




# Compression

**How is it that a 4MP image can be compressed to a few hundred KB without a noticeable change?**



# Lossy Image Compression (JPEG)



Block-based Discrete Cosine Transform (DCT)

# Using DCT in JPEG

- The first coefficient  $B(0,0)$  is the DC component, the average intensity
- The top-left coeffs represent low frequencies, the bottom right – high frequencies



# Image compression using DCT

- Quantize
  - More coarsely for high frequencies (which also tend to have smaller values)
  - Many quantized high frequency values will be zero
- Encode
  - Can decode with inverse dct

Filter responses

$$G = \begin{matrix} & & & \xrightarrow{u} & & & & & \\ \begin{matrix} \downarrow v \\ \end{matrix} & \begin{bmatrix} -415.38 & -30.19 & -61.20 & 27.24 & 56.13 & -20.10 & -2.39 & 0.46 \\ 4.47 & -21.86 & -60.76 & 10.25 & 13.15 & -7.09 & -8.54 & 4.88 \\ -46.83 & 7.37 & 77.13 & -24.56 & -28.91 & 9.93 & 5.42 & -5.65 \\ -48.53 & 12.07 & 34.10 & -14.76 & -10.24 & 6.30 & 1.83 & 1.95 \\ 12.12 & -6.55 & -13.20 & -3.95 & -1.88 & 1.75 & -2.79 & 3.14 \\ -7.73 & 2.91 & 2.38 & -5.94 & -2.38 & 0.94 & 4.30 & 1.85 \\ -1.03 & 0.18 & 0.42 & -2.42 & -0.88 & -3.02 & 4.12 & -0.66 \\ -0.17 & 0.14 & -1.07 & -4.19 & -1.17 & -0.10 & 0.50 & 1.68 \end{bmatrix} \end{matrix}$$

Quantization table

$$Q = \begin{bmatrix} 16 & 11 & 10 & 16 & 24 & 40 & 51 & 61 \\ 12 & 12 & 14 & 19 & 26 & 58 & 60 & 55 \\ 14 & 13 & 16 & 24 & 40 & 57 & 69 & 56 \\ 14 & 17 & 22 & 29 & 51 & 87 & 80 & 62 \\ 18 & 22 & 37 & 56 & 68 & 109 & 103 & 77 \\ 24 & 35 & 55 & 64 & 81 & 104 & 113 & 92 \\ 49 & 64 & 78 & 87 & 103 & 121 & 120 & 101 \\ 72 & 92 & 95 & 98 & 112 & 100 & 103 & 99 \end{bmatrix}$$

Quantized values

$$B = \begin{bmatrix} -26 & -3 & -6 & 2 & 2 & -1 & 0 & 0 \\ 0 & -2 & -4 & 1 & 1 & 0 & 0 & 0 \\ -3 & 1 & 5 & -1 & -1 & 0 & 0 & 0 \\ -3 & 1 & 2 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$


# JPEG Compression Summary

1. Convert image to YCrCb
2. Subsample color by factor of 2
  - People have bad resolution for color
3. Split into blocks (8x8, typically), subtract 128
4. For each block
  - a. Compute DCT coefficients
  - b. Coarsely quantize
    - Many high frequency components will become zero
  - c. Encode (with run length encoding and then Huffman coding for leftovers)

<http://en.wikipedia.org/wiki/YCbCr>

<http://en.wikipedia.org/wiki/JPEG>

# Why do we get different, distance-dependent interpretations of hybrid images?



# Application: Hybrid Images




- A. Oliva, A. Torralba, P.G. Schyns, ["Hybrid Images,"](#) SIGGRAPH 2006



# Clues from Human Perception


- Early processing in humans filters for various orientations and scales of frequency
- Perceptual cues in the mid-high frequencies dominate perception
- When we see an image from far away, we are effectively subsampling it



Early Visual Processing: Multi-scale edge and blob filters


# Campbell-Robson contrast sensitivity curve

---




# Hybrid Image in FFT


Hybrid Image



Low-passed Image



High-passed Image



# Review: Image filtering

$$g[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$f[\cdot, \cdot]$

$h[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	0	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0


$$h[m, n] = \sum_{k, l} f[k, l] g[m+k, n+l]$$

# Image filtering

$$g[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$f[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$

	0	10							

$$h[m, n] = \sum_{k, l} f[k, l] g[m+k, n+l]$$

# Image filtering

$$g[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$f[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$



	0	10	20						

$$h[m, n] = \sum_{k, l} f[k, l] g[m+k, n+l]$$


# Filtering in spatial domain

1	0	-1
2	0	-2
1	0	-1



intensity image



# Filtering in frequency domain





FFT





$\times$

$=$



Inverse FFT






# Review of Filtering

- Filtering in frequency domain
  - Can be faster than filtering in spatial domain (for large filters)
  - Can help understand effect of filter
  - Algorithm:
    1. Convert image and filter to fft (fft2 in matlab)
    2. Pointwise-multiply ffts
    3. Convert result to spatial domain with ifft2


# Review of Filtering

- Linear filters for basic processing
  - Edge filter (high-pass)
  - Gaussian filter (low-pass)


$[-1 \ 1]$



FFT of Gradient Filter



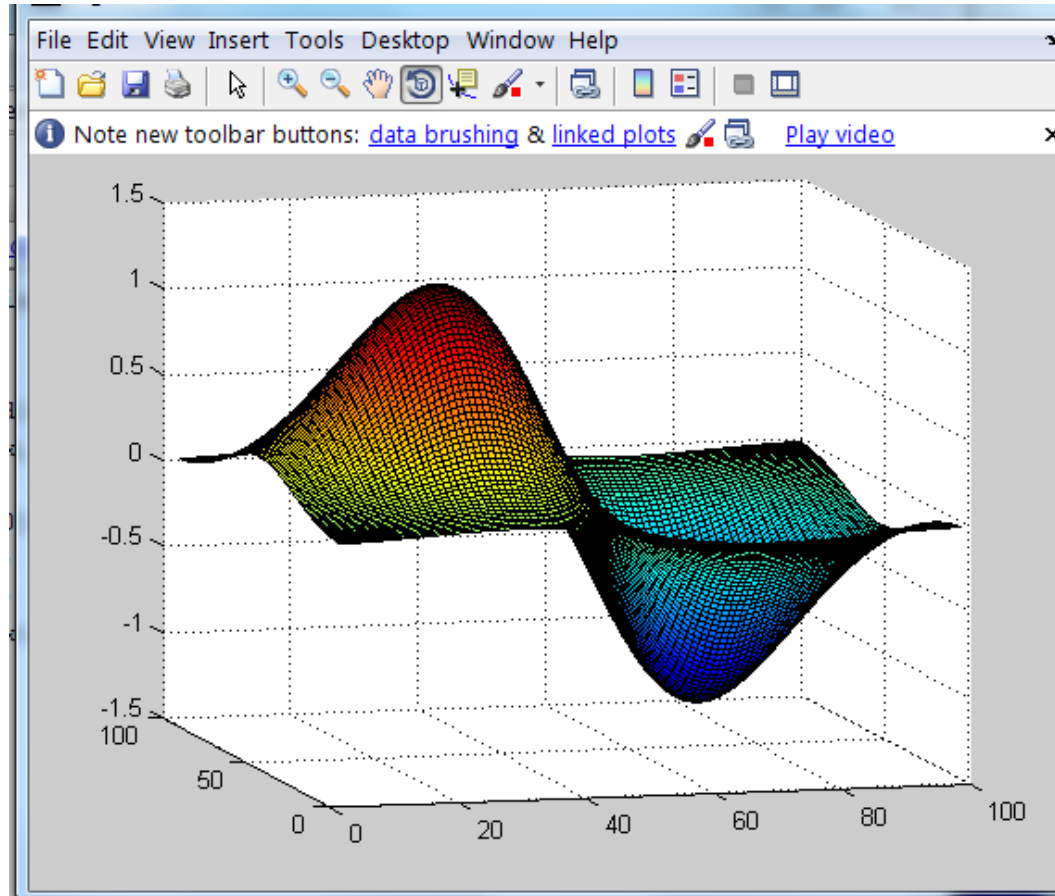
FFT of Gaussian



Gaussian





# Review of Filtering


- Derivative of Gaussian



# Things to Remember

- Sometimes it makes sense to think of images and filtering in the frequency domain
  - Fourier analysis
- Can be faster to filter using FFT for large images ( $N \log N$  vs.  $N^2$  for auto-correlation)
- Images are mostly smooth
  - Basis for compression
- Remember to low-pass before sampling





The blue and green colors are actually the same



<http://blogs.discovermagazine.com/badastronomy/2009/06/24/the-blue-and-the-green/>

# Previous Lectures

- We've now touched on the first three chapters of Szeliski.
  - 1. Introduction
  - 2. Image Formation
  - 3. Image Processing
- Now we're moving on to
  - 4. Feature Detection and Matching
  - Multiple views and motion (7, 8, 11)

# Edge / Boundary Detection

Computer Vision

Szeliski 4.2

James Hays



# 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)




# 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

# 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







# Intensity profile





# With a little Gaussian noise



# Effects of noise

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

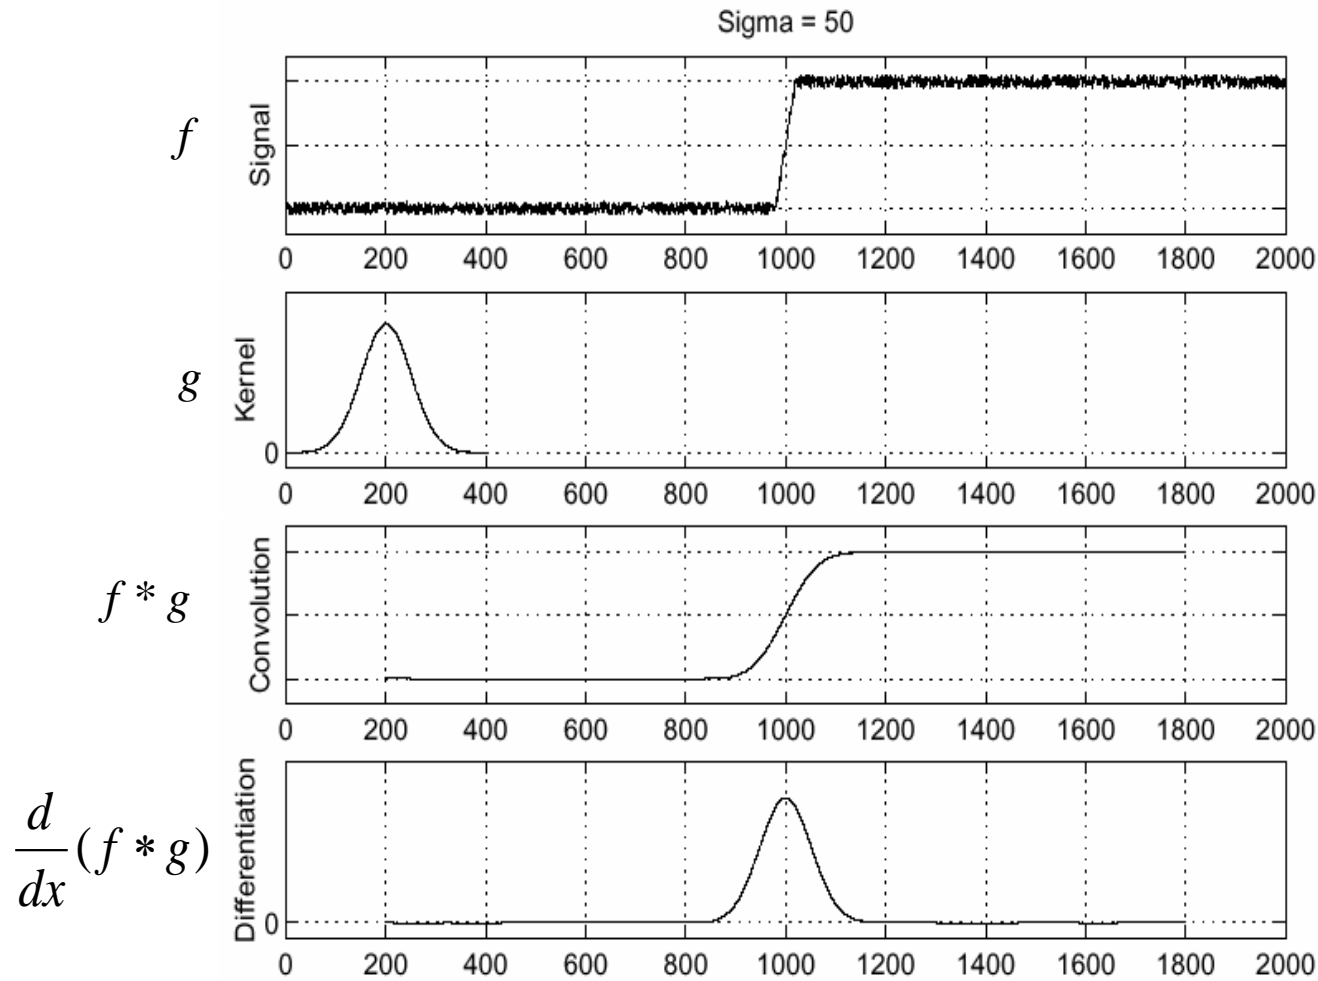


Where is the edge?

# 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:  
$$\frac{d}{dx}(f * g) = f * \frac{d}{dx}g$$
- This saves us one operation:




# Derivative of Gaussian filter



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



# 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.

22,000 citations!


Source: L. Fei-Fei

# Example




original image (Lena)



# Derivative of Gaussian filter



x-direction



y-direction




# Compute Gradients (DoG)



X-Derivative of Gaussian



Y-Derivative of Gaussian



Gradient Magnitude


# Get Orientation at Each Pixel

- Threshold at minimum level
- Get orientation




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

# 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.



# Sidebar: Interpolation options

- `imx2 = imresize(im, 2, interpolation_type)`
- 'nearest'
  - Copy value from nearest known
  - Very fast but creates blocky edges
- 'bilinear'
  - Weighted average from four nearest known pixels
  - Fast and reasonable results
- 'bicubic' (default)
  - Non-linear smoothing over larger area (4x4)
  - Slower, visually appealing, may create negative pixel values





# Before Non-max Suppression



After non-max suppression




# Hysteresis thresholding

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



# 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.



# 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

# Where do humans see boundaries?



image



human segmentation



gradient magnitude

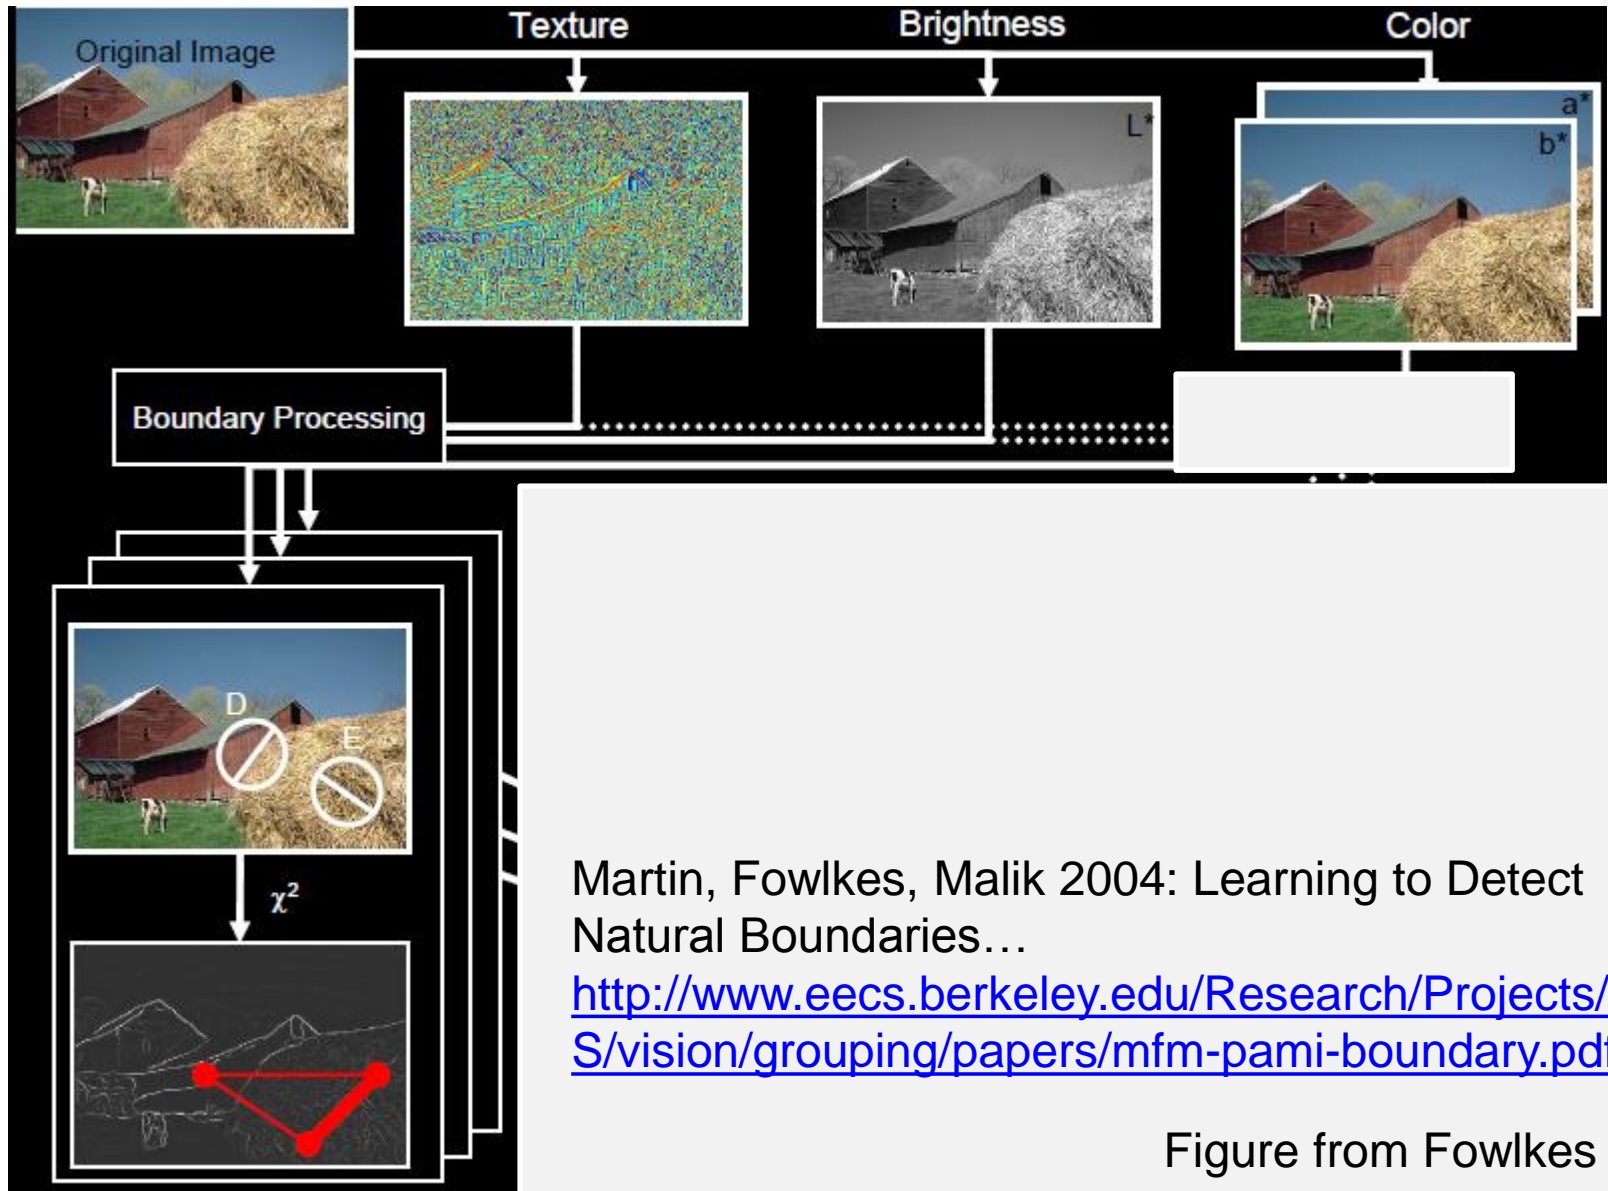


- Berkeley segmentation database:

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


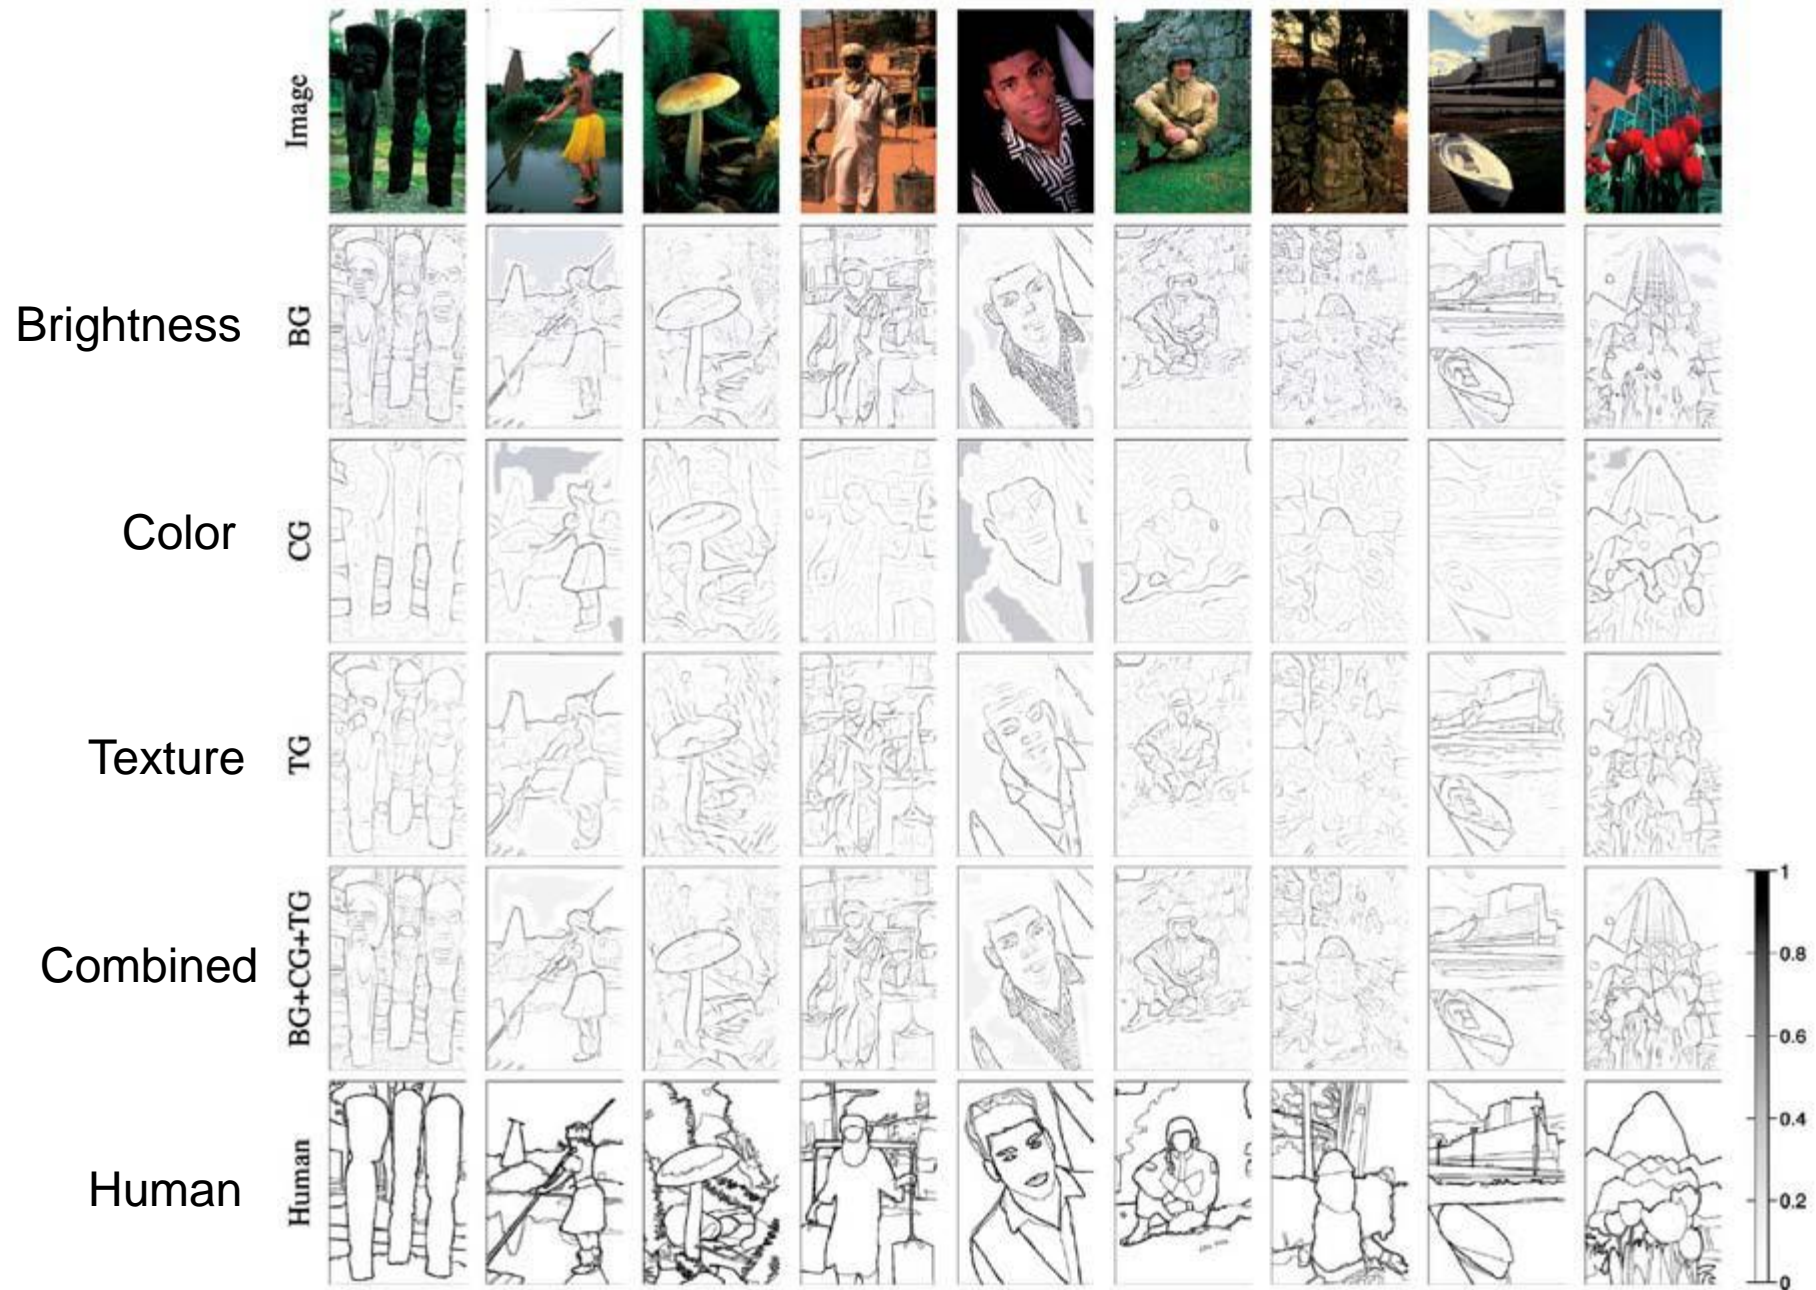


# pB boundary detector





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



# Results




# Results






Human (0.95)



Pb (0.63)






For more:

<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/bsds/bench/html/108082-color.html>

# 45 years of boundary detection



# State of edge detection

- Local edge detection works well
  - But many false positives from illumination and texture edges
- Some methods to take into account longer contours, but could probably do better
- Modern methods that actually “learn” from data.
- Poor use of object and high-level information