

# Edge and Texture

CS 554 – Computer Vision

Pinar Duygulu

Bilkent University

# Filters for features

- Previously, thinking of filtering as a way to remove or reduce **noise**
- Now, consider how filters will allow us to abstract higher-level “**features**”.
  - Map raw pixels to an intermediate representation that will be used for subsequent processing
  - Goal: reduce amount of data, discard redundancy, preserve what’s useful



# Edge detection

- **Goal:** Identify sudden changes (discontinuities) in an image
  - Intuitively, most semantic and shape information from the image can be encoded in the edges
  - More compact than pixels
- **Ideal:** artist's line drawing (but artist is also using object-level knowledge)



# Edge detection

- **Goal:** map image from 2d array of pixels to a set of curves or line segments or contours.
- **Why?**




Figure from J. Shotton et al., PAMI 2007


- **Main idea:** look for strong gradients, post-process

# Why do we care about edges?


- Extract information, recognize objects



- Recover geometry and viewpoint



# Origin of Edges




- Edges are caused by a variety of factors

# What can cause an edge?

Reflectance change:  
appearance  
information, texture


Change in surface  
orientation: shape



Depth discontinuity:  
object boundary

Cast shadows

# Contrast and invariance






# Closeup of edges



# Closeup of edges



# Closeup of edges






# Closeup of edges



# Characterizing edges

- An edge is a place of rapid change in the image intensity function



# Differentiation and convolution

For 2D function,  $f(x,y)$ , the partial derivative is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:

$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x + 1, y) - f(x, y)}{1}$$

To implement above as convolution, what would be the associated filter?



# Partial derivatives of an image



$$\frac{\partial f(x, y)}{\partial x}$$

$$\frac{\partial f(x, y)}{\partial y}$$

-1	1
----	---



-1	?	1
1	or	-1

Which shows changes with respect to x?

(showing flipped filters)


# Assorted finite difference filters

**Prewitt:**  $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

**Sobel:**  $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

**Roberts:**  $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$  ;  $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

```
>> My = fspecial('sobel');  
>> outim = imfilter(double(im), My);  
>> imagesc(outim);  
>> colormap gray;
```








# Image gradient

The gradient of an image:

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

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






The gradient direction (orientation of edge normal) is given by:

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

The *edge strength* is given by the gradient magnitude


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

# Intensity




Source: D. Hoiem

# With a little Gaussian noise





Source: Hays, Brown



Source: D. Hoiem

# Effects of noise

- Consider a single row or column of the image
  - Plotting intensity as a function of position gives a signal

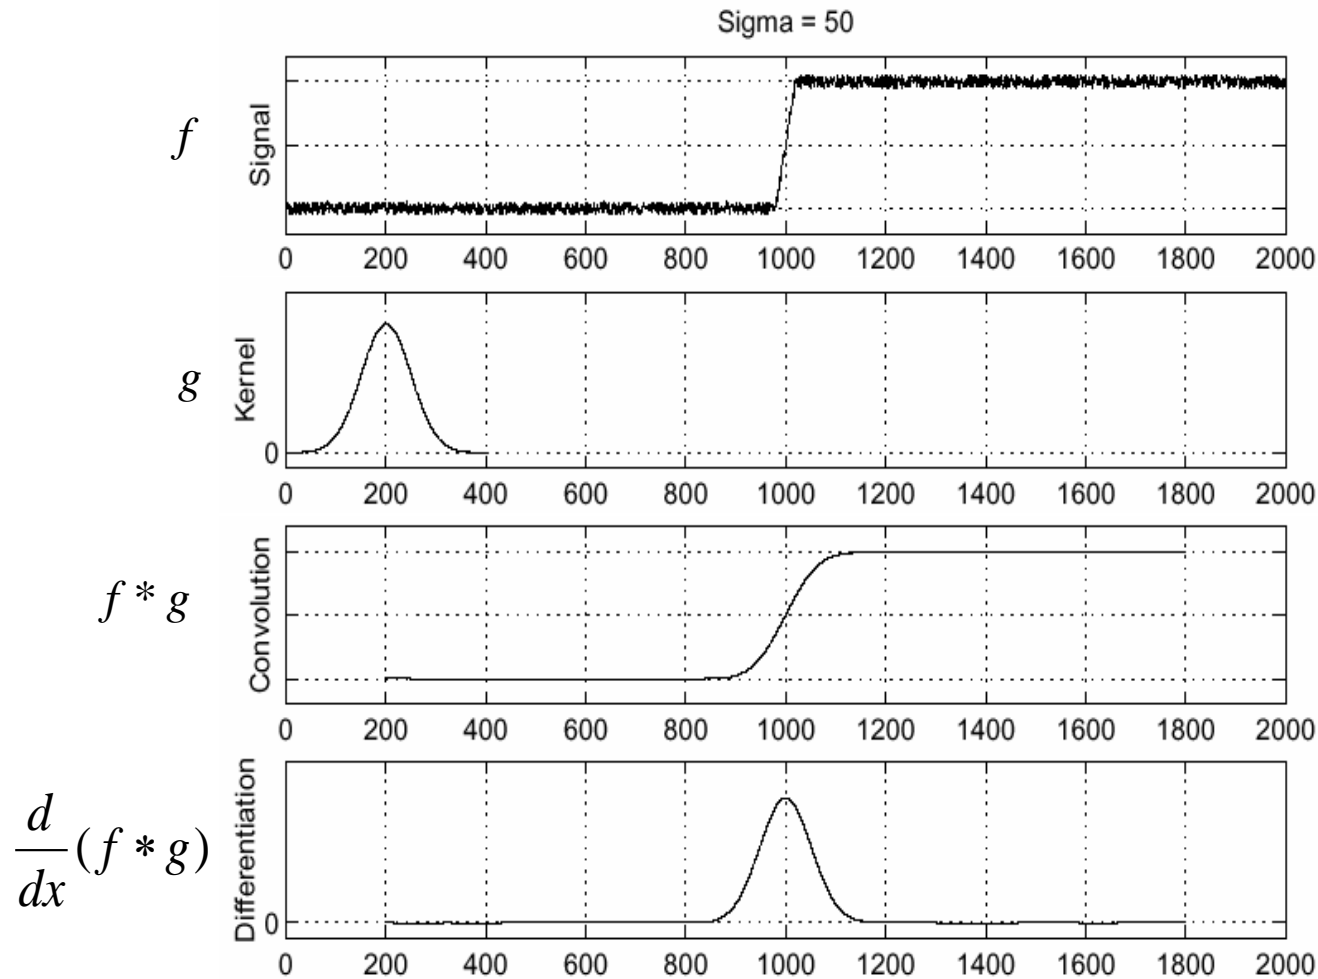


Where is the edge?

# Effects of noise

- Difference filters respond strongly to noise
  - Image noise results in pixels that look very different from their neighbors
  - Generally, the larger the noise the stronger the response
- What can we do about it?

# Solution: smooth first




- To find edges, look for peaks in  $\frac{d}{dx}(f * g)$


# Derivative theorem of convolution

- Differentiation is convolution, and convolution is associative:


- This saves us



# Derivative of Gaussian filter




$$* [1 \ -1] =$$








# Derivative of Gaussian filters



x-direction




y-direction




# Laplacian of Gaussian

Consider  $\frac{\partial^2}{\partial x^2}(h \star f)$


$f$



$\frac{\partial^2}{\partial x^2}h$




$(\frac{\partial^2}{\partial x^2}h) \star f$



Where is the edge?


Zero-crossings of bottom graph

# 2D edge detection filters




Gaussian

$$h_{\sigma}(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$



derivative of Gaussian

$$\frac{\partial}{\partial x} h_{\sigma}(u, v)$$



Laplacian of Gaussian



$$\nabla^2 h_{\sigma}(u, v)$$

- $\nabla^2$  is the Laplacian operator:


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

# Smoothing with a Gaussian

Recall: parameter  $\sigma$  is the “scale” / “width” / “spread” of the Gaussian kernel, and controls the amount of smoothing.




...



# Effect of $\sigma$ on derivatives



$\sigma = 1$  pixel



$\sigma = 3$  pixels

The apparent structures differ depending on Gaussian's scale parameter.

Larger values: larger scale edges detected  
Smaller values: finer features detected



# So, what scale to choose?

It depends what we're looking for.



Too fine of a scale...can't see the forest for the trees.

Too coarse of a scale...can't tell the maple grain from the cherry.

# Thresholding


- Choose a threshold value  $t$
- Set any pixels less than  $t$  to zero (off)
- Set any pixels greater than or equal to  $t$  to one (on)

# Original image





# Gradient magnitude image




# Thresholding gradient with a lower threshold



# Thresholding gradient with a higher threshold



## Tradeoff between smoothing and localization



1 pixel

3 pixels

7 pixels

- Smoothed derivative removes noise, but blurs edge. Also finds edges at different “scales”.

# Designing an edge detector

- Criteria for a good edge detector:
  - **Good detection:** the optimal detector should find all real edges, ignoring noise or other artifacts
  - **Good localization**
    - the edges detected must be as close as possible to the true edges
    - the detector must return one point only for each true edge point
- Cues of edge detection
  - Differences in color, intensity, or texture across the boundary
  - Continuity and closure
  - High-level knowledge

# Canny edge detector

- This is probably the most widely used edge detector in computer vision
- Theoretical model: step-edges corrupted by additive Gaussian noise
- Canny has shown that the first derivative of the Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio* and localization


J. Canny, [A Computational Approach To Edge Detection](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.

# Example





original image (Lena)


# Derivative of Gaussian filter



x-direction



y-direction





# The Canny edge detector



original image (Lena)


# Compute Gradients (DoG)



X-Derivative of Gaussian



Y-Derivative of Gaussian



Gradient Magnitude

# The Canny edge detector





norm of the gradient

# The Canny edge detector





thresholding

# The Canny edge detector



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

# Non-maximum suppression



Check if pixel is local maximum along gradient direction, select single max across width of the edge  
– requires checking interpolated pixels  $p$  and  $r$


# Get Orientation at Each Pixel

- Threshold at minimum level




$$\text{theta} = \text{atan2}(\text{gy}, \text{gx})$$

# Non-maximum suppression for each orientation



At  $q$ , we have a maximum if the value is larger than those at both  $p$  and at  $r$ . Interpolate to get these values.






# Before Non-max Suppression



# After non-max suppression



# The Canny edge detector




Problem:  
pixels along  
this edge  
didn't survive  
the  
thresholding

thinning  
(non-maximum suppression)

# Hysteresis thresholding

- Check that maximum value of gradient value is sufficiently large
  - drop-outs? use **hysteresis**
    - use a high threshold to start edge curves and a low threshold to continue them.



# Hysteresis thresholding

- Threshold at low/high levels to get weak/strong edge pixels
- Do connected components, starting from strong edge pixels



# Hysteresis thresholding



original image



high threshold  
(strong edges)



low threshold  
(weak edges)



hysteresis threshold

# Final Canny Edges



# Canny edge detector

1. Filter image with x, y derivatives of Gaussian
  2. Find magnitude and orientation of gradient
  3. Non-maximum suppression:
    - Thin multi-pixel wide “ridges” down to single pixel width
  4. Thresholding and linking (hysteresis):
    - Define two thresholds: low and high
    - Use the high threshold to start edge curves and the low threshold to continue them
- MATLAB: `edge(image, 'canny')`



# Effect of $\sigma$ (Gaussian kernel spread/size)



original



Canny with  $\sigma = 1$




Canny with  $\sigma = 2$

The choice of  $\sigma$  depends on desired behavior

- large  $\sigma$  detects large scale edges
- small  $\sigma$  detects fine features

# Object boundaries vs. edges



Background

Texture




Shadows

# Edge detection is just the beginning...

image

human segmentation

gradient magnitude



Berkeley segmentation database:

<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/>



***Much more on segmentation later in term...***




# Representing Texture



# Texture and Material



# Texture and Orientation



# Texture and Scale



# What is texture?

Regular or stochastic patterns caused by bumps, grooves, and/or markings




# How can we represent texture?

- Compute responses of blobs and edges at various orientations and scales

# Overcomplete representation: filter banks



LM Filter Bank



Code for filter banks: [www.robots.ox.ac.uk/~vgg/research/texclass/filters.html](http://www.robots.ox.ac.uk/~vgg/research/texclass/filters.html)

# Filter banks

- Process image with each filter and keep




# How can we represent texture?

- Measure responses of blobs and edges at various orientations and scales
- Idea 1: Record simple statistics (e.g., mean, std.) of absolute filter responses


# Can you match the texture to the response?

Filters


response?




1



2




3



Mean abs responses

A



B





C



# Representing texture

- Idea 2: take vectors of filter responses at each pixel and cluster them, then take histograms.




# Building Visual Dictionaries

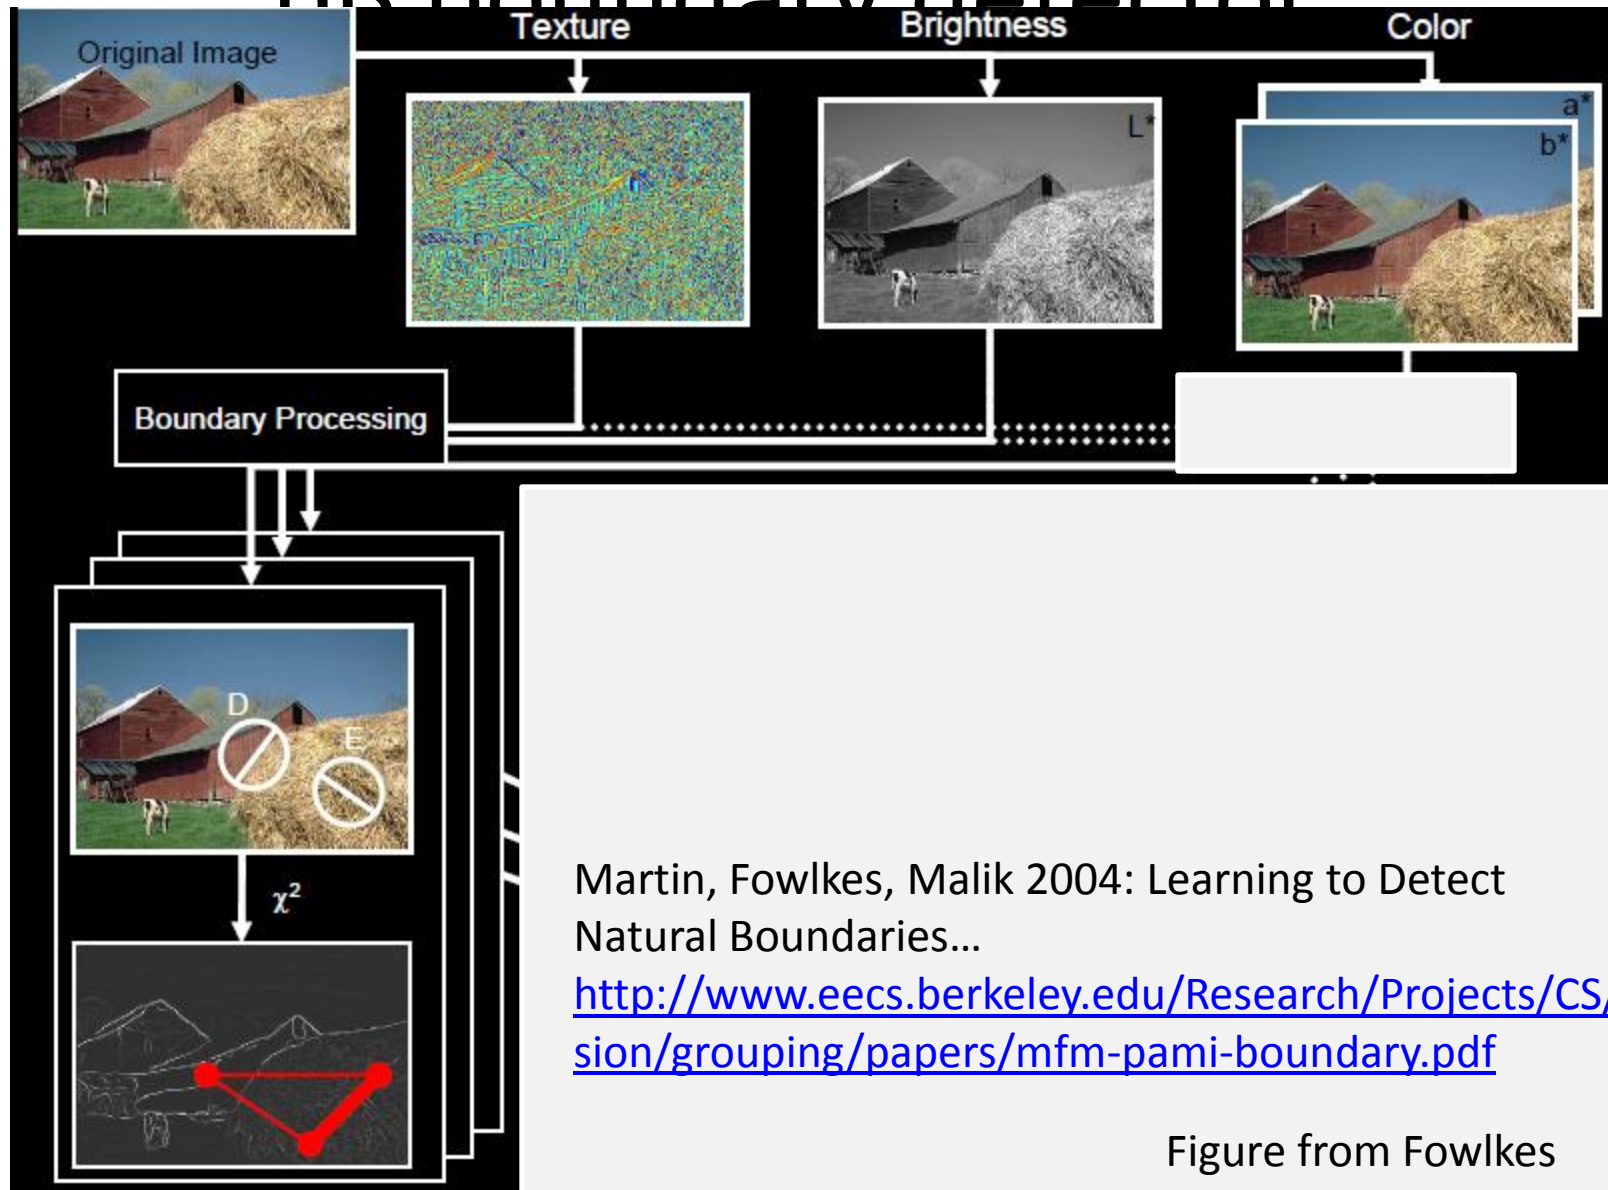
1. Sample patches from a database
  - E.g., 128 dimensional SIFT vectors



2. Cluster the patches
  - Cluster centers are the dictionary
3. Assign a codeword (number) to each new patch, according to the nearest cluster



# nB boundary detector



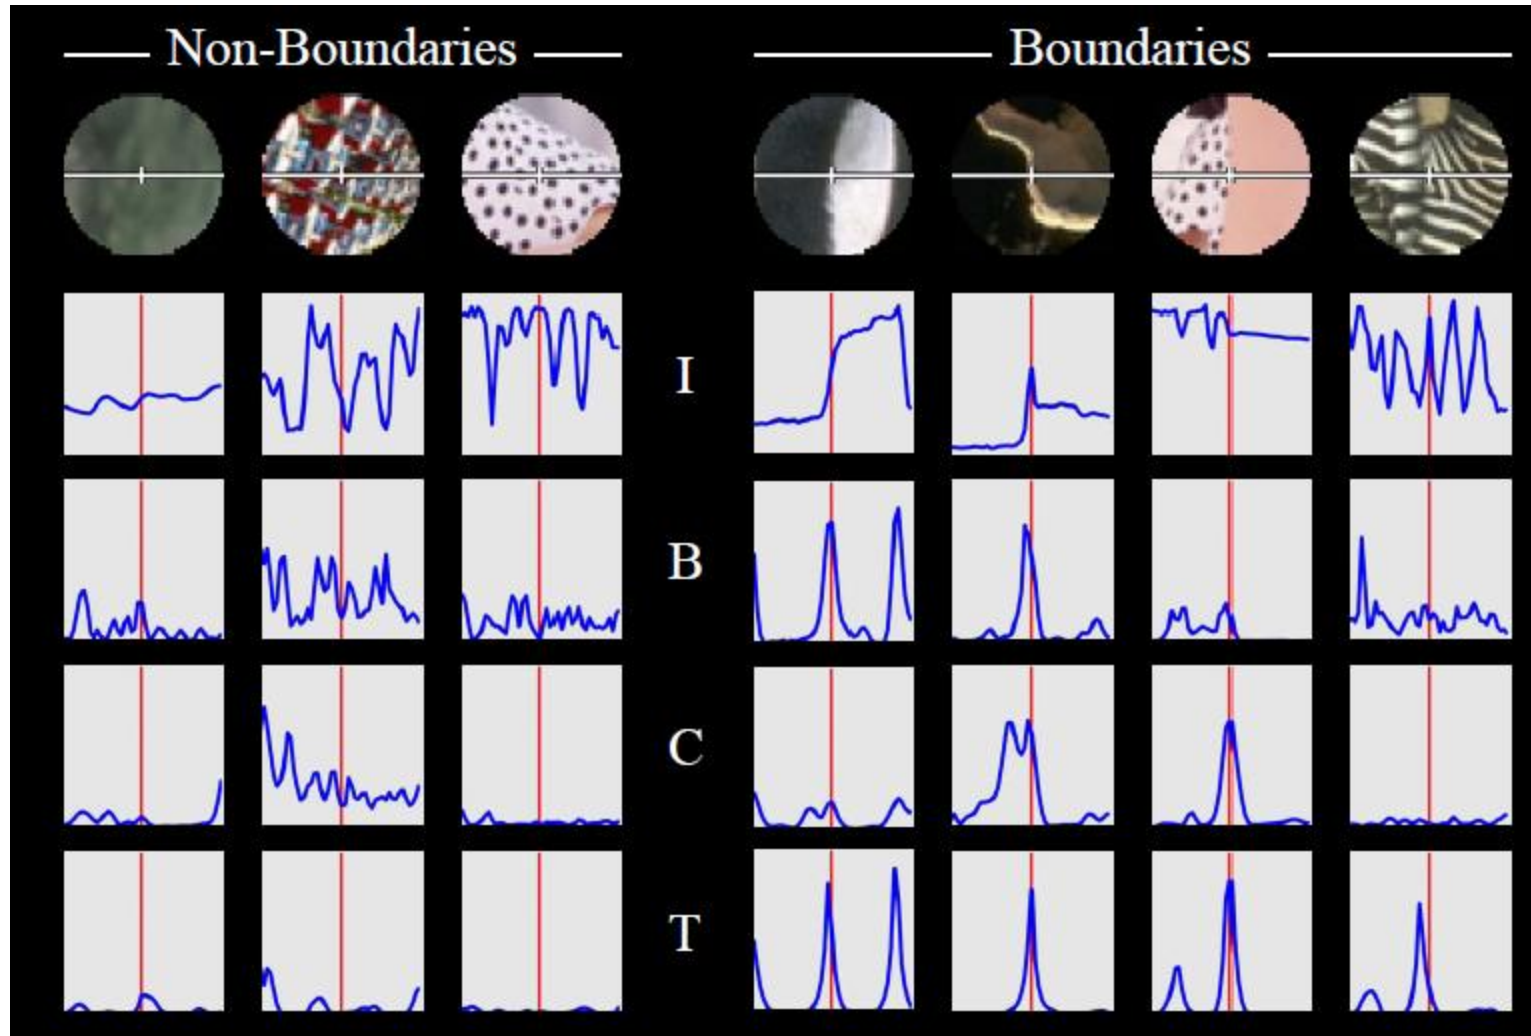
Martin, Fowlkes, Malik 2004: Learning to Detect Natural Boundaries...

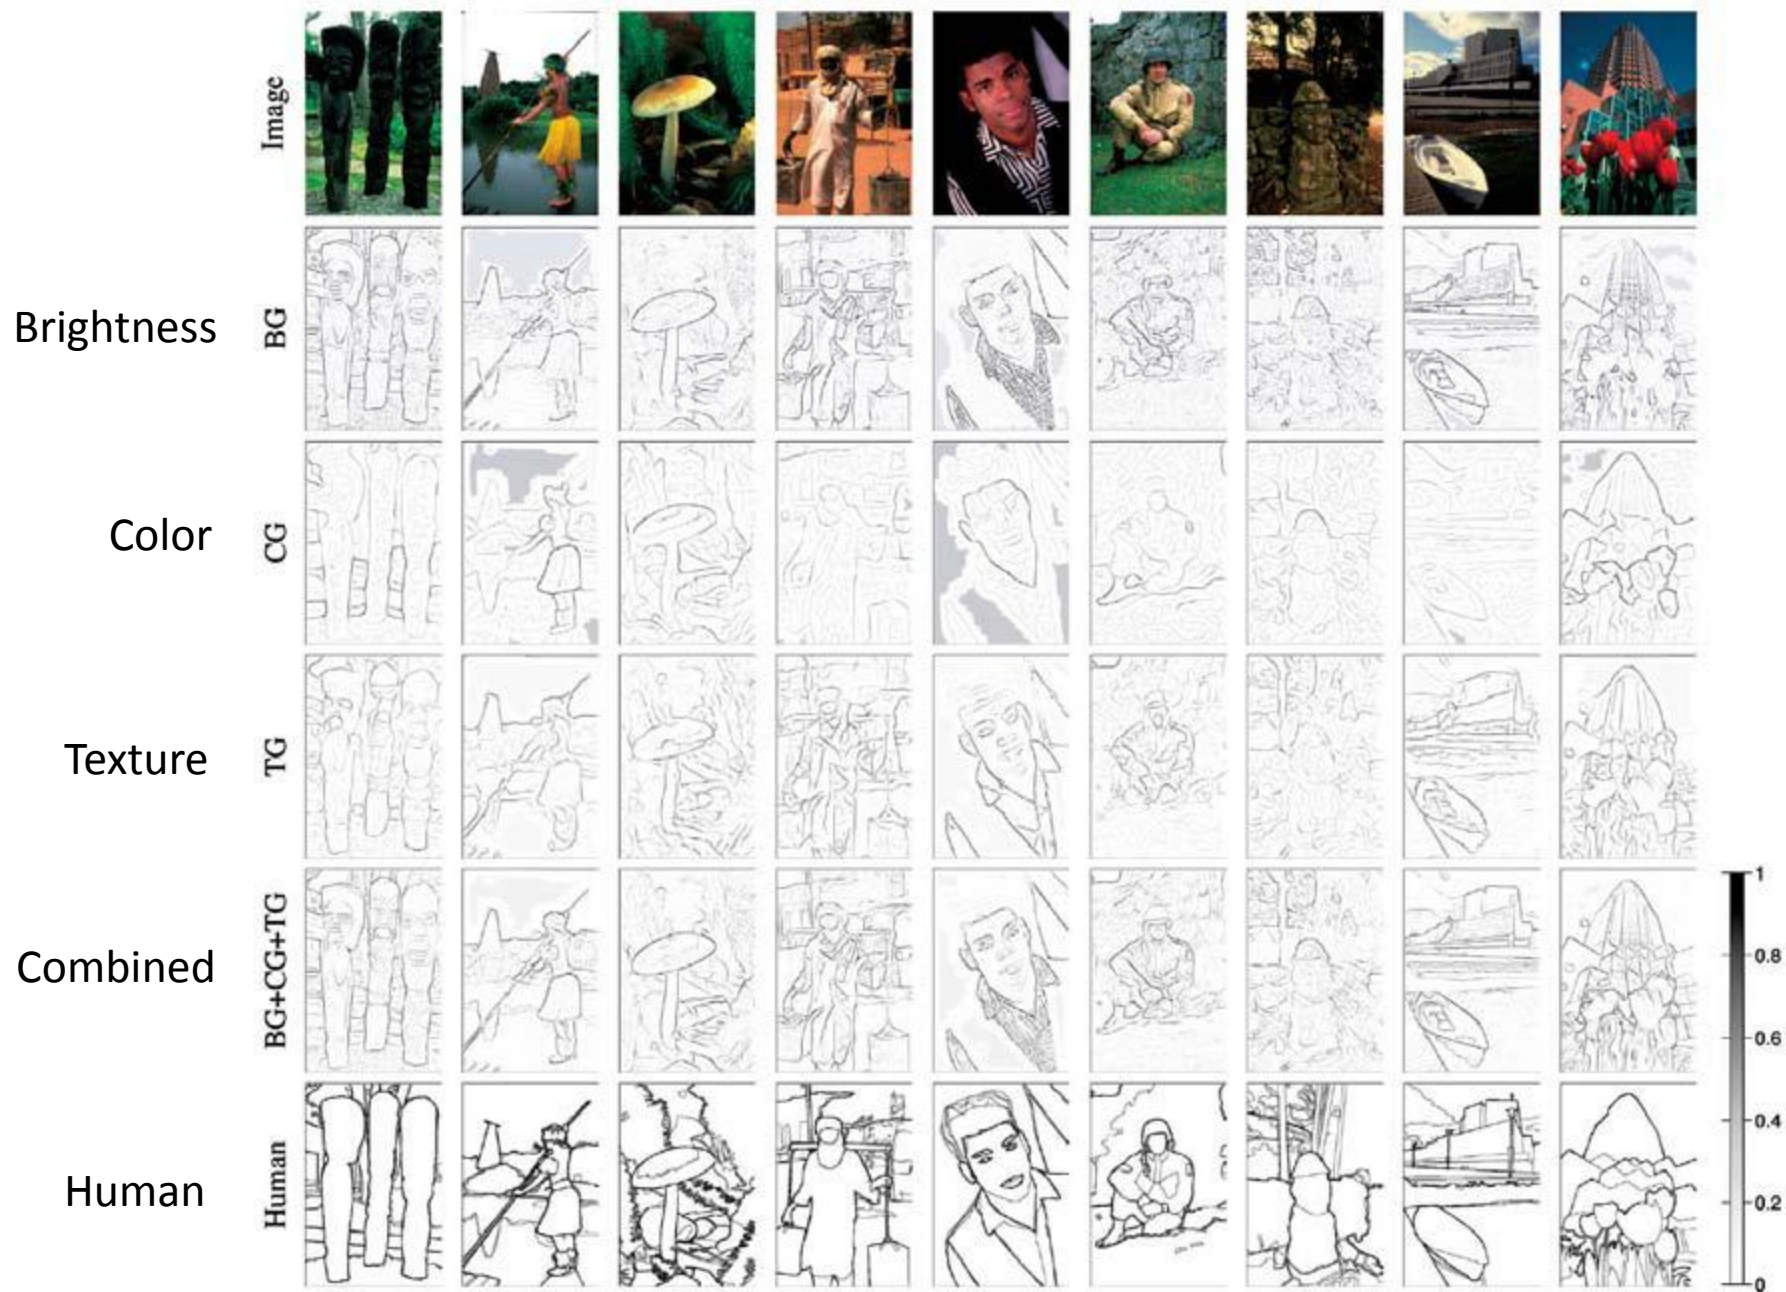
<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/papers/mfm-pami-boundary.pdf>

Figure from Fowlkes

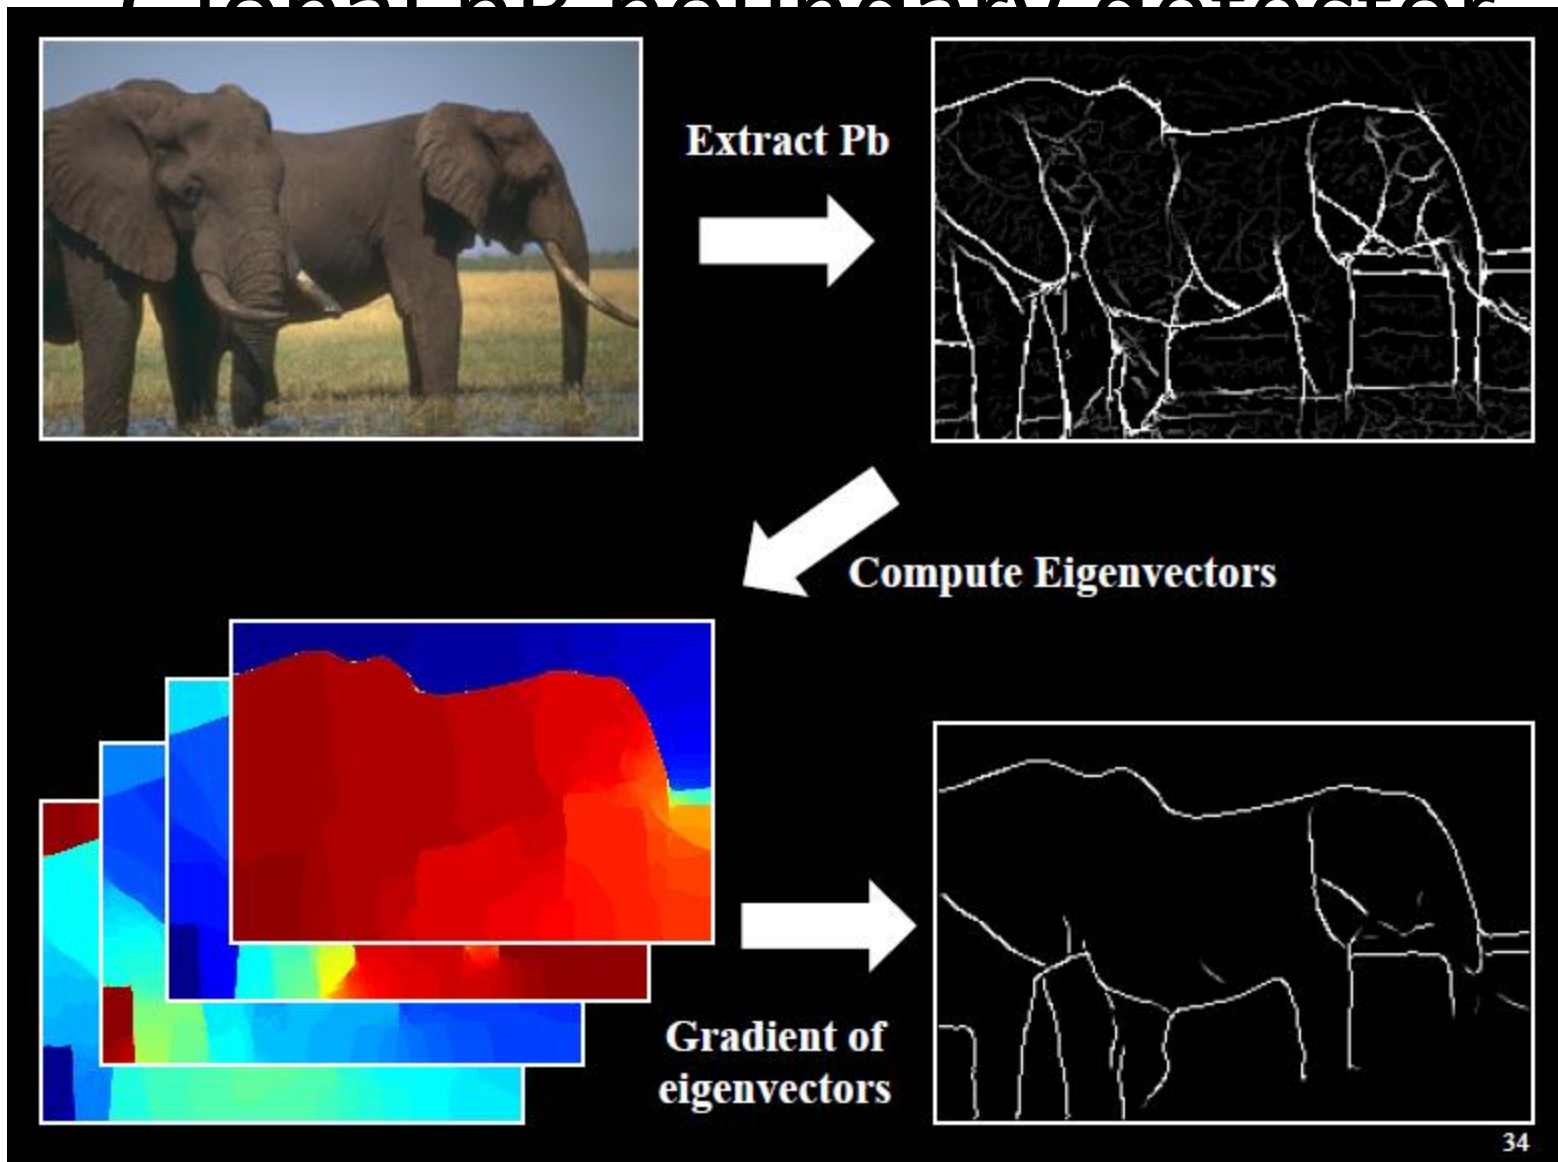


# pB Boundary Detector





# Global $pD$ boundary detector



# 45 years of boundary detection

