



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

Szeliski 4.2

Computer Vision

CS 143, Brown

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

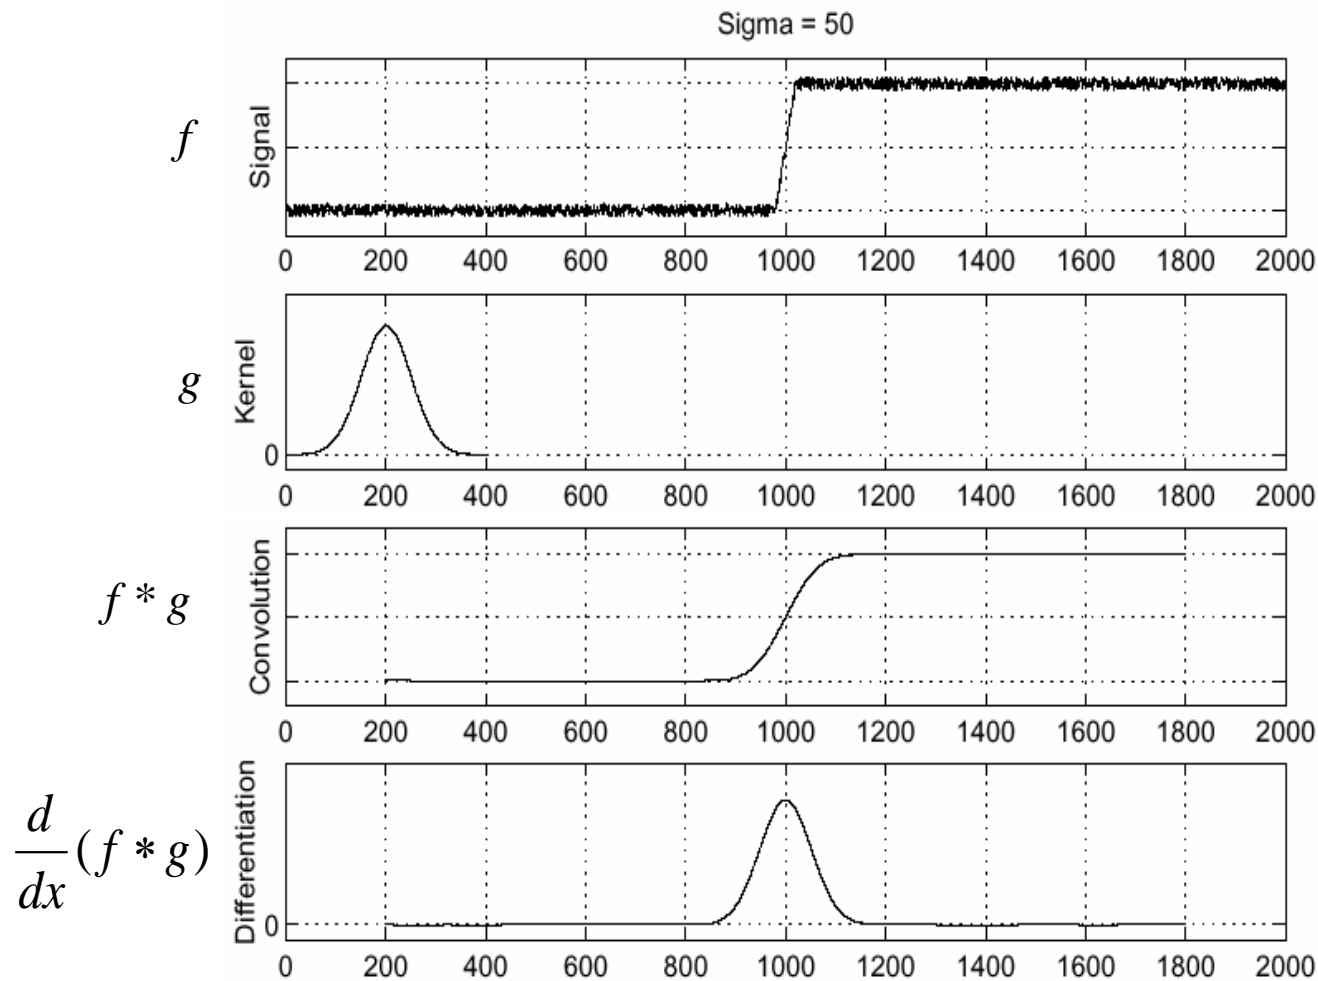


Where is the edge?

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.

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 σ (Gaussian kernel spread/size)



original



Canny with $\sigma = 1$



Canny with $\sigma = 2$

The choice of σ depends on desired behavior


- large σ detects large scale edges
- small σ detects fine features

Where do humans see boundaries?



image



human segmentation



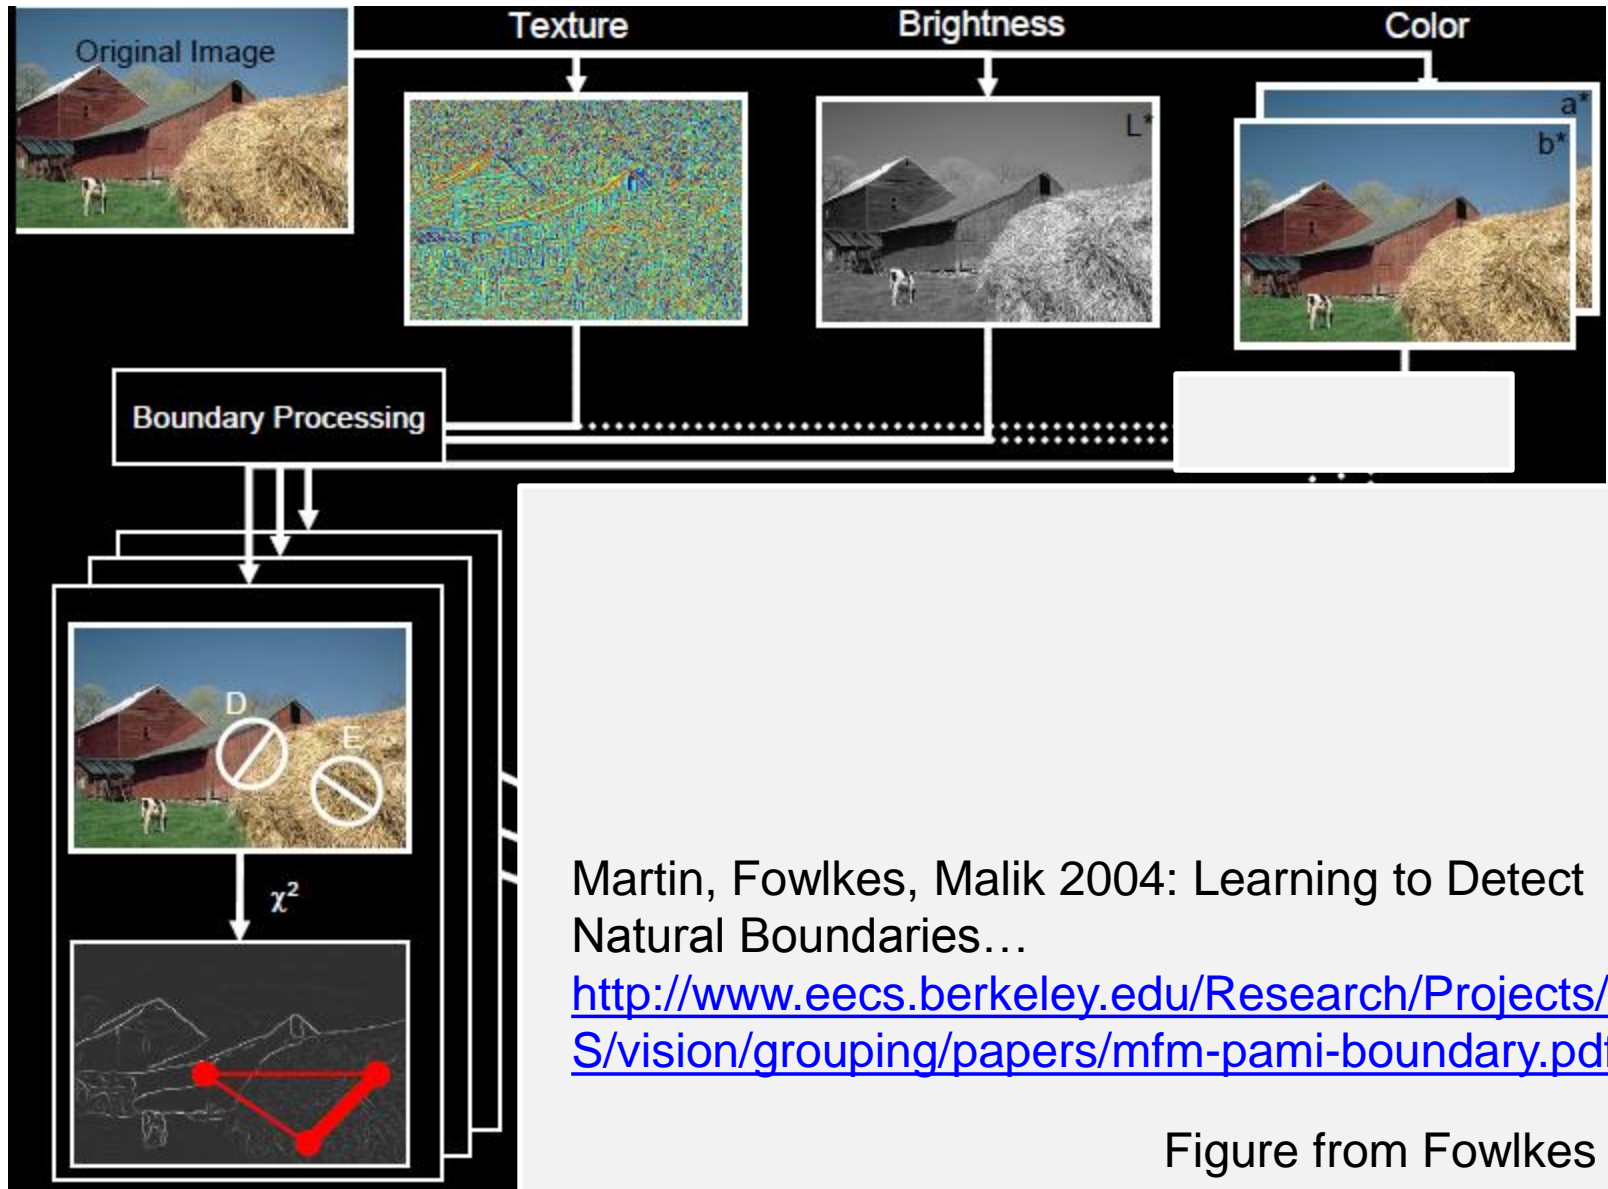
gradient magnitude



- 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

pB Boundary Detector

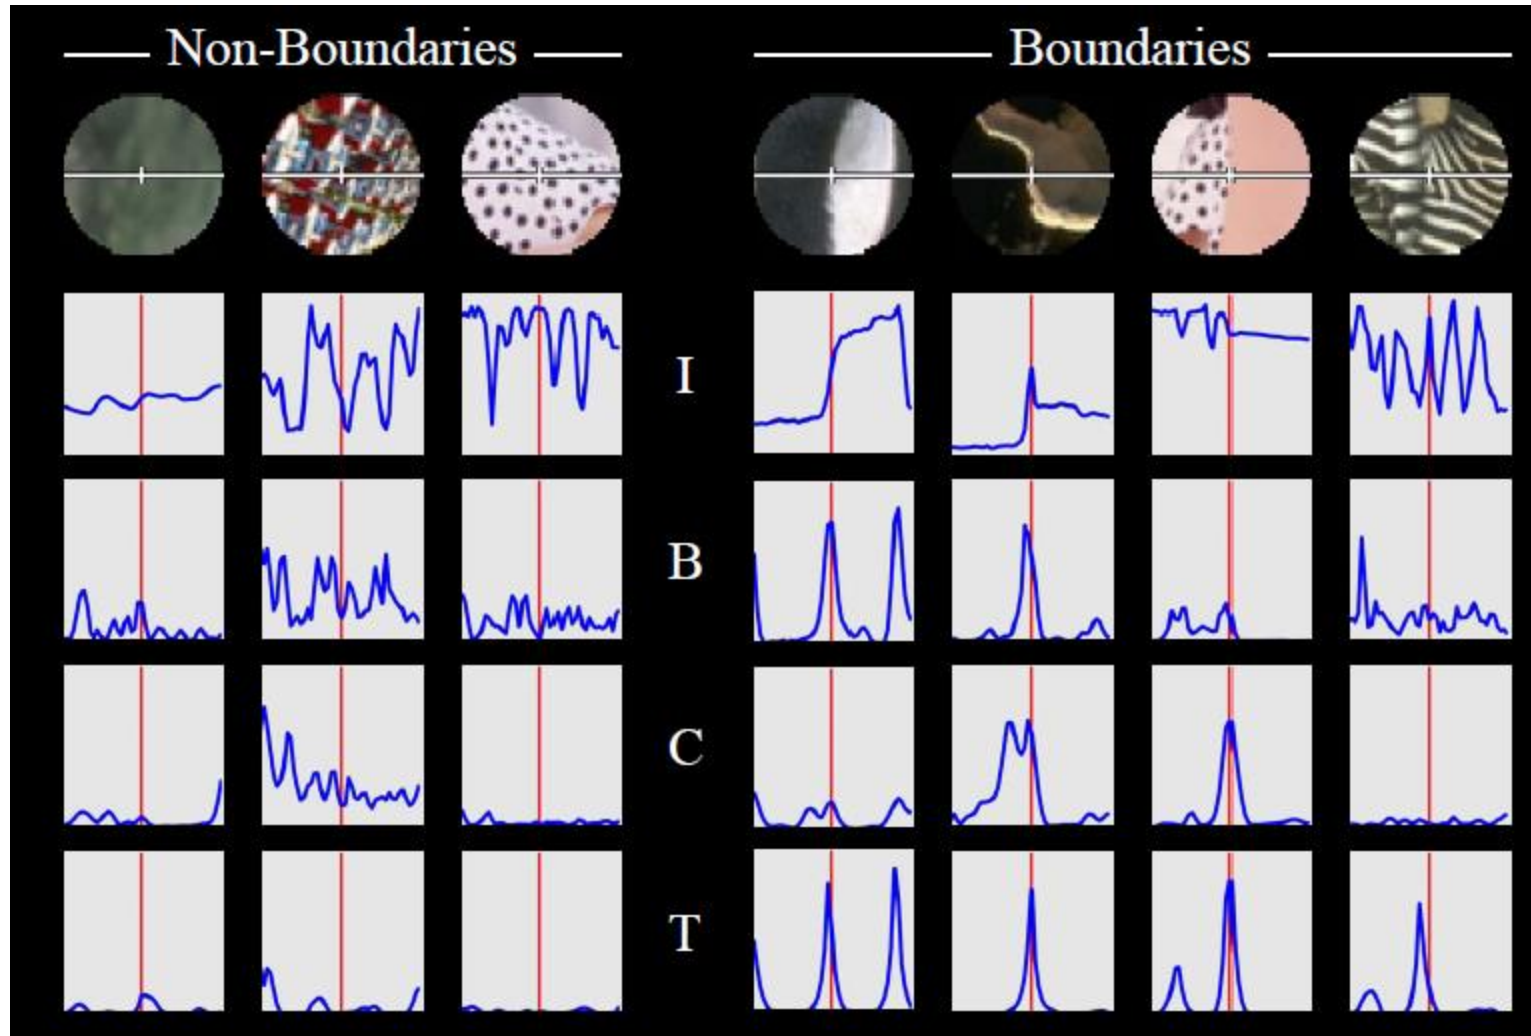

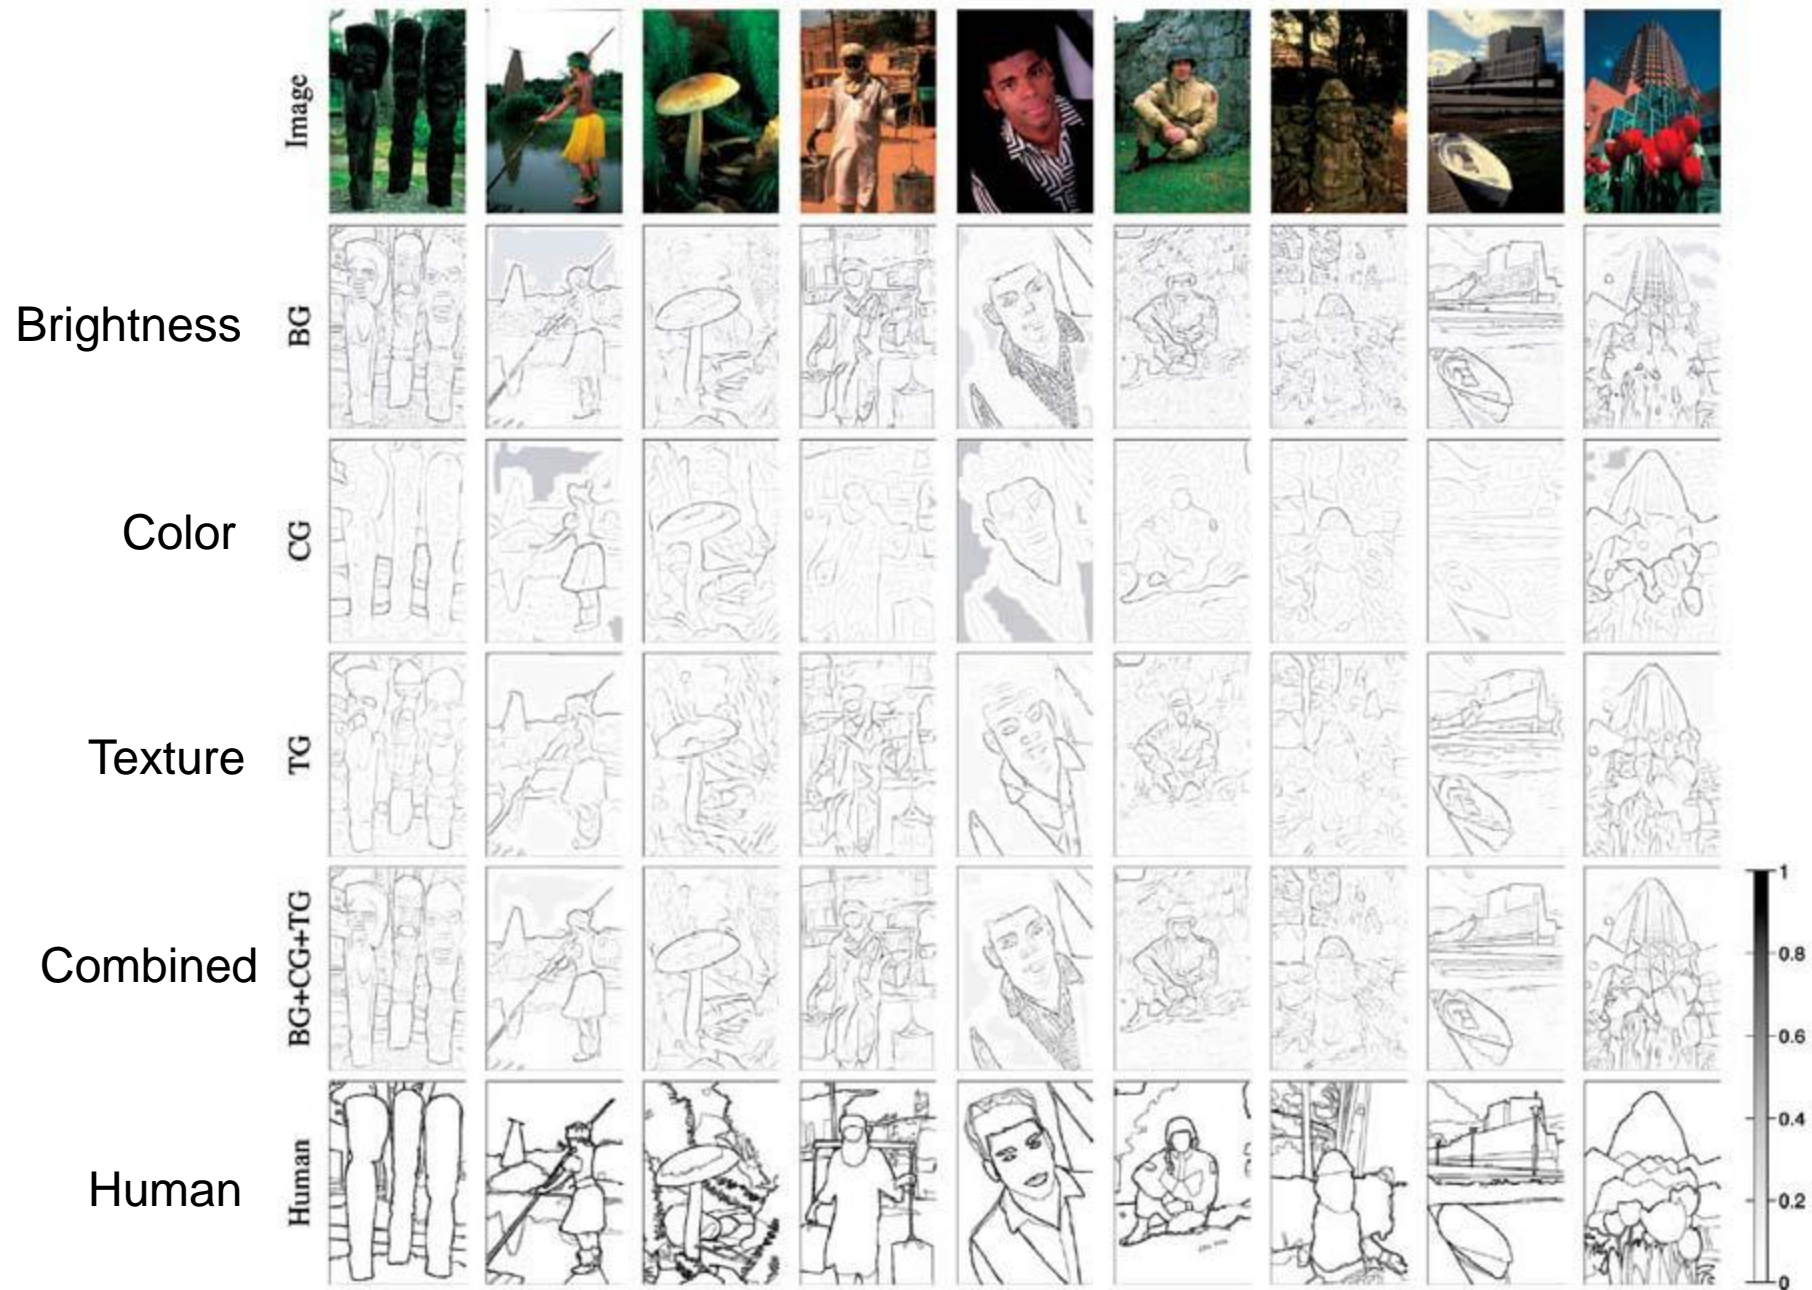





Figure from Fowlkes







Results



Results








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
- Few methods that actually “learn” from data. Your project 5, Sketch Tokens, will do so.
- Poor use of object and high-level information

Style and abstraction in portrait sketching, Berger et al. SIGGRAPH 2013



- Learn from artist's strokes so that edges are more likely in certain parts of the face.

Questions