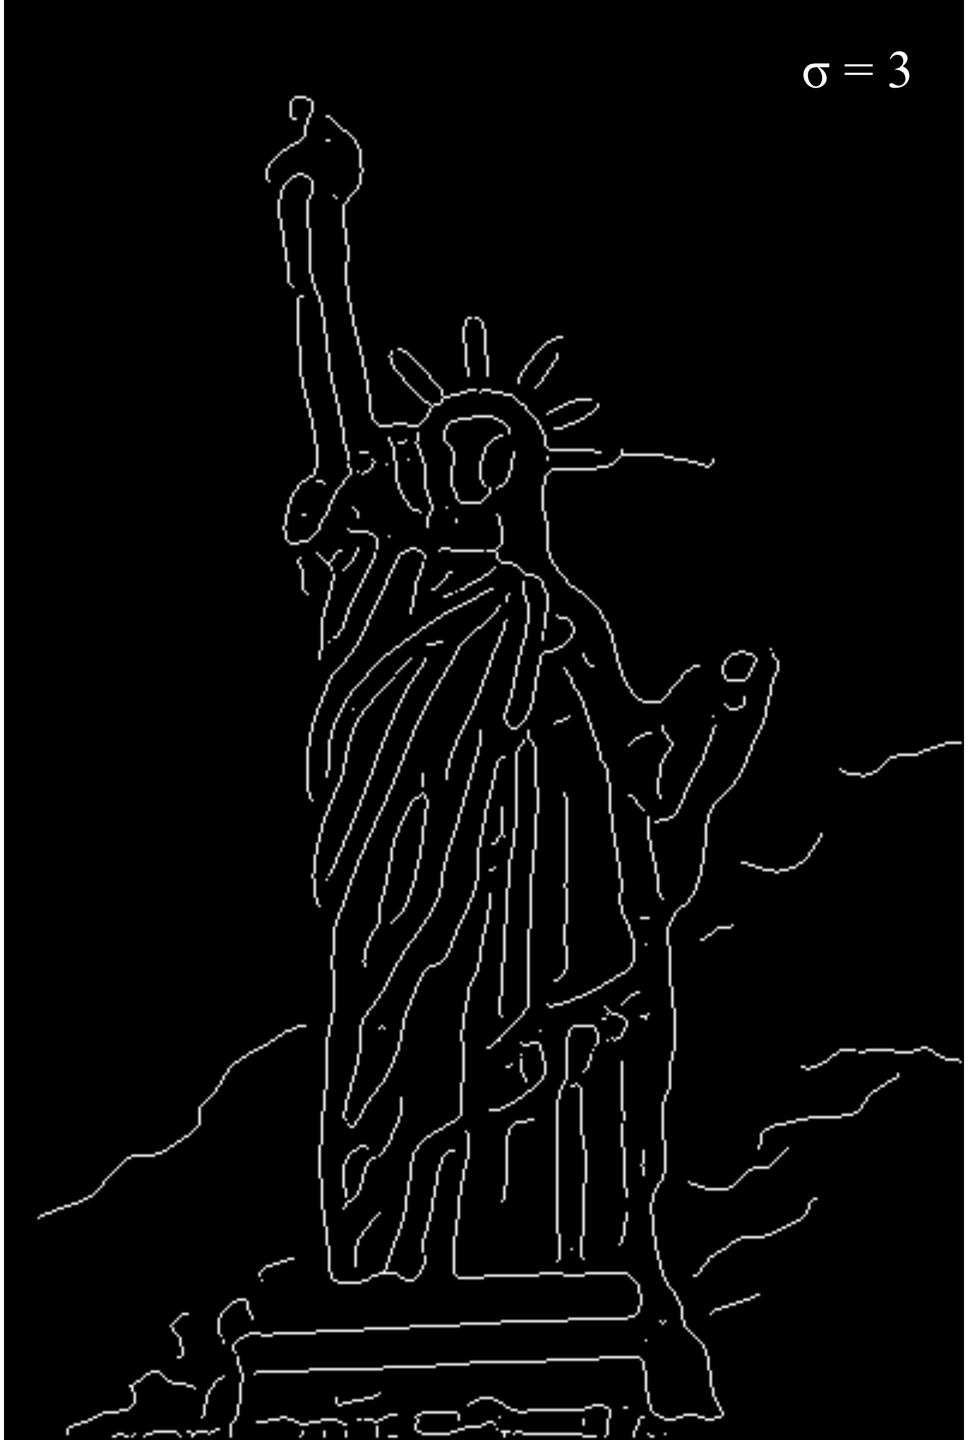
σ = 3
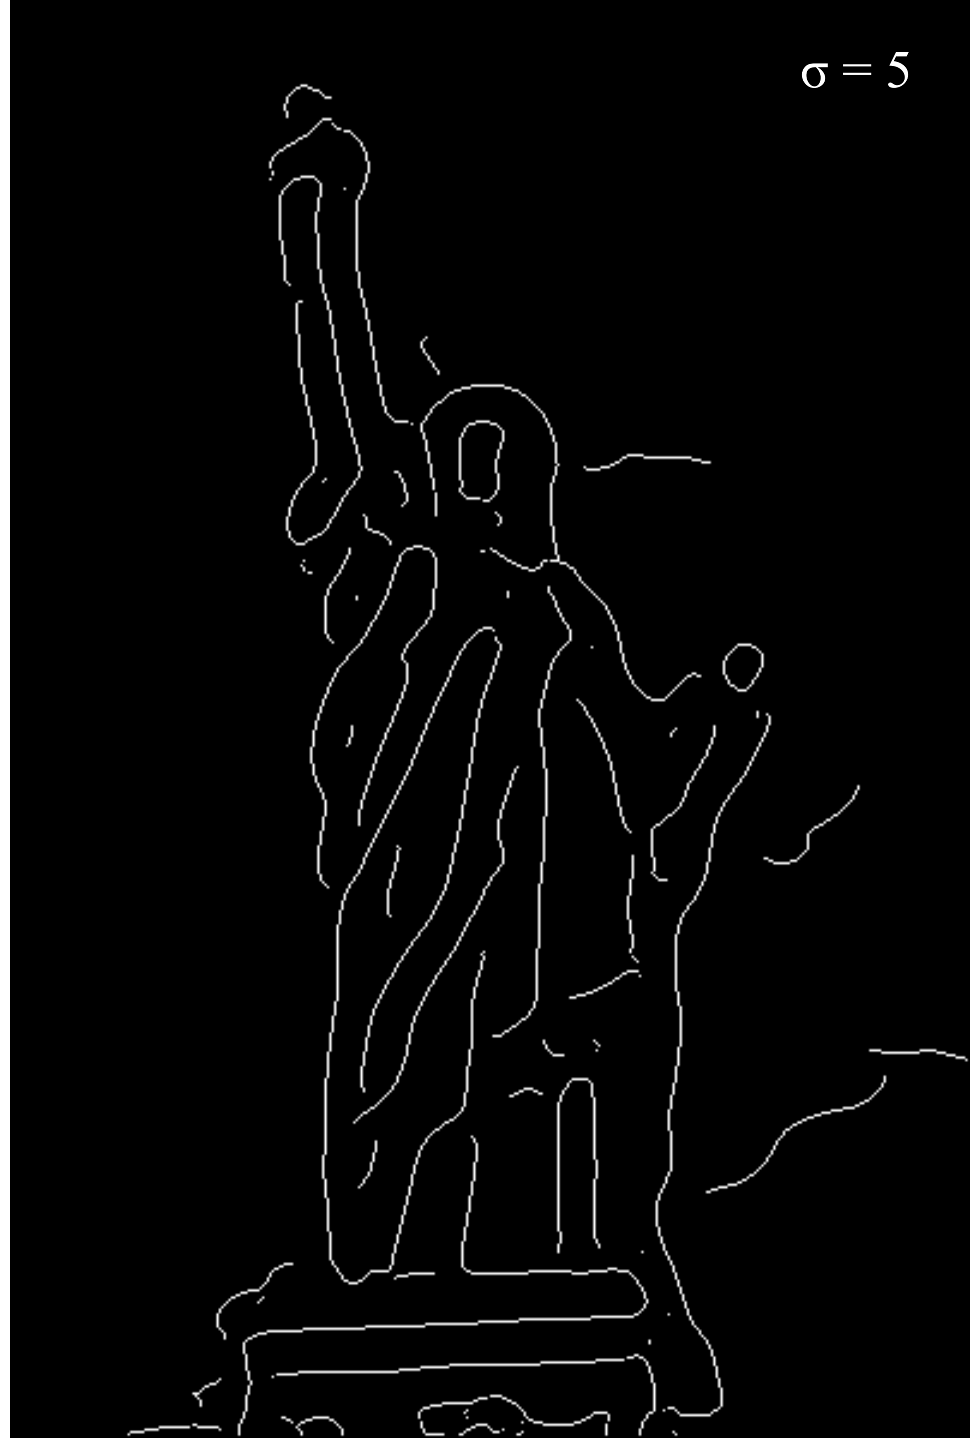
σ = 5

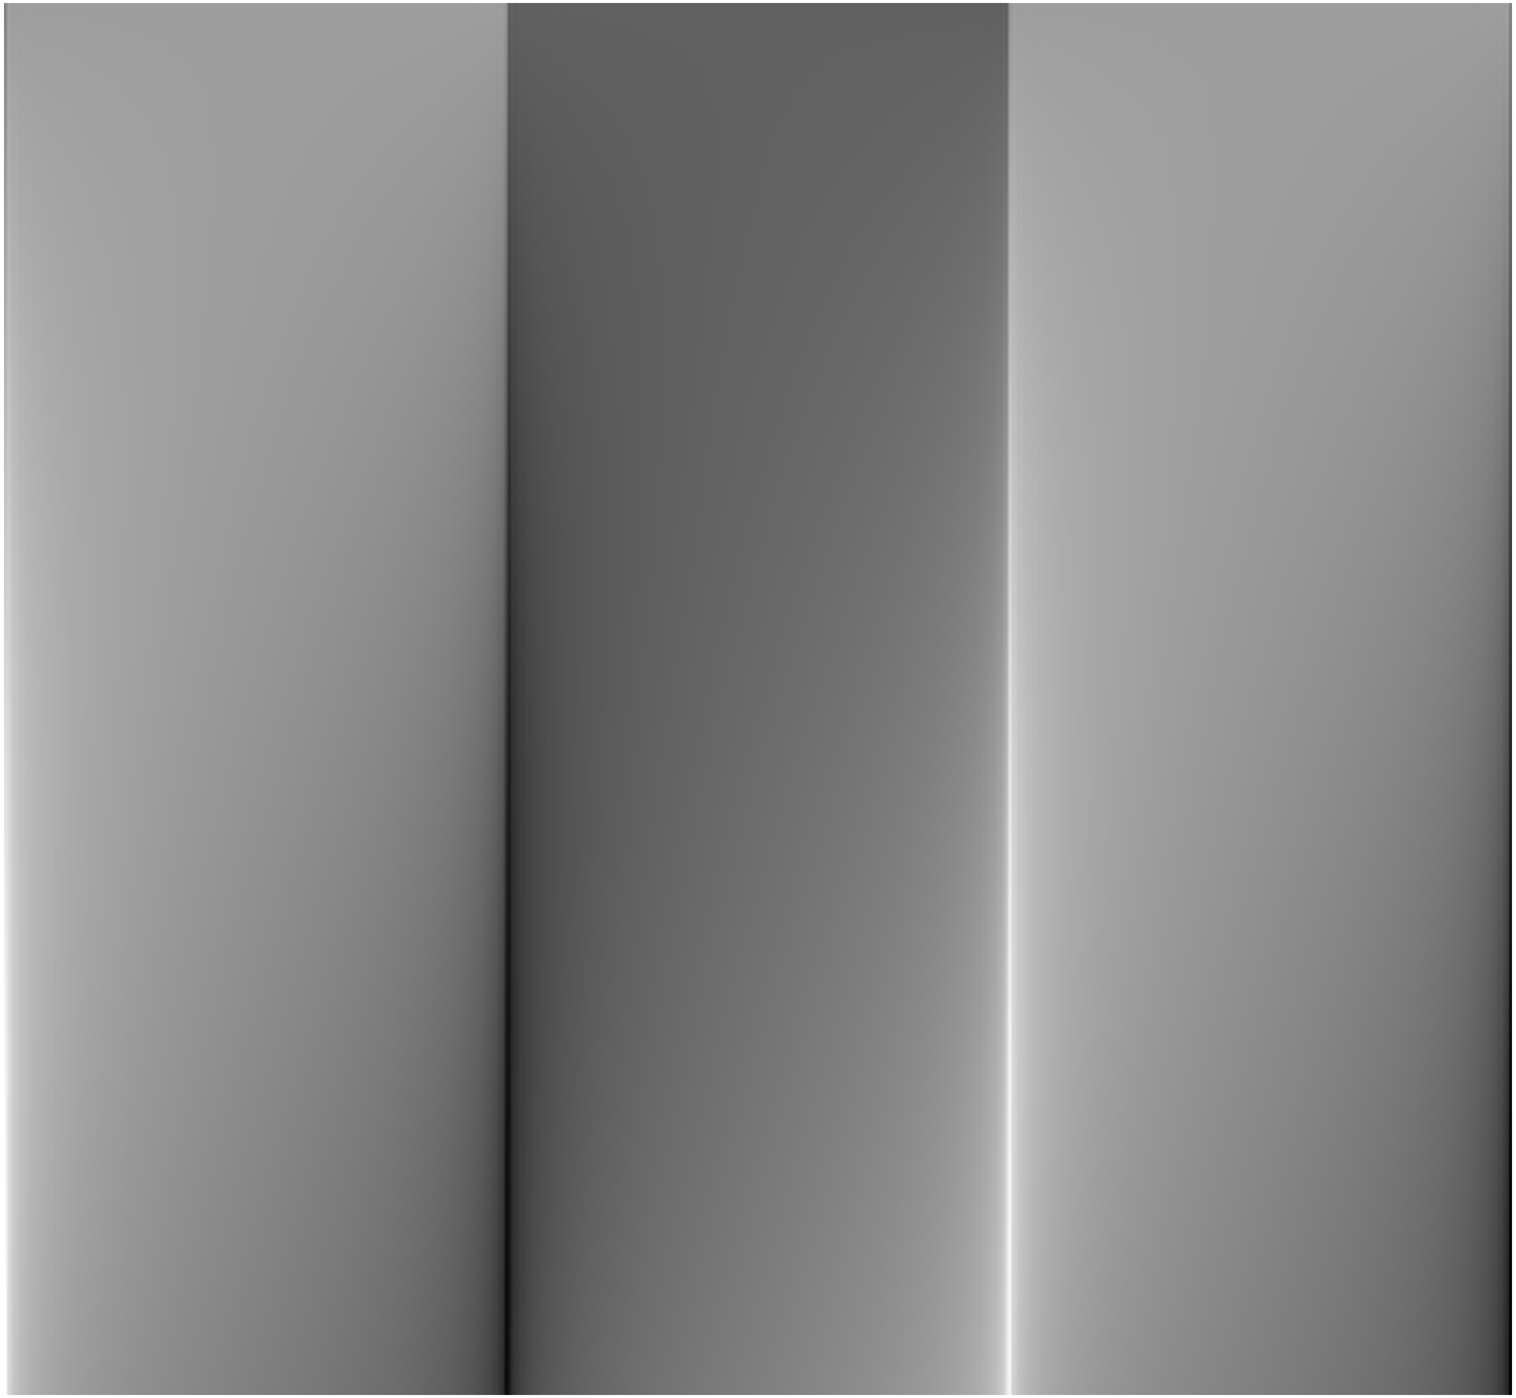

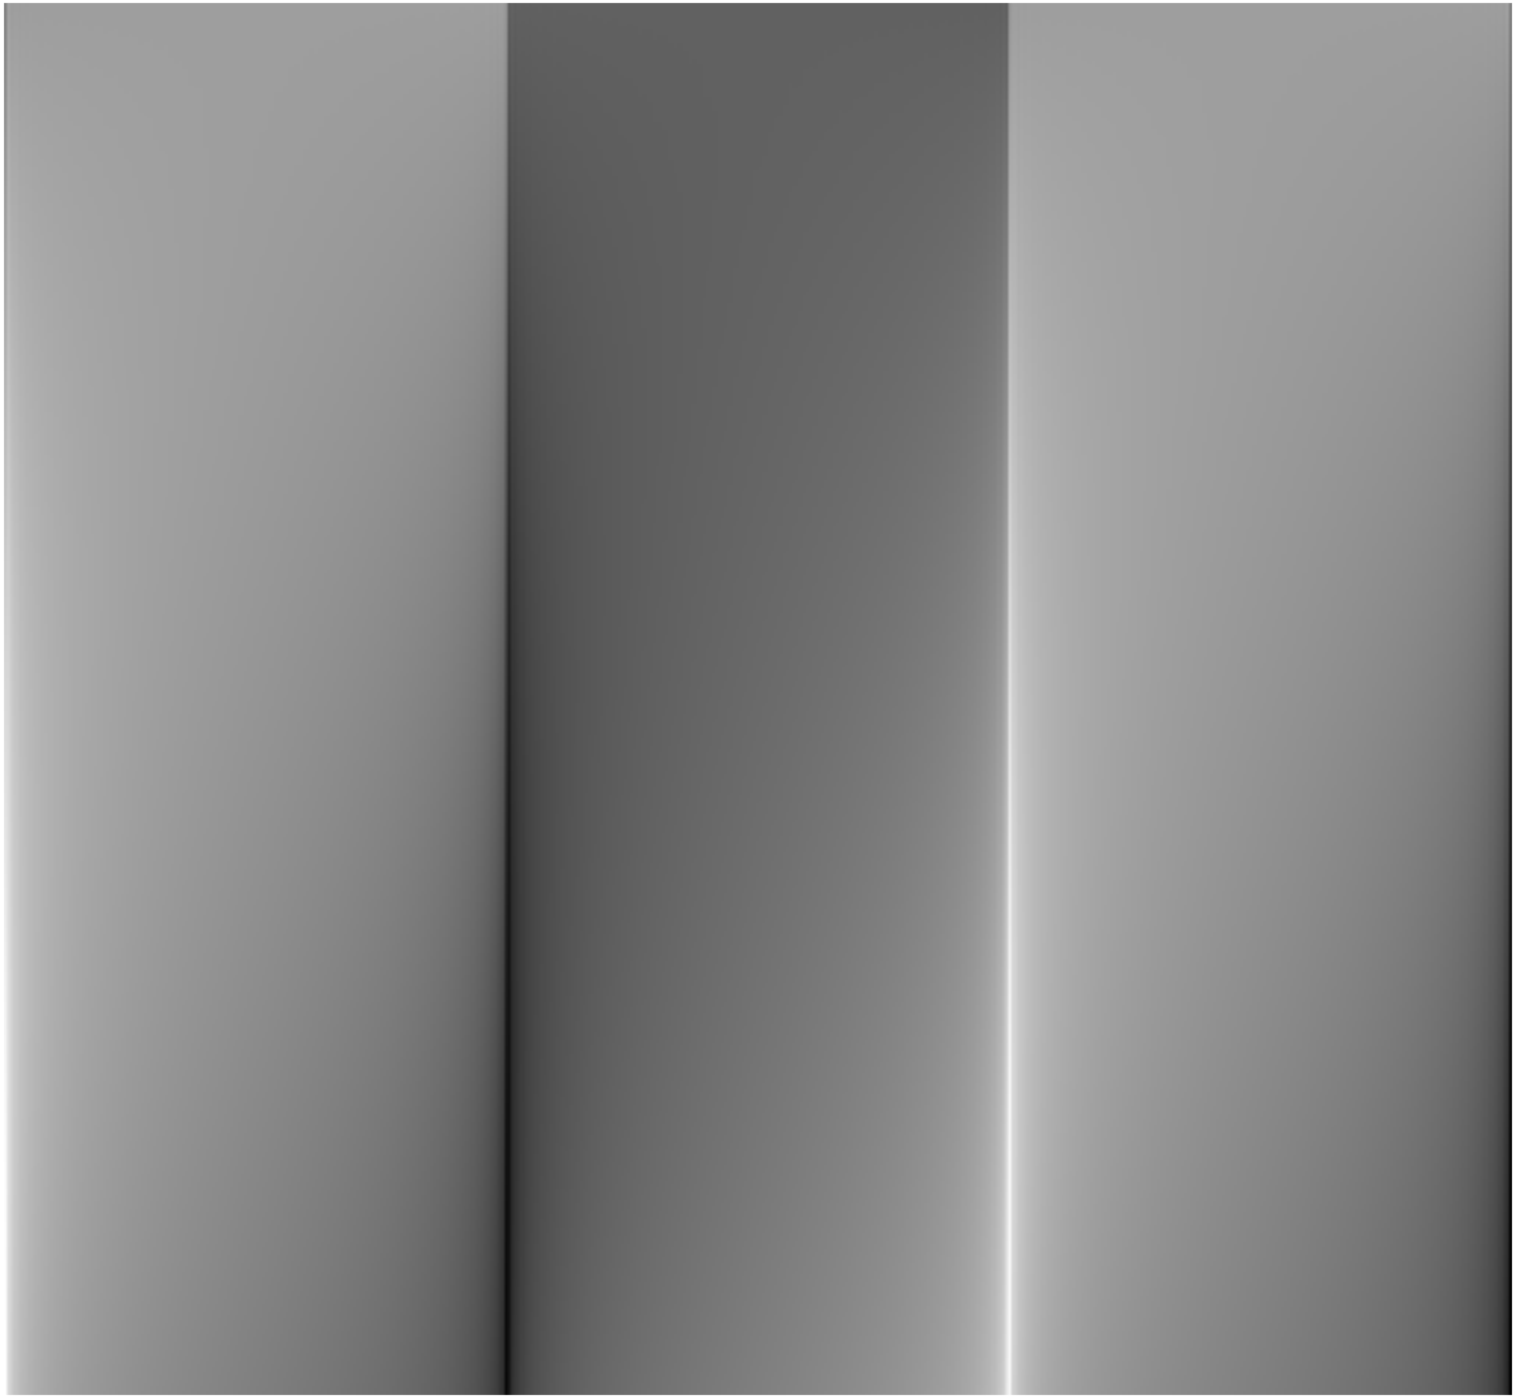

Canny

$\sigma = 1$

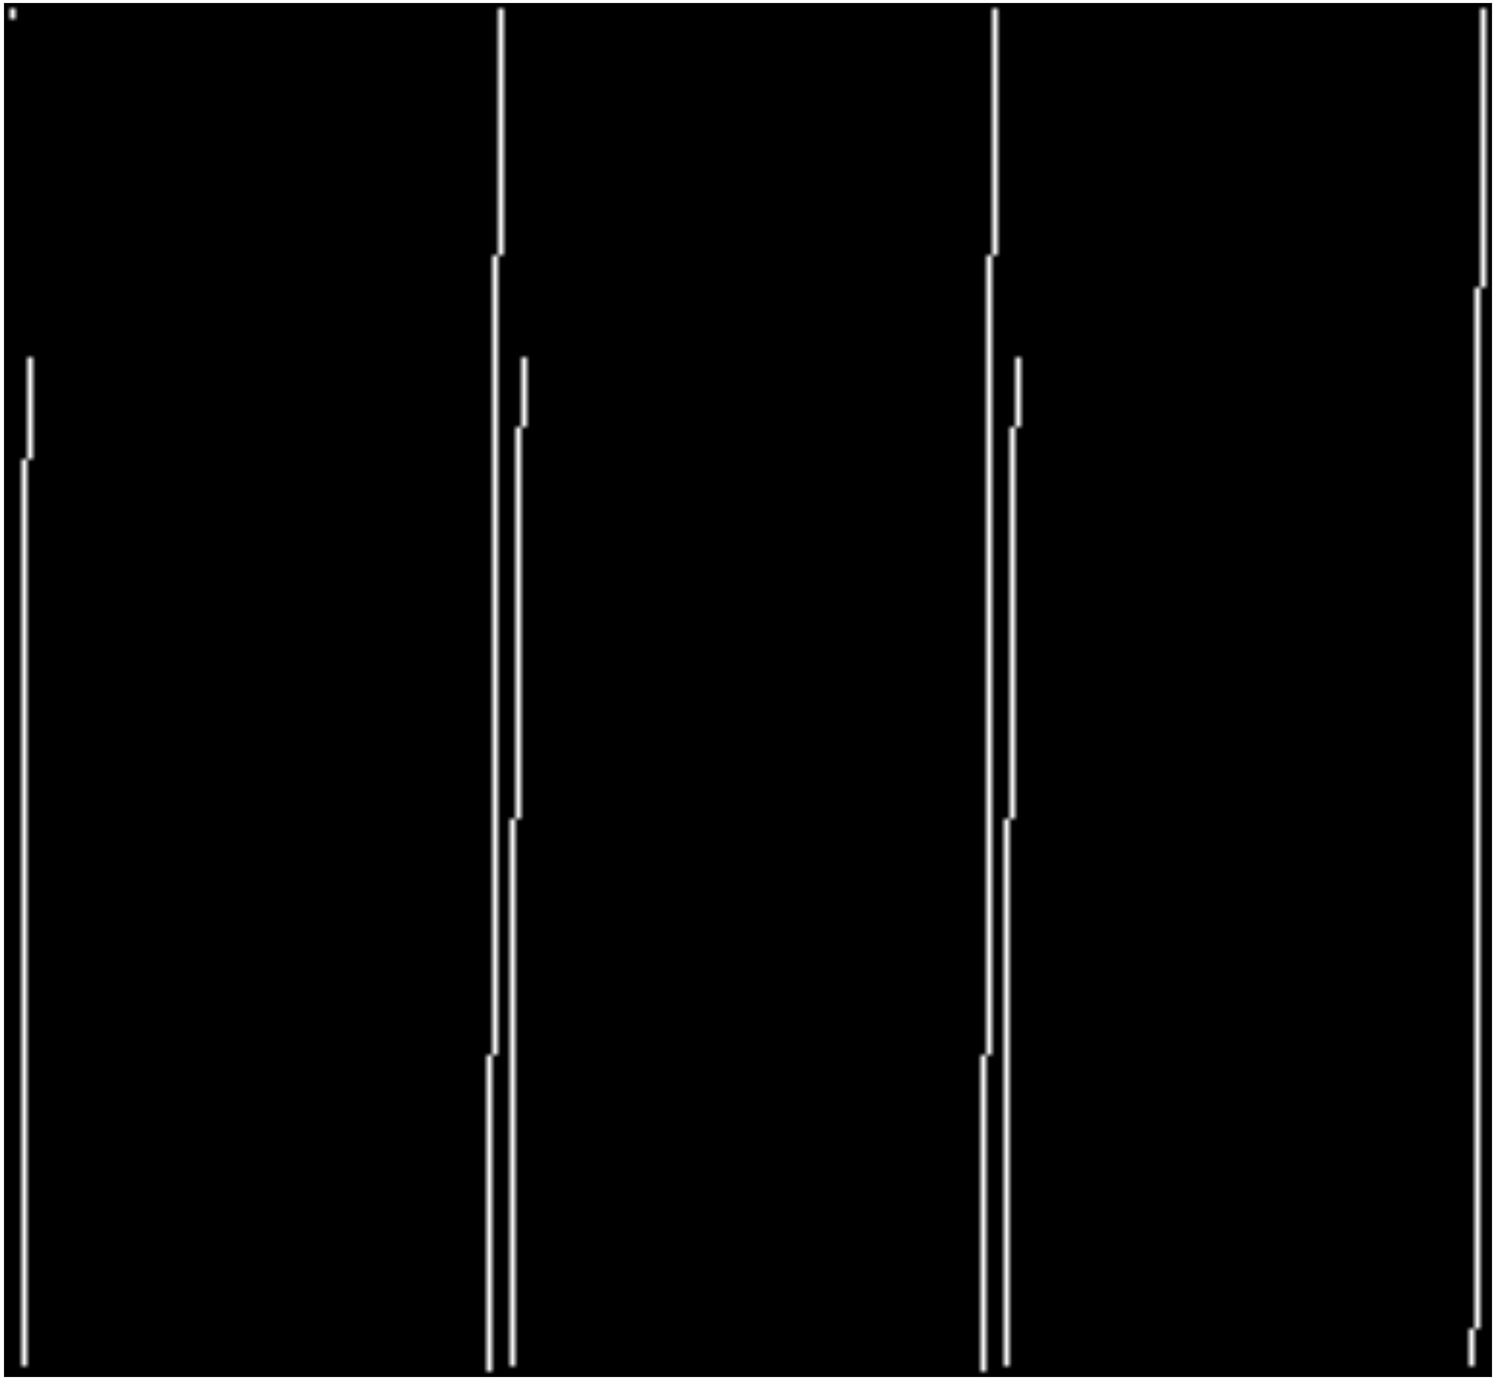

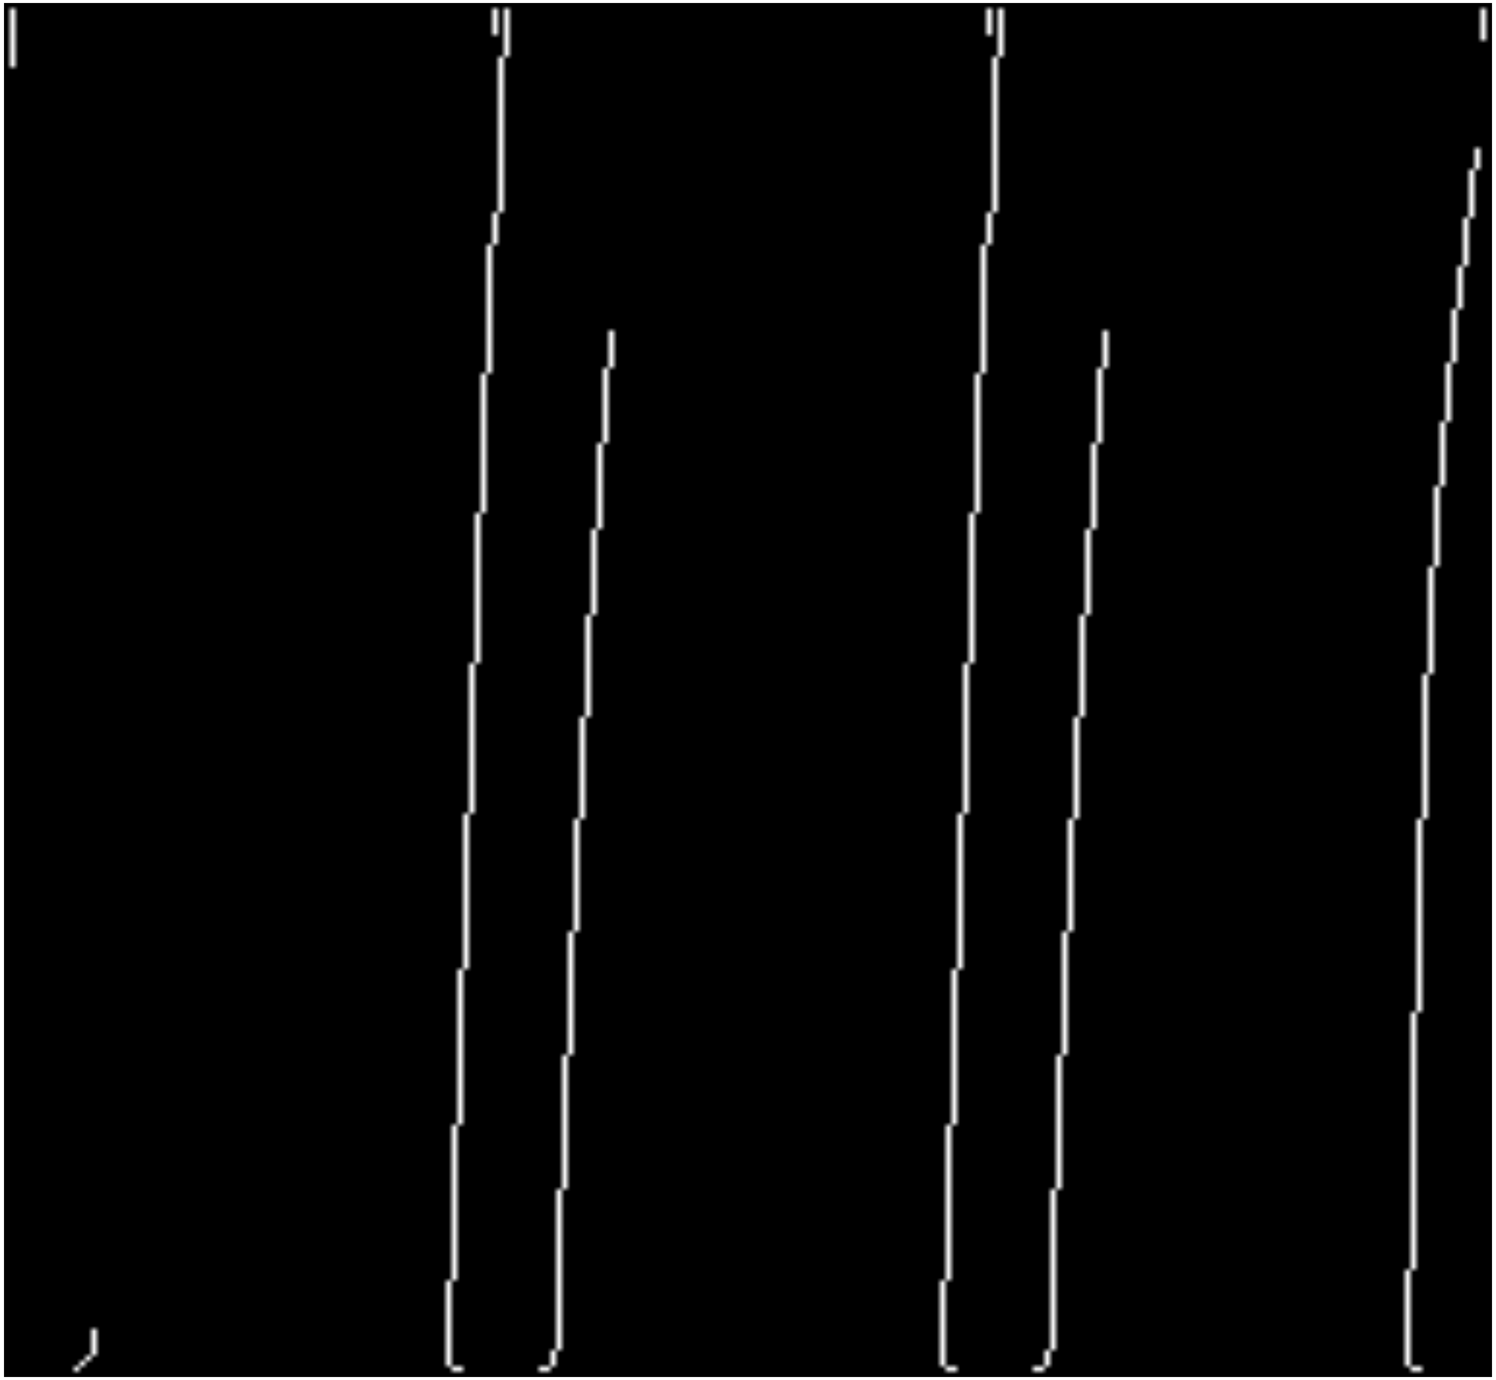

Canny
$\sigma = 5$

Much of computer vision depends on the ability to correctly detect, localize and match local features

But…

- Edges, corners and other features are *not* simple step changes in luminance.
- Gradient based operators do not correctly detect and localize many image features.
- The localization of gradient based features varies with scale of analysis.
- Thresholds are sensitive to image illumination and contrast.
- Despite the obvious importance of edges we do not really know how to *use* them.

Corners

Advances in reconstruction algorithms have renewed our interest in, and reliance on, the detection of feature points or ‘corners’.

Pollefeys’ castle

What's the Problem?

- Current corner detectors do not work reliably with images of varying lighting and contrast.

- Localization of features can be inaccurate and depends on the scale of analysis.

Harris Operator

Form the gradient covariance matrix

$$G = \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$

where I_x and I_y are image gradients in x and y

A corner occurs when eigenvalues are similar and large.
Corner strength is given by

$$R = \det(G) - k(tr(G))^2 \qquad (k = 0.04)$$

Note that R has units (intensity gradient)^4

The Harris operator is very sensitive to local contrast

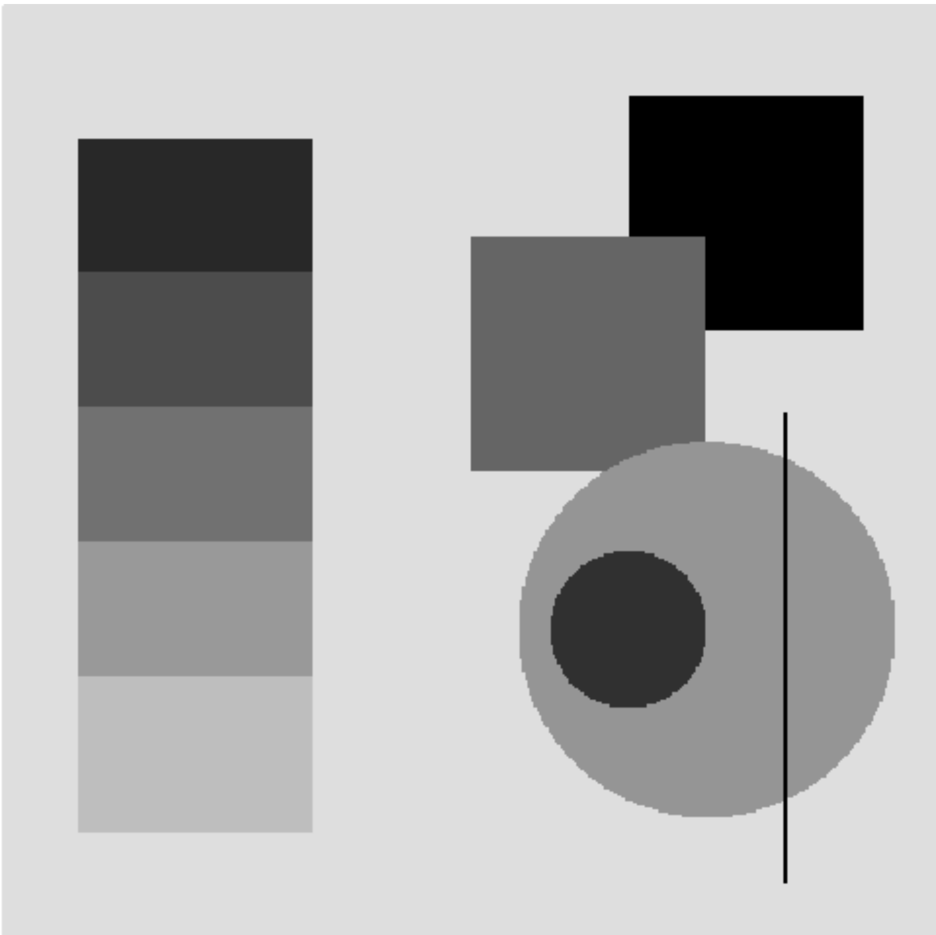

Test image

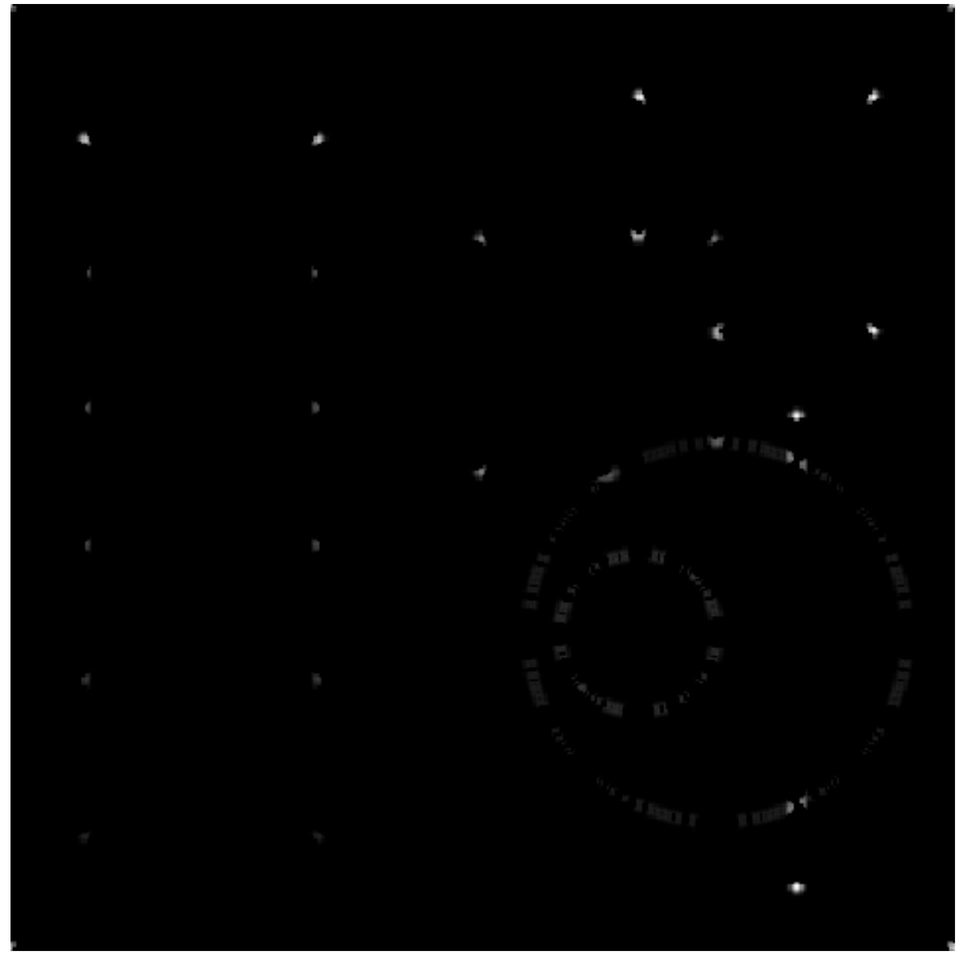

Fourth root of Harris corner strength
(Max value is ~1.5×10^{10})

The location of Harris corners is sensitive to scale

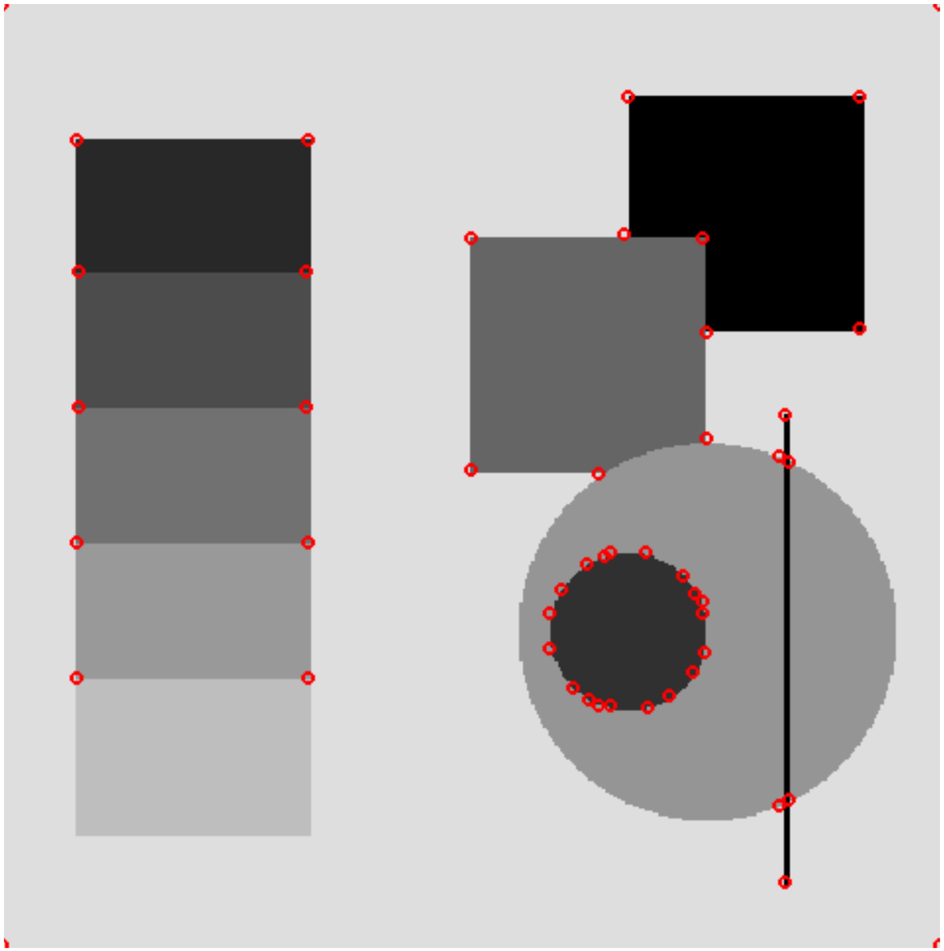

Harris corners, $\sigma = 1$

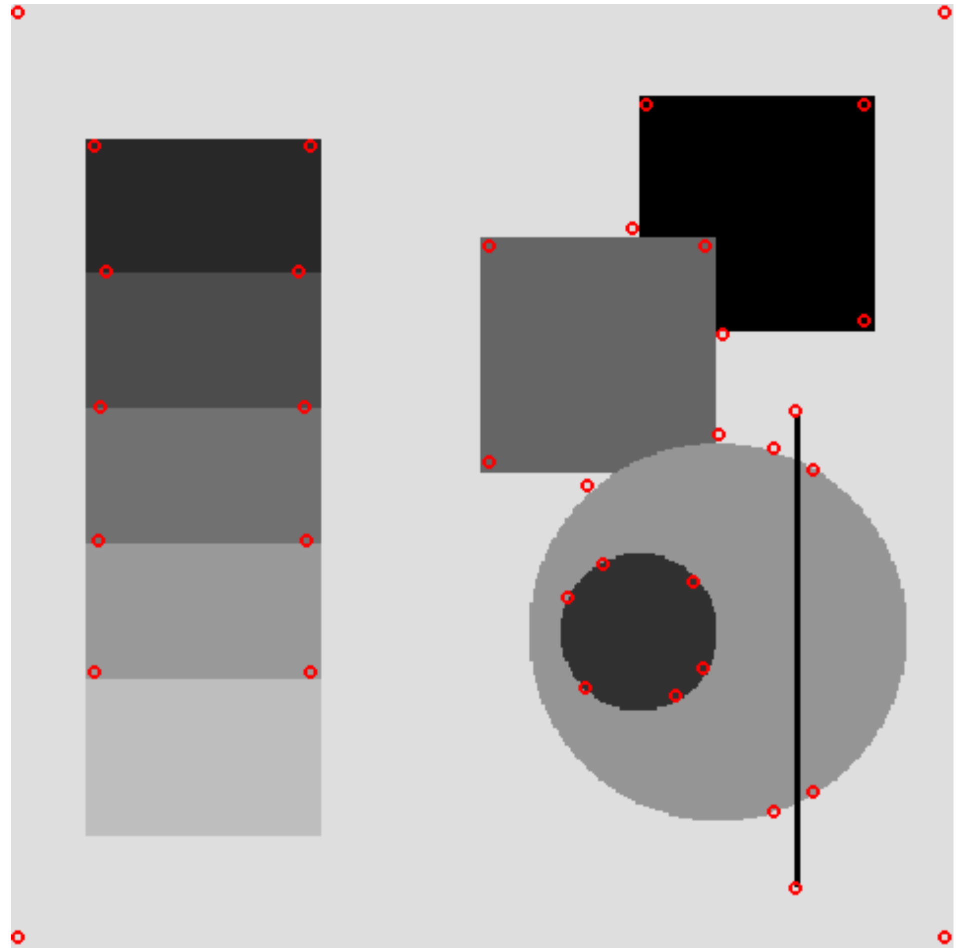

Harris corners, $\sigma = 7$

Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.

As a consequence, successful 3D reconstruction from matched corner points requires extensive outlier removal and the application of robust estimation techniques.

Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.

As a consequence, successful 3D reconstruction from matched corner points requires extensive outlier removal and the application of robust estimation techniques.

To minimize these problems we need a feature operator that is maximally invariant to illumination and scale.

Features are Perceived at Points of Phase Congruency

- Don't think of features in terms of derivatives!
- Think of features in terms of local frequency components.

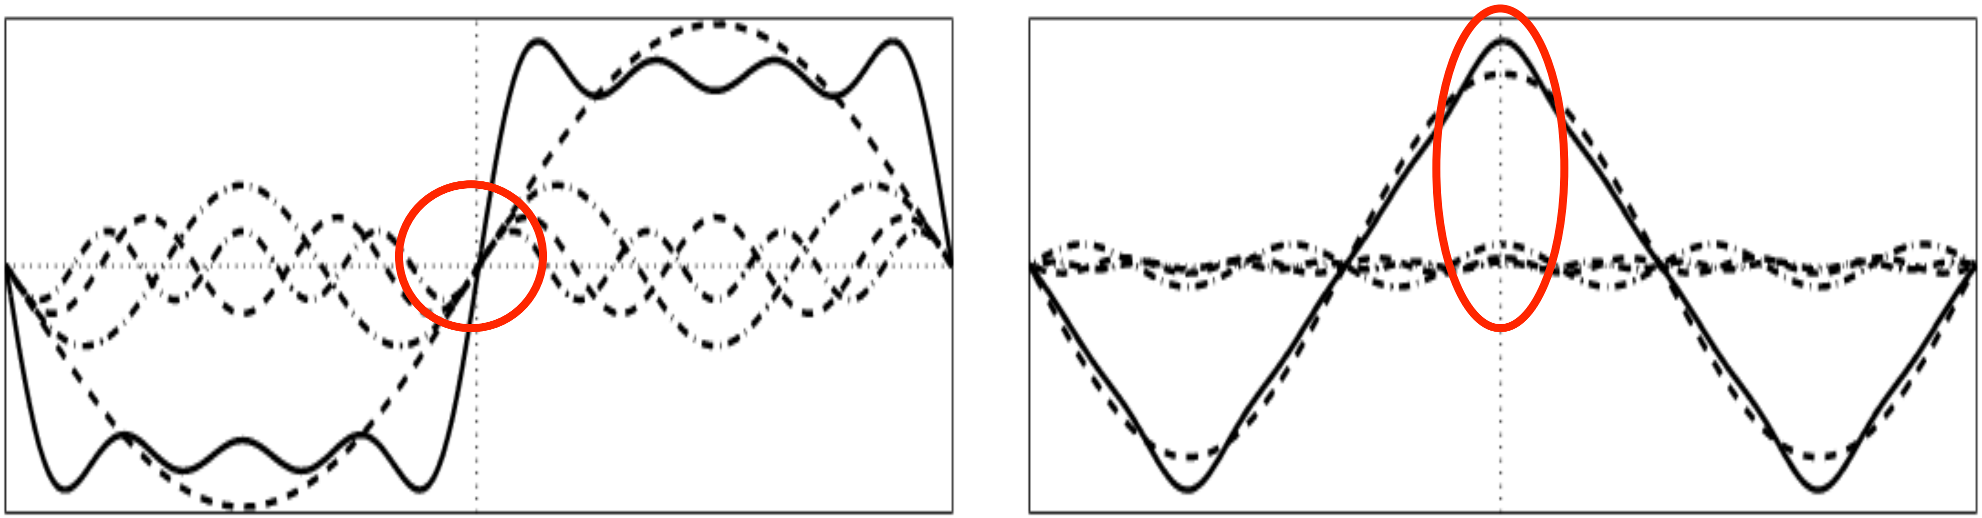

- The Fourier components are all in phase at the point of the step in the square wave, and at the peaks and troughs of the triangular wave.
- This property is stable over scale.

(Morrone & Owens 1987)

Phase is important!

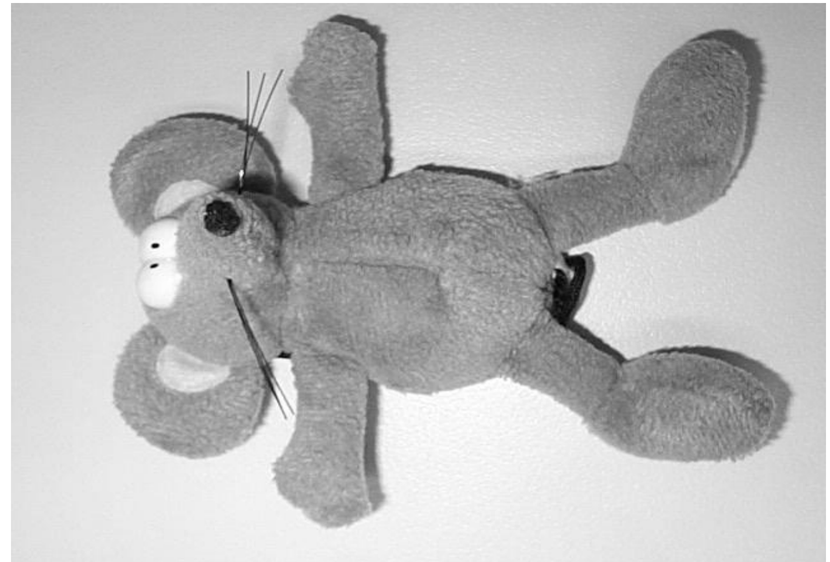

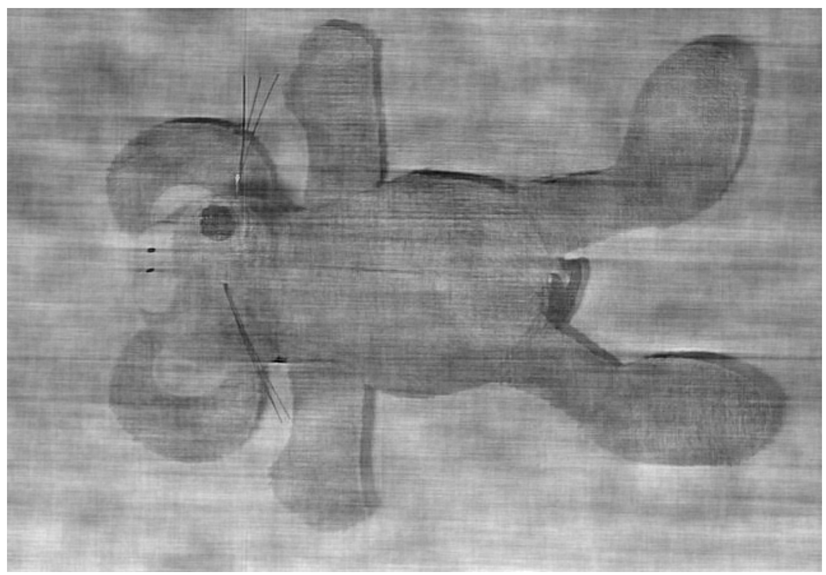

Amplitude of Peter, Phase of Mouse...

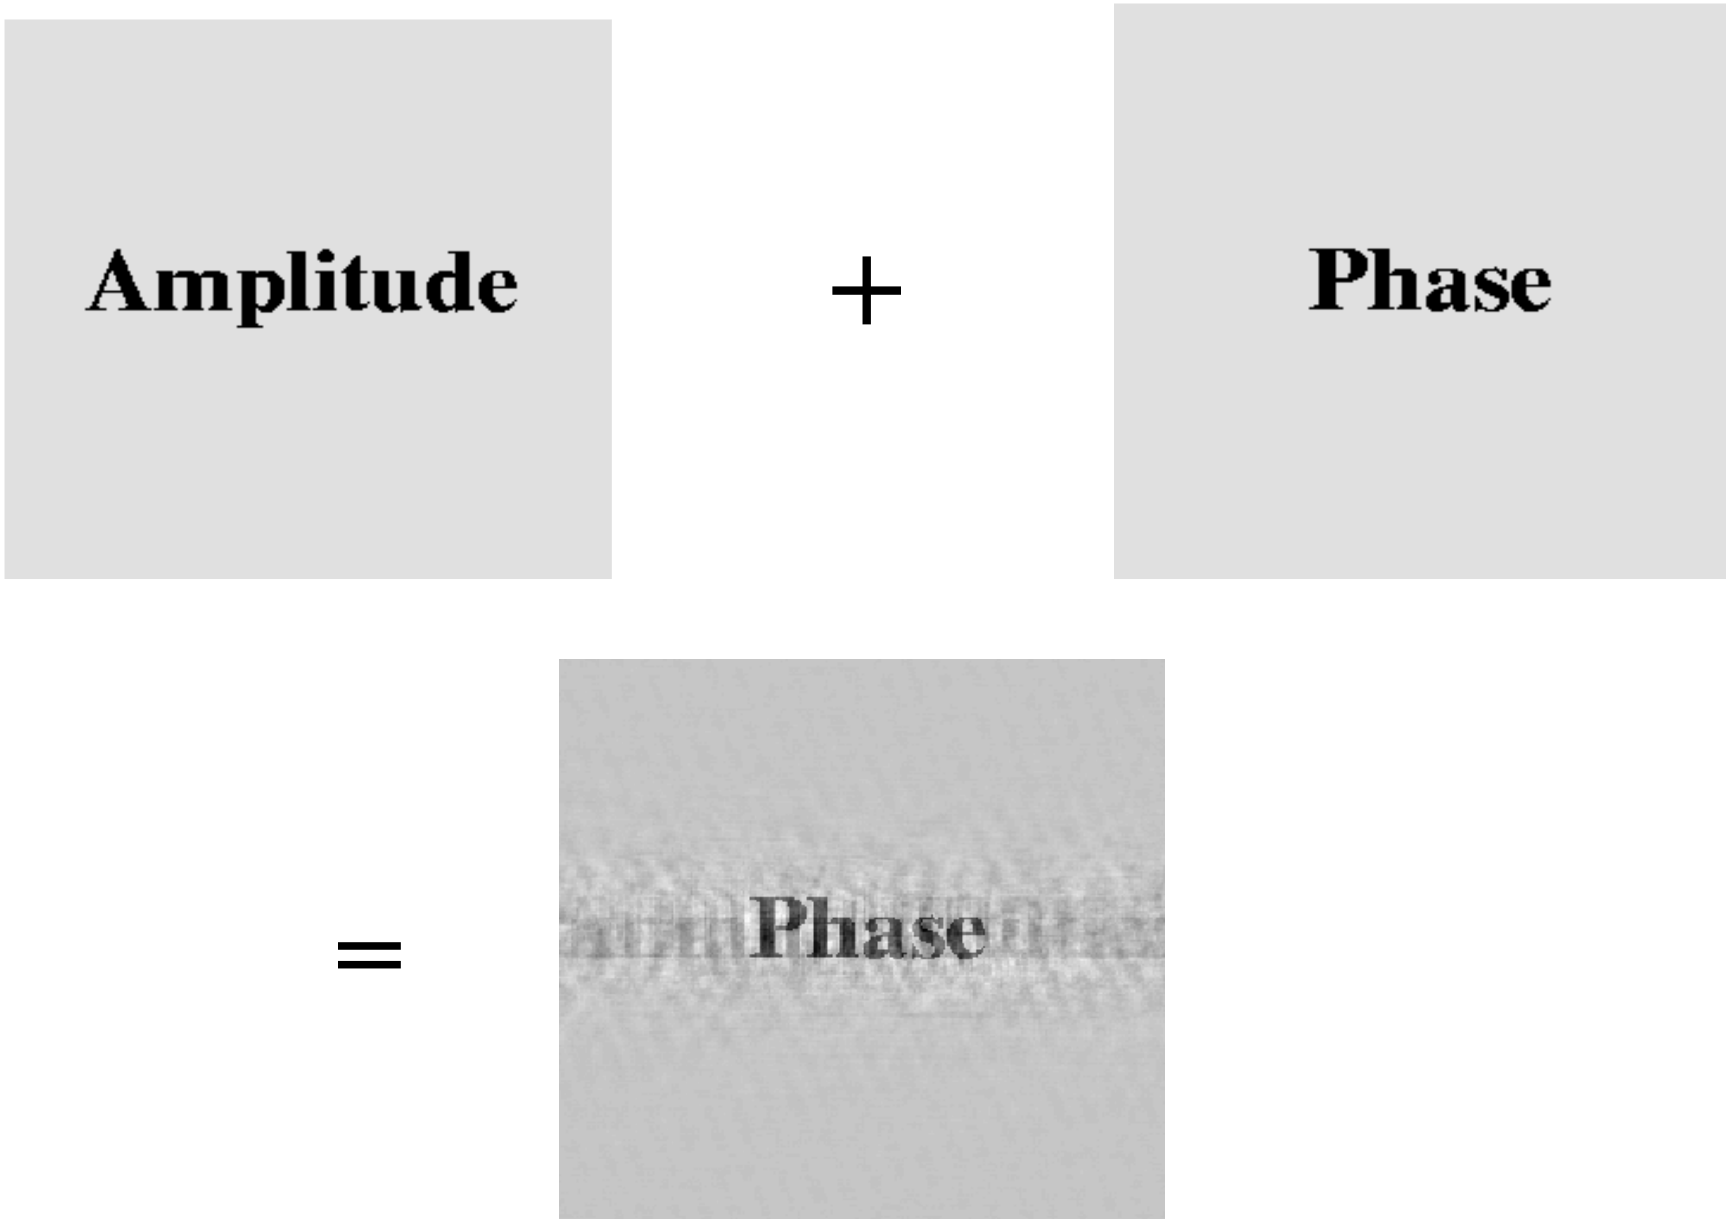

Phase and amplitude mixed image

(Oppenheim and Lim 1981)

Phase Congruency Models a Wide Range of Feature Types

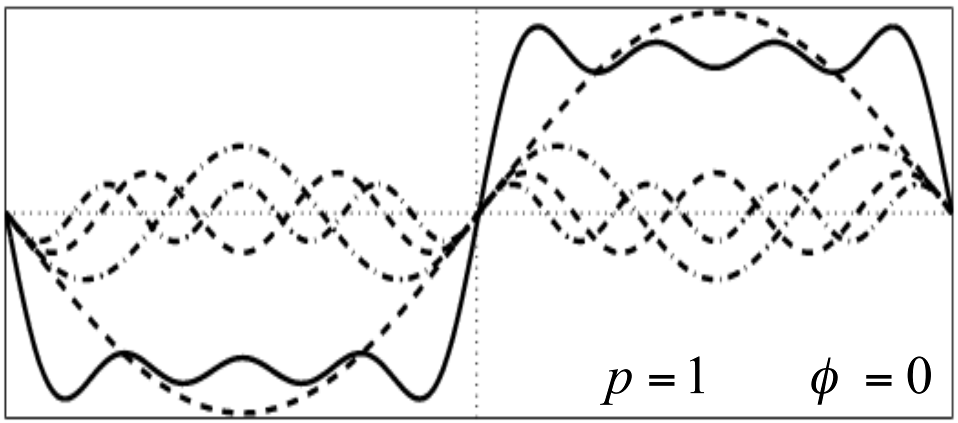

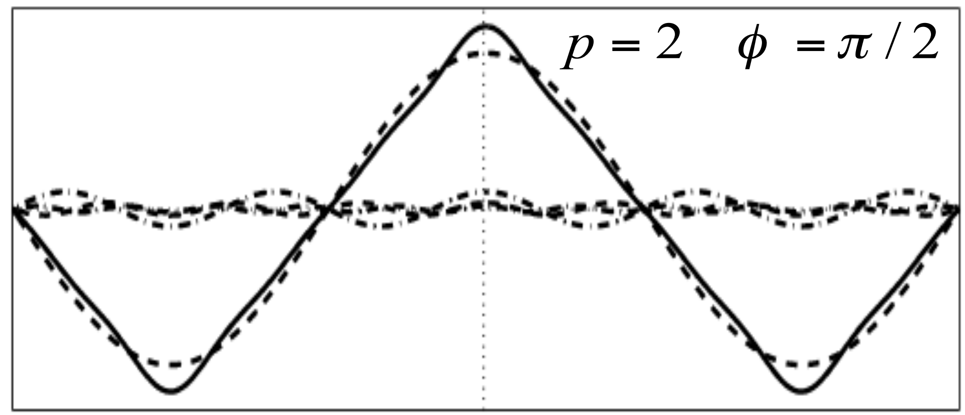

$$f(x) = \sum_{n=0}^{\infty} \frac{1}{(2n+1)^p} \sin[(2n+1)\, x + \phi\ ]$$

phase offset

amplitude decay with frequency

A continuum of feature types from step to line can be obtained by varying the phase offset. Sharpness is controlled by amplitude decay.

Congruency of Phase at any Angle Produces a Feature

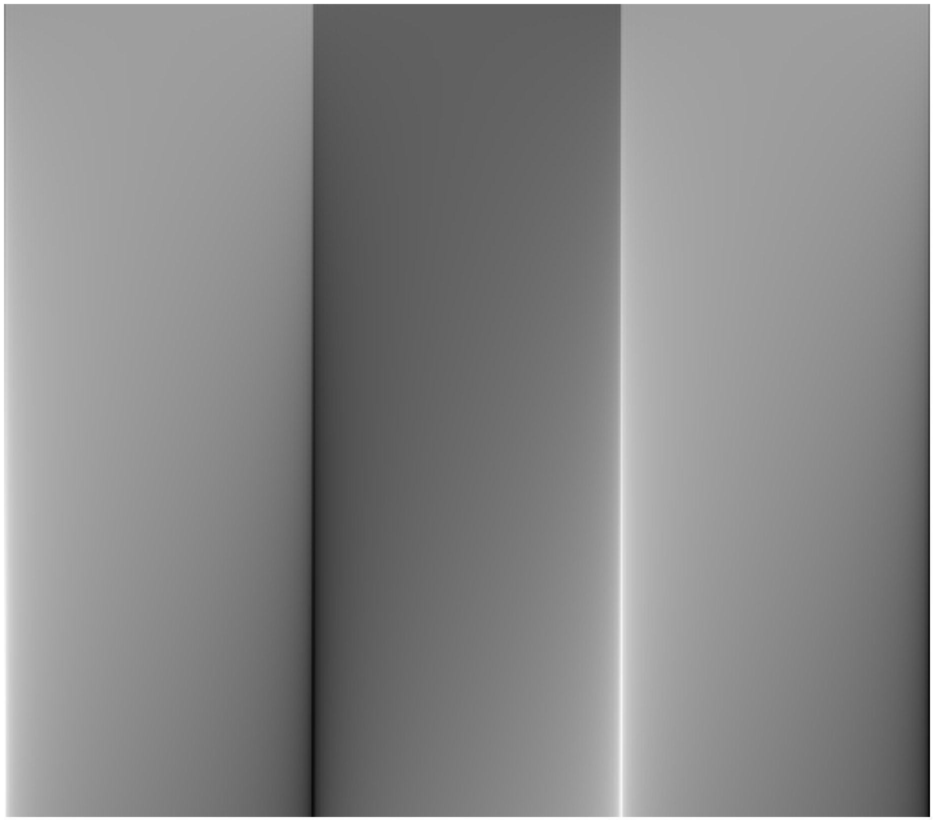

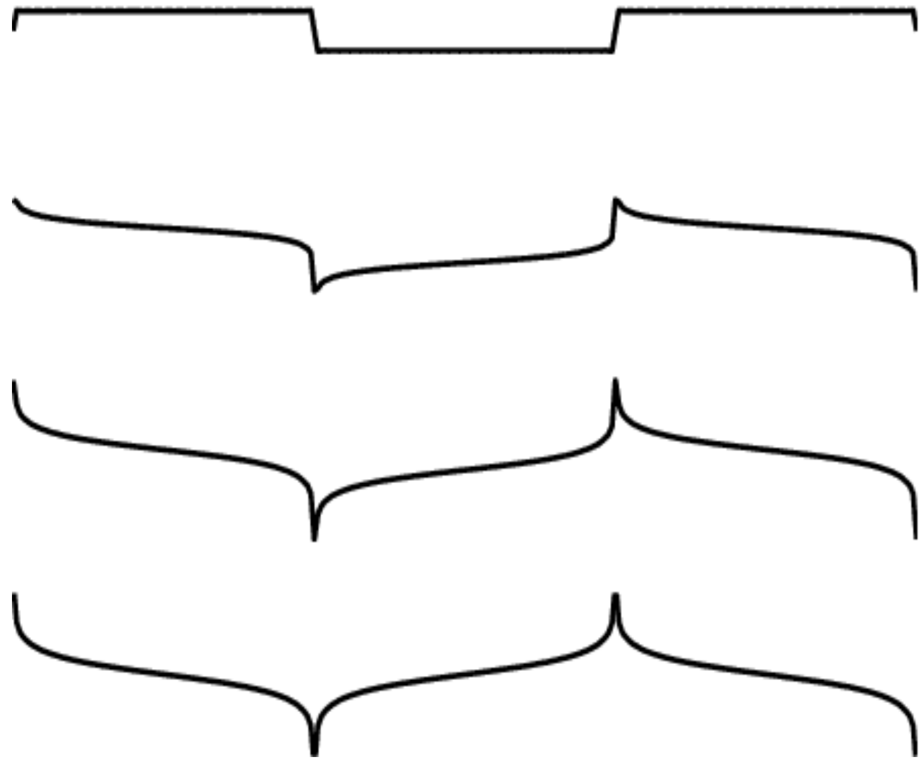

Interpolation of a step to a line by varying ϕ from 0 at the top to $\pi / 2$ at the bottom

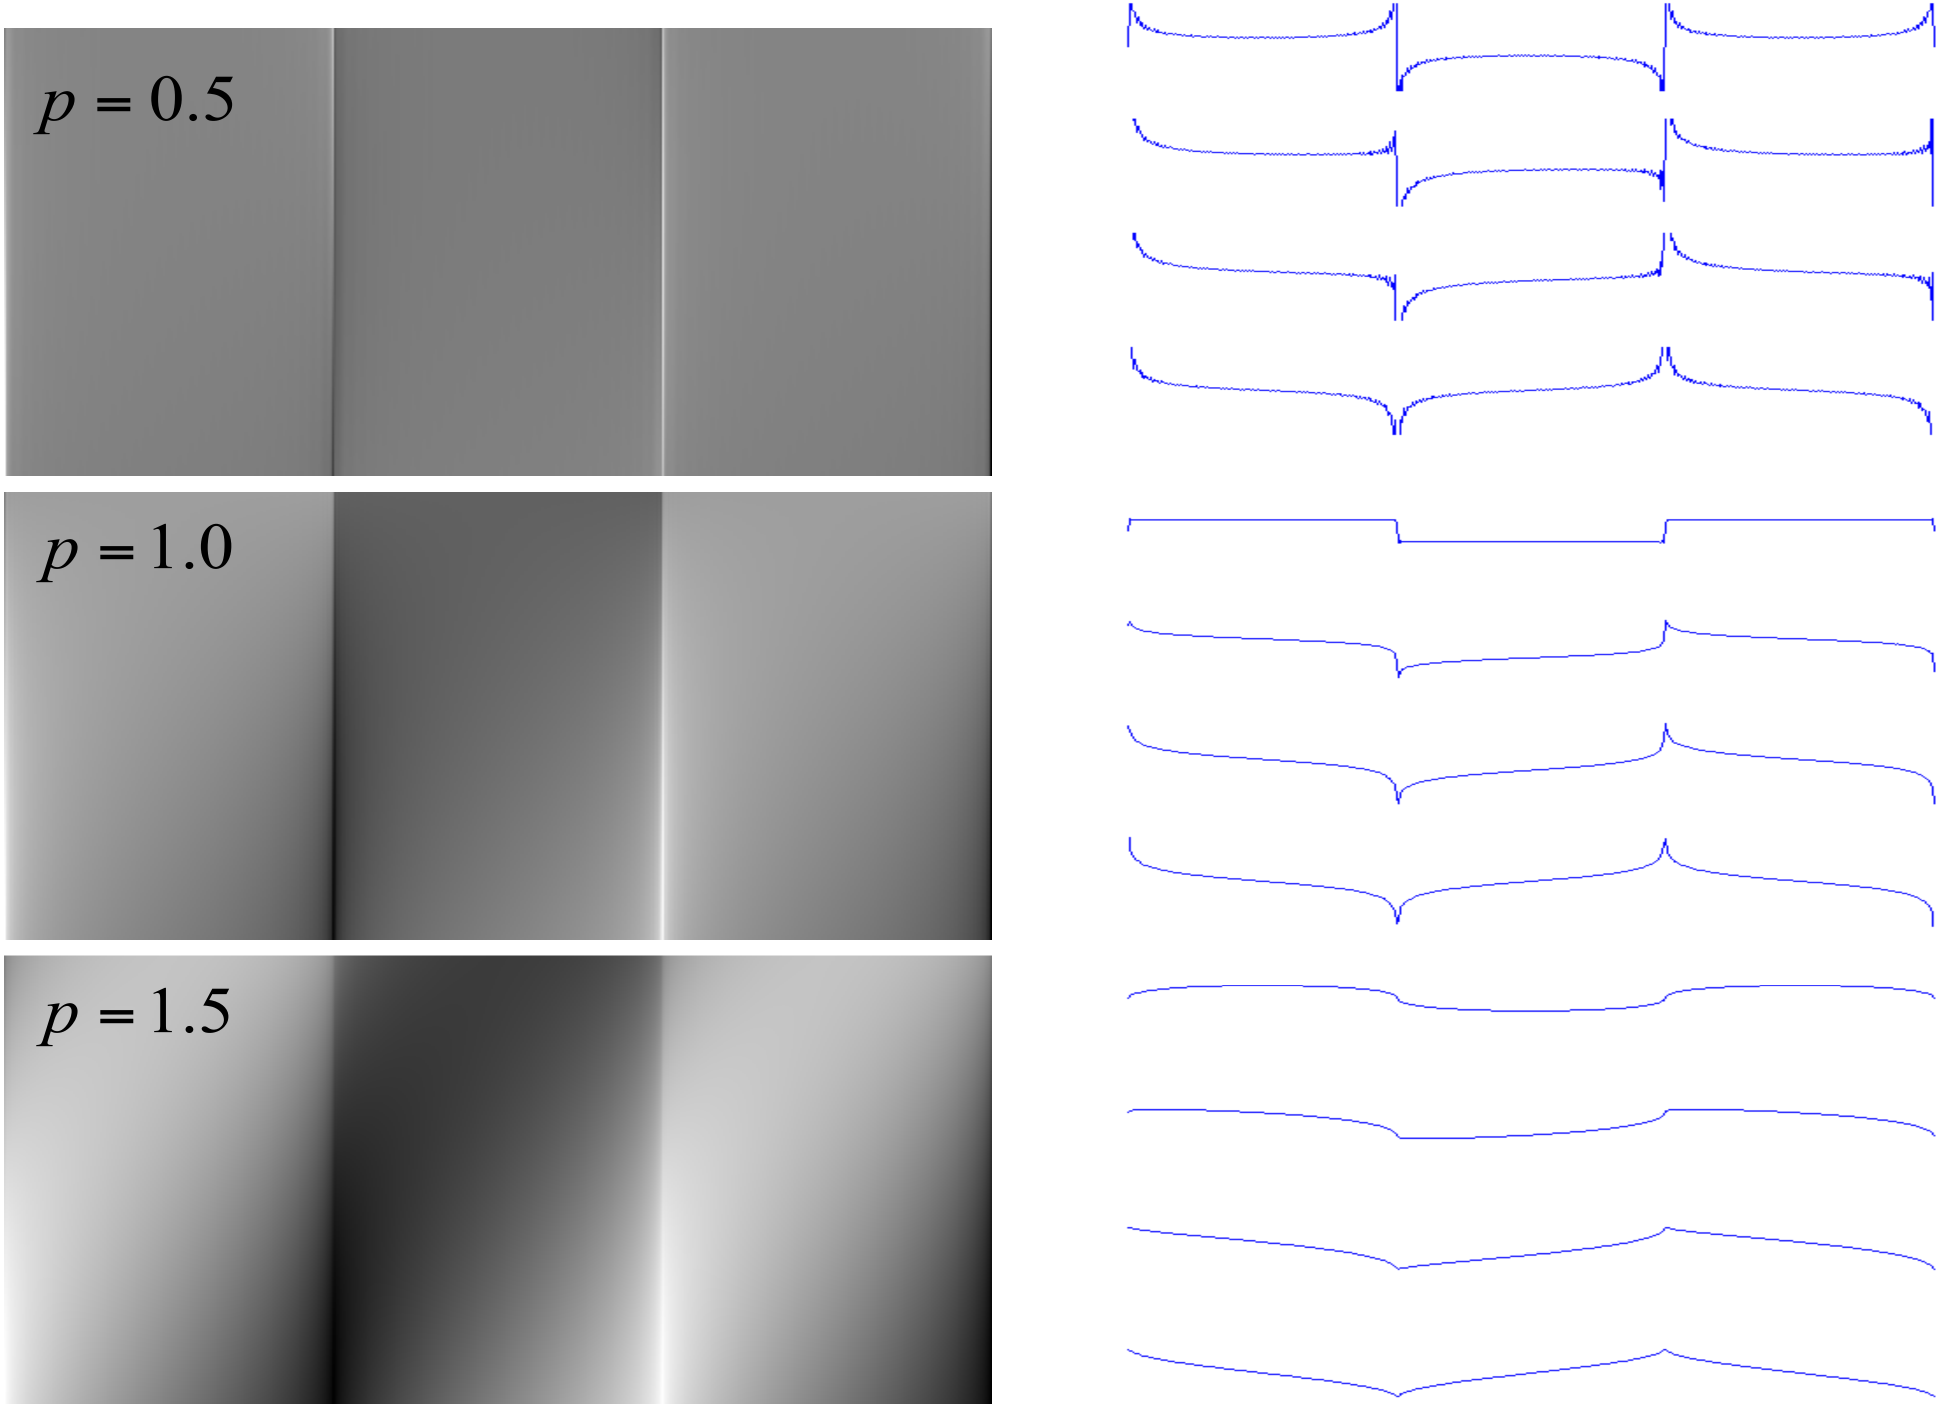

$p = 0.5$
$p = 1.0$
$p = 1.5$

Measuring Phase Congruency

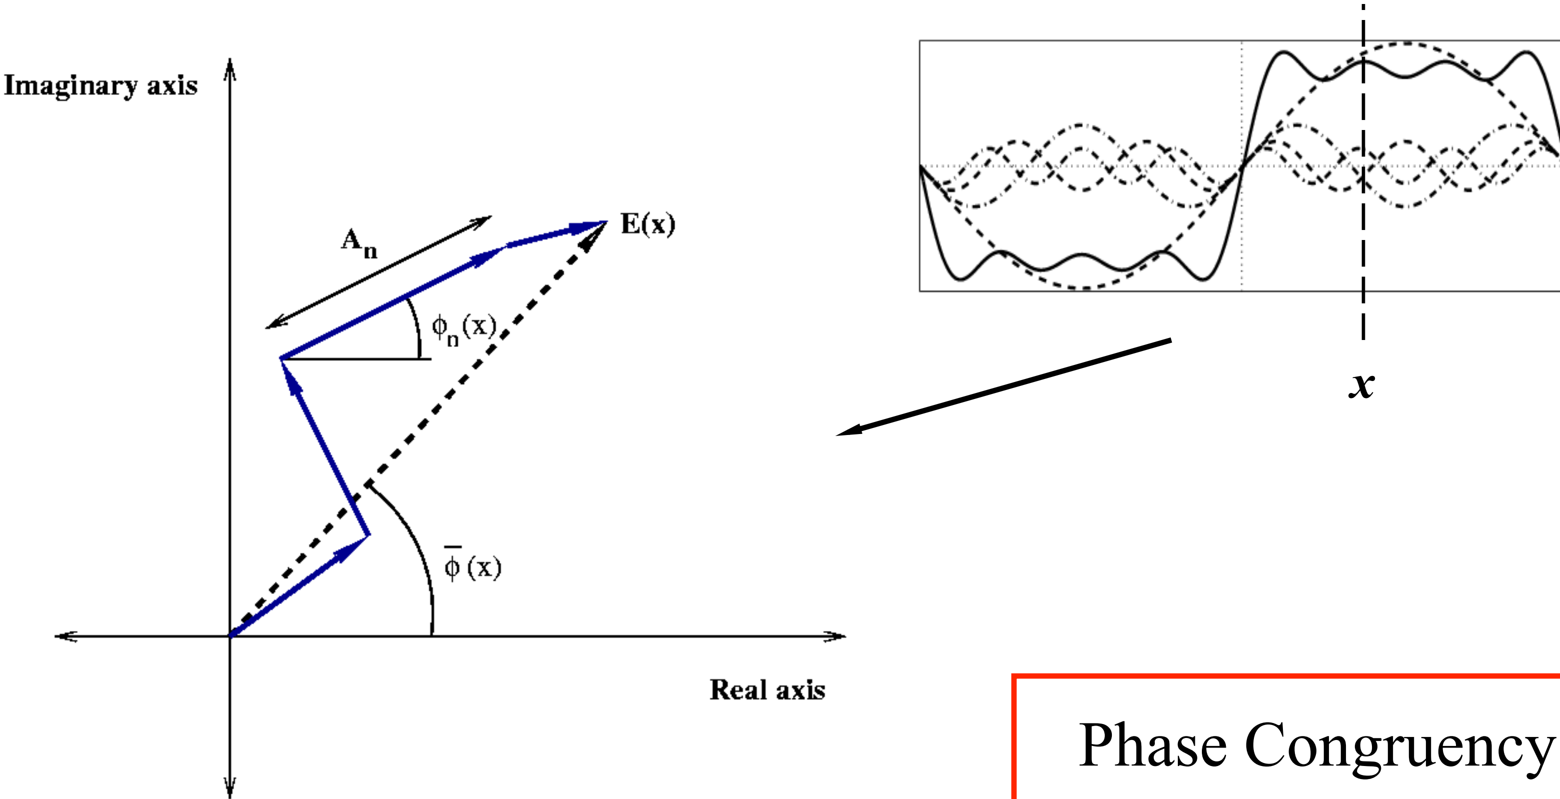

Polar diagram of local Fourier components at a location x plotted head to tail:

Amplitude $A_n(x)$
Phase $\phi_n(x)$
Local Energy $E(x)$

Phase Congruency is the ratio

$$PC(x) = \frac{|E(x)|}{\sum_n A_n(x)}$$

$$0 \le PC(x) \le 1$$

Failure of a gradient operator on a simple synthetic image…

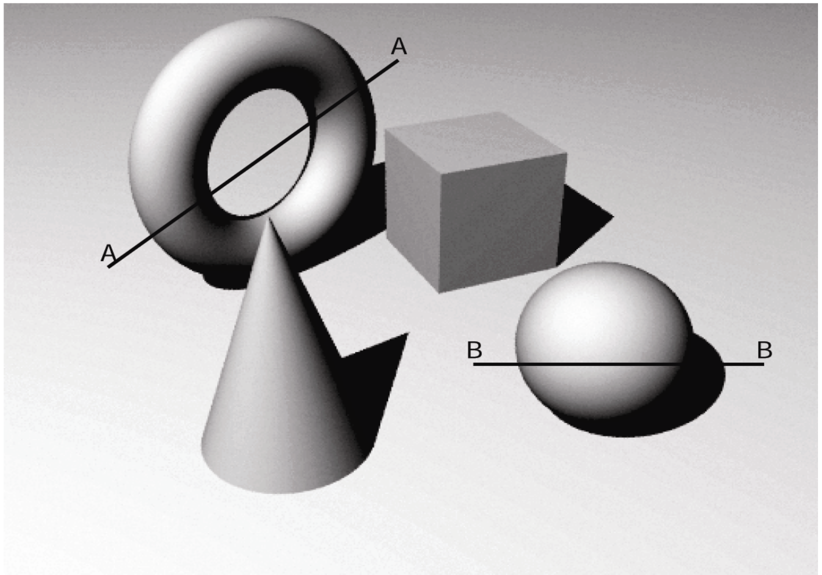

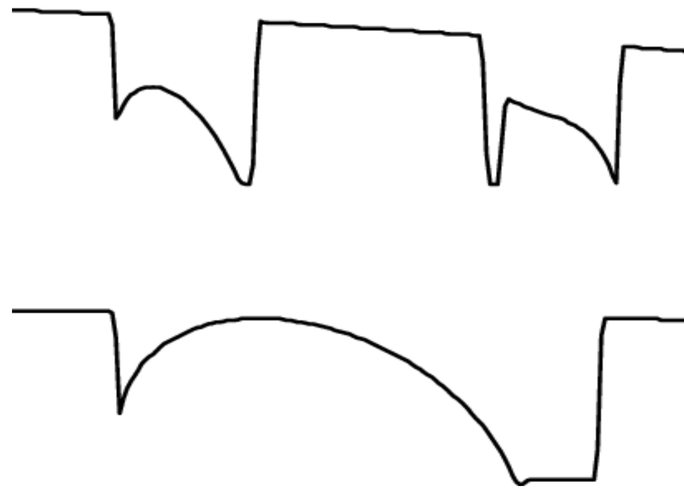

Section A-A

Section B-B

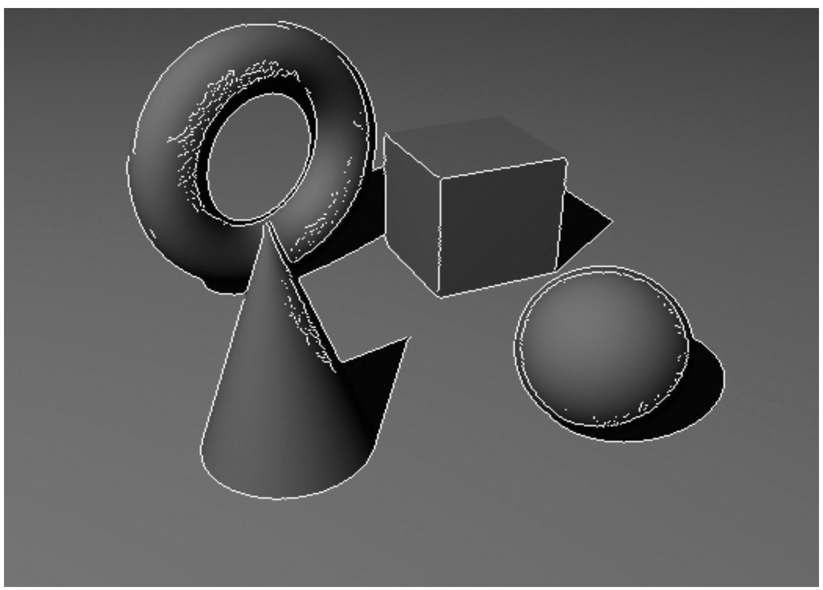

Canny edge map

A gradient based operator merely marks points of maximum gradient.

Note the doubled response around the sphere and the confused torus boundary.

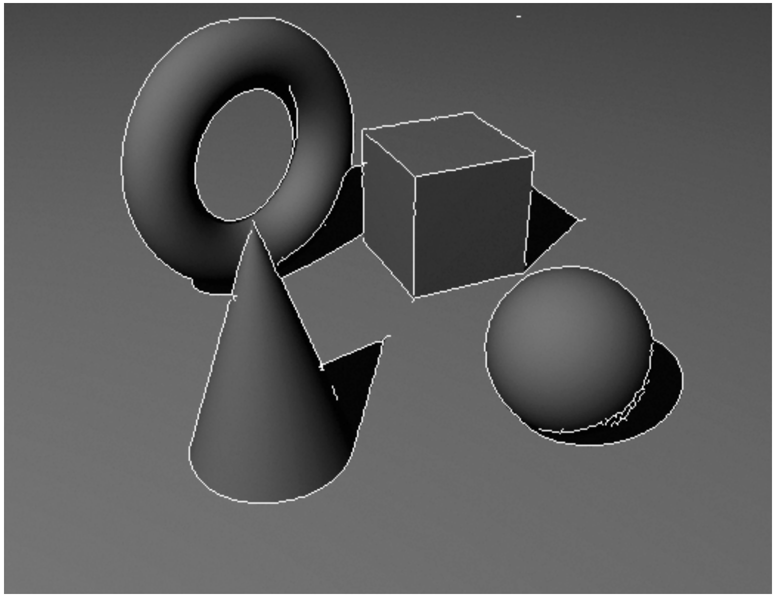

Phase Congruency edge map

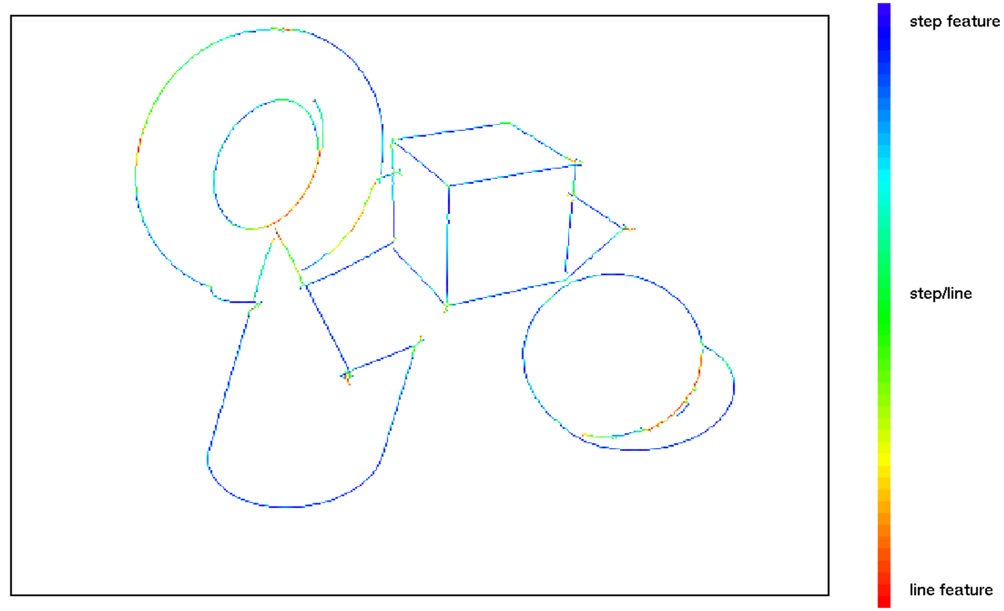

Feature classifications

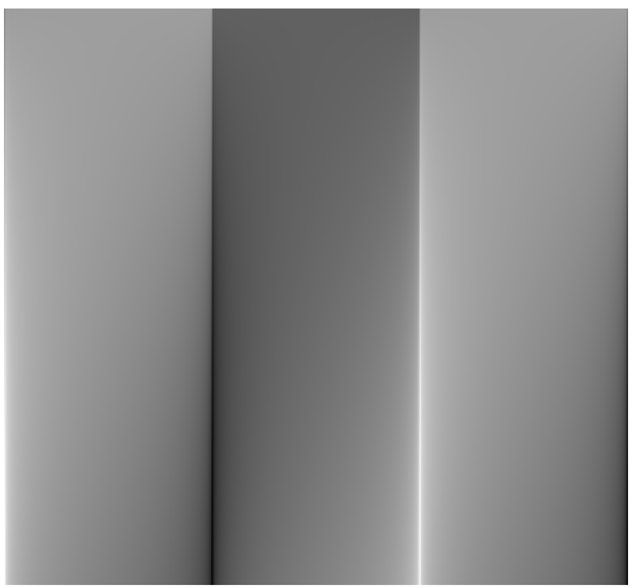

Test grating

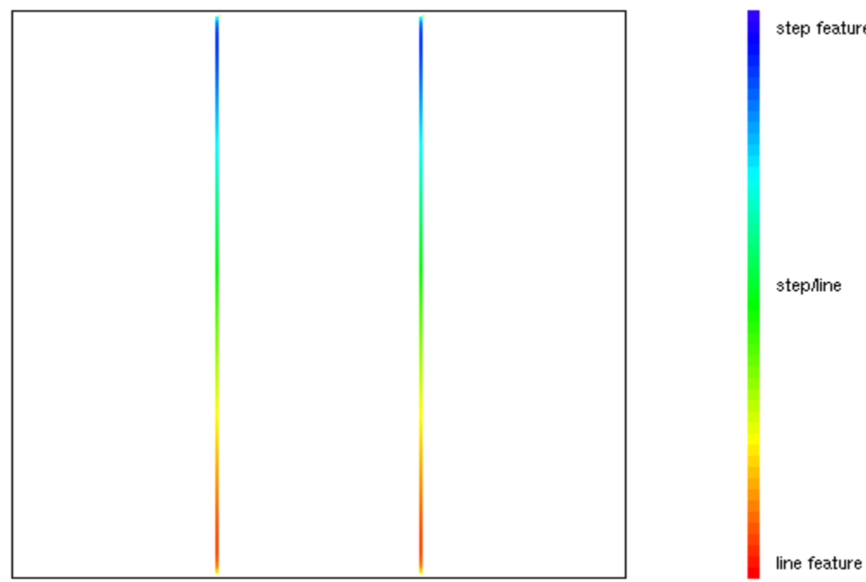

Feature classifications

Implementation

Calculate local frequency information by convolving the image with banks of quadrature pairs of log-Gabor wavelets (Field 1987).

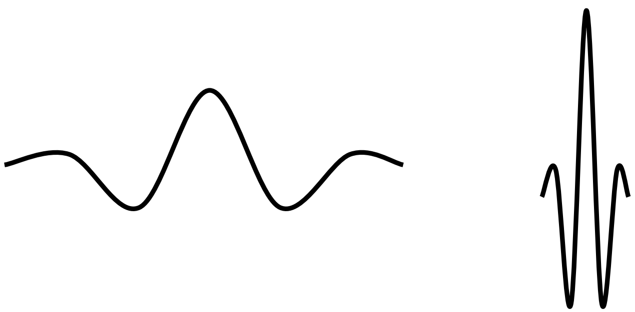

even-symmetric wavelets
(real-valued)

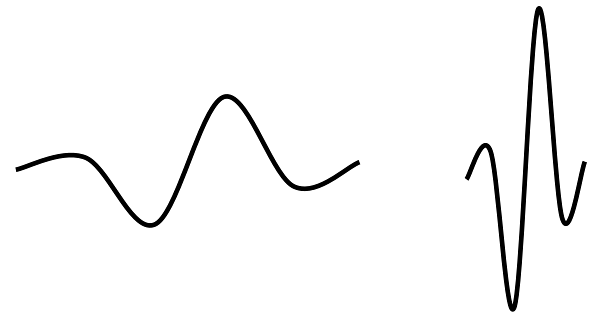

odd-symmetric wavelets
(imaginary-valued)

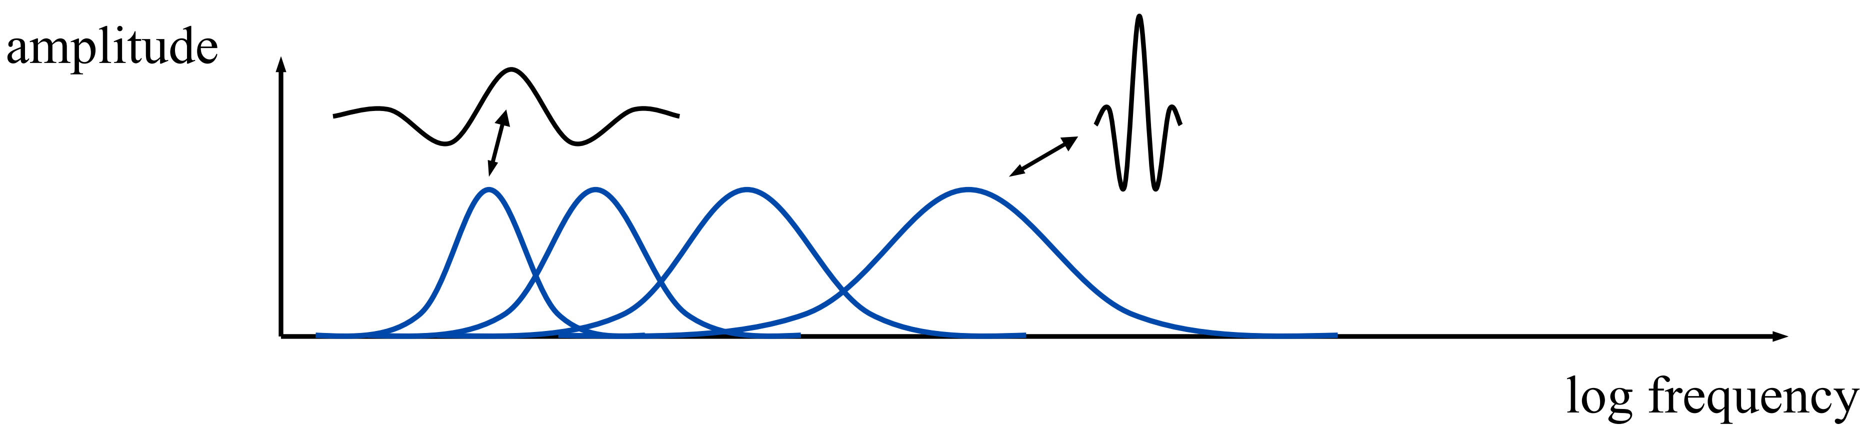

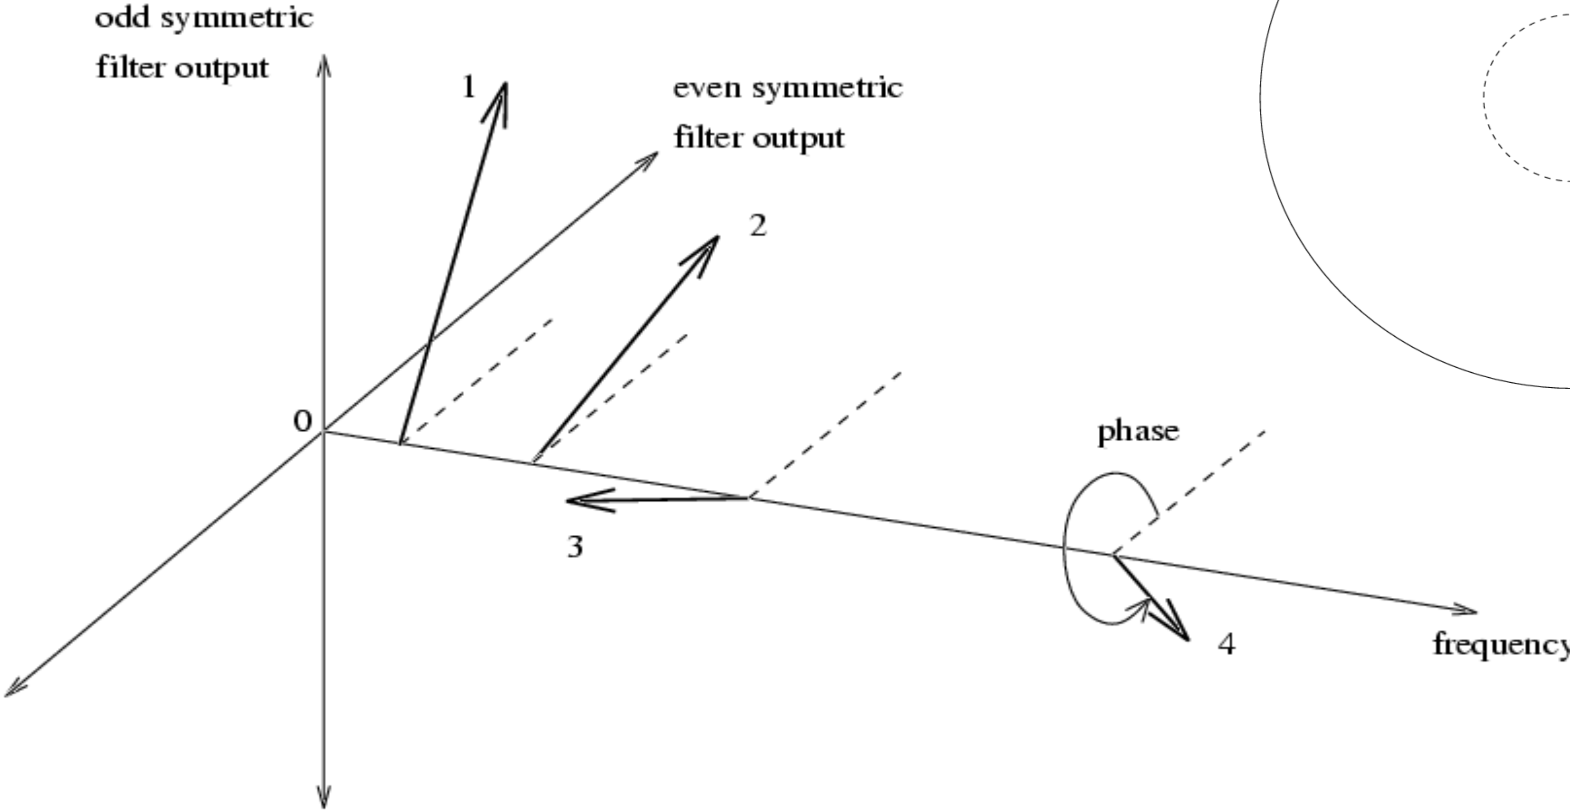

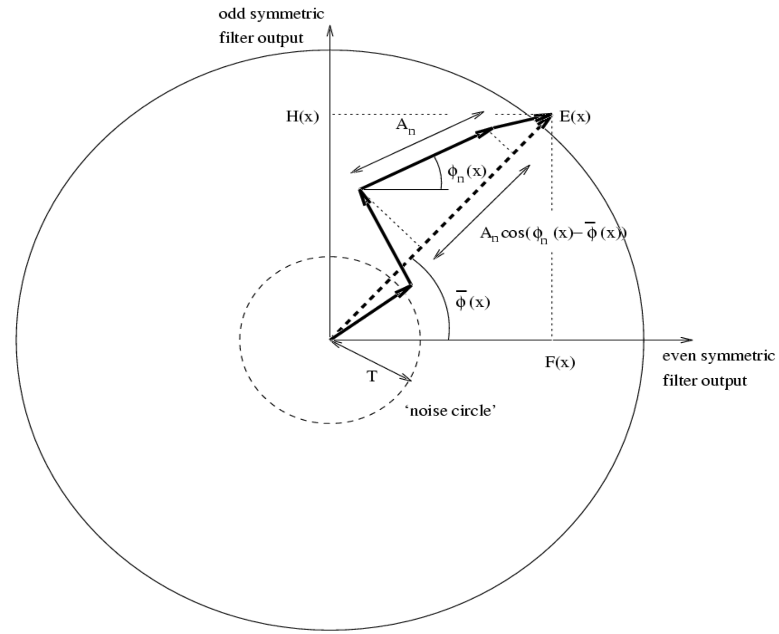

For each point in the signal the responses from the quadrature pairs of filters at different scales will form response vectors that encode phase and amplitude.

Implementation

Local frequency information is obtained by applying quadrature pairs of log-Gabor filters over six orientations and 3-4 scales (typically).

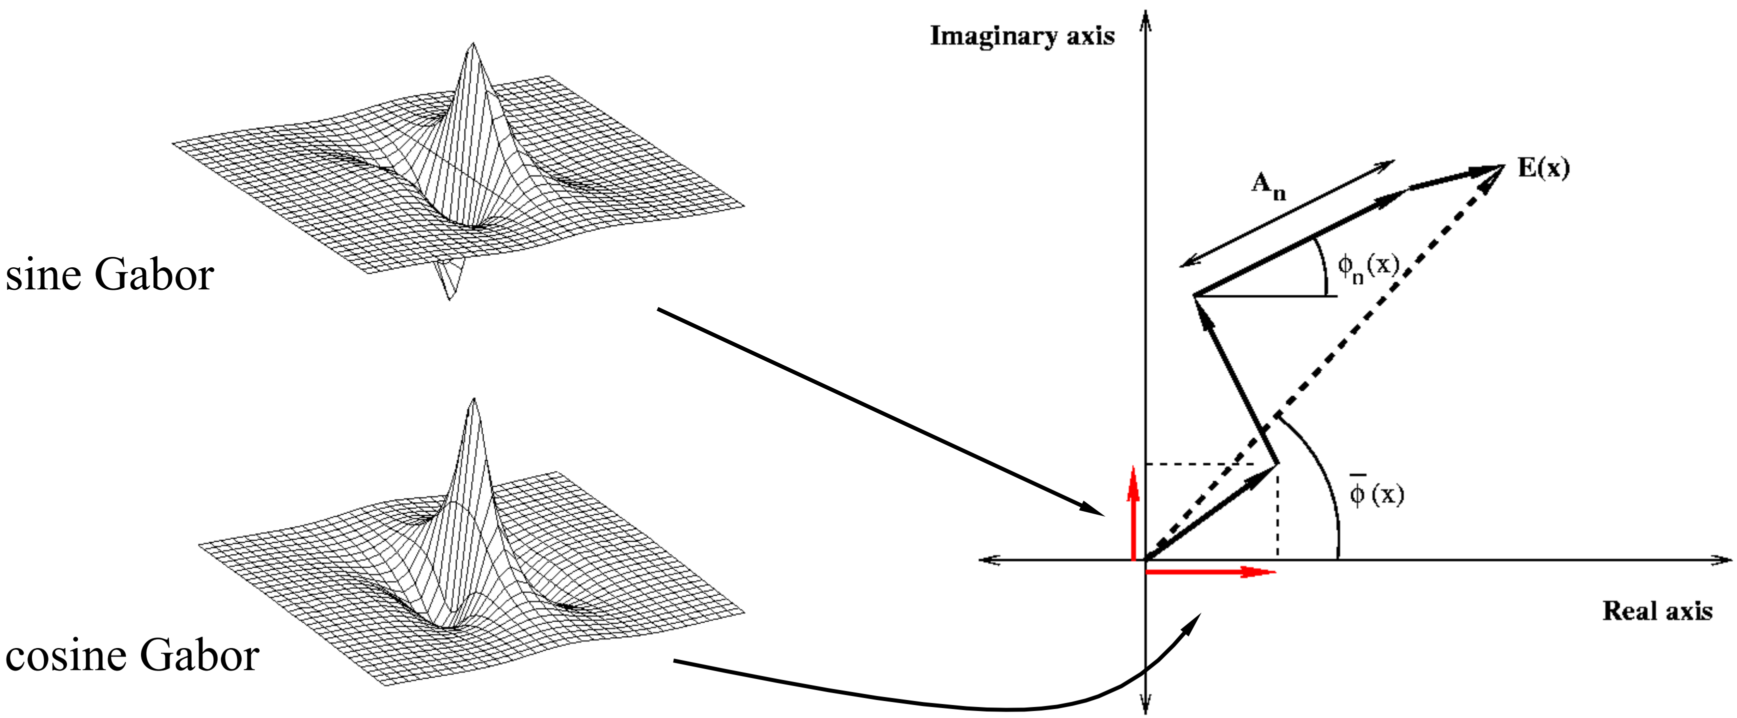

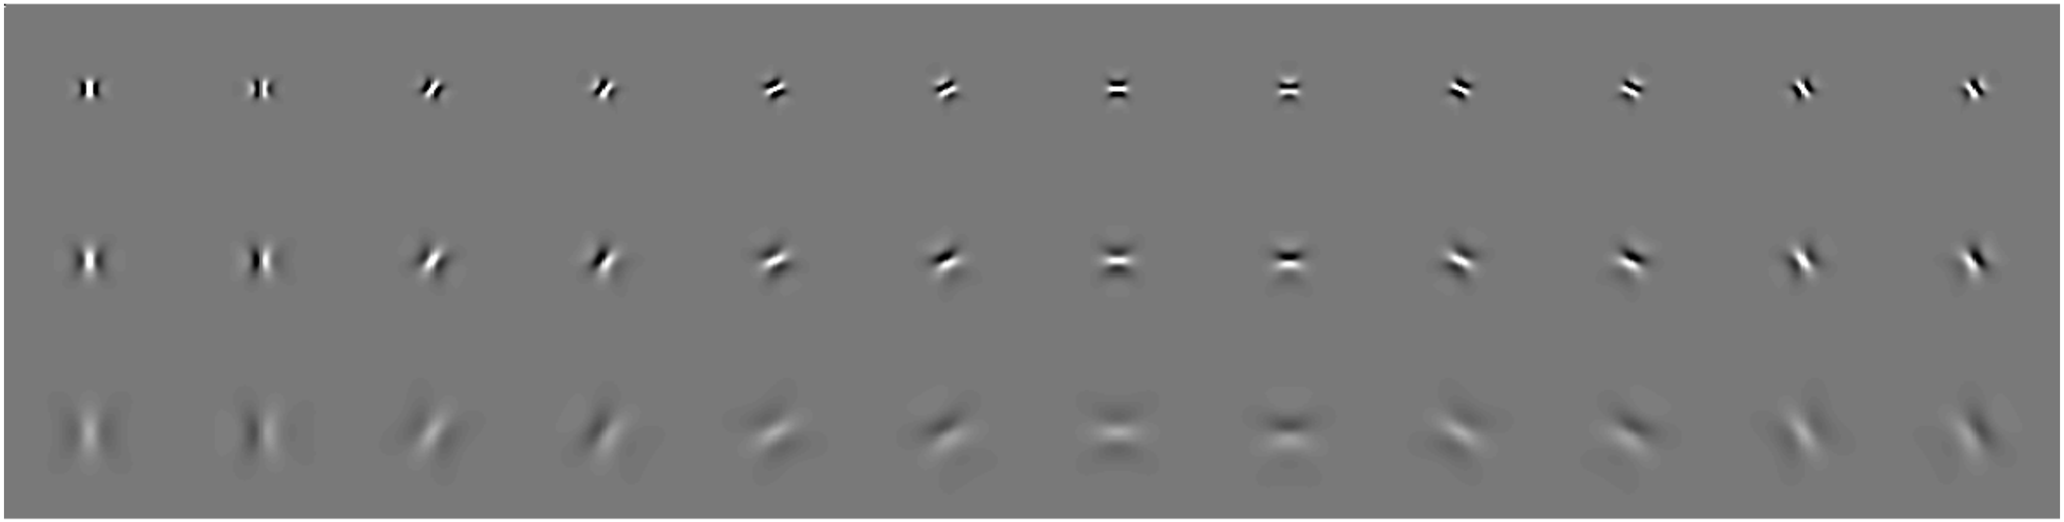

Typical Filter Bank

Implementation

Local frequency information is obtained by applying quadrature pairs of log-Gabor filters over six orientations and 3-4 scales (typically).

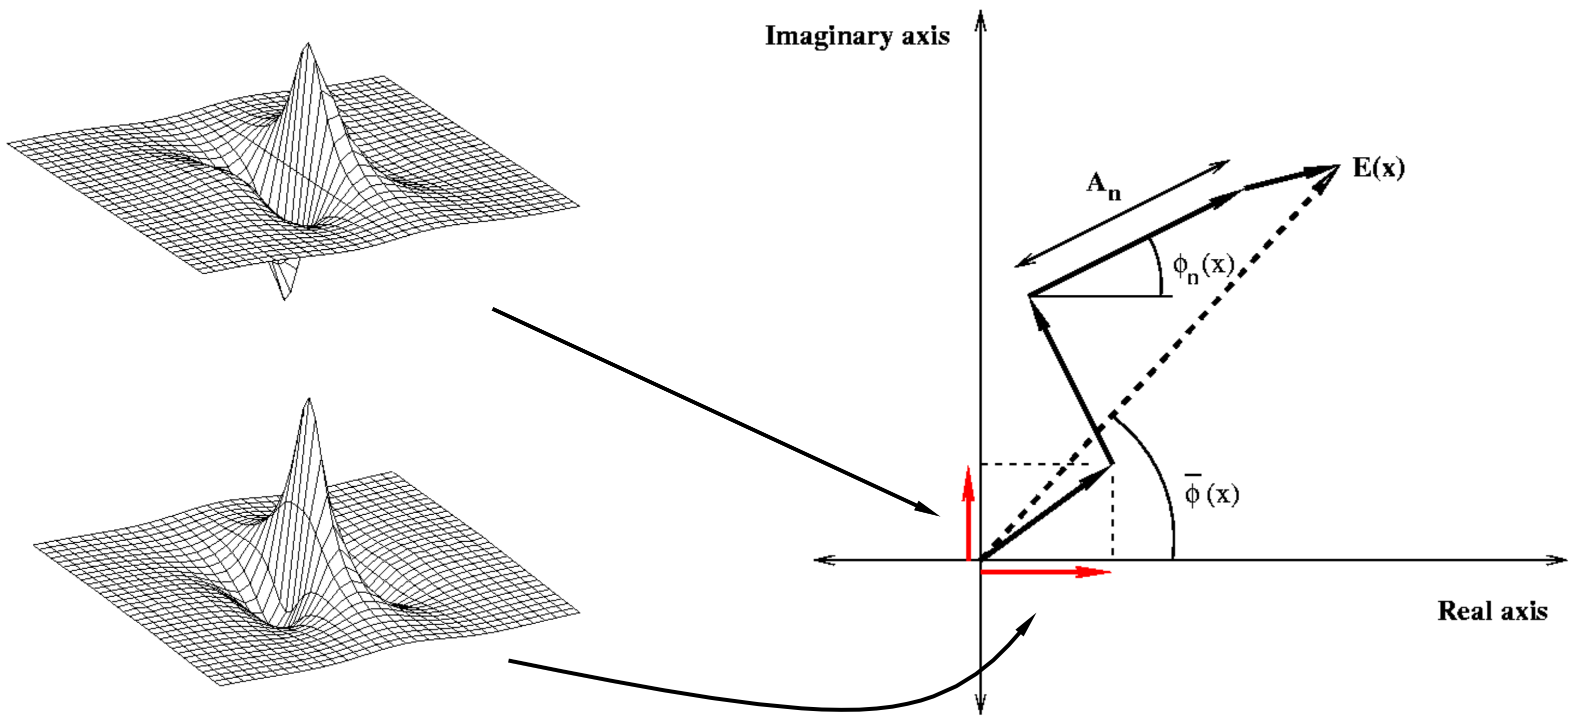

Phase Congruency values are calculated for every orientation
- how do we combine this information?

Detecting Corners and Edges

At each point in the image compute the Phase Congruency covariance matrix.

$$G = \begin{bmatrix} \sum PC_x^2 & \sum PC_x PC_y \\ \sum PC_x PC_y & \sum PC_y^2 \end{bmatrix}$$

where PC_x and PC_y are the x and y components of Phase Congruency for each orientation.

The minimum and maximum singular values correspond to the minimum and maximum moments of Phase Congruency.

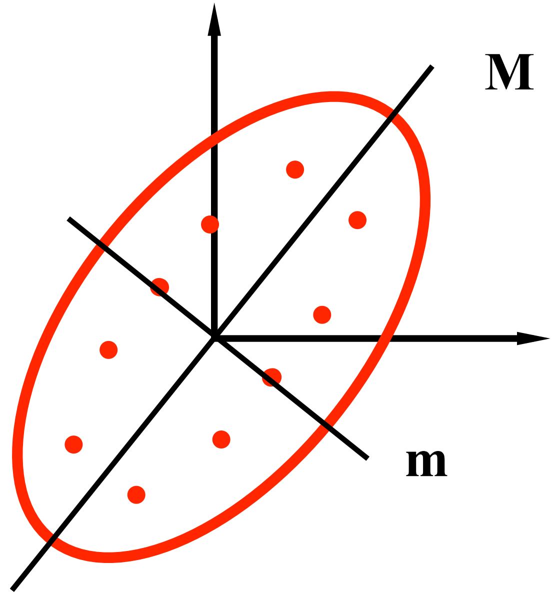

- The magnitude of the maximum moment, **M**, gives an indication of the significance of the feature.

- If the minimum moment, **m**, is also large we have an indication that the feature has a strong 2D component and can be classified as a corner.

- The principal axis, about which the moment is minimized, provides information about the orientation of the feature.

Comparison with Harris Operator

- The minimum and maximum values of the Phase Congruency singular values/moments are bounded to the range 0-1 and are dimensionless. This provides invariance to illumination and contrast.

- Phase Congruency moment values can be used directly to determine feature significance. Typically a threshold of 0.3 - 0.4 can be used.

- Eigenvalues/Singular values of a gradient covariance matrix are unbounded and have units $(\text{gradient})^4$ - the cause of the difficulties with the Harris operator.

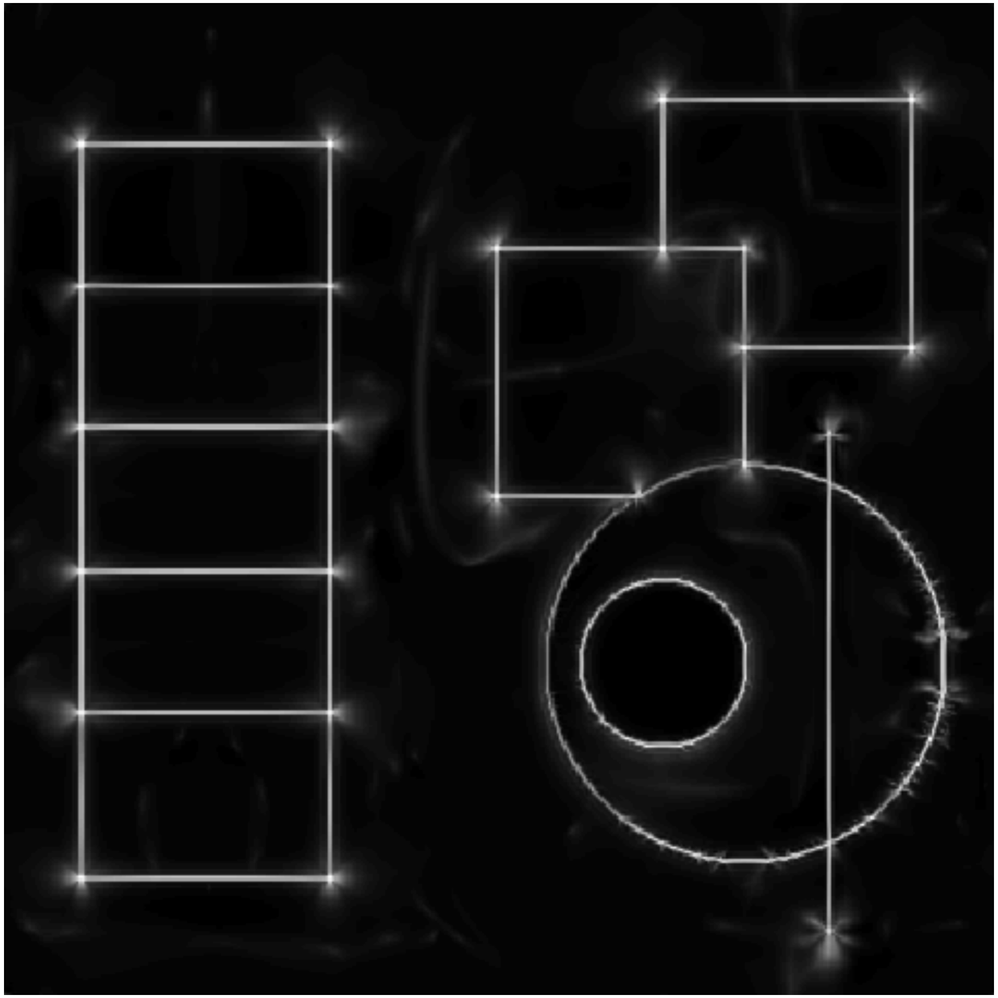

Edge strength
- given by magnitude of <u>maximum</u>
Phase Congruency moment.

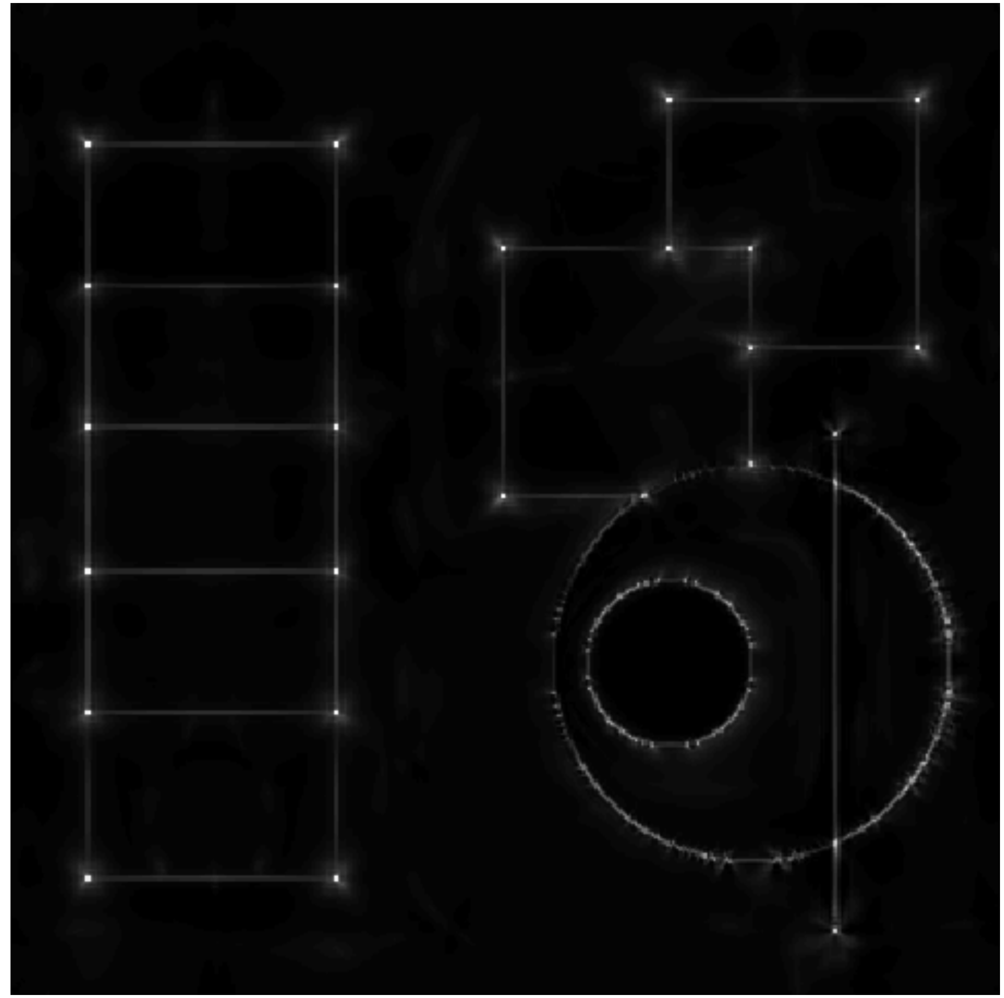

Corner strength
- given by magnitude of <u>minimum</u>
Phase Congruency moment.

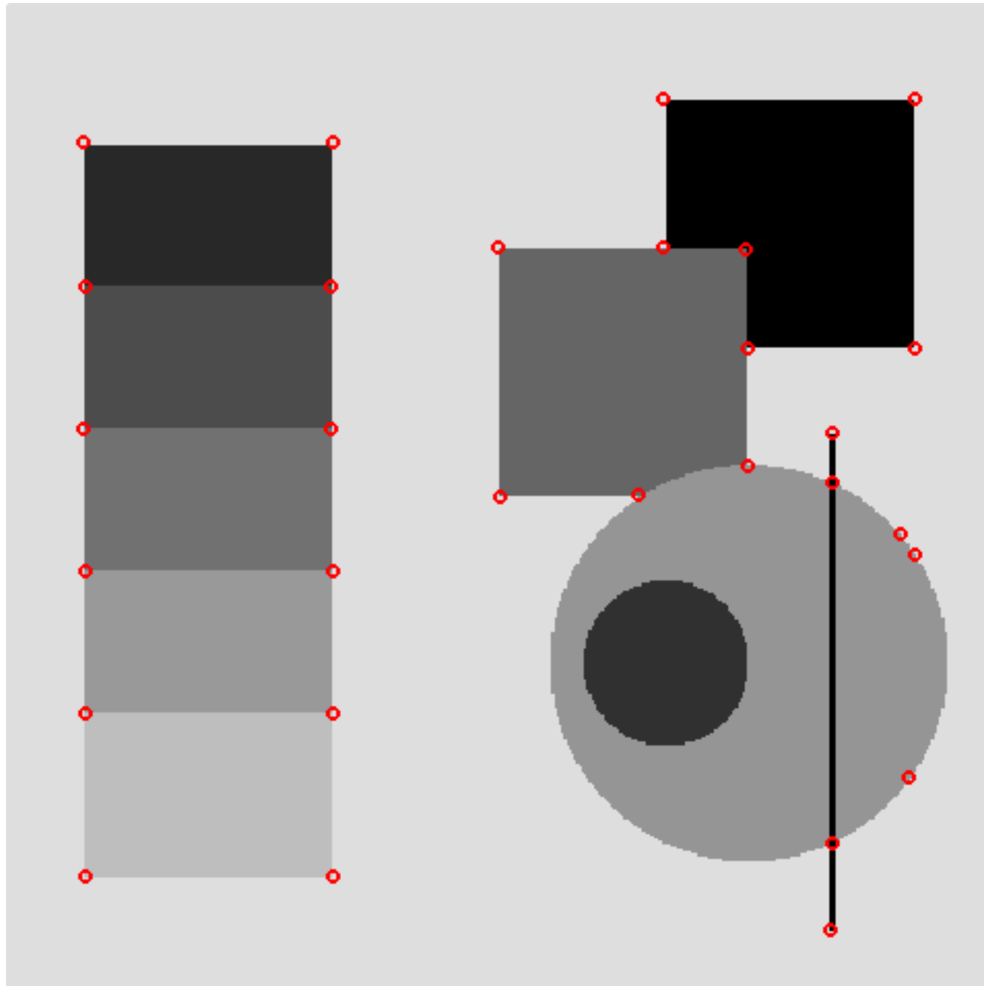

Phase Congruency corners thresholded at 0.4

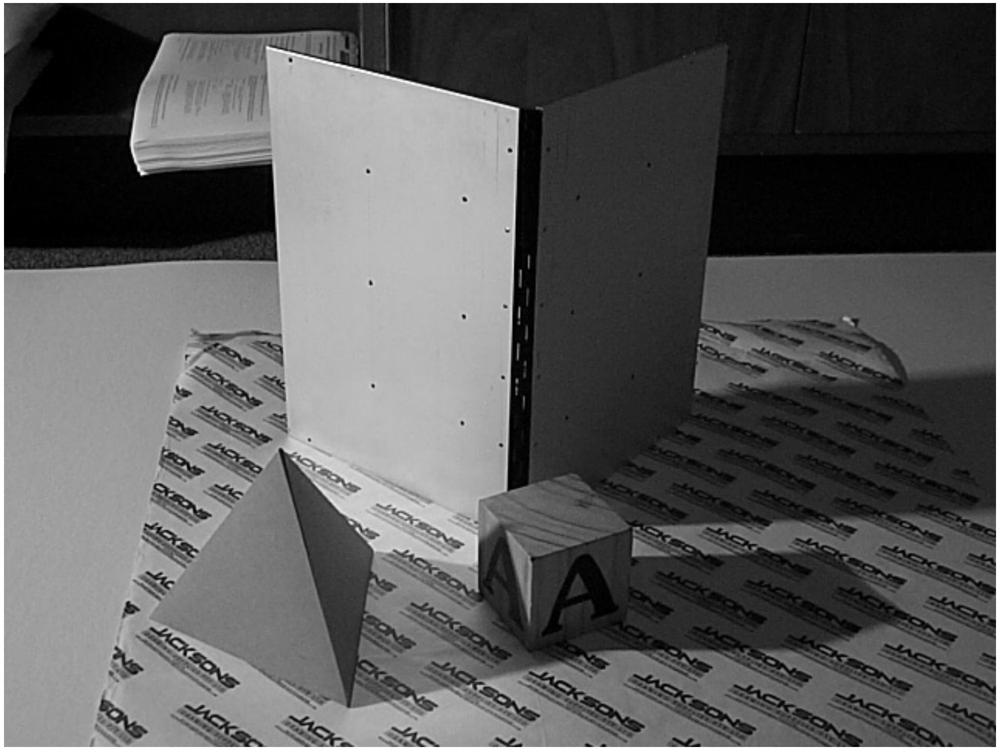

Image with strong shadows

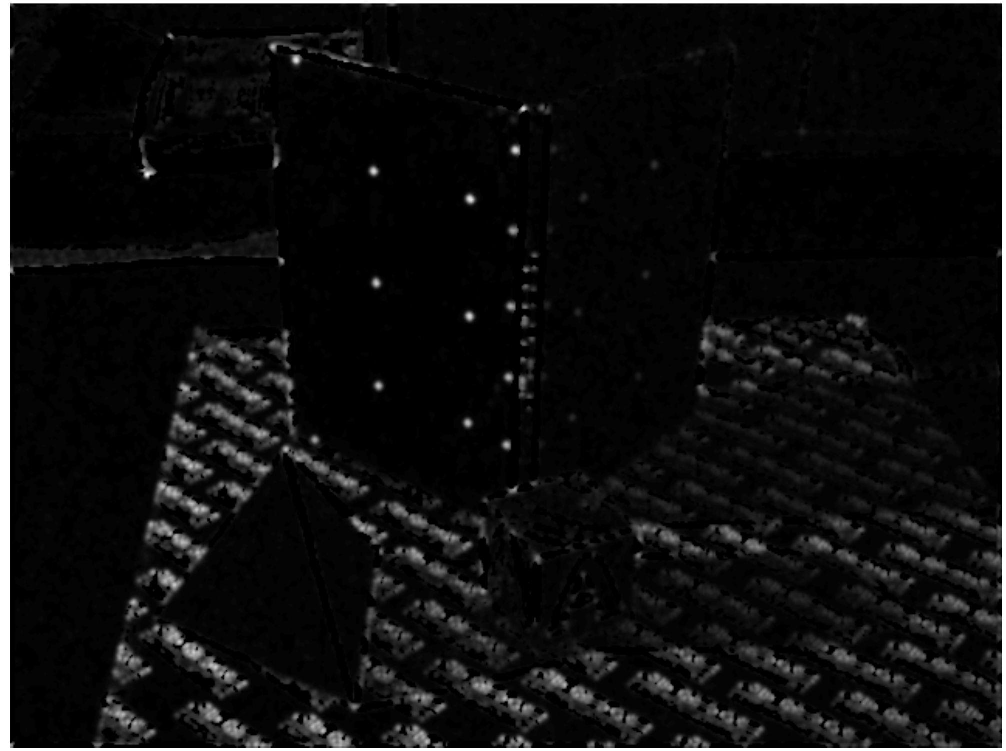

Fourth root of Harris corner strength (max value ~1.25×10^{10})

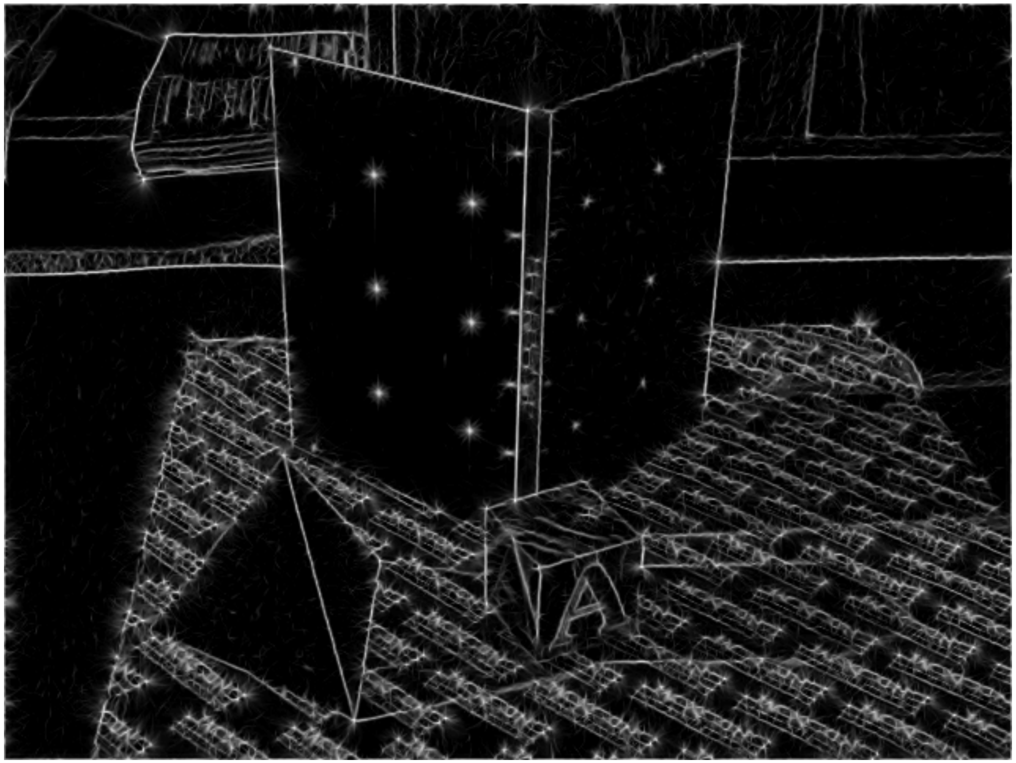

Phase Congruency edge strength
(max possible value = 1.0)

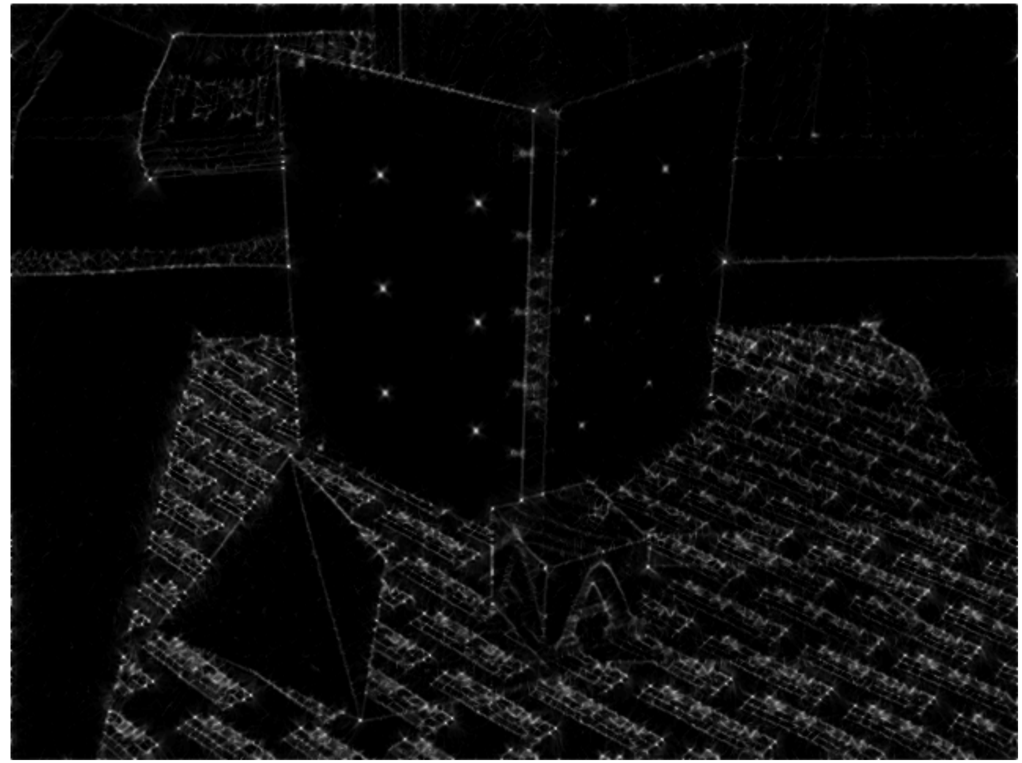

Phase Congruency corner strength
(max possible value = 1.0)

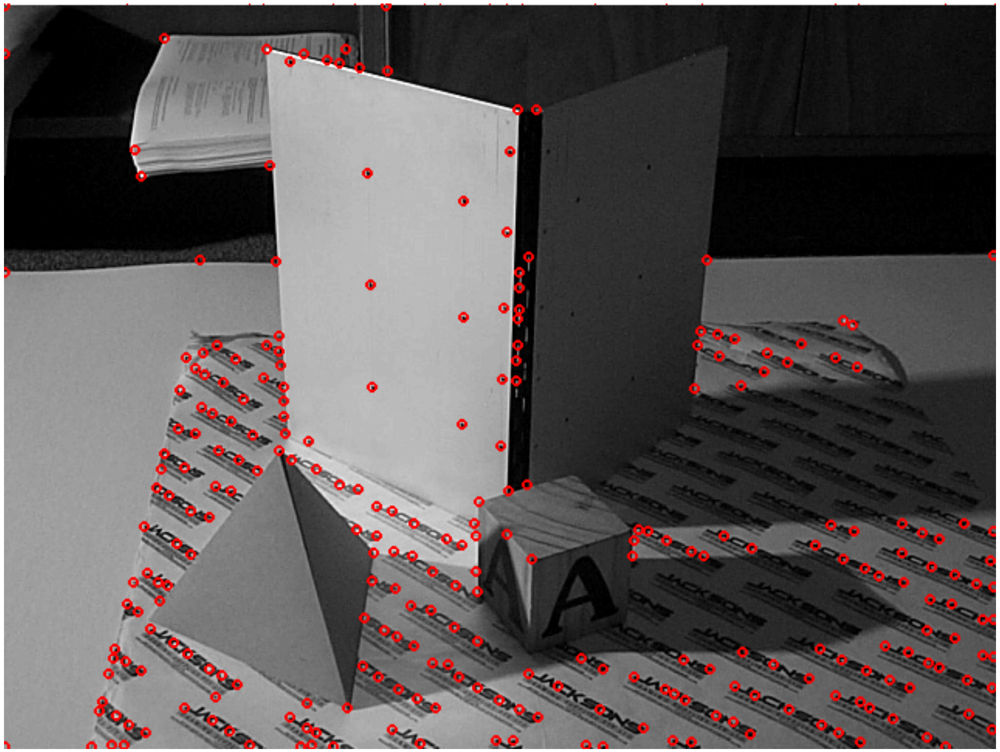

Harris corners thresholded at 10^8
(max corner strength $\sim 1.25 \times 10^{10}$)

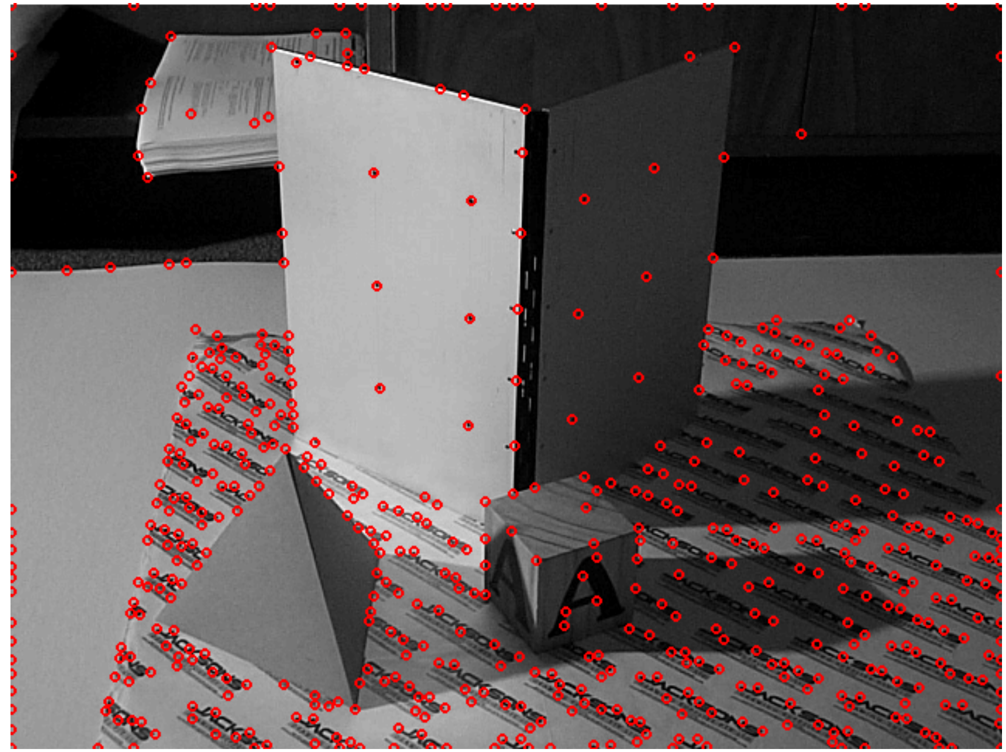

Phase Congruency corners thresholded at 0.4


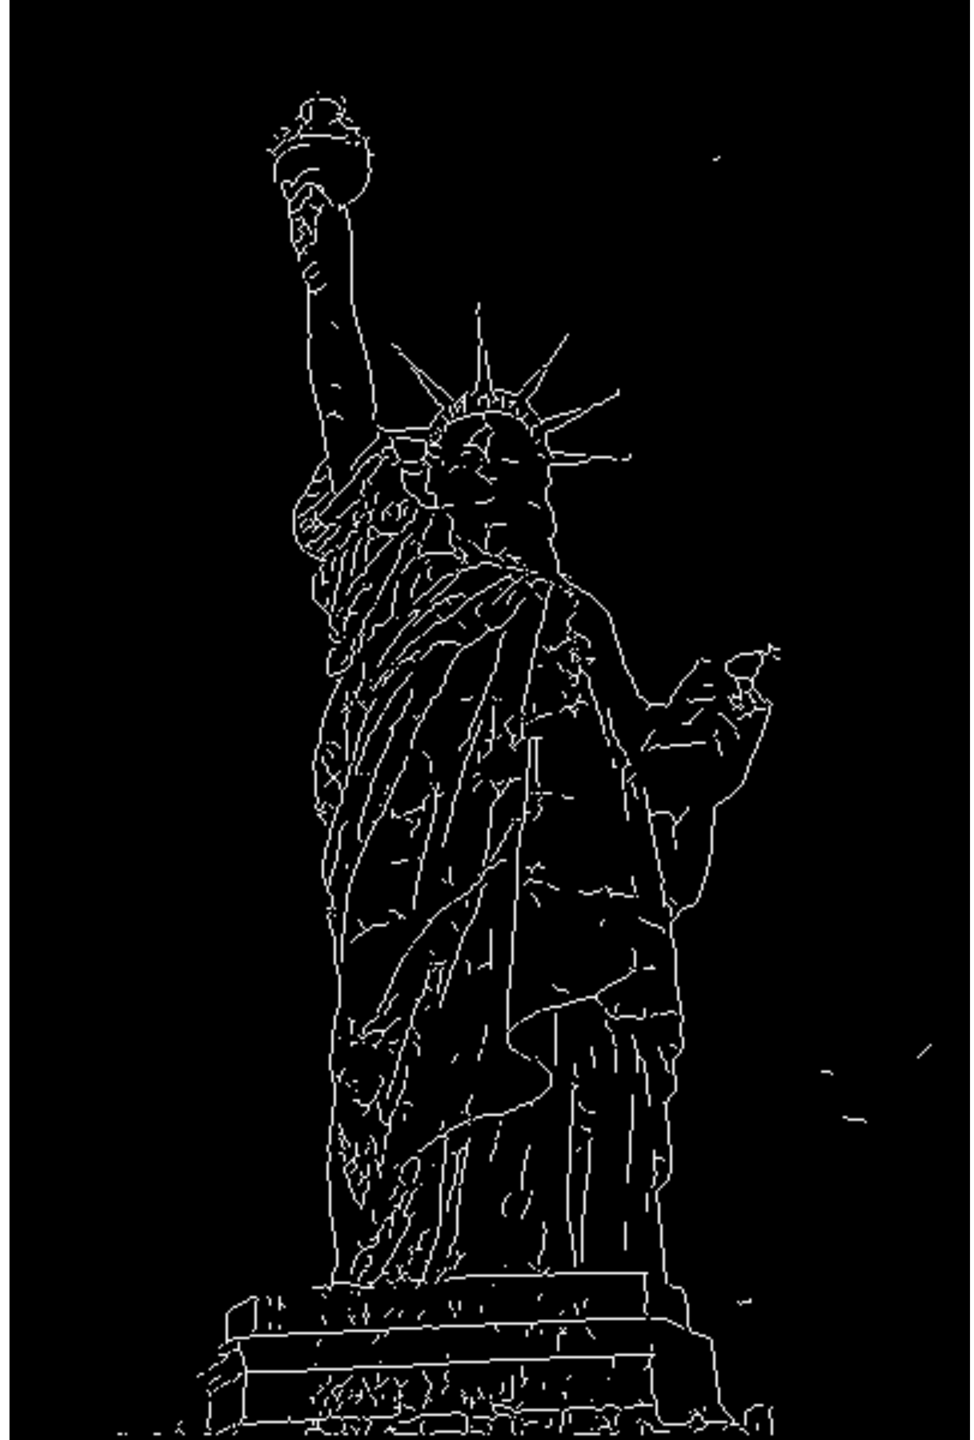


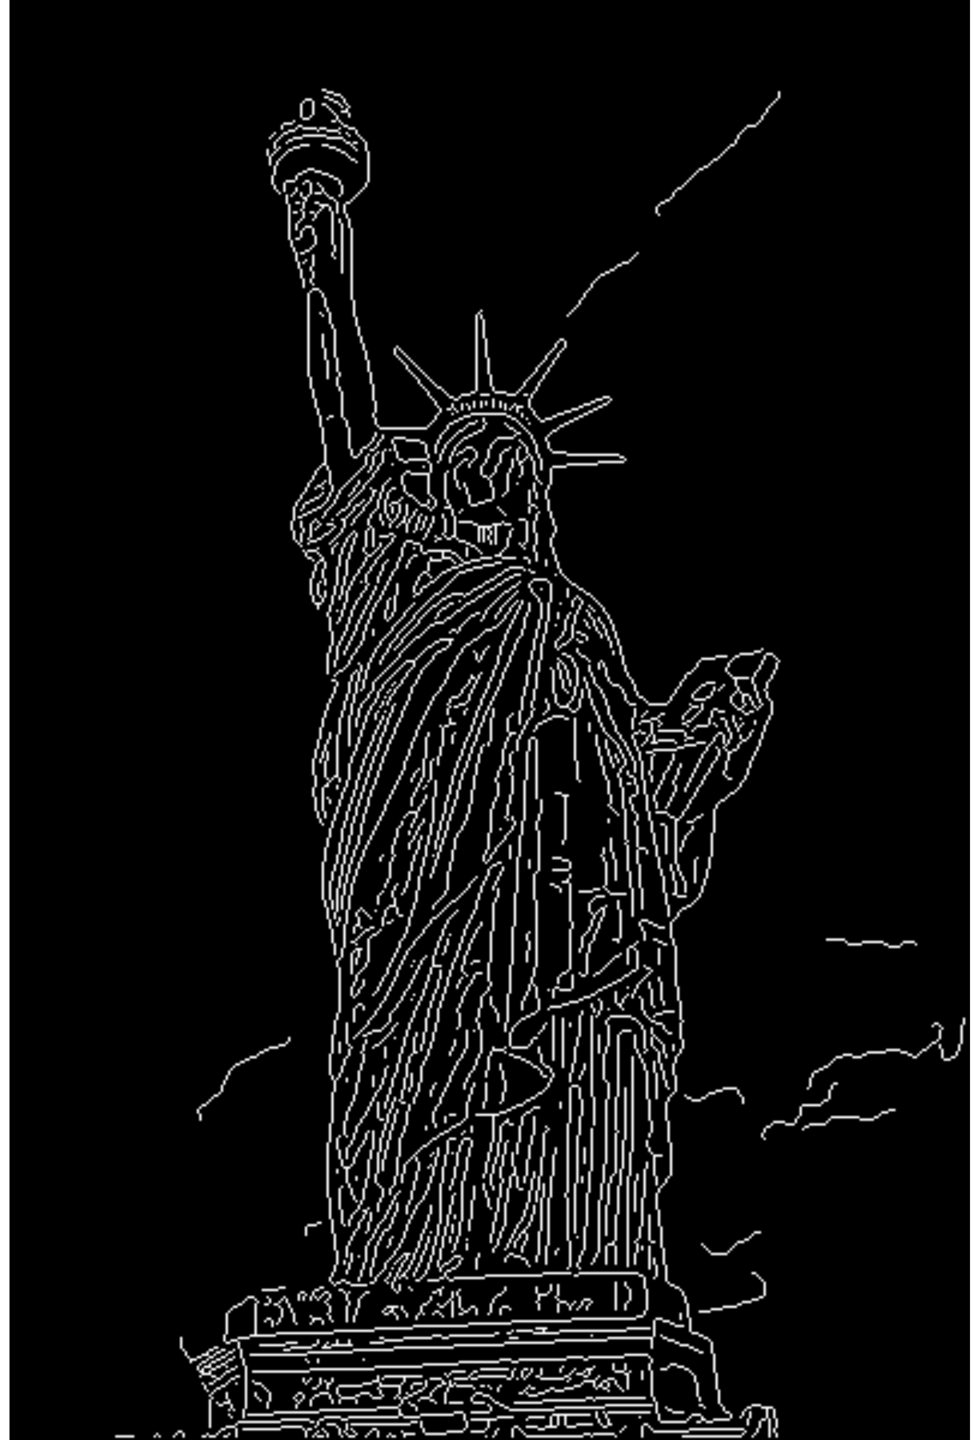

Good Things:

- Phase Congruency is a dimensionless quantity.

- Invariant to contrast and scale.

- Value ranges between 0 and 1.

$$0 \longleftrightarrow 1$$

no congruency   perfect congruency

- Threshold values can be fixed for wide classes of images.

- Provides classification of features.

Problems:

- Degenerates when Fourier components are very small.
- Degenerates when there is only one Fourier component.
- Sensitive to noise.
- Localization is not sharp.

Degeneracy with a single frequency component

- Congruency of phase is only significant if it occurs over a distribution of frequencies.
- What is a good distribution?
- Natural images have an amplitude spectrum that decay at $1/\omega$ (Field 1987)

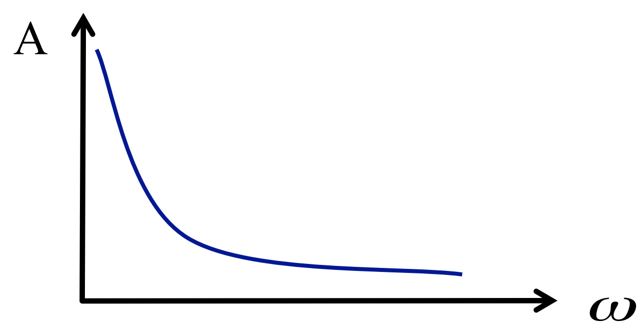

weighting function for frequency spread

$$PC(x) = \frac{W(x)\,|E(x)|}{\sum_n A_n(x) + \varepsilon}$$

small constant to avoid division by 0

Noise

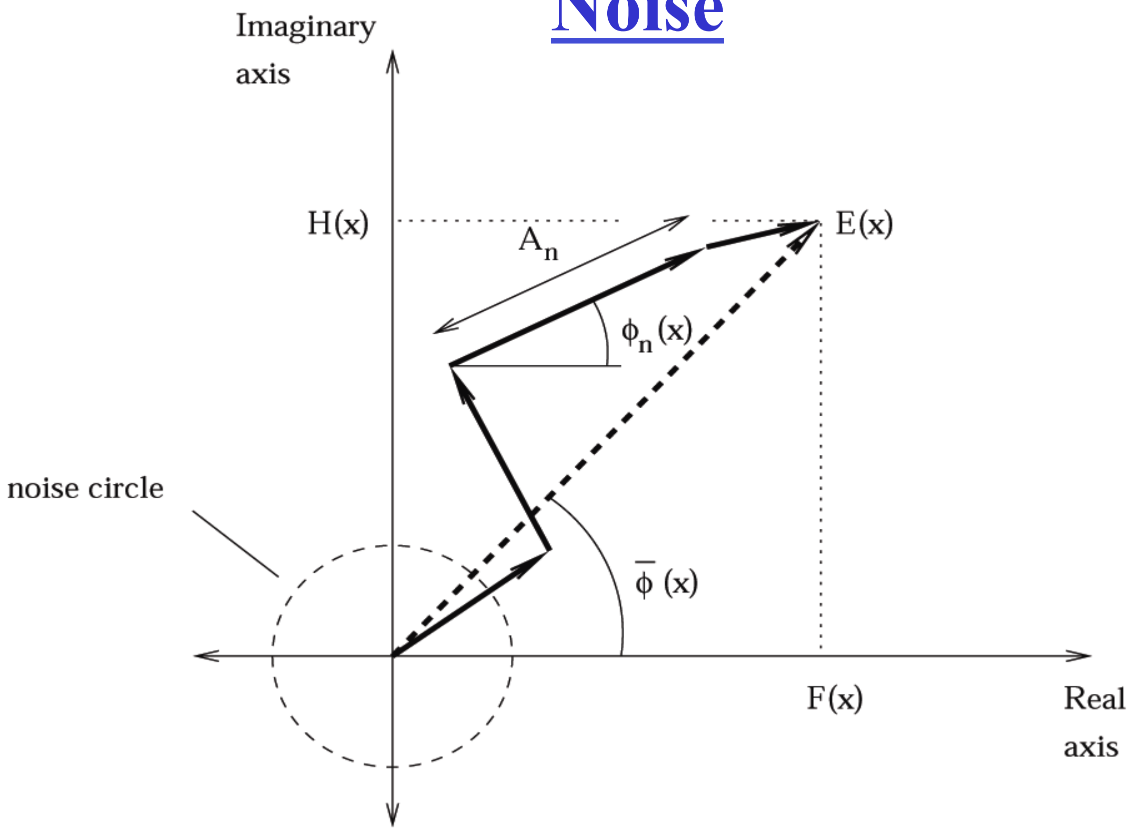

- Being a normalized quantity Phase Congruency is sensitive to noise.
- The radius of the noise circle represents the value of |***E(x)***| one can expect from noise.
- If ***E(x)*** falls within this circle our confidence in any Phase Congruency value falls to 0.

Noise Compensation

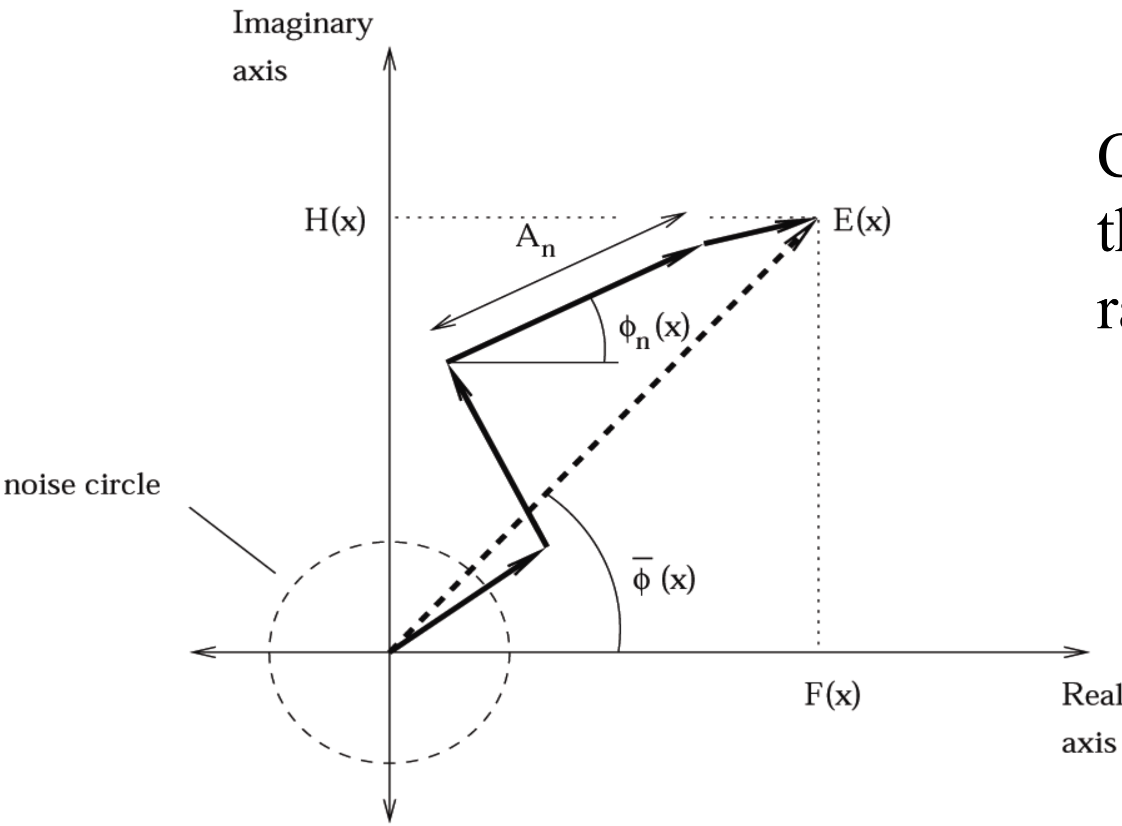

Calculate Phase Congruency using the amount that |***E(x)*** | exceeds the radius of the noise circle, ***T.***

$$PC(x) = \frac{W(x)\lfloor |E(x)| - T \rfloor}{\sum_n A_n(x) + \varepsilon}$$

Radius of the Noise Circle…

The radius of the noise circle is the expected distance from the origin if we take an n-step random walk in the plane

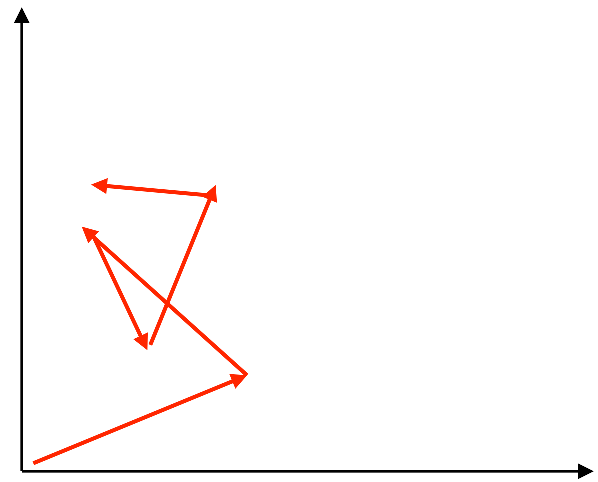

- The size of the steps correspond to the expected noise response of the filters at each scale.

- The direction of the steps are random.

Radius of the Noise Circle…

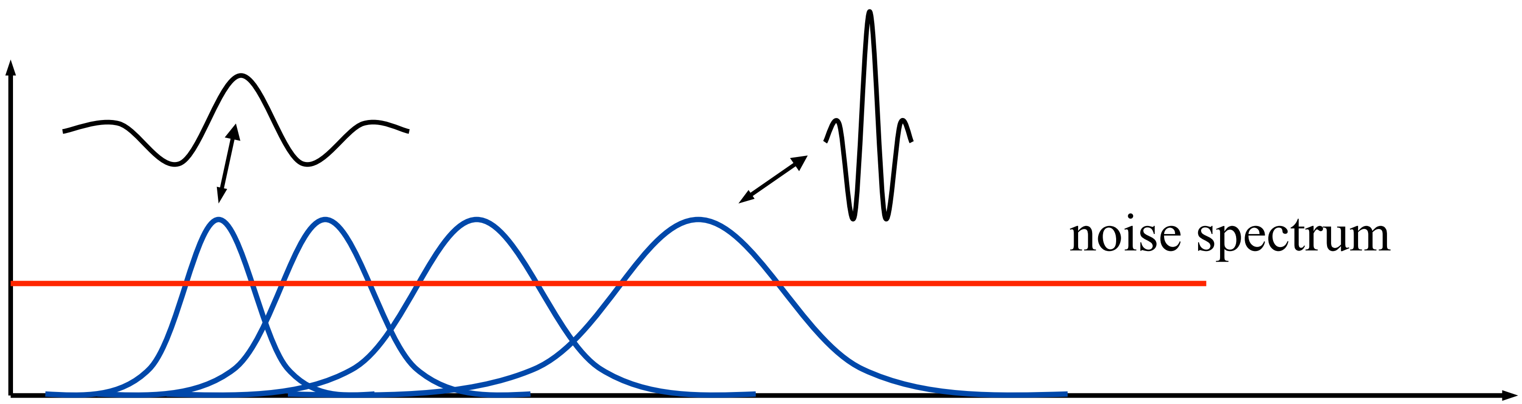

If we assume the noise spectrum is flat filters will gather energy from the noise as a function of their bandwidth (which is proportional to centre frequency).

Smallest wavelet has largest bandwidth
$\Rightarrow$ gets the most energy from noise.

Smallest wavelet has the most local response to features in the signal
$\Rightarrow$ most of the time it is only responding to noise.

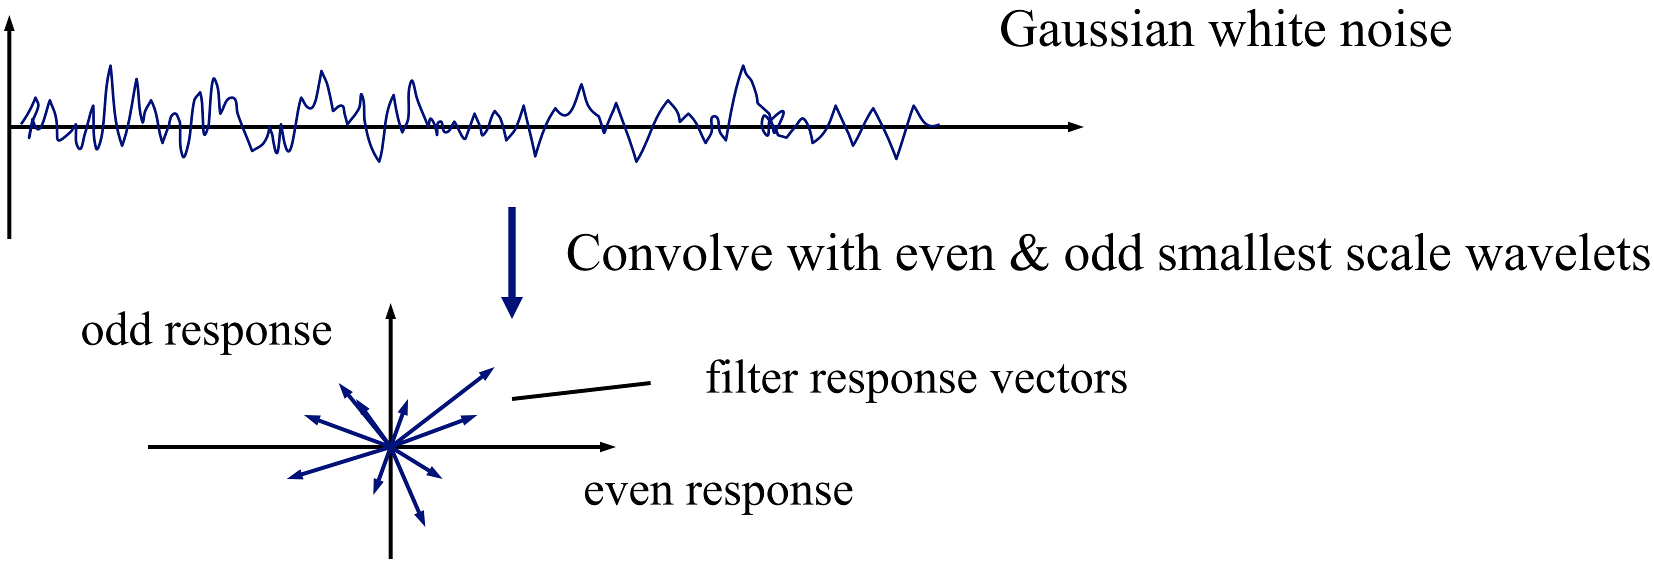

The distribution of the positions of the response vectors will be a 2D Gaussian (+ some contamination from feature responses).

We are interested in the distribution of the <u>magnitude</u> of the responses
- This will be a Rayleigh Distribution.

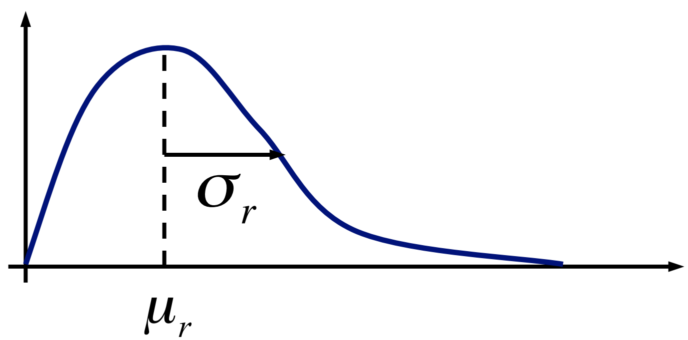

Mean μ_r Variance σ_r^2
are described by <u>one</u> parameter

- To obtain a robust estimate of μ_r for the smallest scale filter we find the median response of the filter over the whole image. This minimizes the influence of any contamination to the distribution from responses to features.

- μ_r and the median have a fixed relationship.

- Get median → get μ_r → get σ_r

- Expected noise responses at other scales are determined by the bandwidths of the other filters relative to the smallest scale filter pair.

- Given the expected filter responses we can solve for the expected distribution of the sum of the responses (which will be another Rayleigh distribution). Set noise threshold in terms of the overall distribution.

$$T = \mu_r + \overset{2\text{ - }3}{k}\,\sigma_r$$

Noise Compensation Example

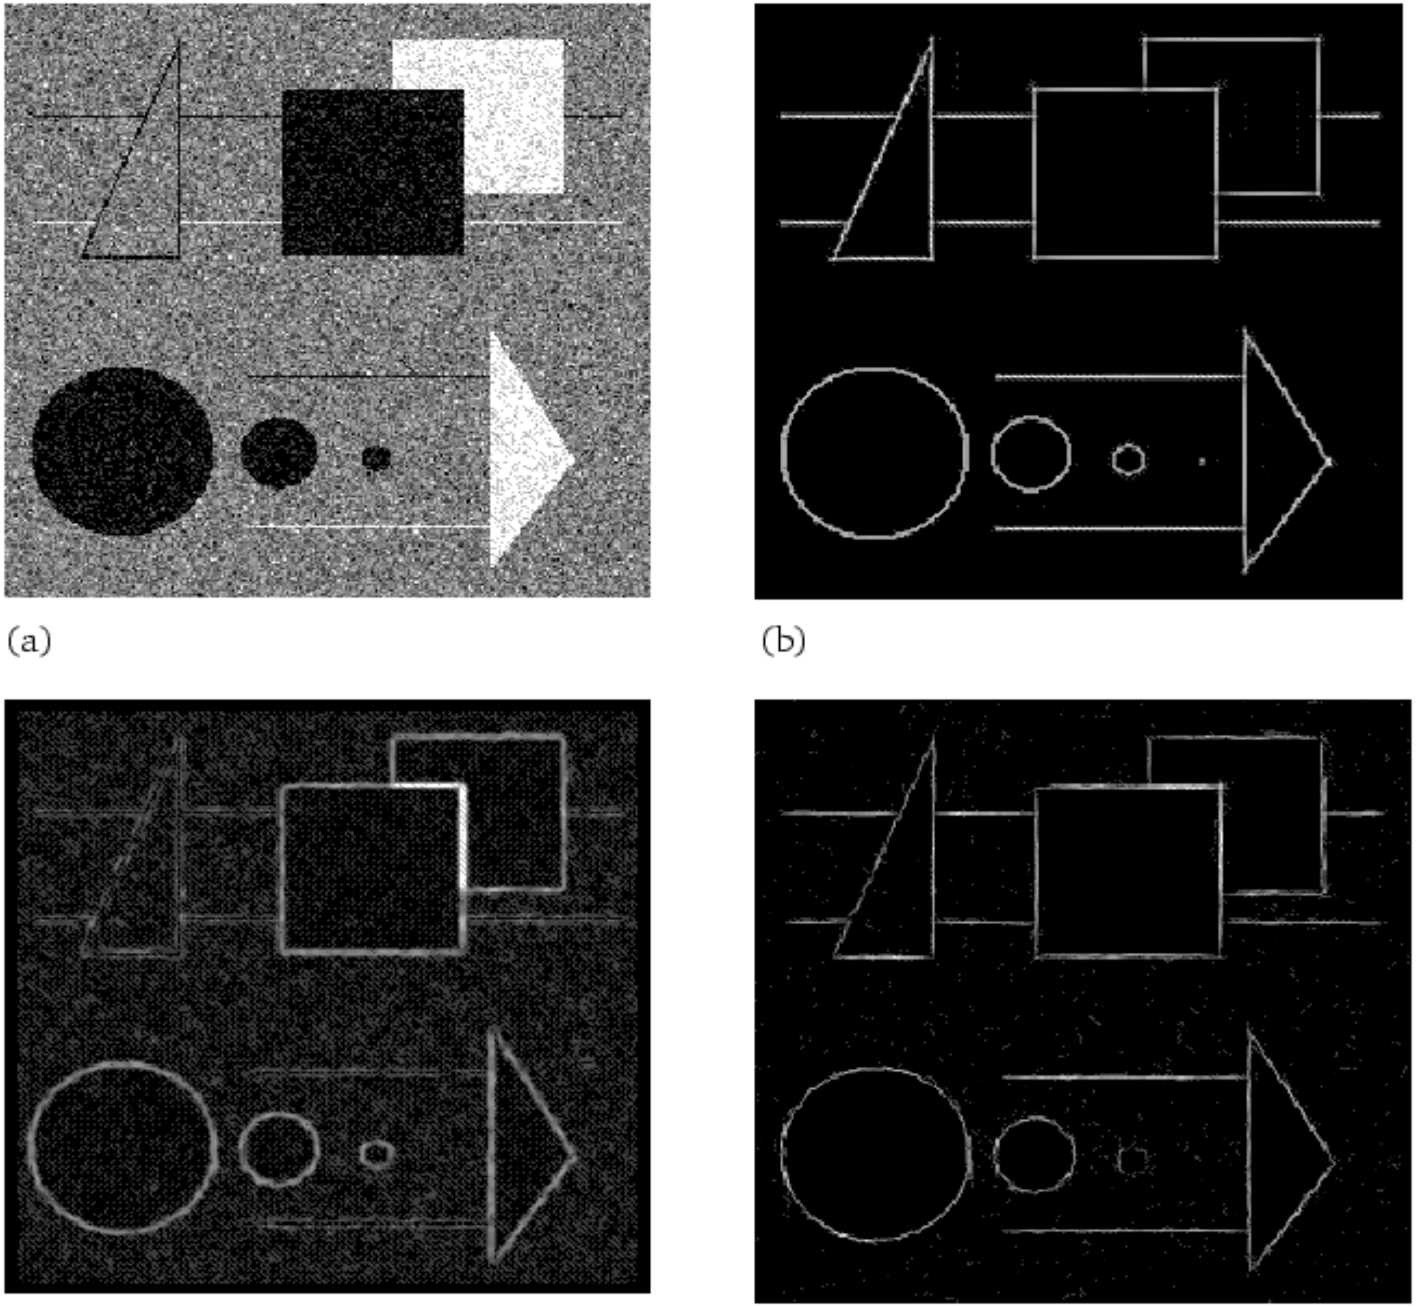

(a) (b) (c) (d)

a) Test image with 256 grey levels + Gaussian noise with standard deviation of 40 grey values.
b) Raw Phase Congruency response on noise-free image.
c) Raw Canny edge strength on noisy image.
d) Raw Phase Congruency on noisy image.

The Localization Problem

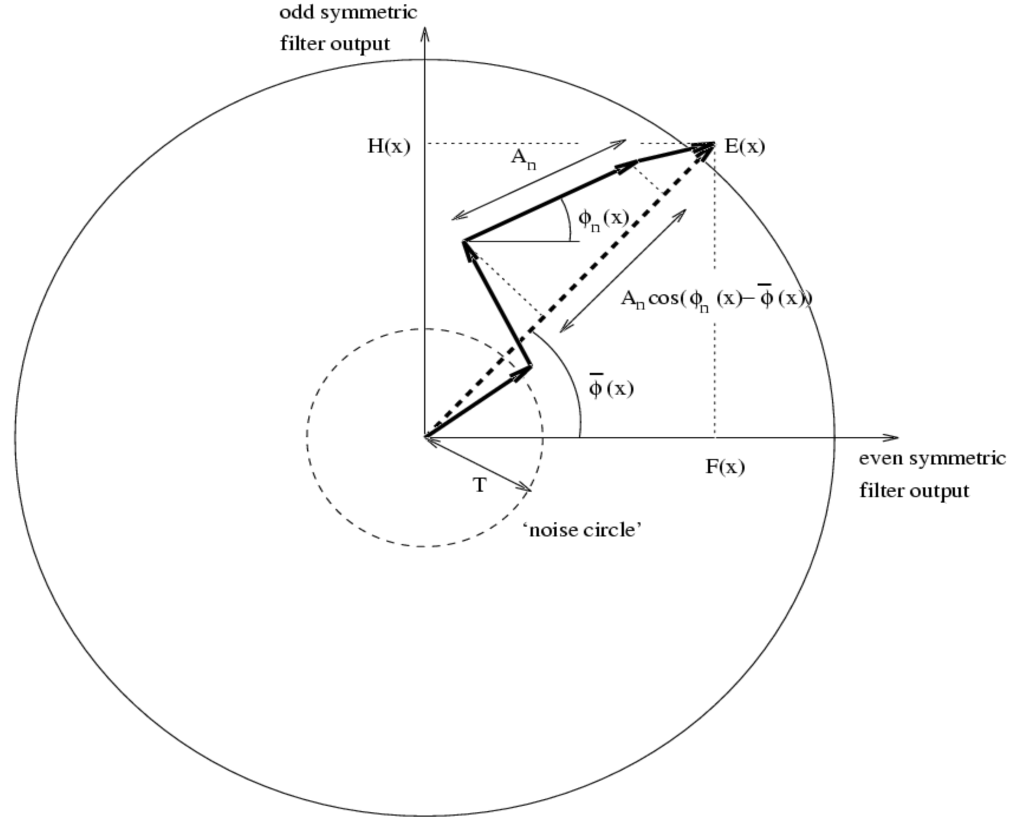

$$PC_1(x) = \frac{|E(x)|}{\sum_n A_n(x)}$$

$$= \frac{\sum_n A_n \cos(\phi_n(x) - \overline{\phi}(x))}{\sum_n A_n(x)}$$

This measure of Phase Congruency is a function of the *cosine* of the phase deviation.

The cosine function has zero slope at the origin.

Improving Localization

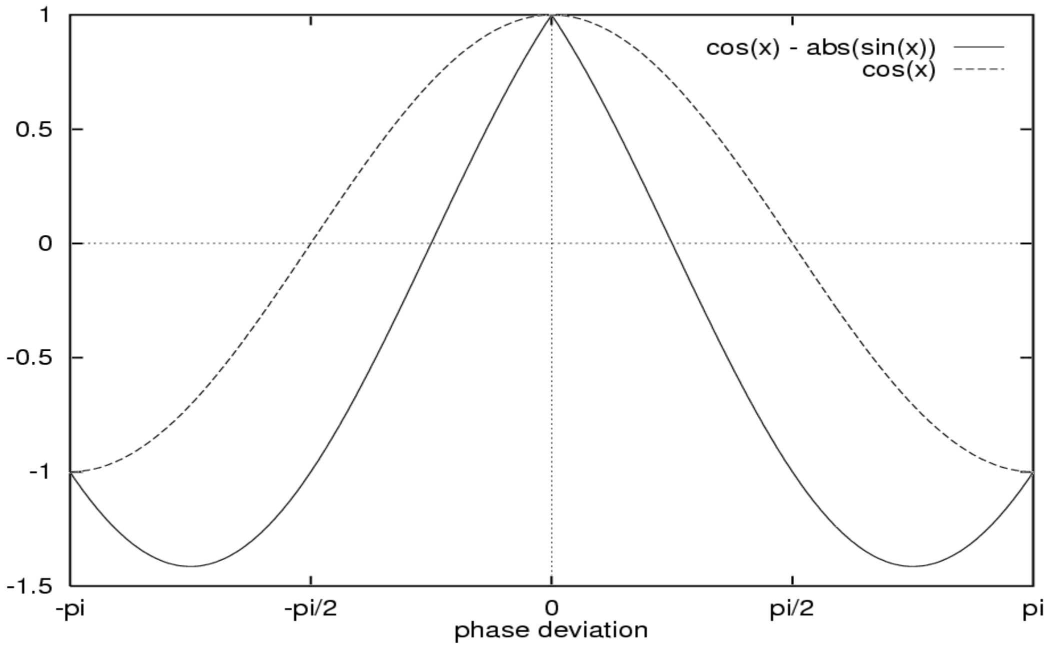

cosine - |sine| varies nearly linearly as one moves away from the origin.

$$PC_2(x) = \frac{\sum_n A_n[\cos(\phi_n(x) - \overline{\phi}(x)) - |\sin(\phi_n(x) - \overline{\phi}(x))|]}{\sum_n A_n(x)}$$

(easily calculated with dot and cross products)

The Final Equation

frequency spread weighting

noise threshold

$$PC_2(x) = \frac{W(x)\left\lfloor \sum_n \overbrace{A_n[\cos(\phi_n(x) - \overline{\phi}(x)) - |\sin(\phi_n(x) - \overline{\phi}(x))|]}^{\text{‘energy’}} - T \right\rfloor}{\sum_n A_n(x) + \varepsilon}$$

small constant to avoid division by 0

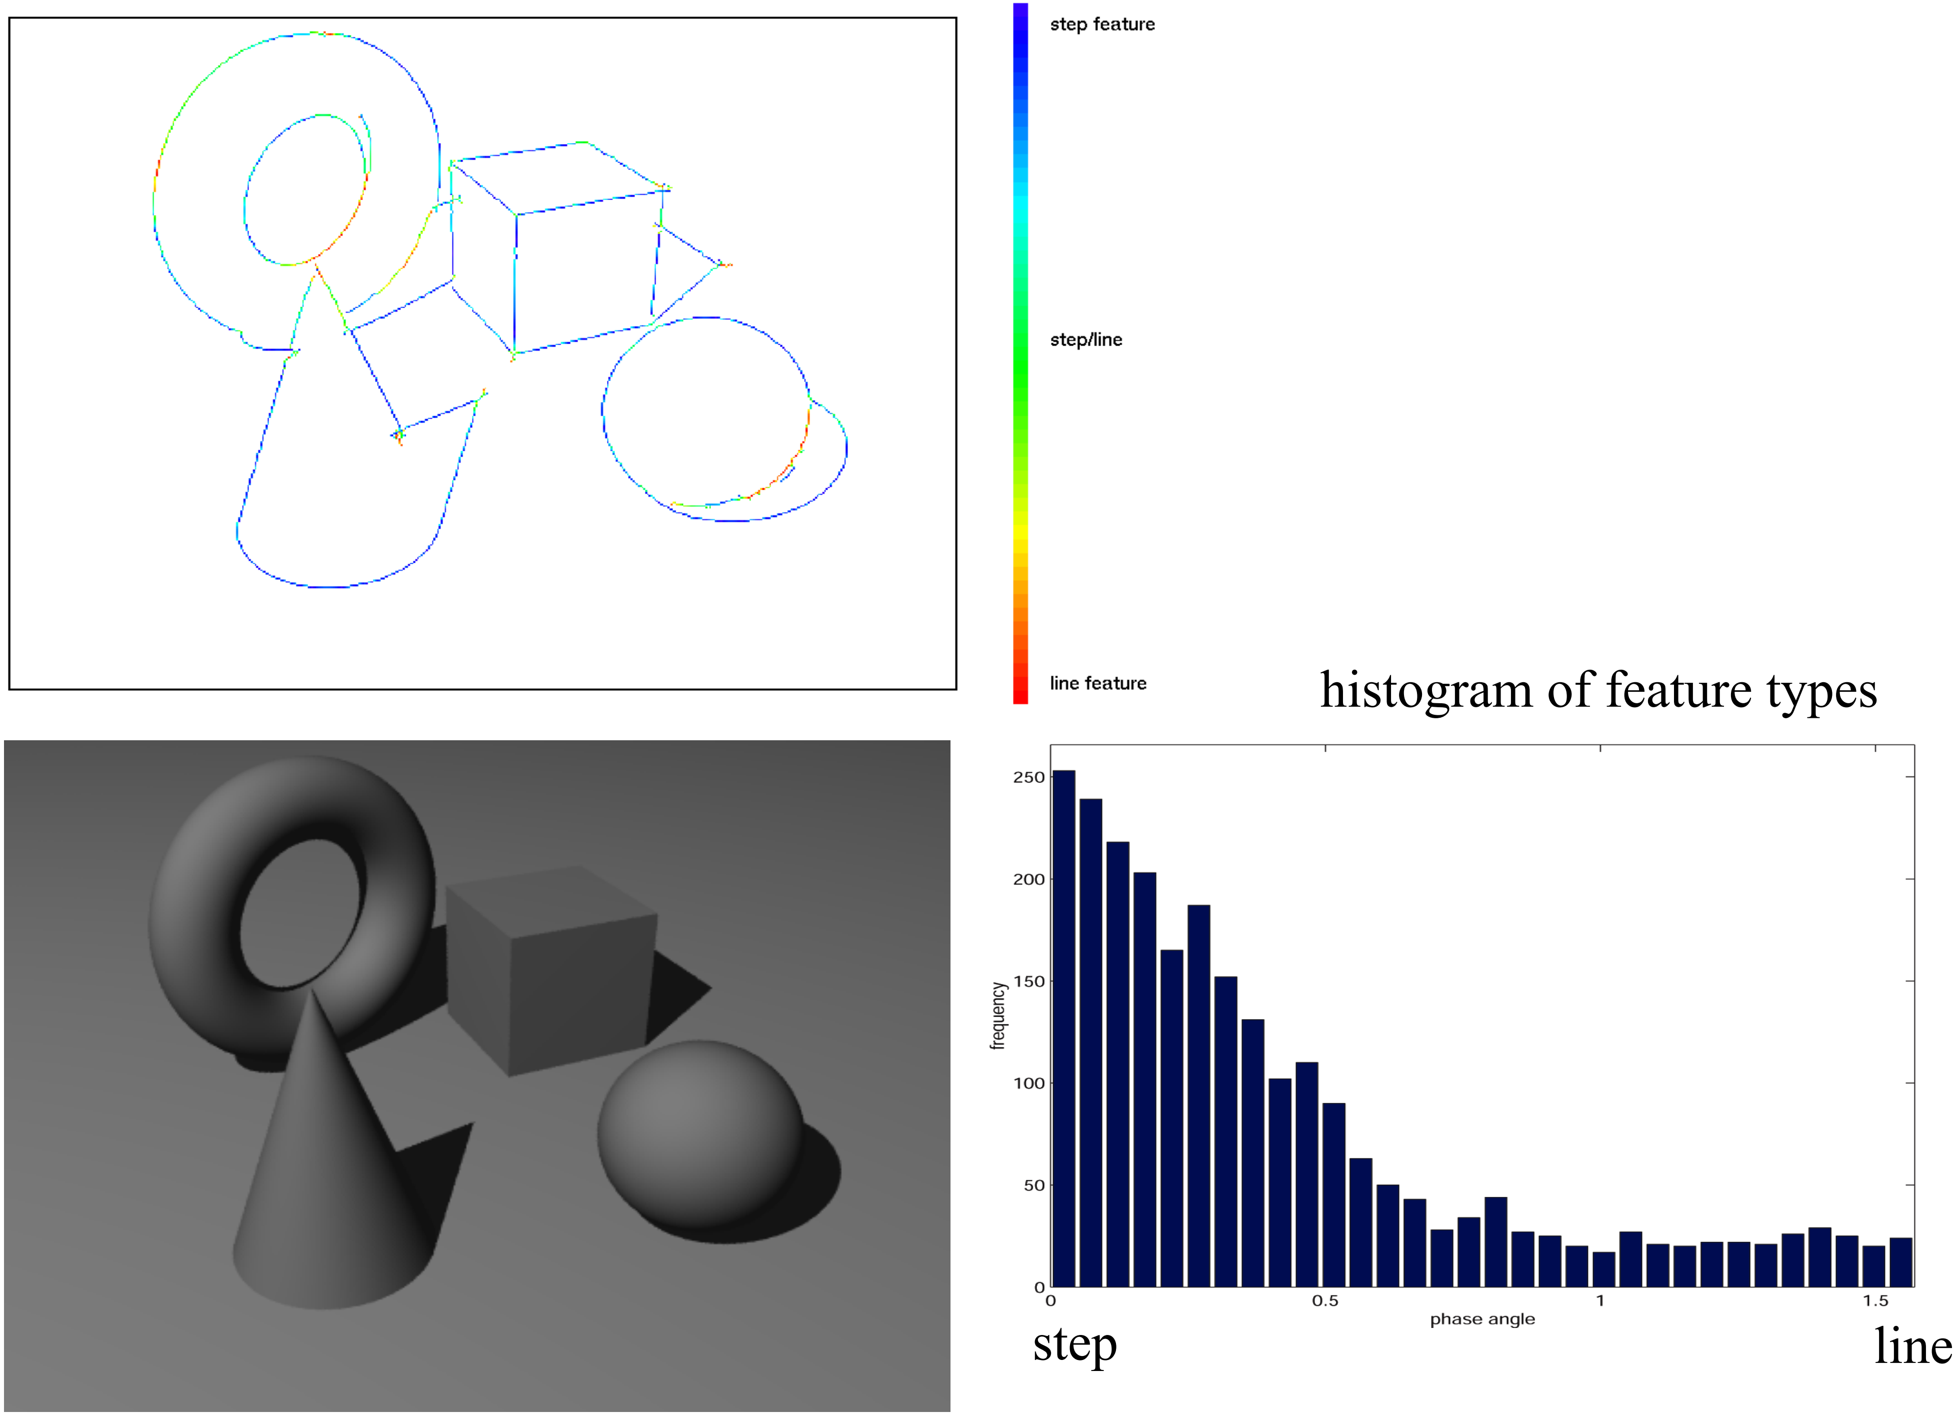
step feature
step/line
line feature
histogram of feature types
frequency
0
50
100
150
200
250
0
0.5
1
1.5
phase angle
step
line

Canny

Phase Congruency

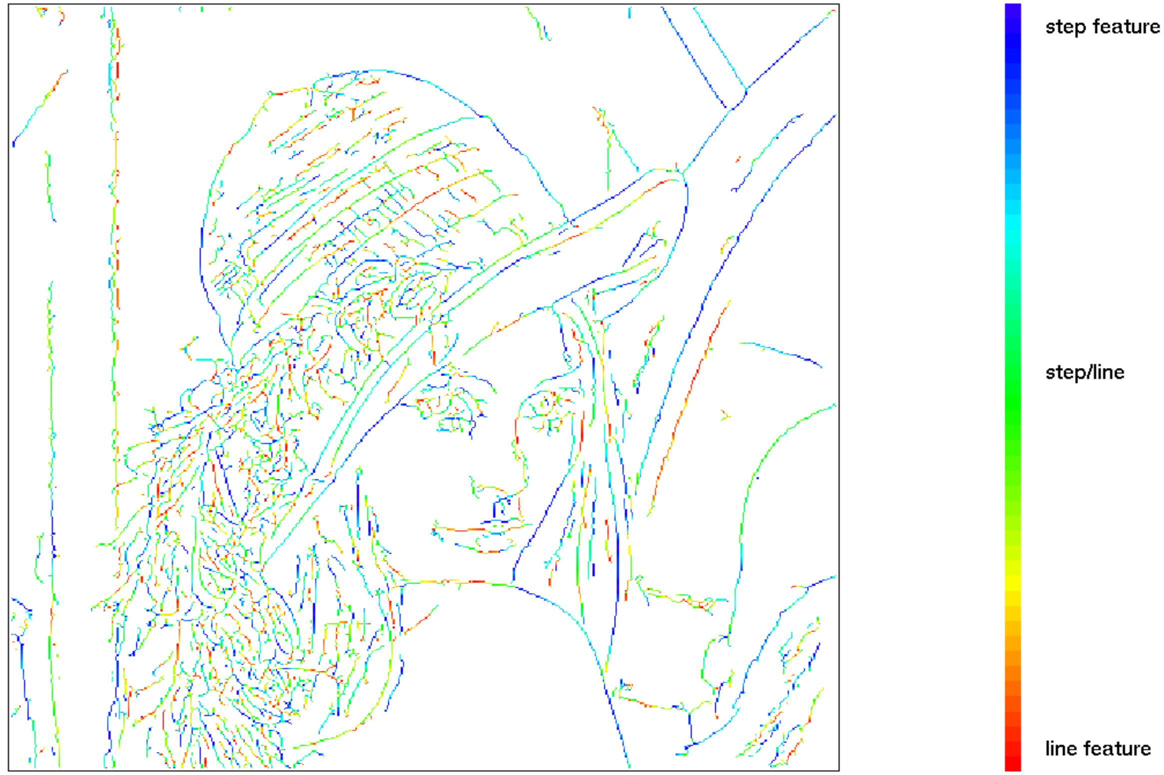

histogram of feature types

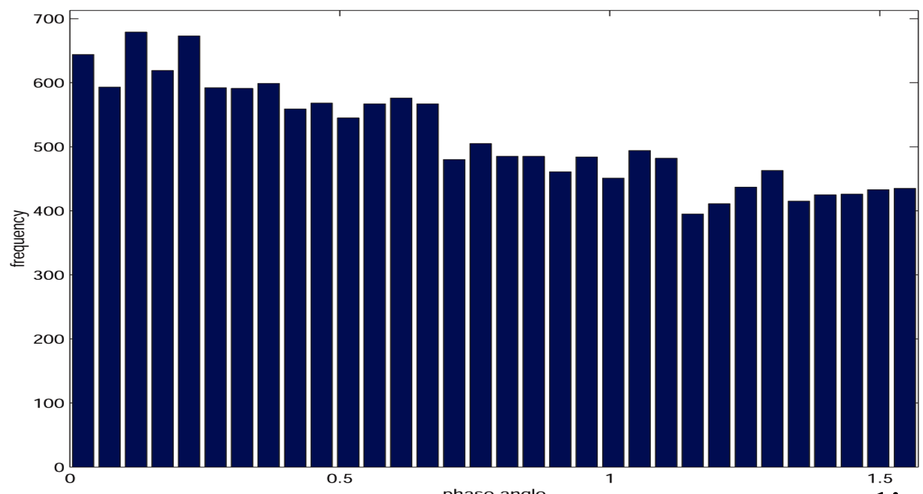

step line

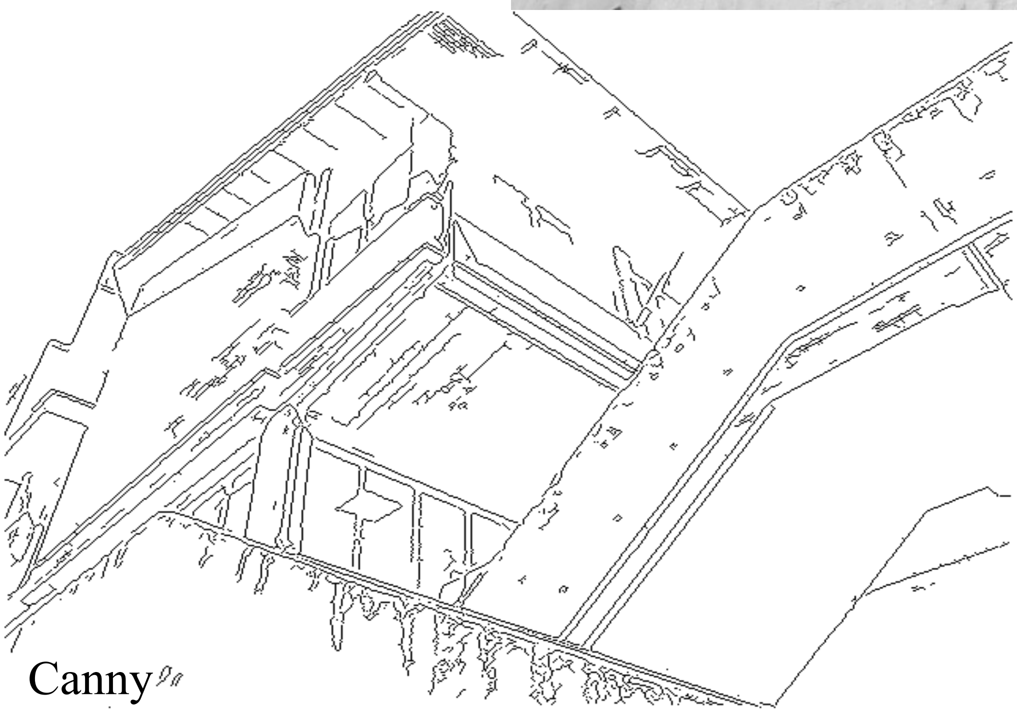

Canny

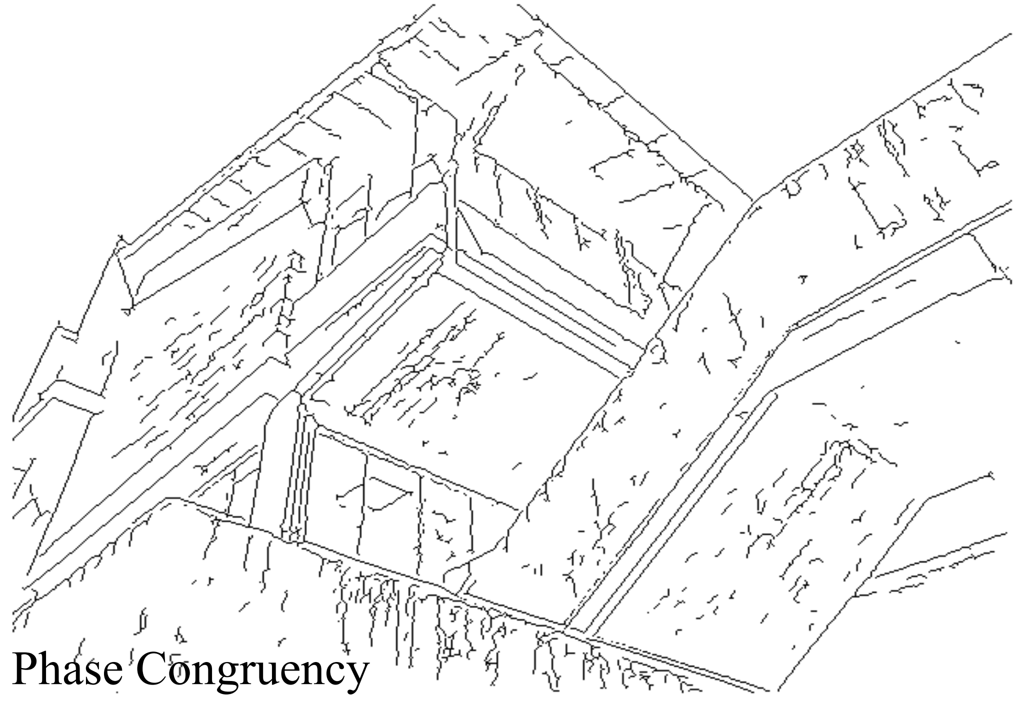

Phase Congruency

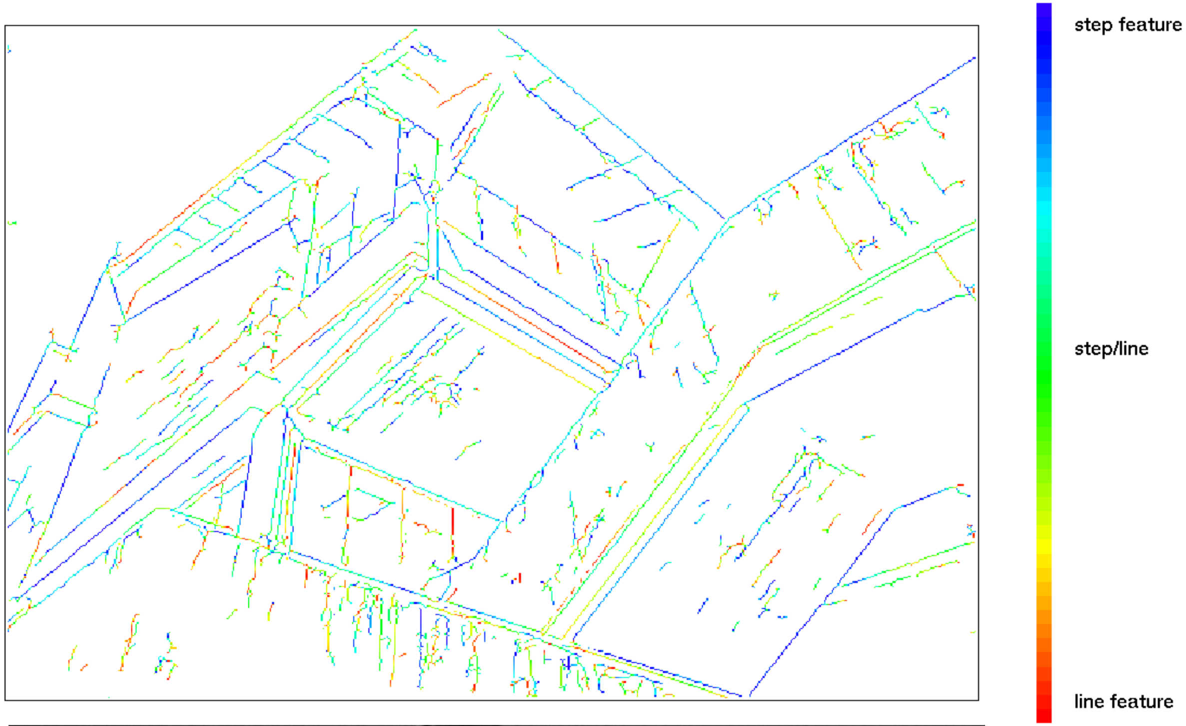

histogram of feature types

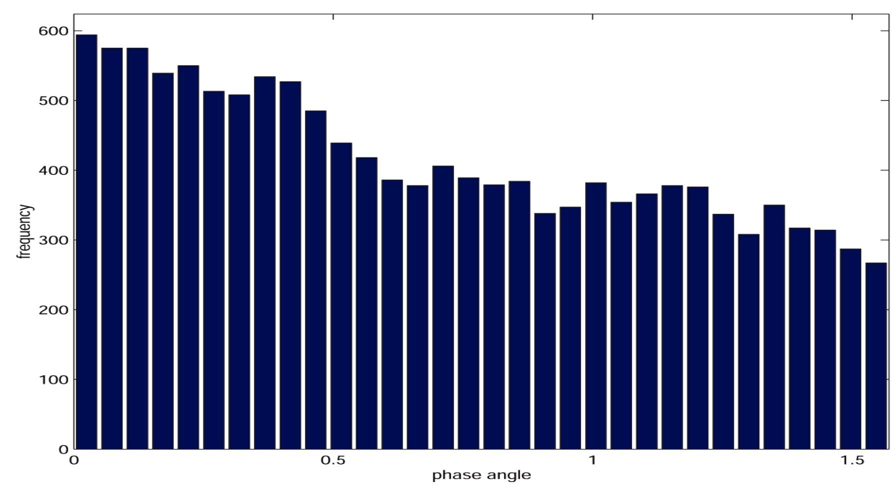

step

line

Phase Symmetry

- At a point of symmetry the Fourier components are at a maxima or minima (at the symmetric points of their cycle)

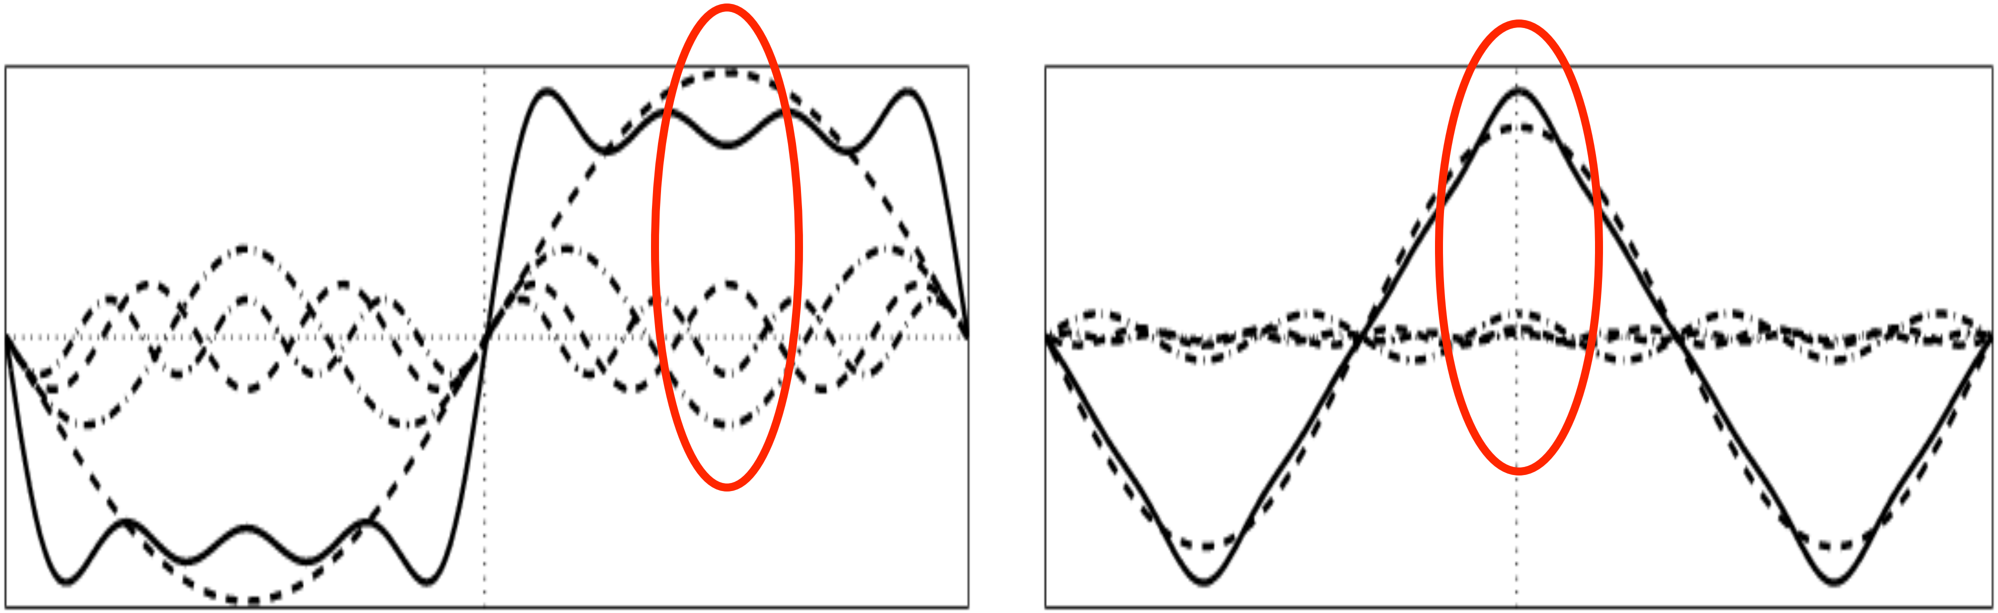

- At a point of asymmetry the Fourier components are at the asymmetric points of their cycle.

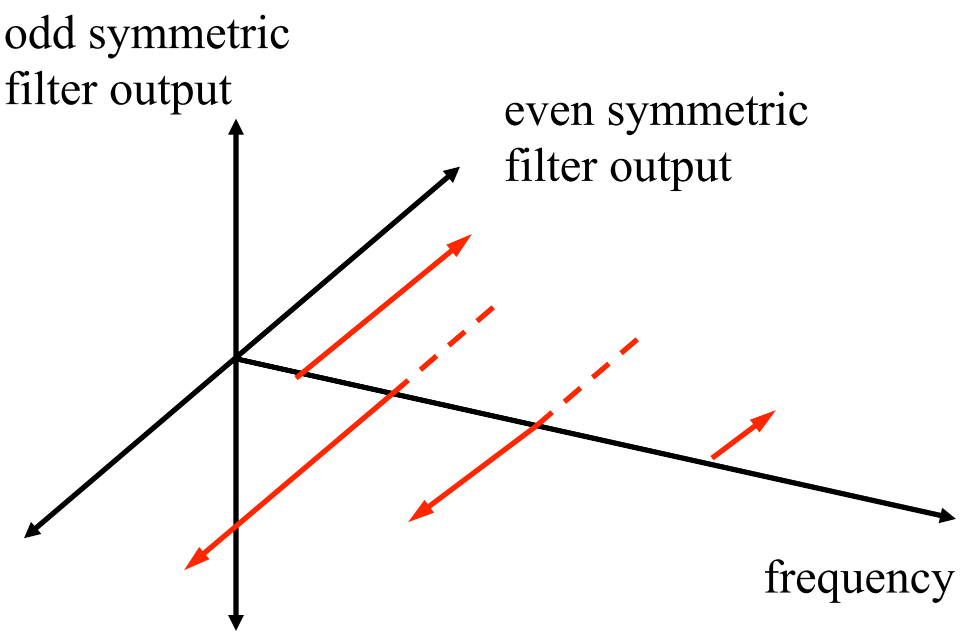

Local phase pattern at a point of symmetry

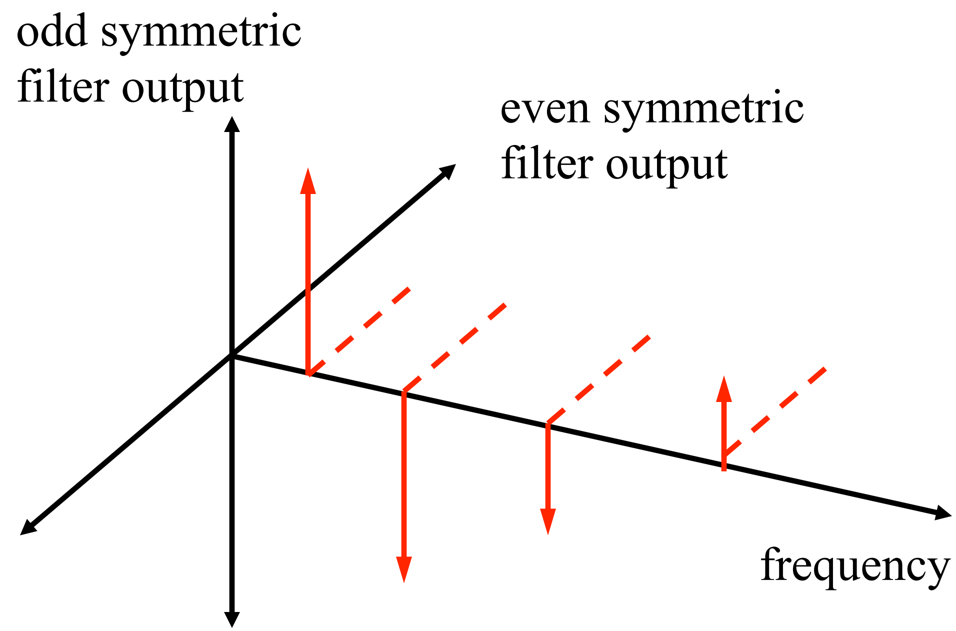

Local phase pattern at a point of asymmetry

$$\text{Phase Symmetry} = \frac{\sum_n A_n\left(\left|\cos(\phi_n(x)\right| - \left|\sin(\phi_n(x)\right|\right)}{\sum_n A_n(x)}$$

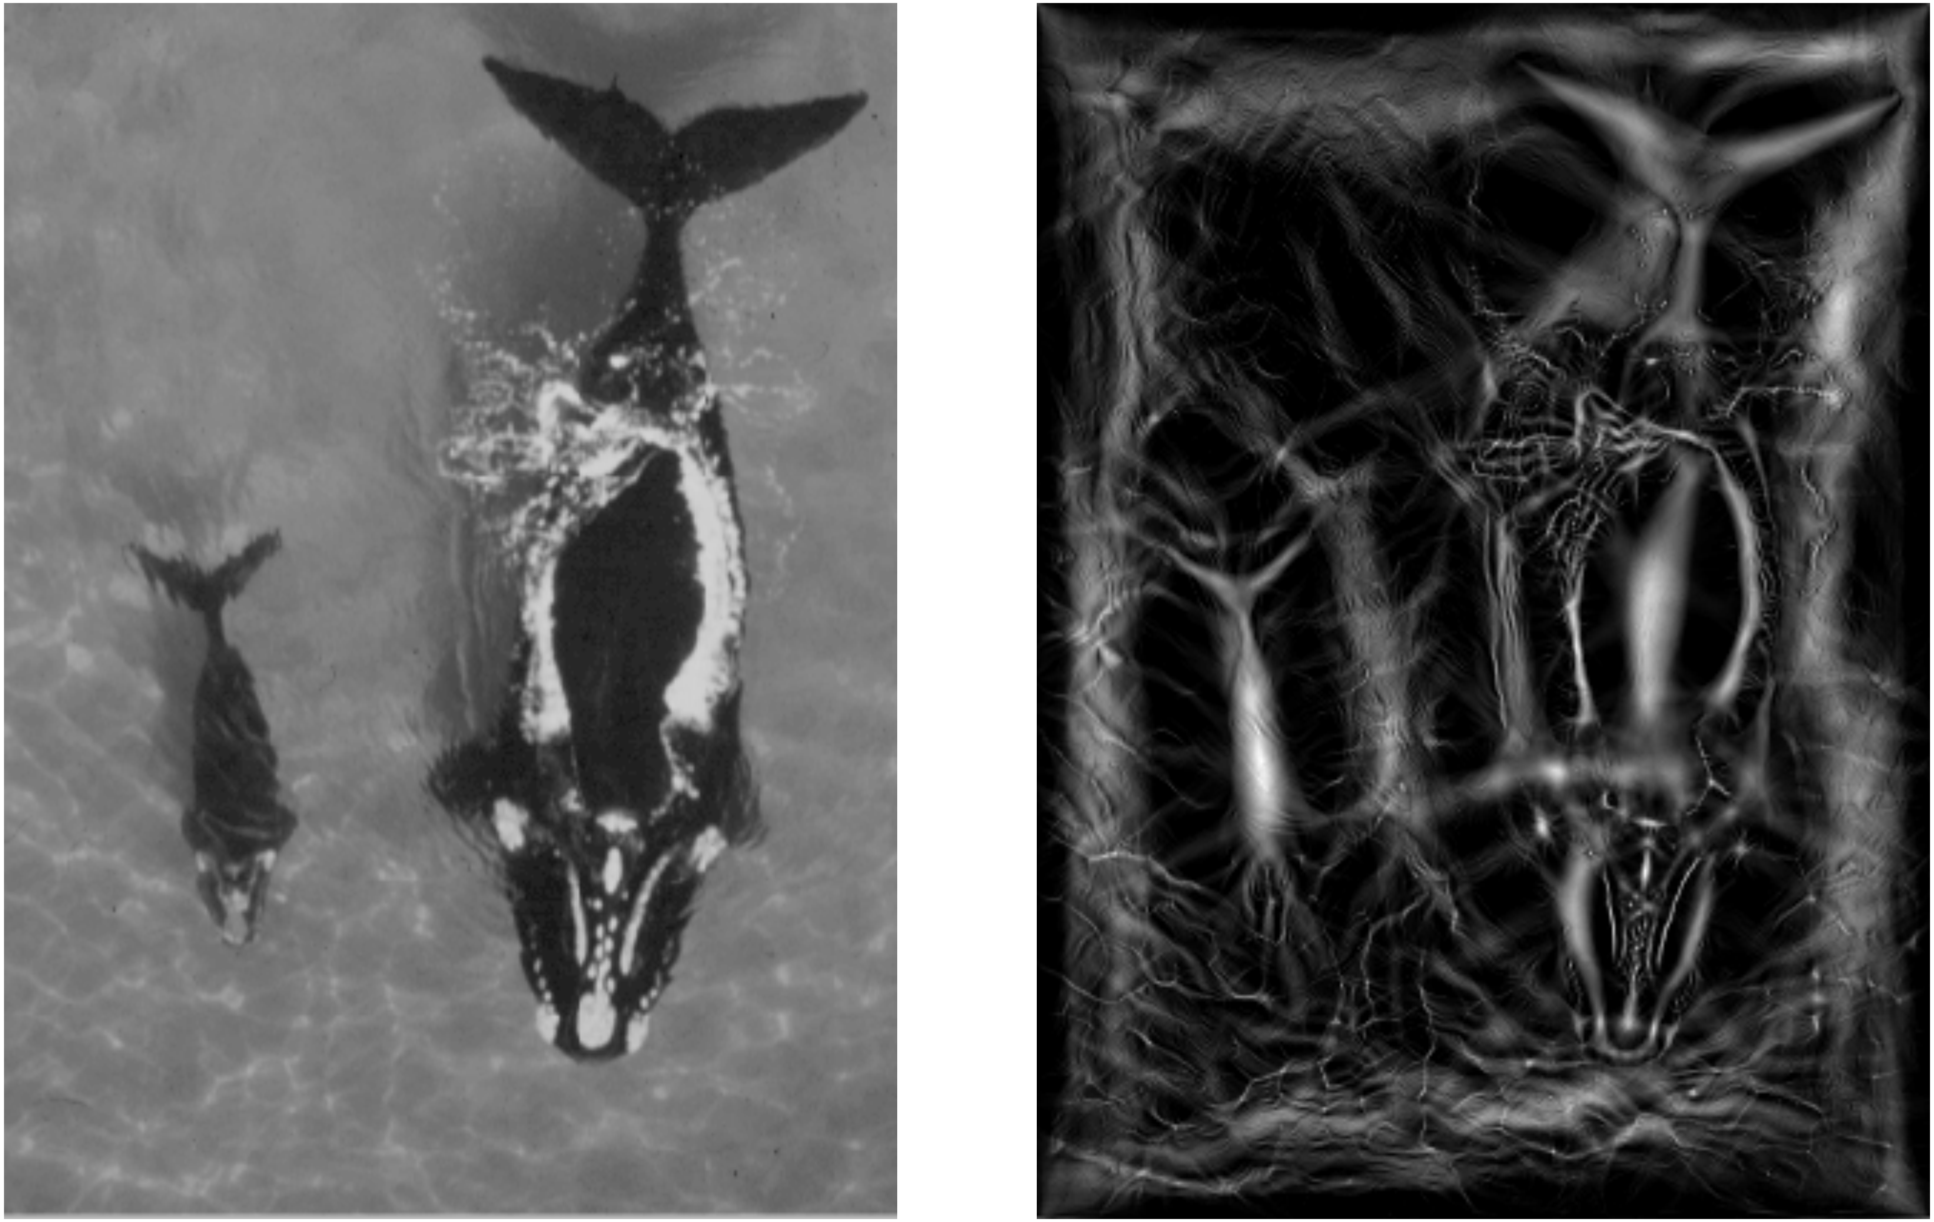

Measures symmetry independently of image contrast

Detection of Regions of Magnetic Discontinuity

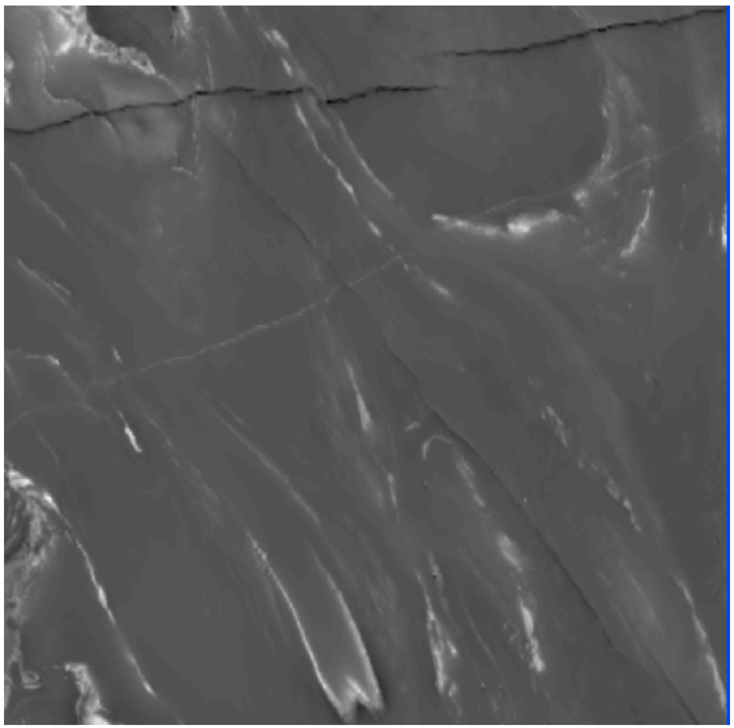

Aeromagnetic RTP image

(Data belongs to Fugro Airborne Surveys Pty Ltd.)

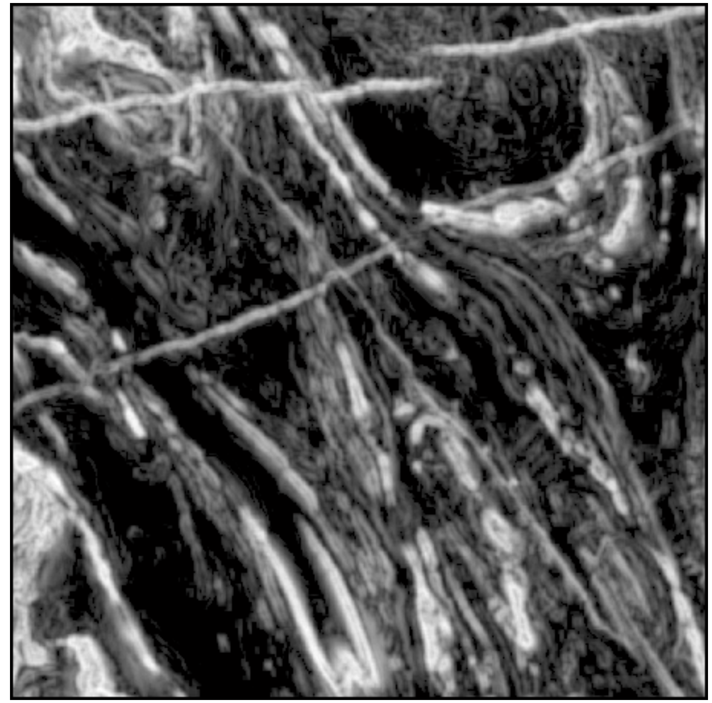

Entropy image

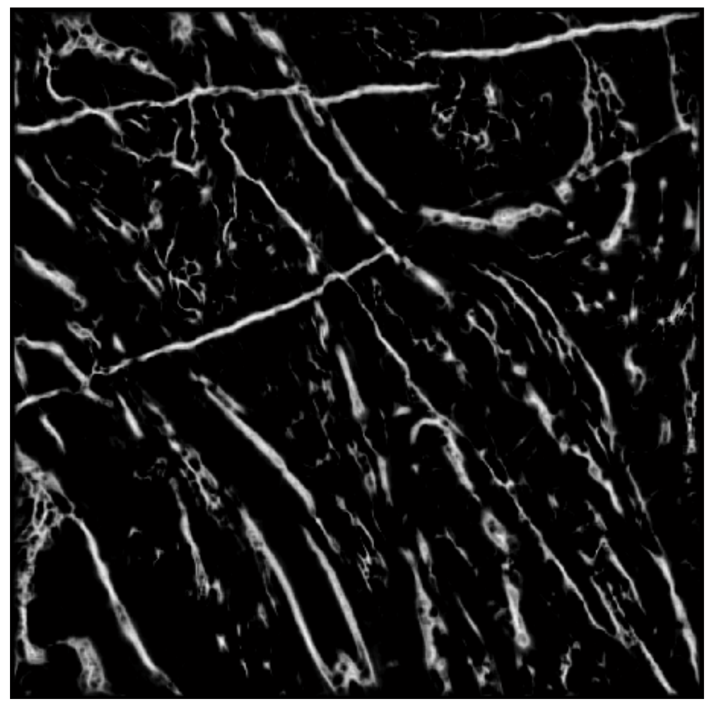

Phase Symmetry

(Only marking +ve features)

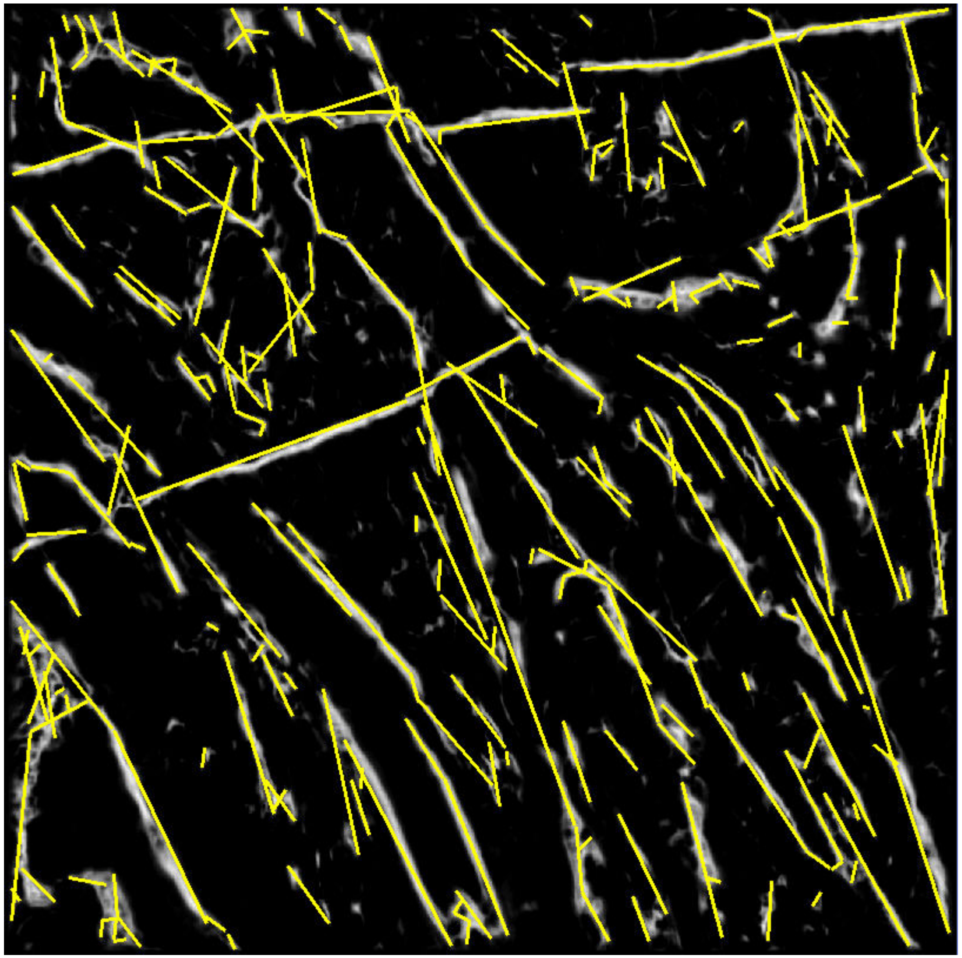

Fitted line segments

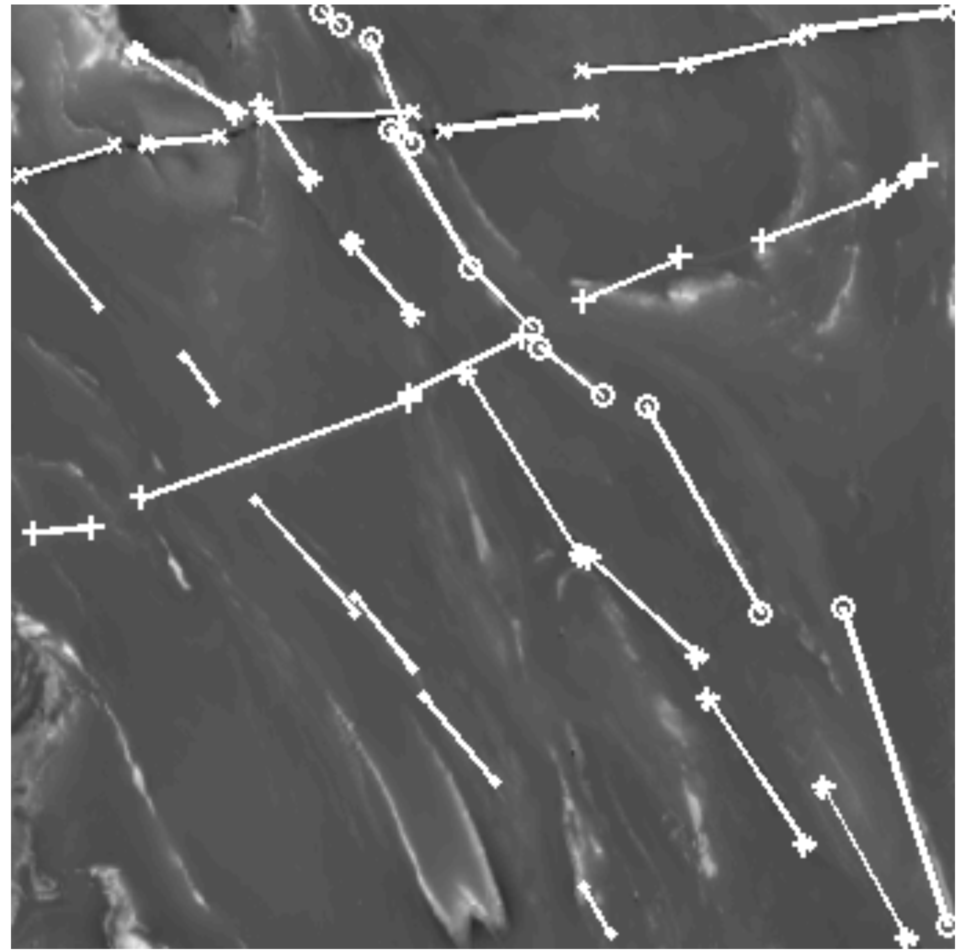

Line segments associated with major linear structures

Phase Based Feature Descriptors ?

Lowe's SIFT descriptor based on histograms of gradient orientations has been very successful, and also influential in the design of other descriptors.

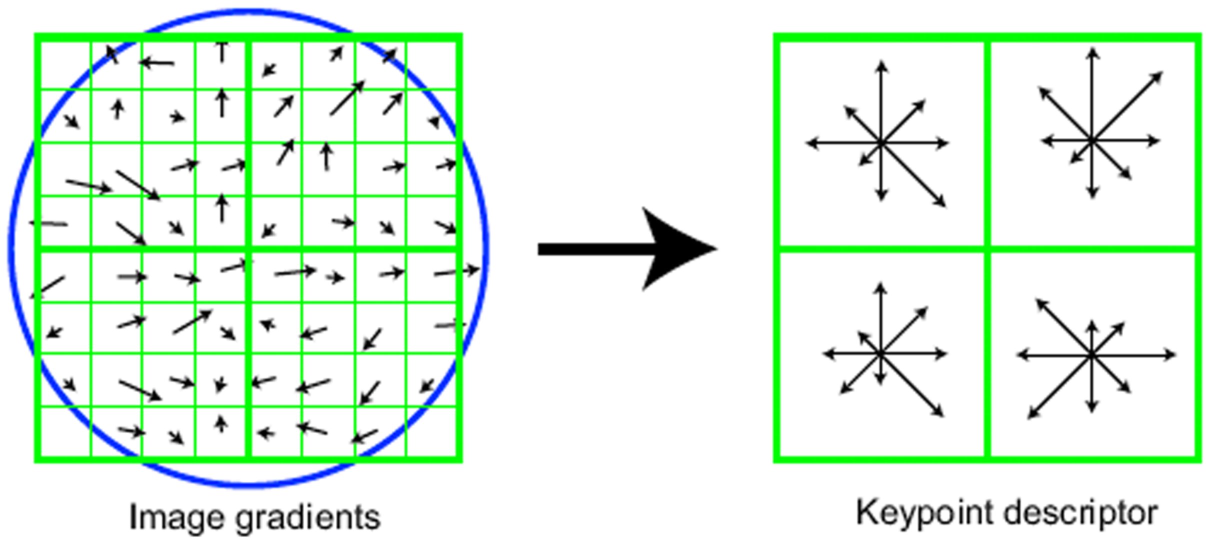

(Lowe 1999, 2004)

Why not build a descriptor based on phase angle and orientation…

Monogenic Filters

Allow efficient calculation of phase ϕ

and orientation θ

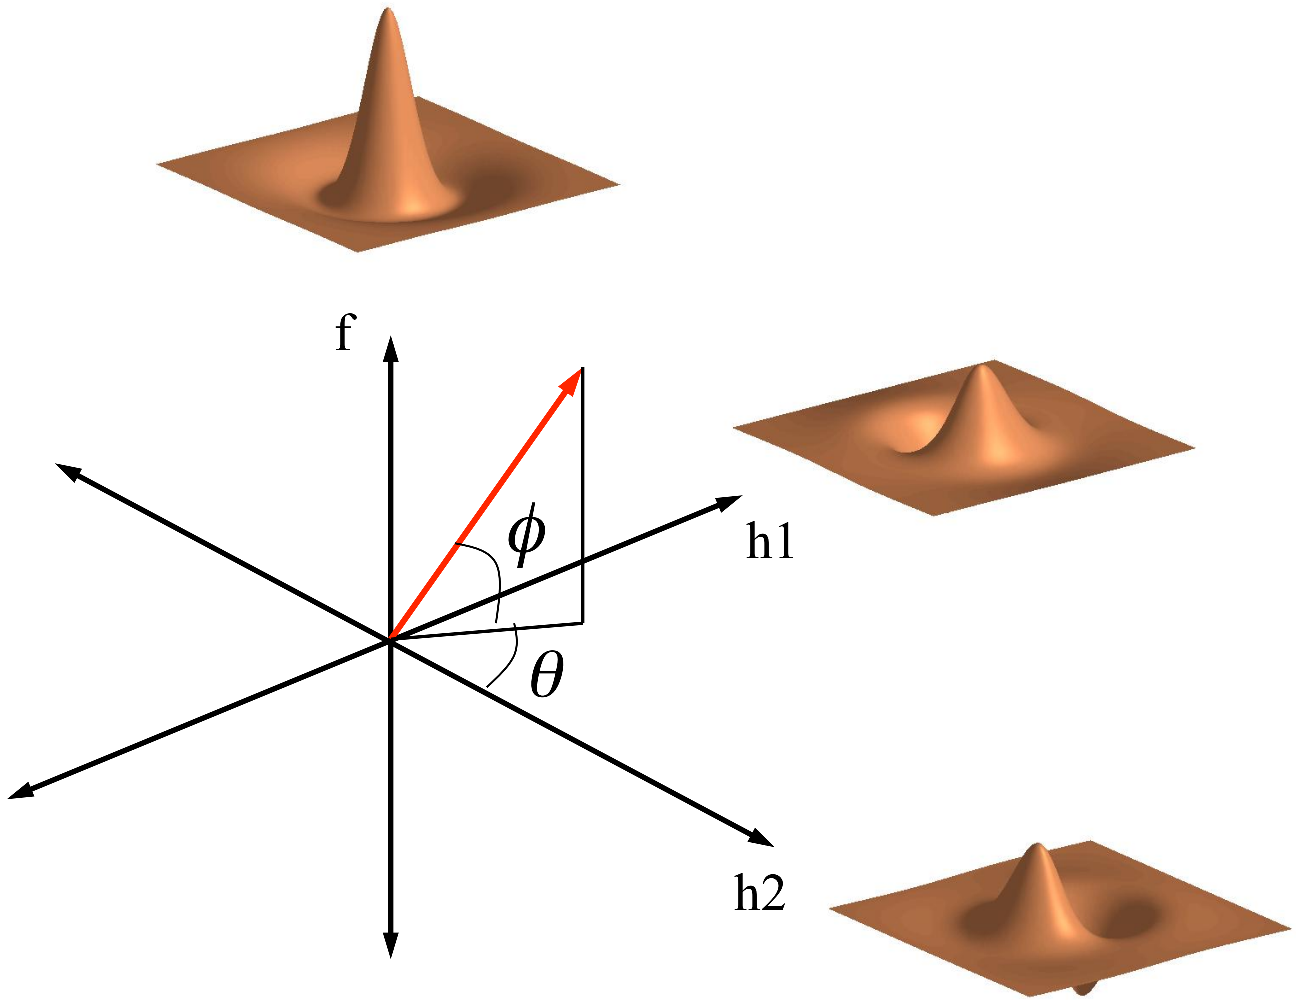

(Michael Felsberg and Gerald Sommer 2000, 2001)

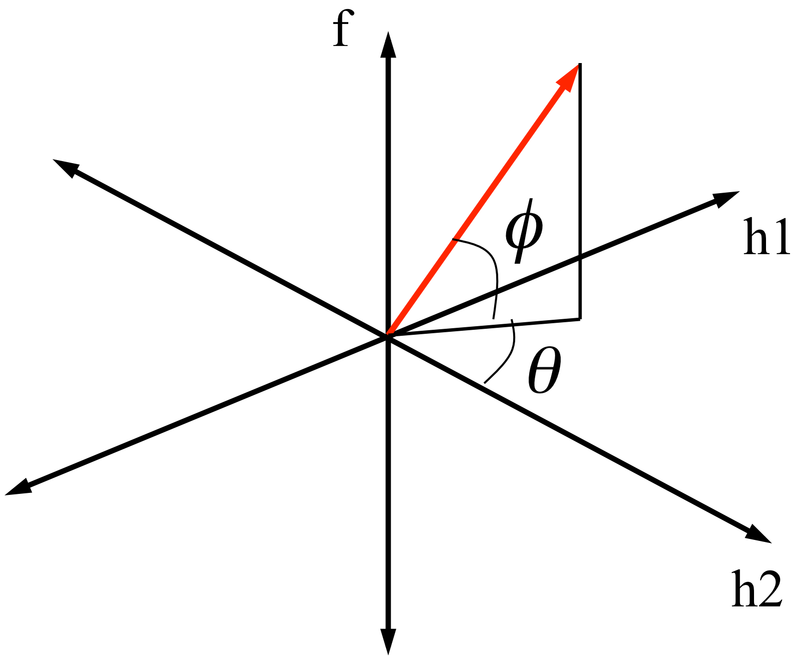

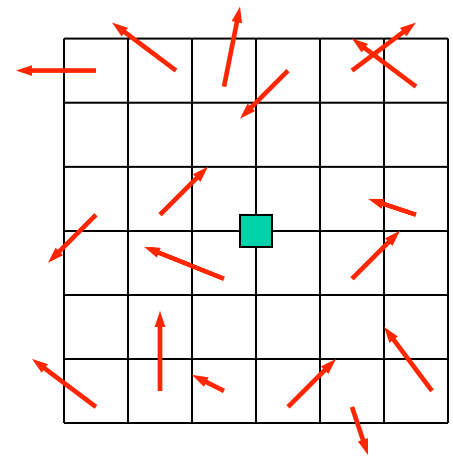

Descriptor can be constructed from phase and orientation values in a region surrounding the feature point.

Good results have been obtained with phase and orientation quantized to just 8 quadrants. This can be encoded with 3 bits allowing very efficient comparison.

Similarity from (local) phase mutual information

(Zhang & Brady, MICCAI 2007)

Pre-Treatment Image with Manually Identified control points

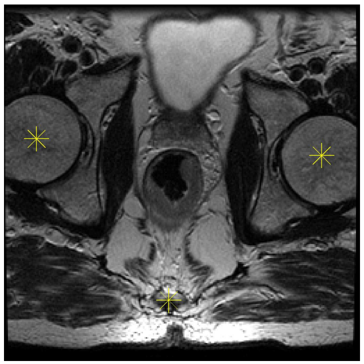

Local Phase (Pre)

Warped Post-Treatment Image

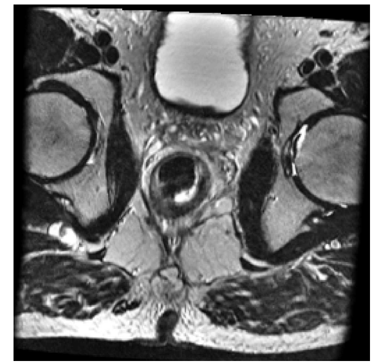

Post-Treatment Image with Manually Identified control points

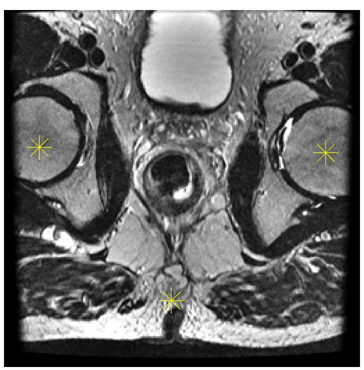

Local Phase (Post)

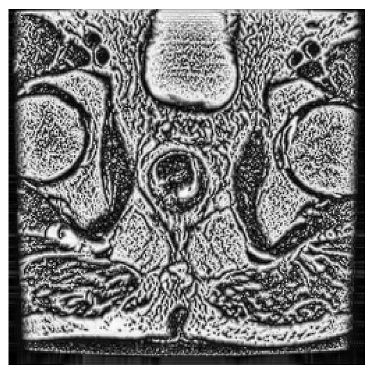

Grid

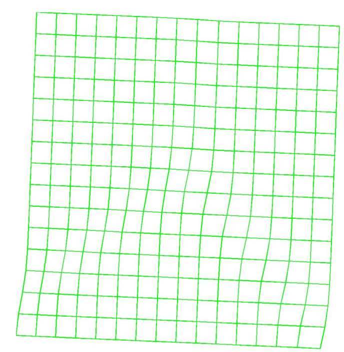

Intensity PDF vs Phase PDF

Pre-Treatment

Intensity PDF: pre

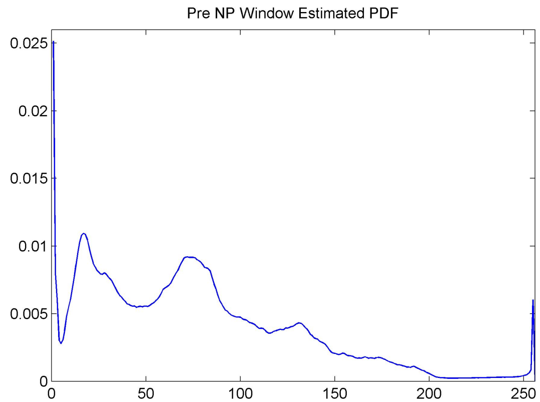

Local Phase PDF: pre

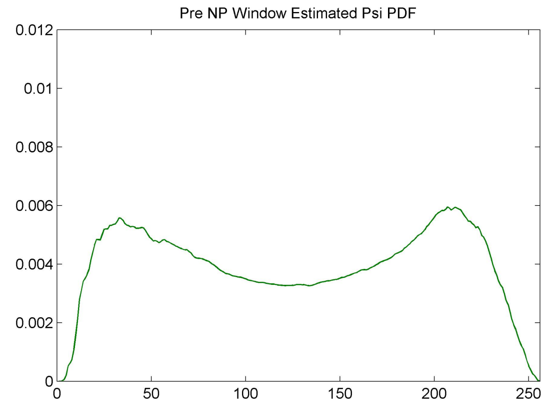

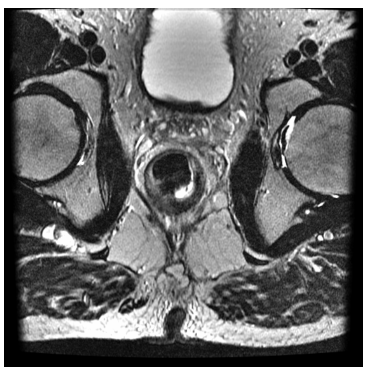

Post-Treatment

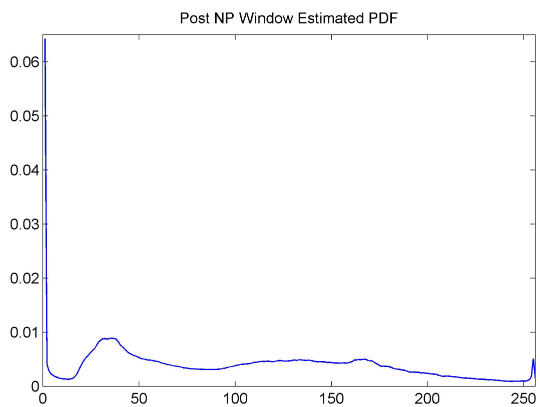

Intensity PDF: post

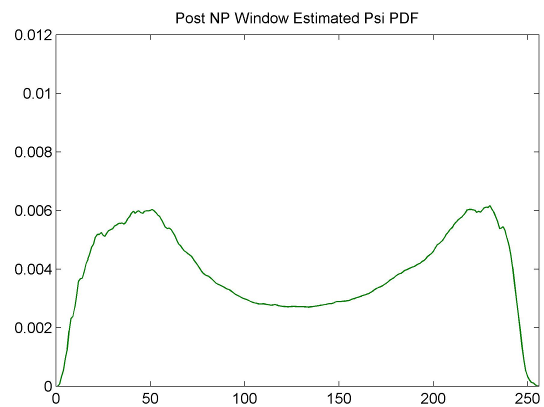

Local Phase PDF: post

How can we use edges?

How can we match edges?

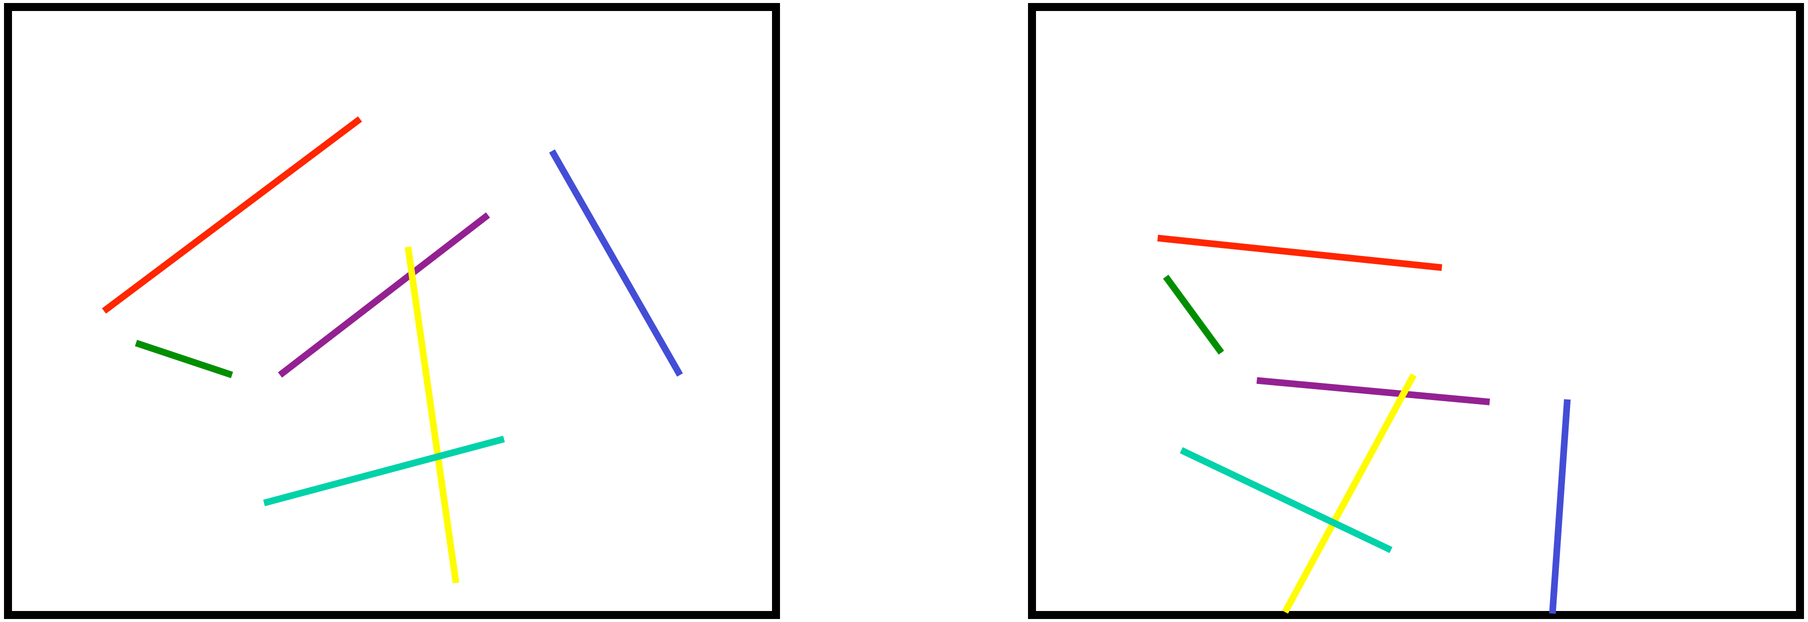

Very little literature on these topics…

What to do with Edges?

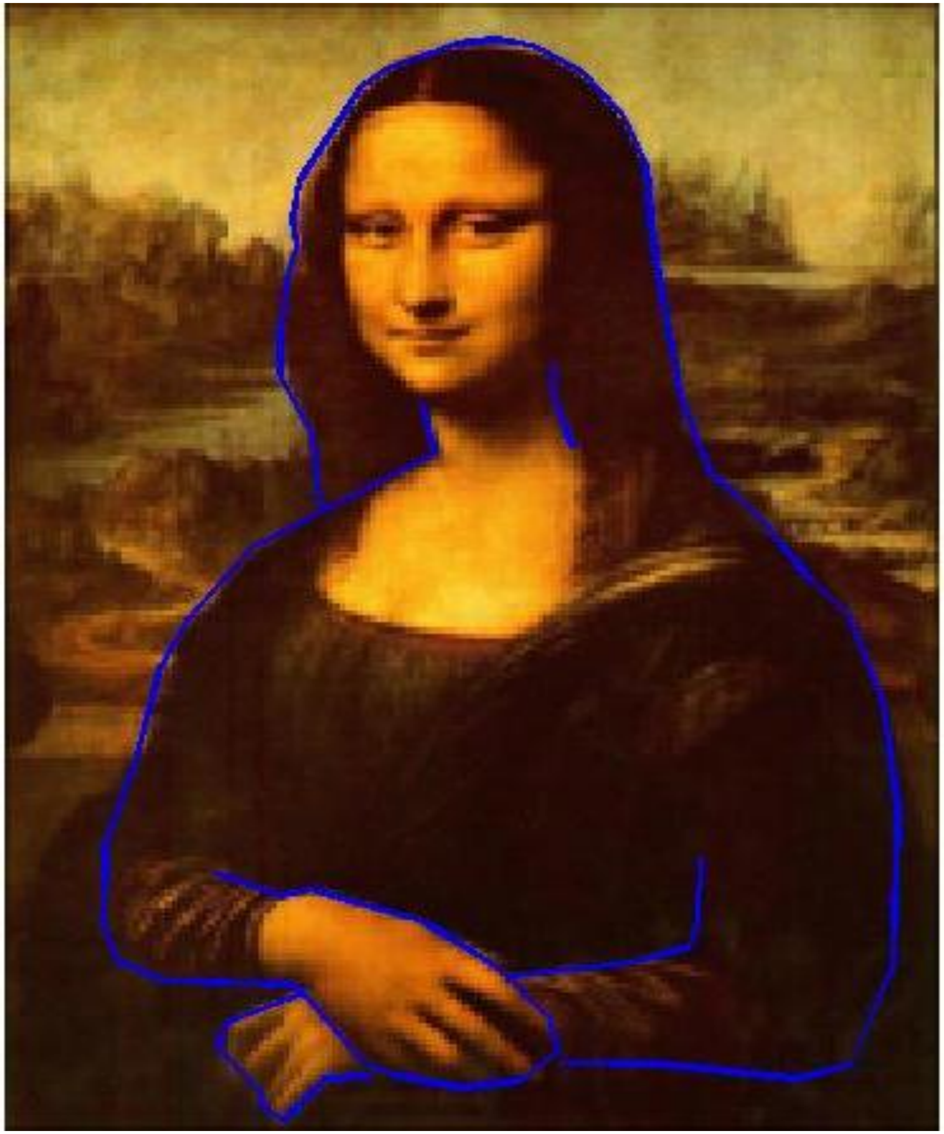

Manually marked contours

Shape from Occluding Contours

Surface normals at occluding contours are perpendicular to viewing direction.

This is a powerful cue to shape.

Shapelets Correlated with Surface Normals Produce Surfaces, ICCV 2005

A crude approximation of the surface normals in the scene…

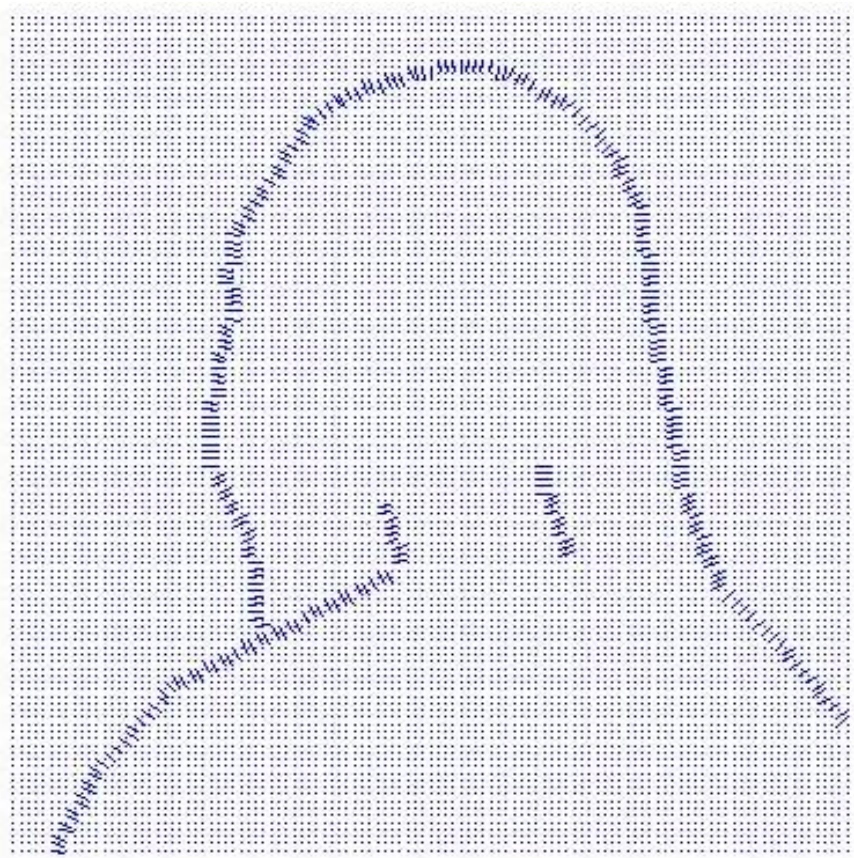

At all points slant is set to 0, except along contours where slant is set to $\pi/2$, and tilt is set to be perpendicular to contour.

This gradient field is not integrable. Project onto nearest integrable field using Shapelets or the Frankot Chellappa algorithm.

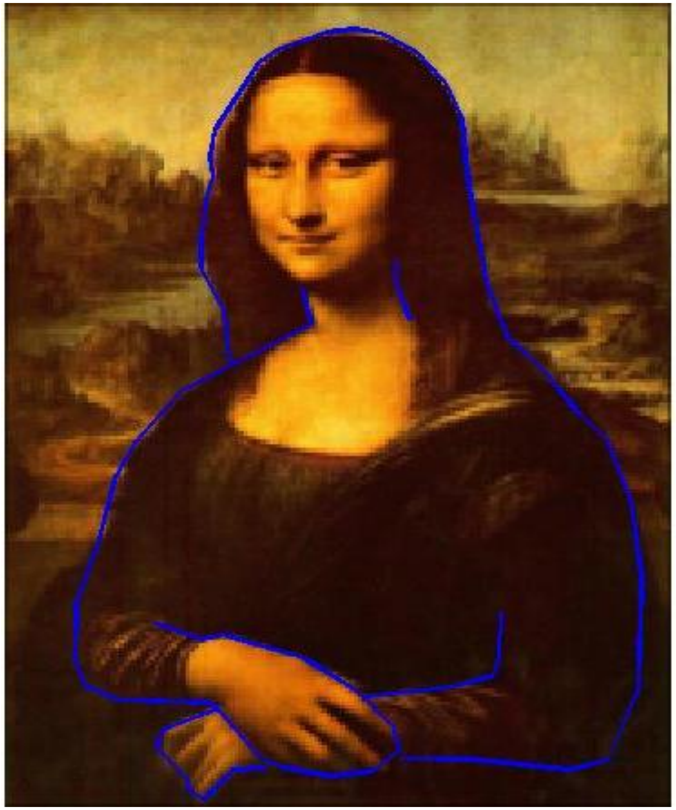

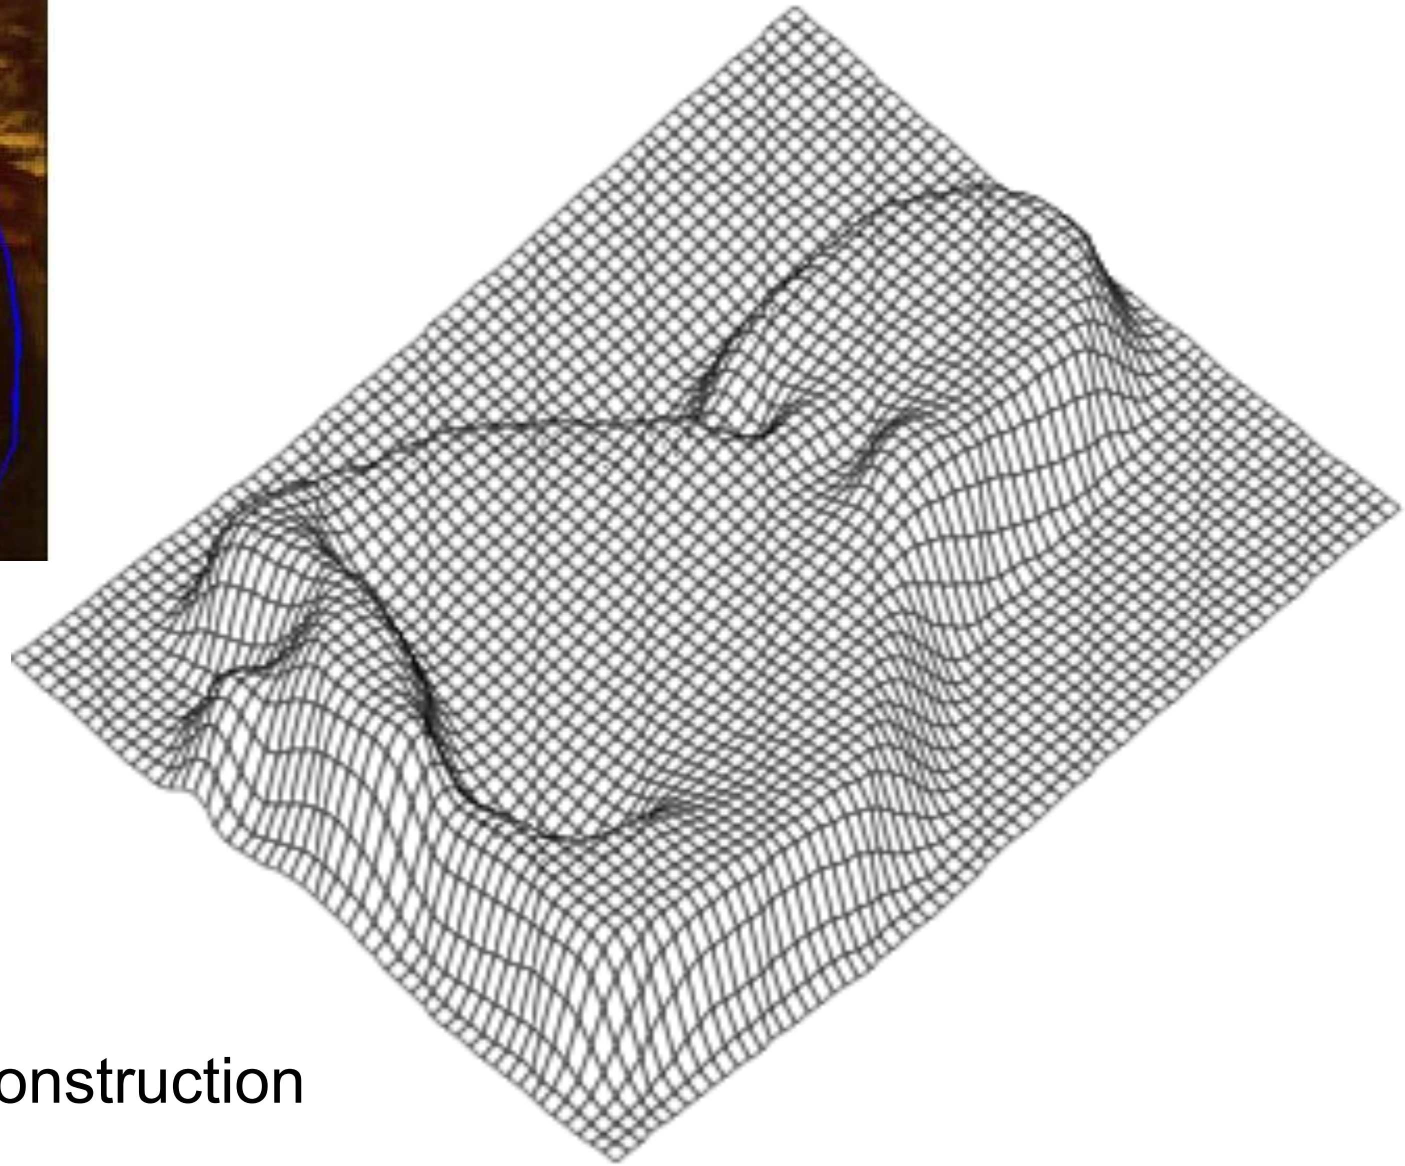

Shapelet reconstruction

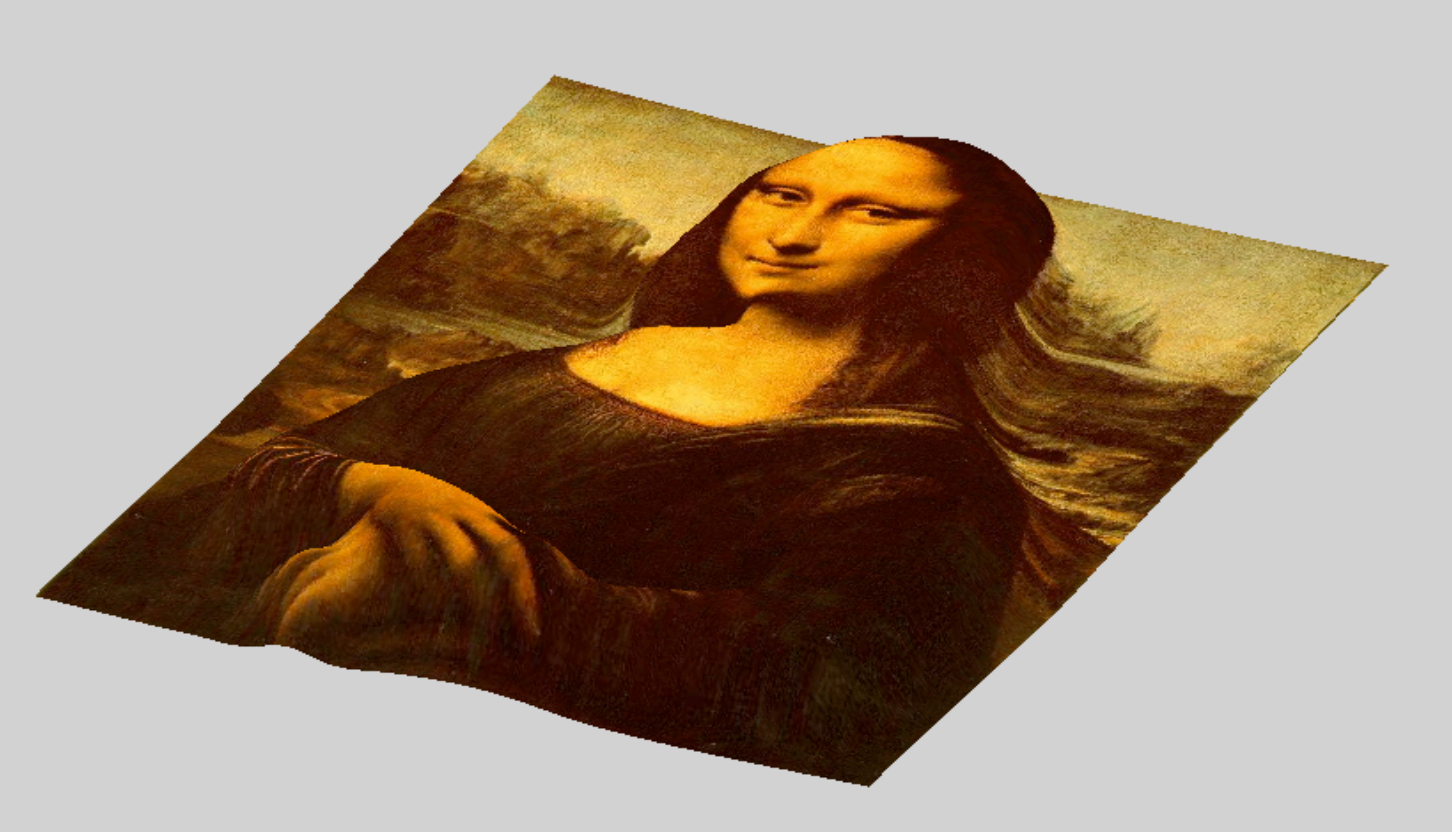

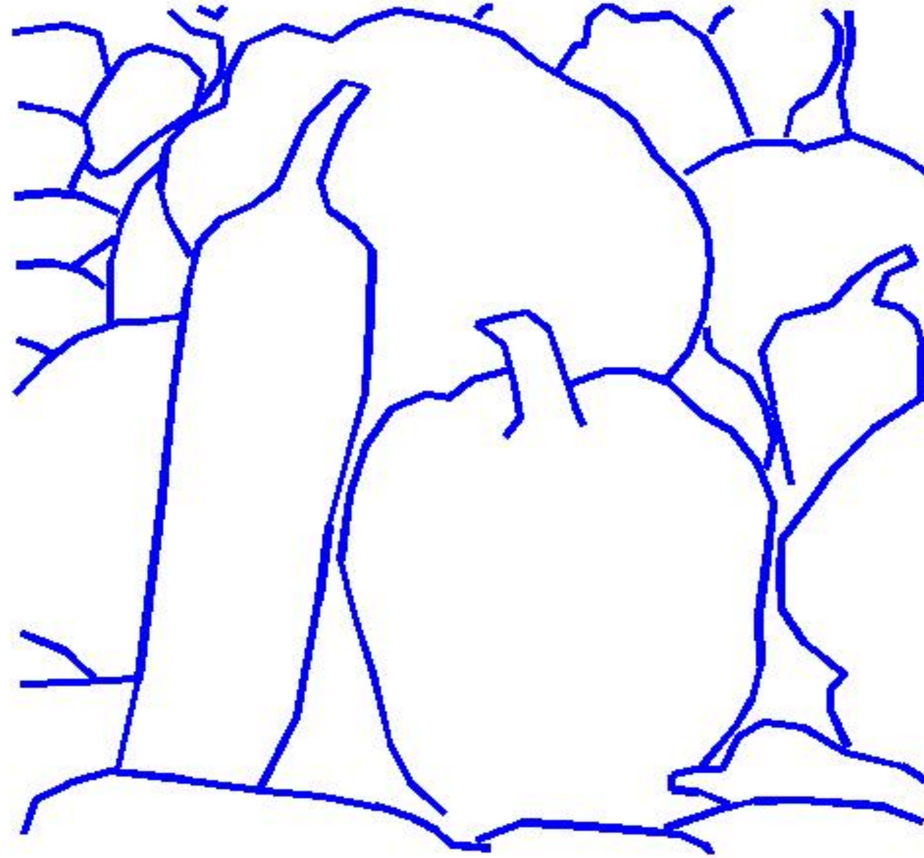

Manually drawn contours

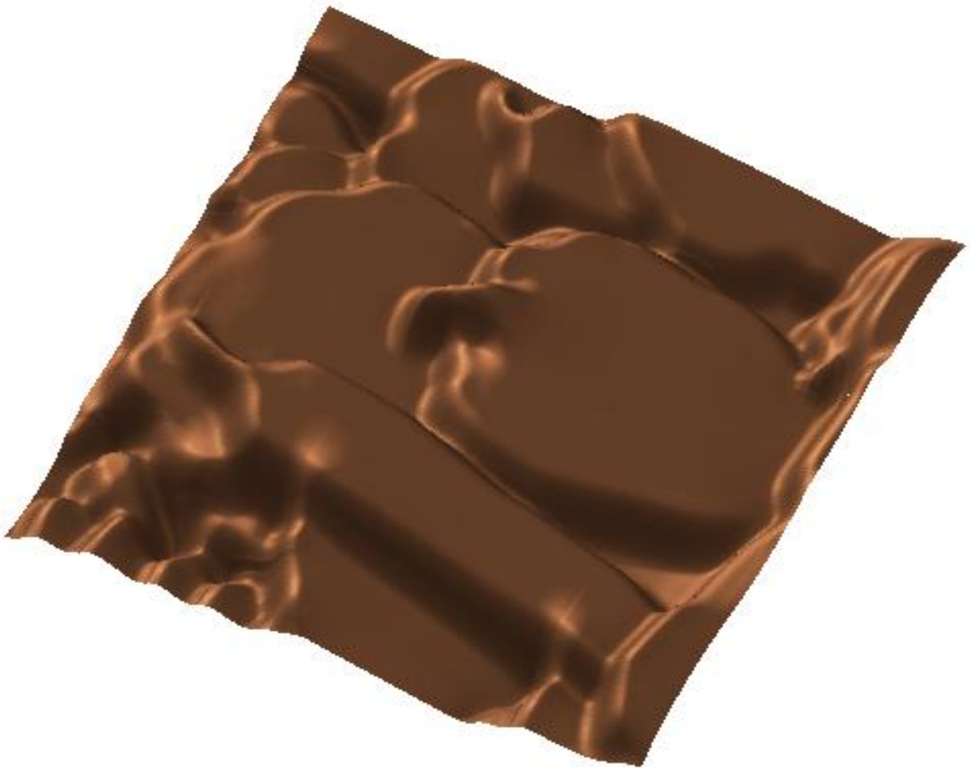


Conclusions

- Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.

Conclusions

- Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.

- Phase Congruency allows a wide range of feature types to be detected within the framework of a single model.

Conclusions

- Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.
- Phase Congruency allows a wide range of feature types to be detected within the framework of a single model.
- Phase Congruency is a dimensionless quantity, invariant to contrast. Features can be tracked over extended sequences more reliably.

Conclusions

- Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.
- Phase Congruency allows a wide range of feature types to be detected within the framework of a single model.
- Phase Congruency is a dimensionless quantity, invariant to contrast. Features can be tracked over extended sequences more reliably.
- The Phase Congruency corner map is a strict subset of the edge map. This simplifies the integration of data computed from edge and corner information.

Conclusions

- Gradient based operators are sensitive to illumination variations and do not localize accurately or consistently.

- Phase Congruency allows a wide range of feature types to be detected within the framework of a single model.

- Phase Congruency is a dimensionless quantity, invariant to contrast. Features can be tracked over extended sequences more reliably.

- The Phase Congruency corner map is a strict subset of the edge map. This simplifies the integration of data computed from edge and corner information.

- Phase is an underused local image attribute that can have many applications.