C280, Computer Vision

Prof. Trevor Darrell
trevor@eecs.berkeley.edu

Lecture 23: Segmentation II &
Computational Photography Teaser

Two presentations today:

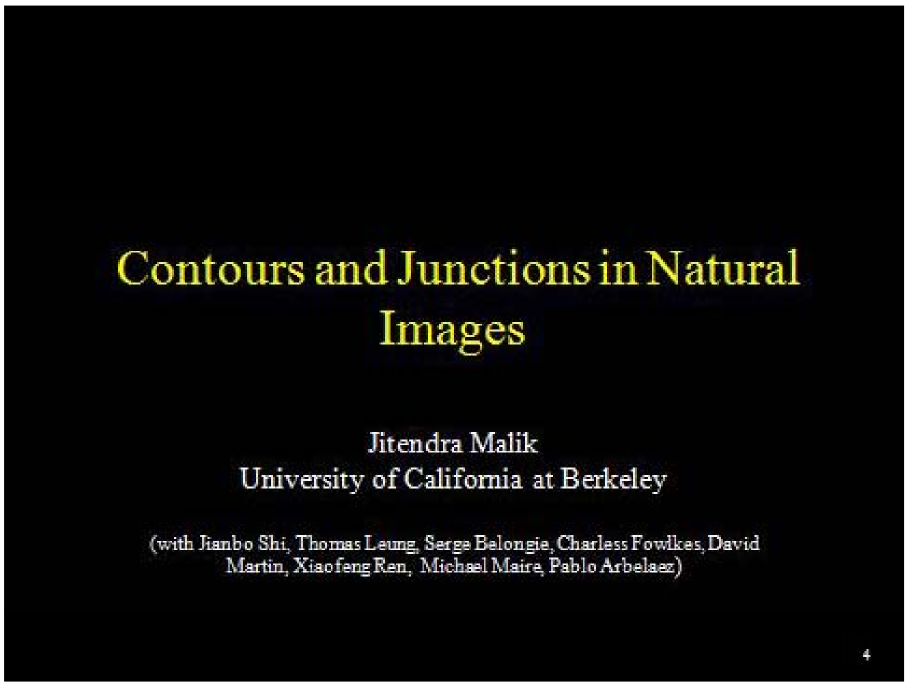

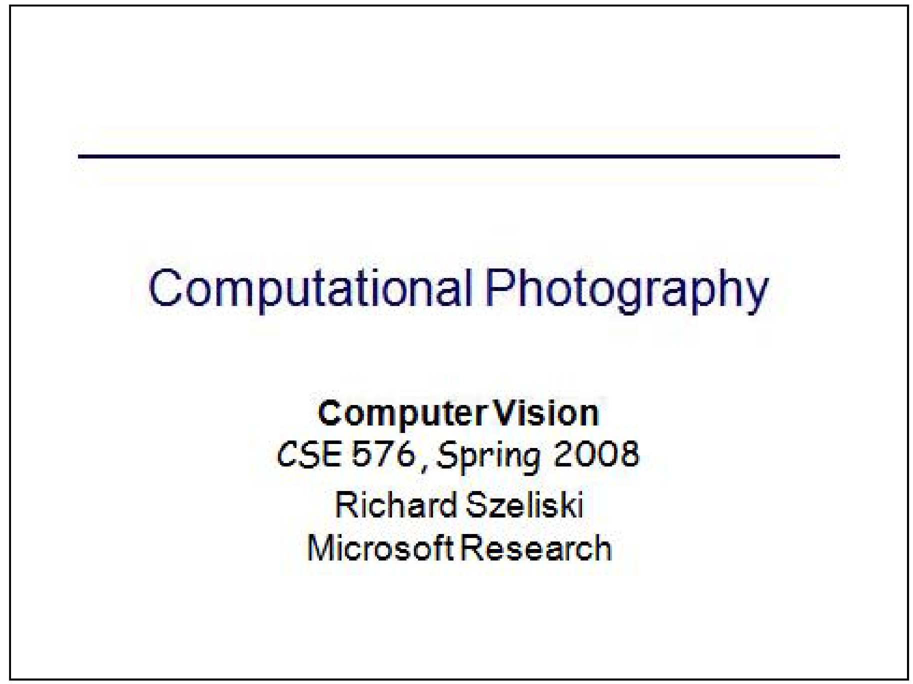

Contours and Junctions in Natural Images

Jitendra Malik
University of California at Berkeley

(with Jianbo Shi, Thomas Leung, Serge Belongie, Charless Fowlkes, David Martin, Xiaofeng Ren, Michael Maire, Pablo Arbelaez)

From Pixels to Perception

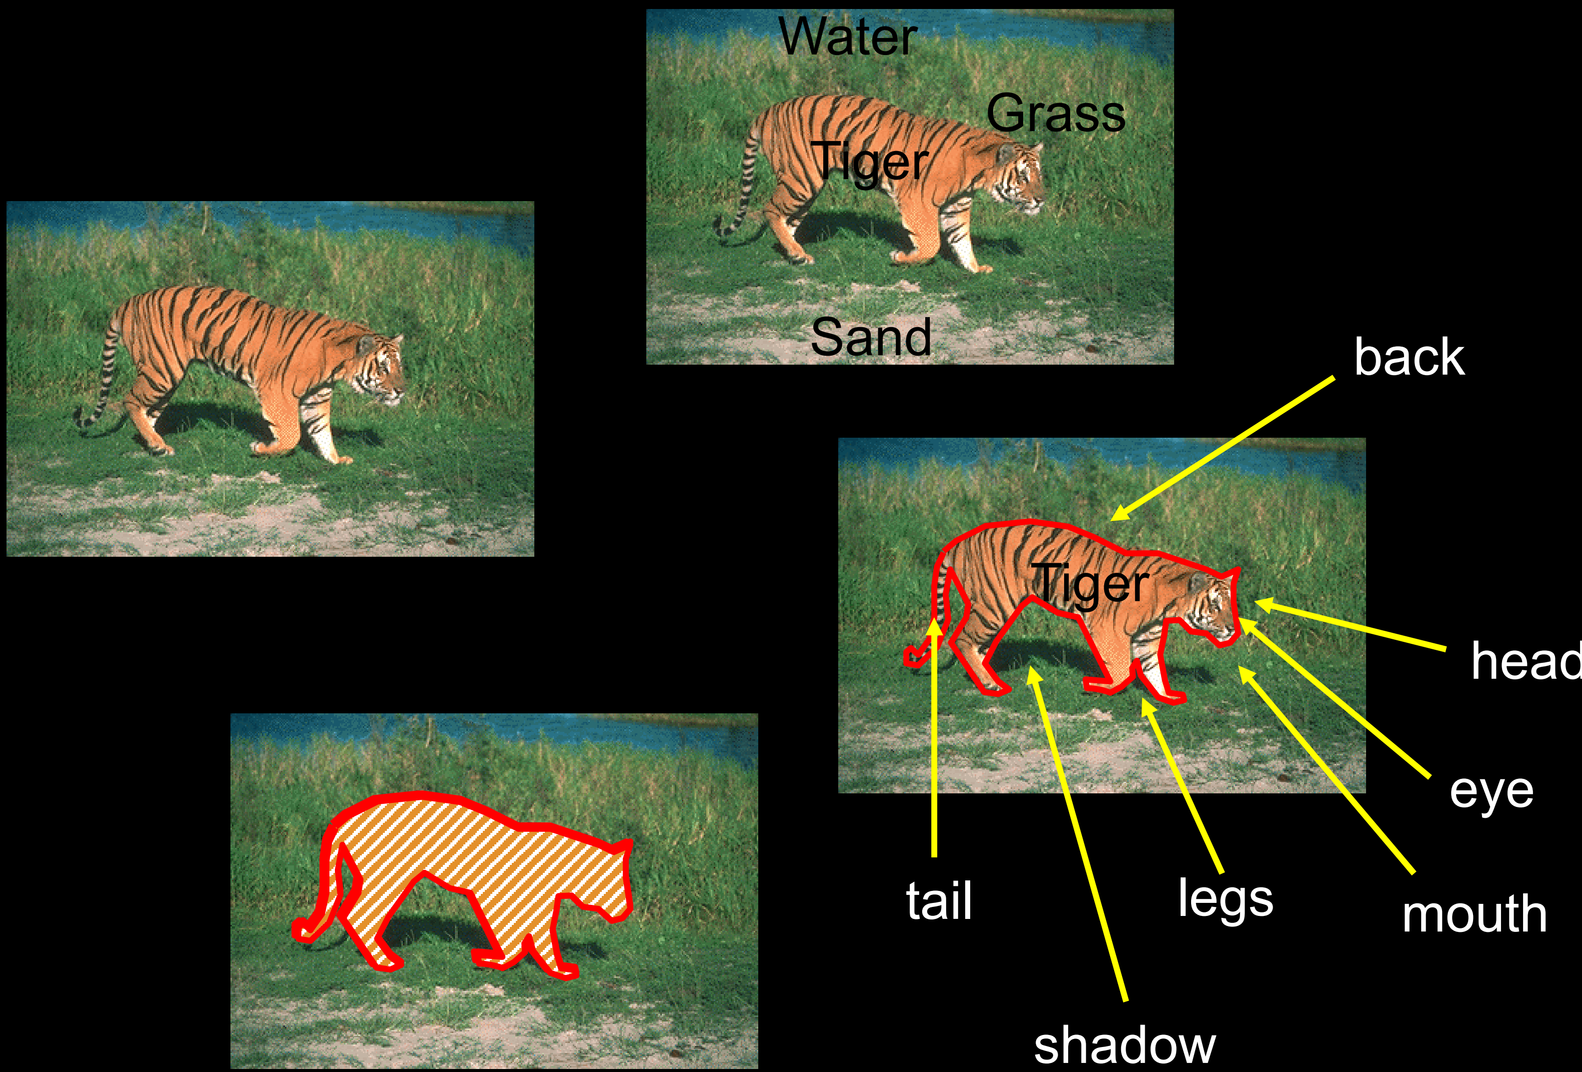

I stand at the window and see a house, trees, sky. Theoretically I might say there were 327 brightnesses and nuances of colour.

Do I have "327"?

No. I have sky, house, and trees.

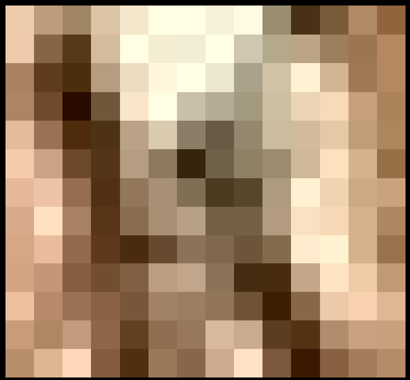

Max Wertheimer

Perceptual Organization

Grouping

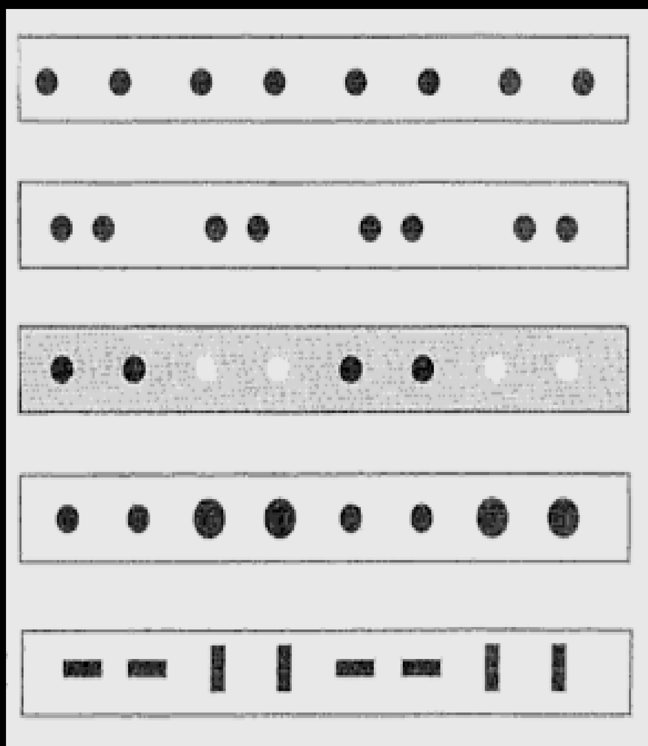

Figure/Ground

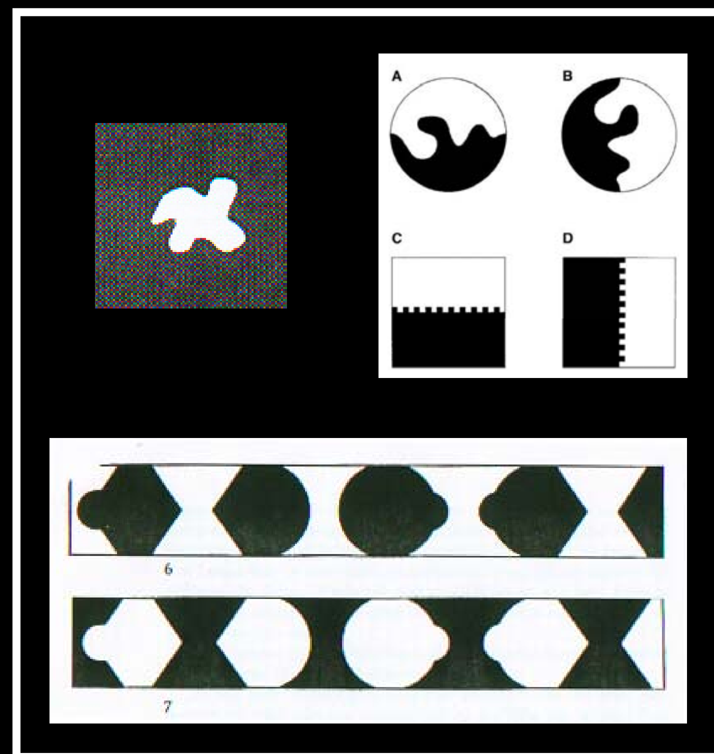

Key Research Questions in Perceptual Organization

- Predictive power
  - Factors for complex, natural stimuli ?
  - How do they interact ?
- Functional significance
  - Why should these be useful or confer some evolutionary advantage to a visual organism?
- Brain mechanisms
  - How are these factors implemented given what we know about V1 and higher visual areas?

Attneave's Cat (1954)
Line drawings convey most of the information

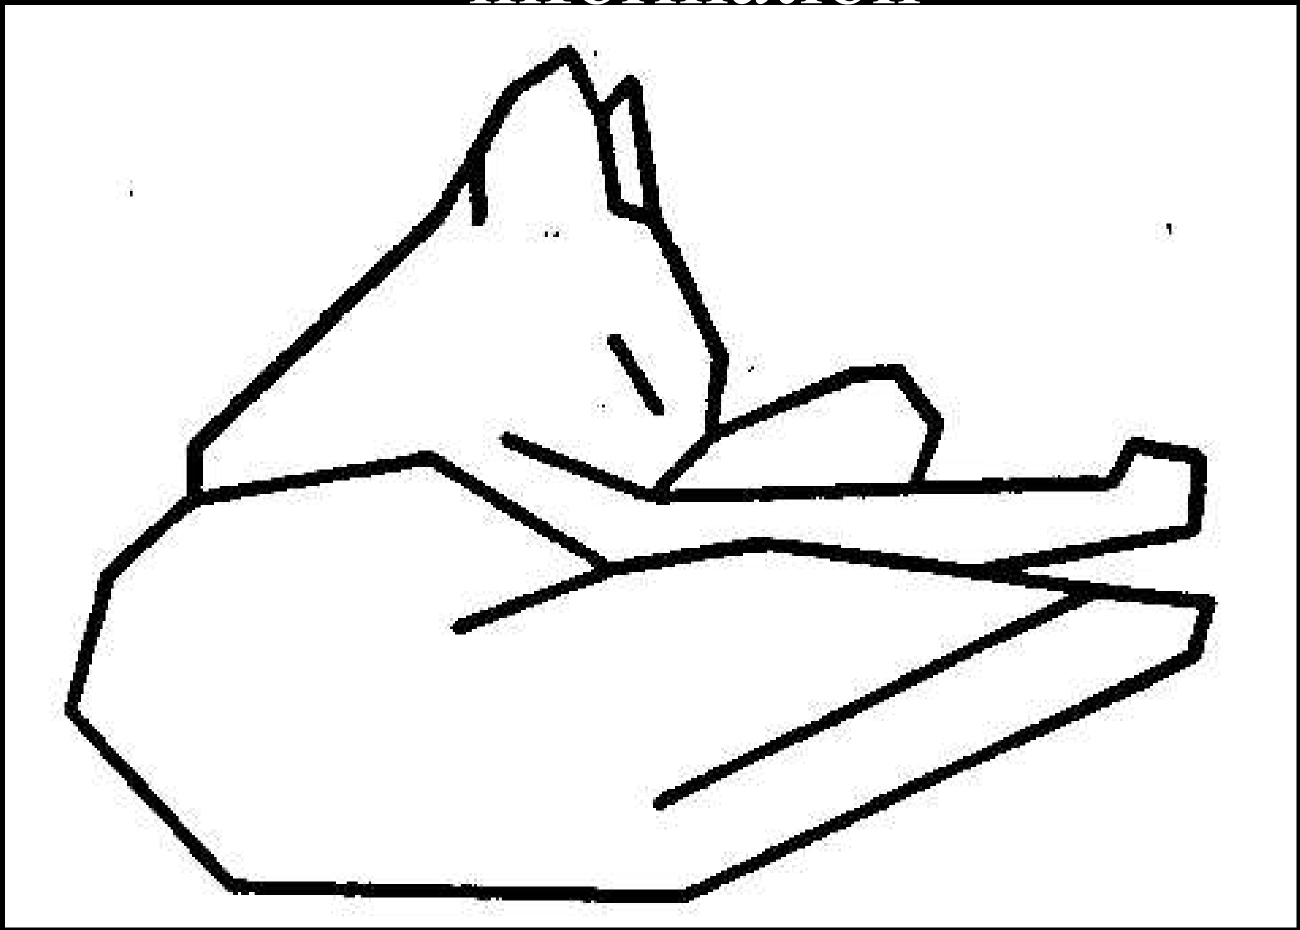

Contours and junctions are fundamental…

- Key to recognition, inference of 3D scene properties, visually- guided manipulation and locomotion…
- This goes beyond local, V1-like, edge-detection. Contours are the result of perceptual organization, grouping and figure/ground processing

Some computer vision history…

- Local Edge Detection was much studied in the 1970s and early 80s (Sobel, Rosenfeld, Binford-Horn, Marr-Hildreth, Canny …)
- Edge linking exploiting curvilinear continuity was studied as well (Rosenfeld, Zucker, Horn, Ullman …)
- In the 1980s, several authors argued for perceptual organization as a precursor to recognition (Binford, Witkin and Tennebaum, Lowe, Jacobs …)

However in the 90s …

1. We realized that there was more to images than edges
   - Biologically inspired filtering approaches (Bergen & Adelson, Malik & Perona..) √
   - Pixel based representations for recognition (Turk & Pentland, Murase & Nayar, LeCun …)
2. We lost faith in the ability of bottom-up vision
   - Do minimal bottom up processing , e.g. tiled orientation histograms don't even assume that linked contours or junctions can be extracted **?**
   - Matching with memory of previously seen objects then becomes the primary engine for parsing an image.

At Berkeley, we took a contrary view…

1. Collect Data Set of Human segmented images
2. Learn Local Boundary Model for combining brightness, color and texture
3. Global framework to capture closure, continuity
4. Detect and localize junctions
5. Integrate low, mid and high-level information for grouping and figure-ground segmentation

Berkeley Segmentation DataSet [BSDS]

D. Martin, C. Fowlkes, D. Tal, J. Malik. "A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics", ICCV, 2001

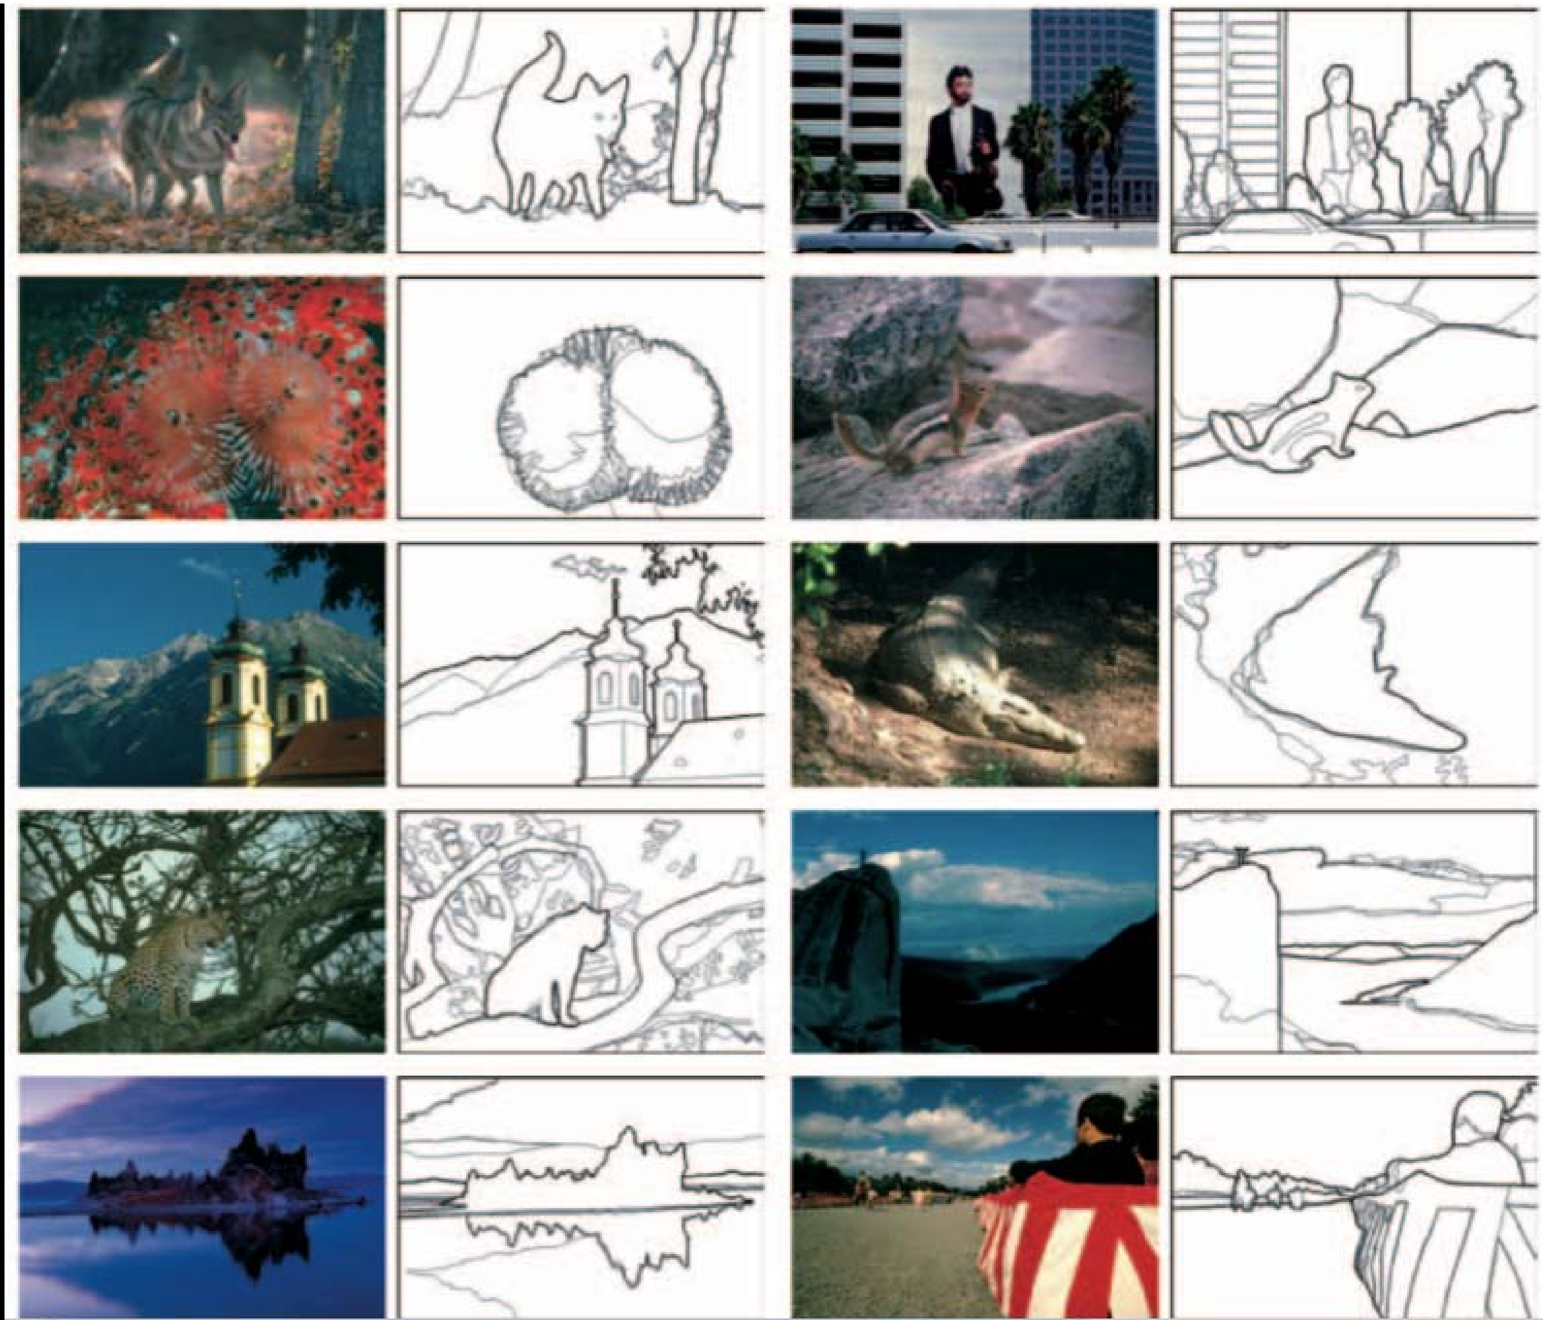

Contour detection ~1970

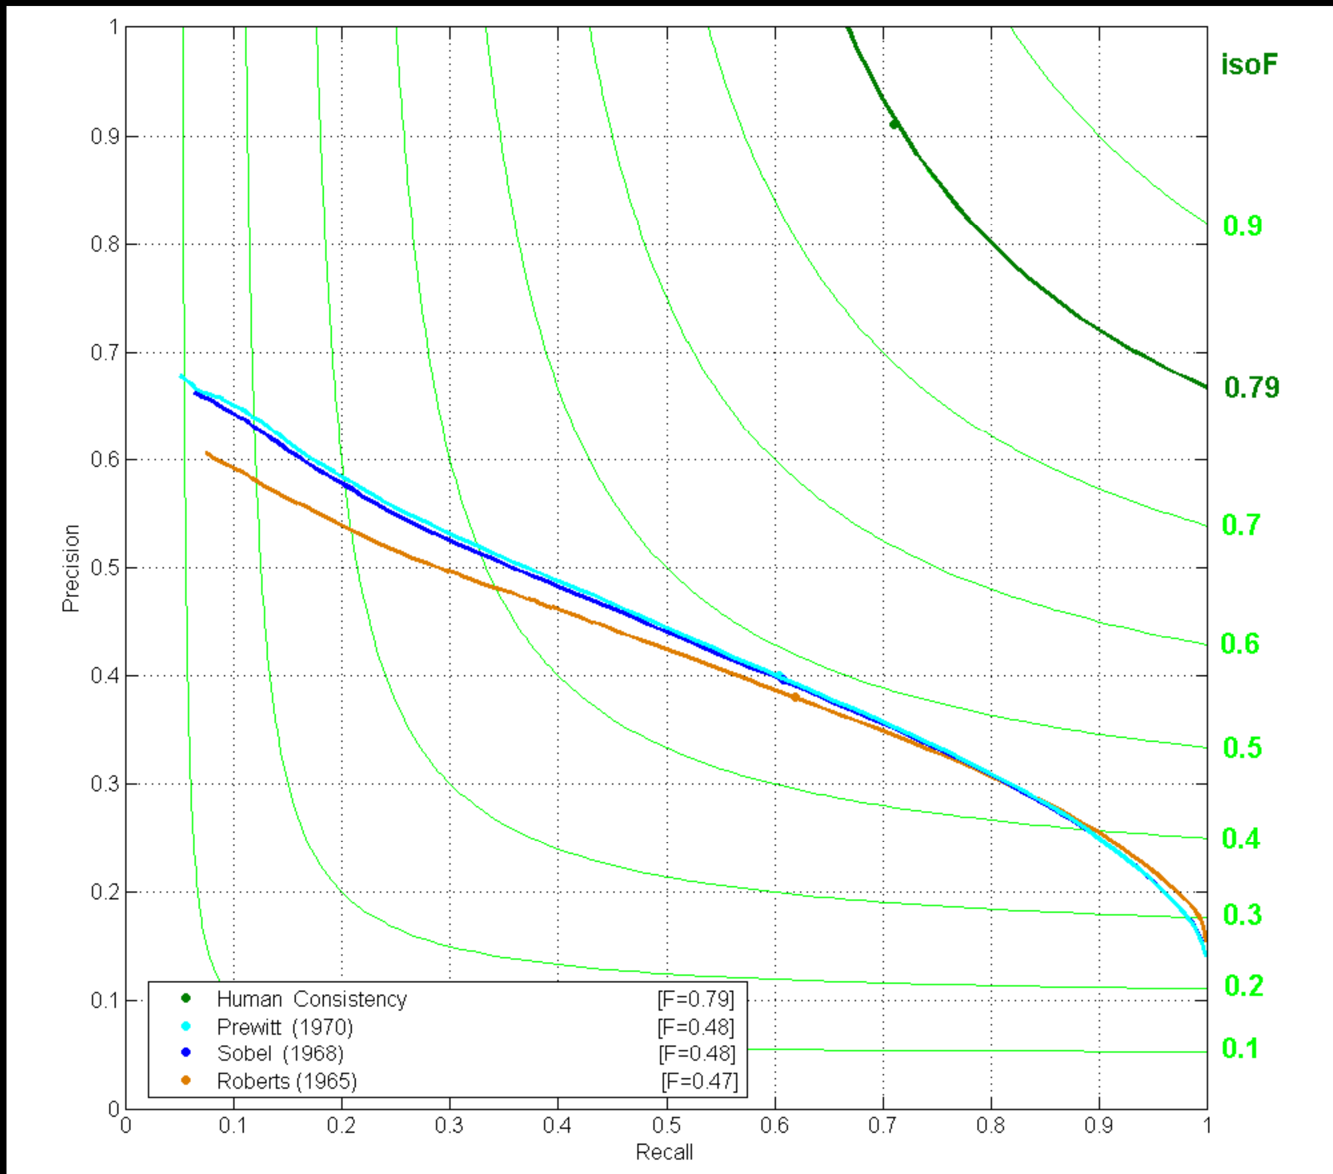

Contour detection ~1990

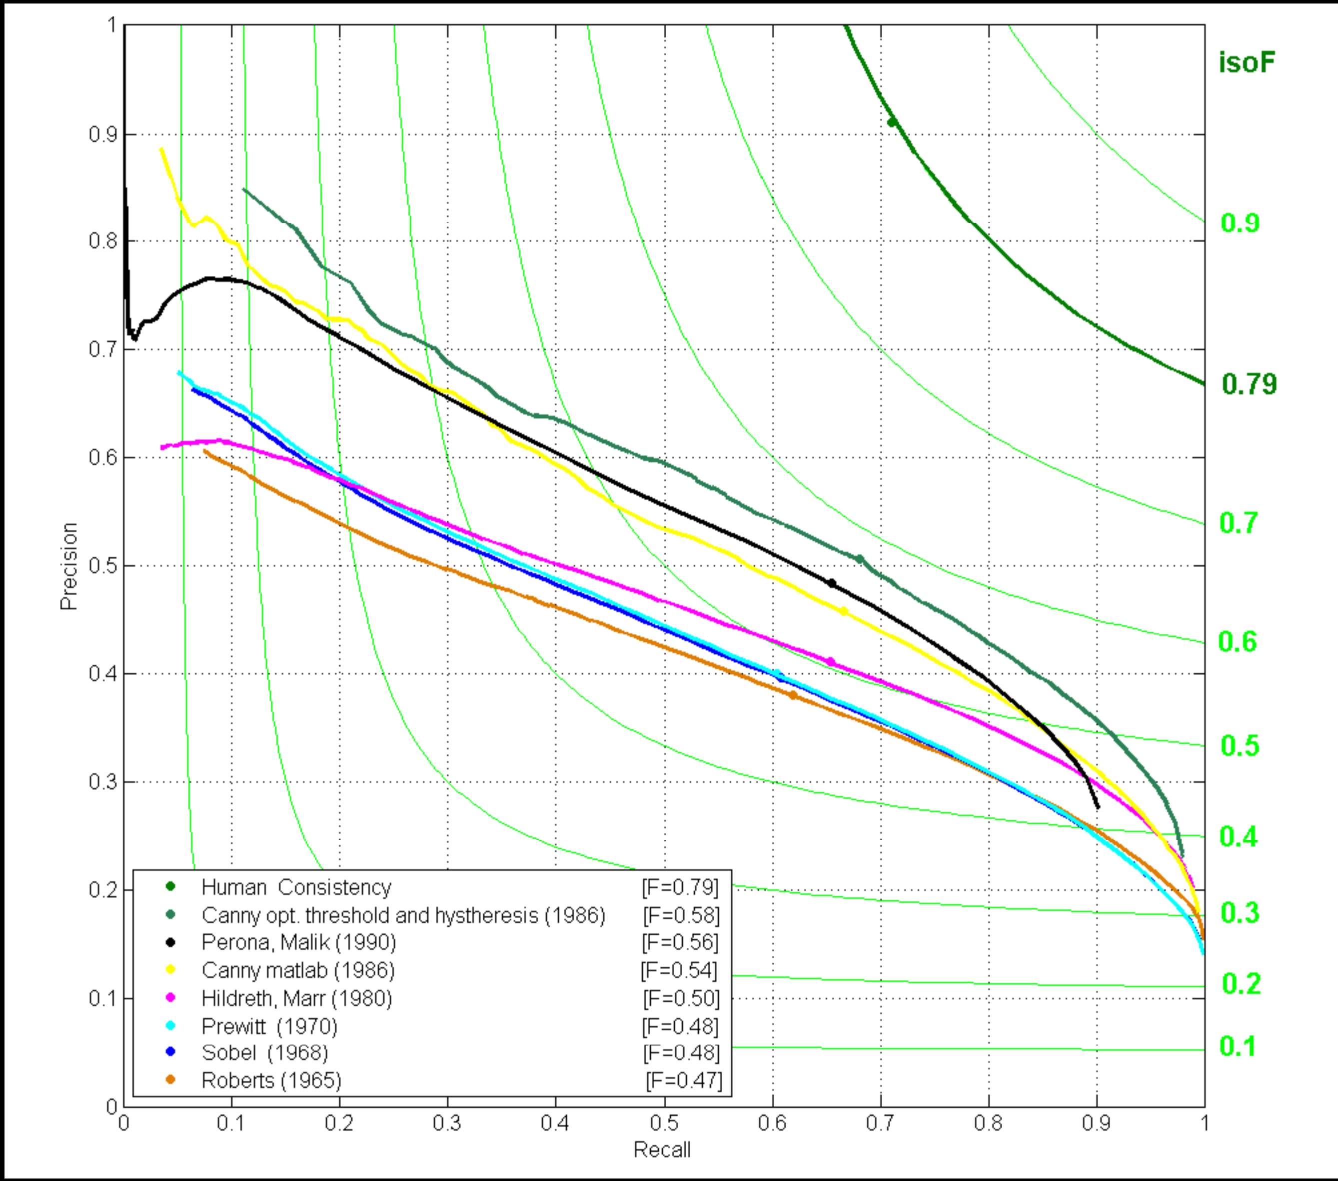

Contour detection ~2004

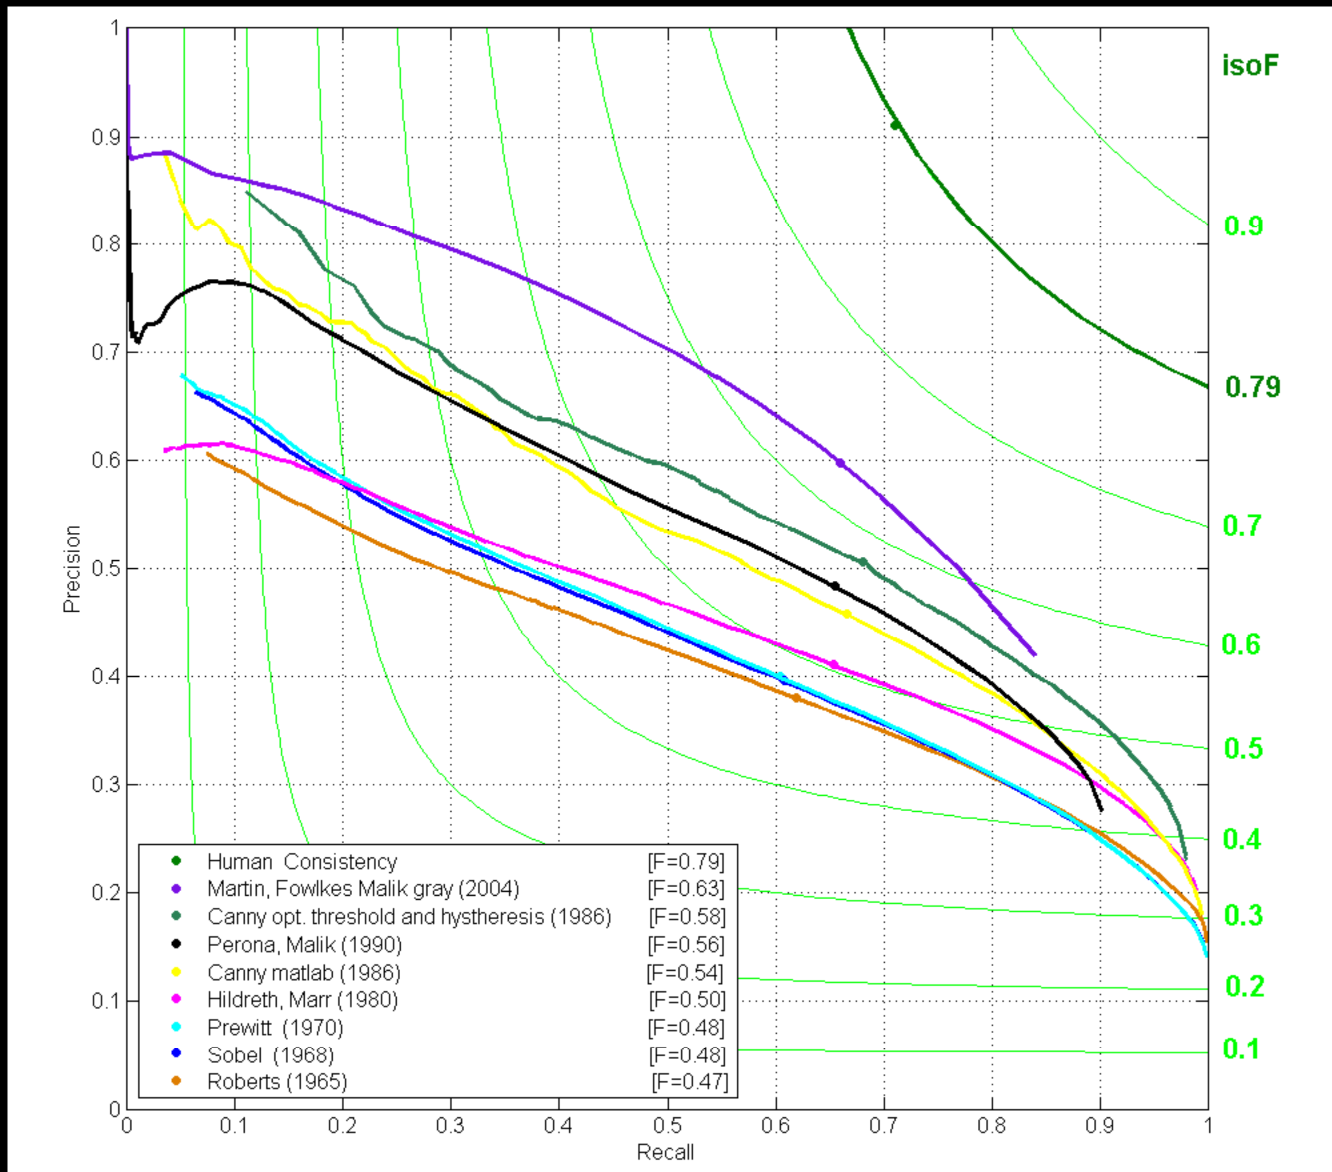

Contour detection ~2008 (gray)

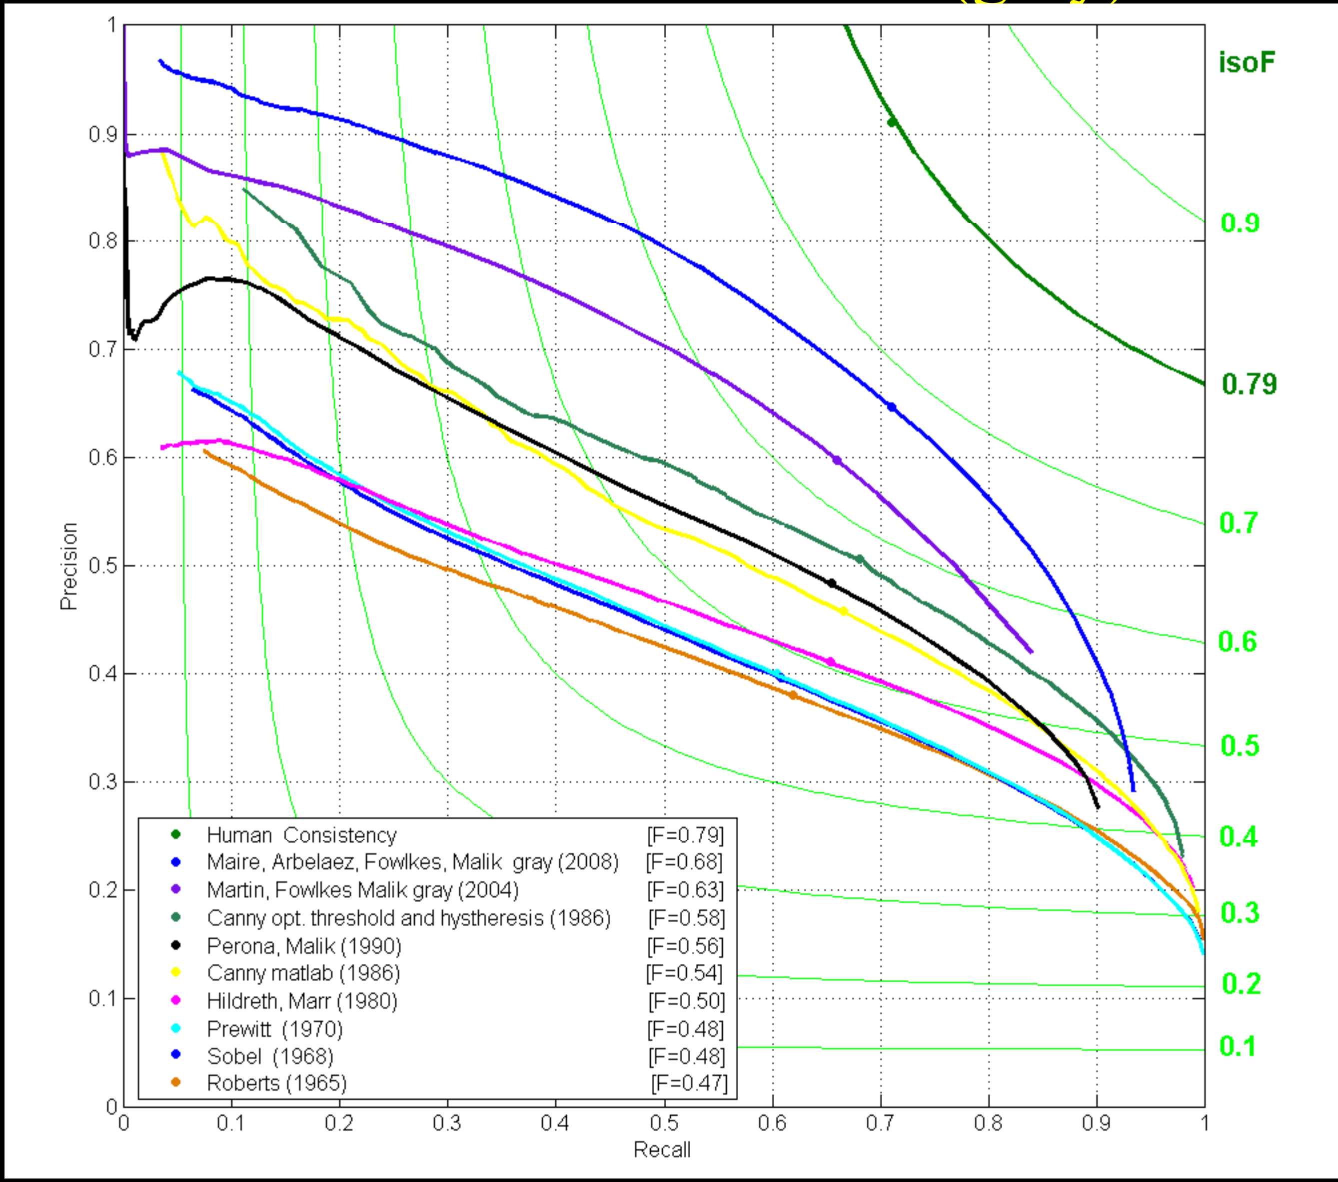

Contour detection ~2008 (color)

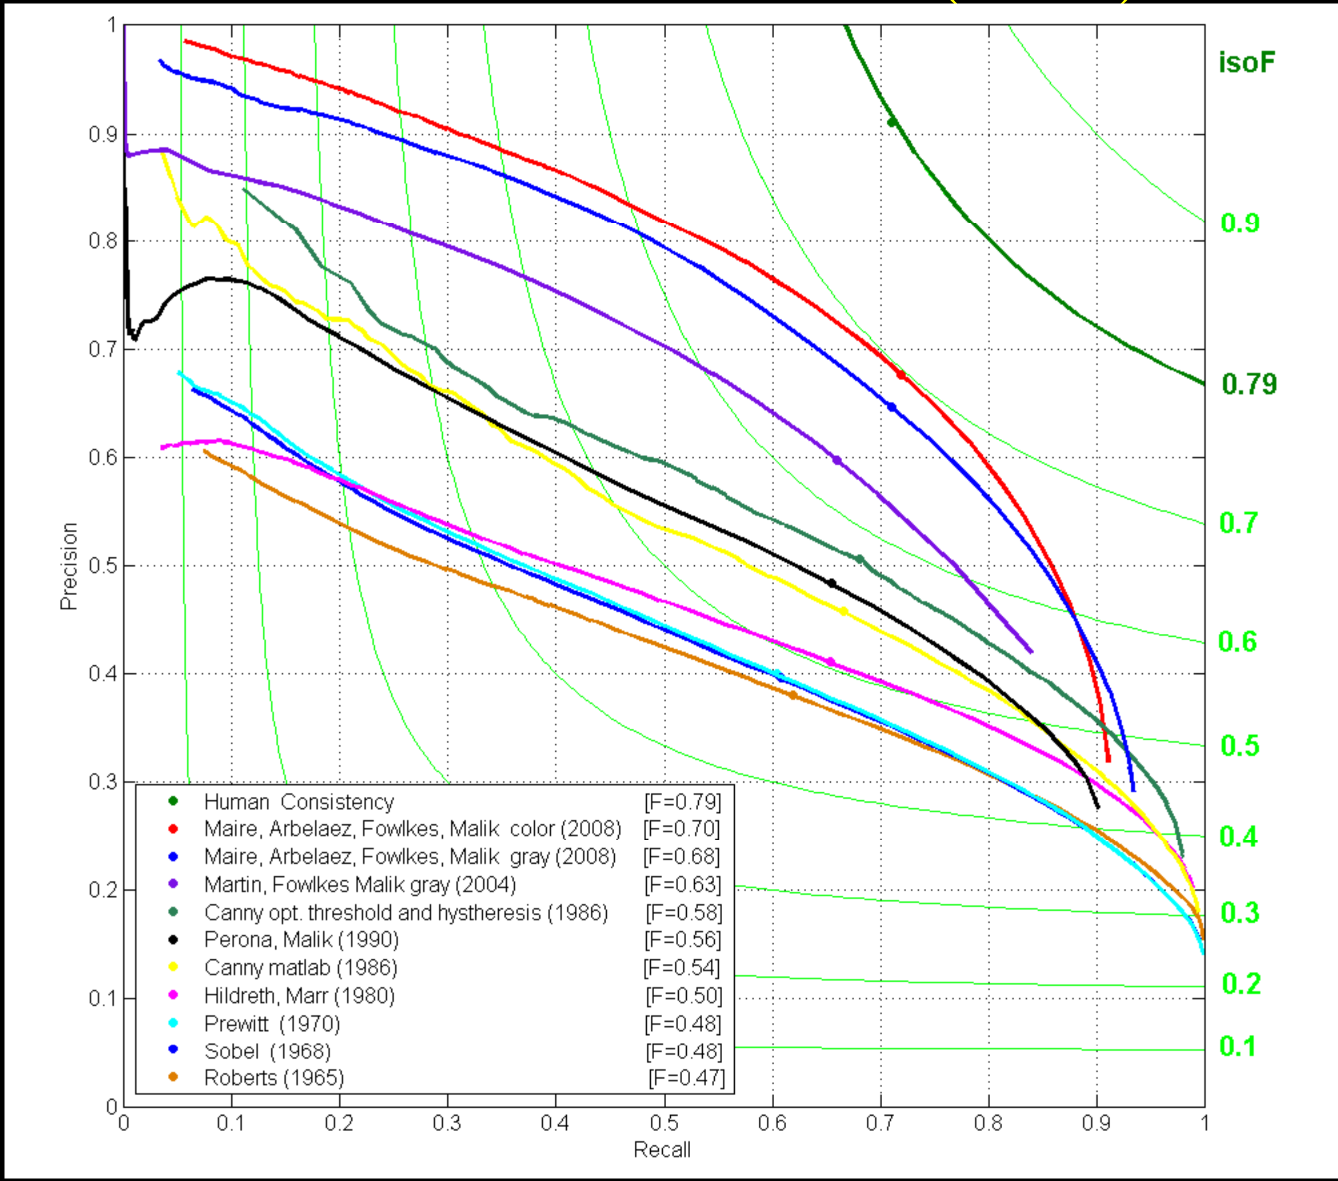

Outline

1. Collect Data Set of Human segmented images
2. Learn Local Boundary Model for combining brightness, color and texture
3. Global framework to capture closure, continuity
4. Detect and localize junctions
5. Integrate low, mid and high-level information for grouping and figure-ground segmentation

Contours can be defined by any of a number of cues (P. Cavanagh)

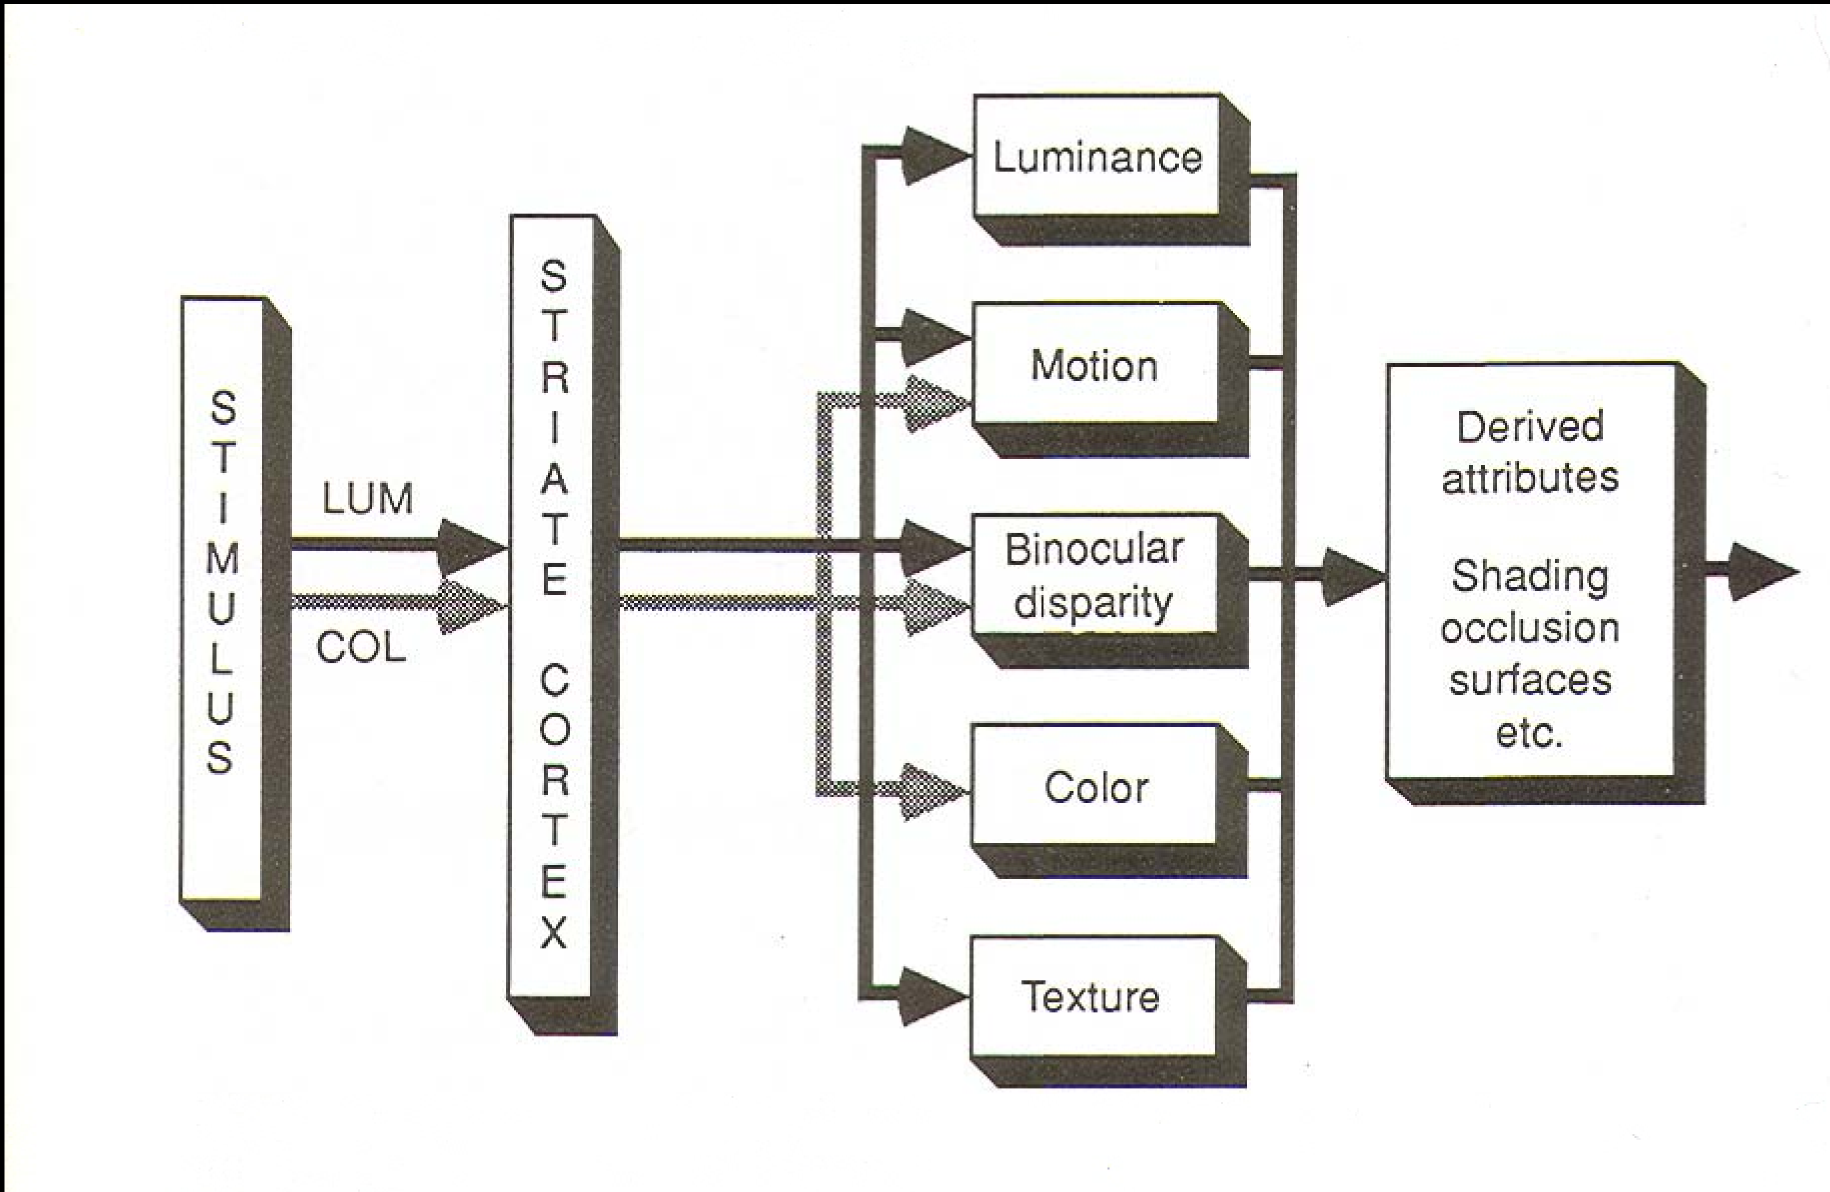

Cue-Invariant Representations

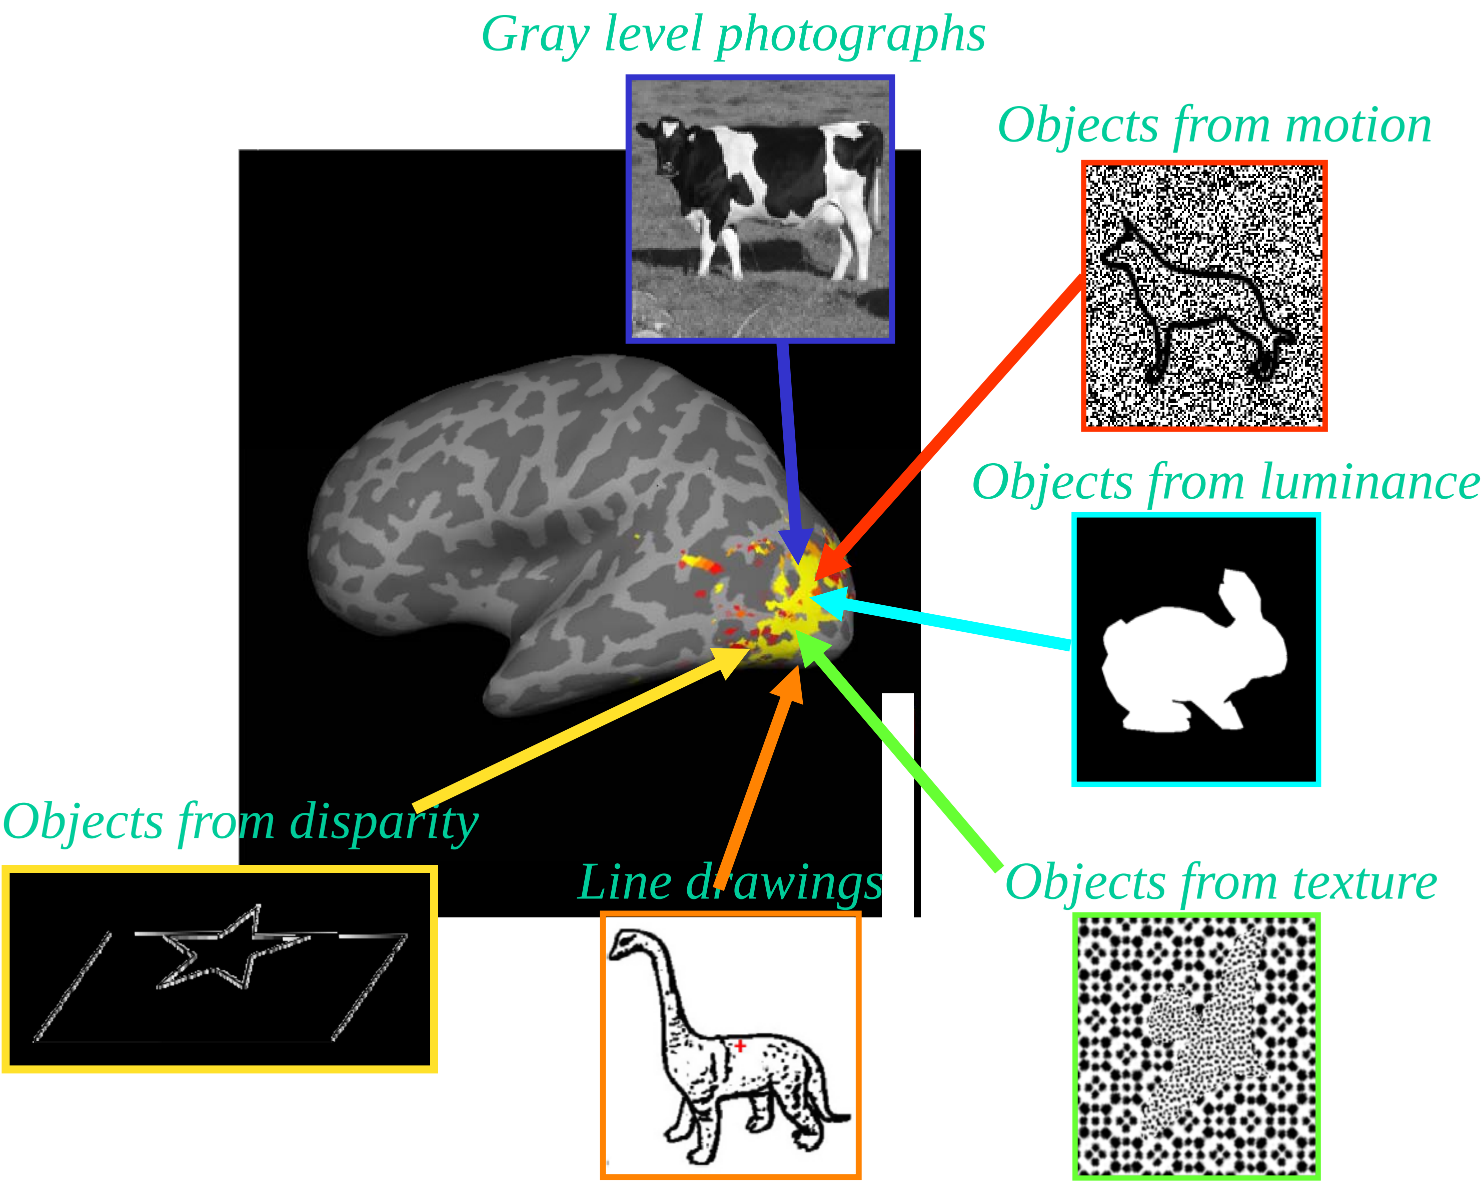

Grill-Spector et al. , Neuron 1998

Martin, Fowlkes, Malik PAMI 04

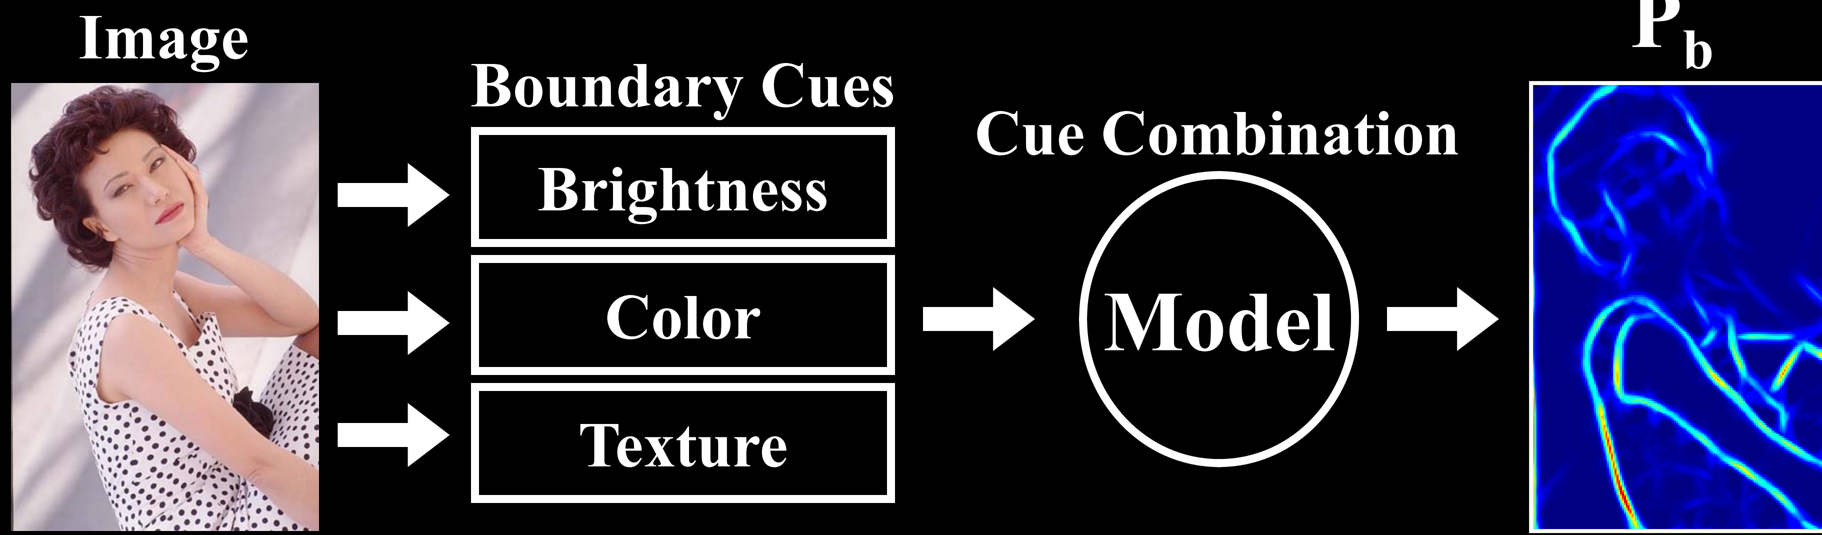

Challenges: texture cue, cue combination

Goal: learn the posterior probability of a boundary $P_b(x,y,\theta)$ from local information only

Individual Features

- 1976 CIE L*a*b* colorspace
- Brightness Gradient BG(x,y,r,θ)
  - Difference of L* distributions
- Color Gradient CG(x,y,r,θ)
  - Difference of a*b* distributions

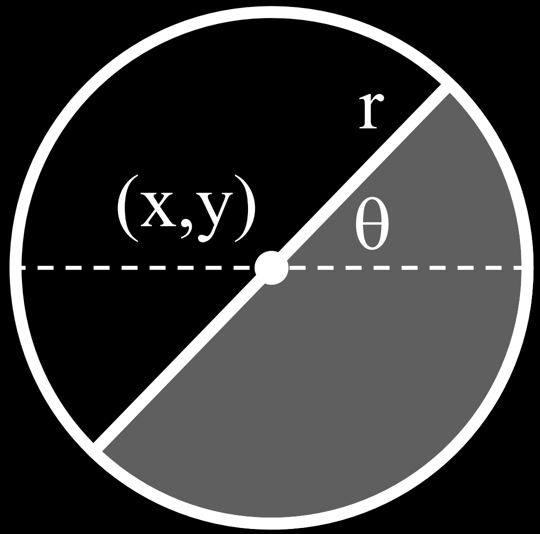

- Texture Gradient TG(x,y,r,θ)
  - Difference of distributions of V1-like filter responses

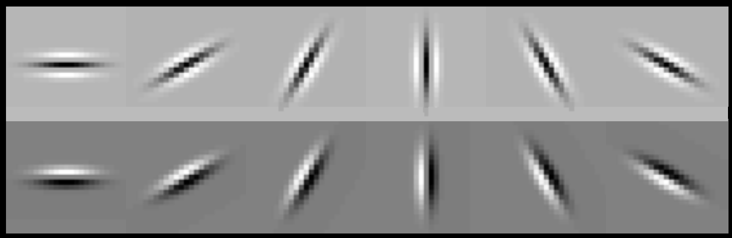

These are combined using logistic regression

Various Cue Combinations

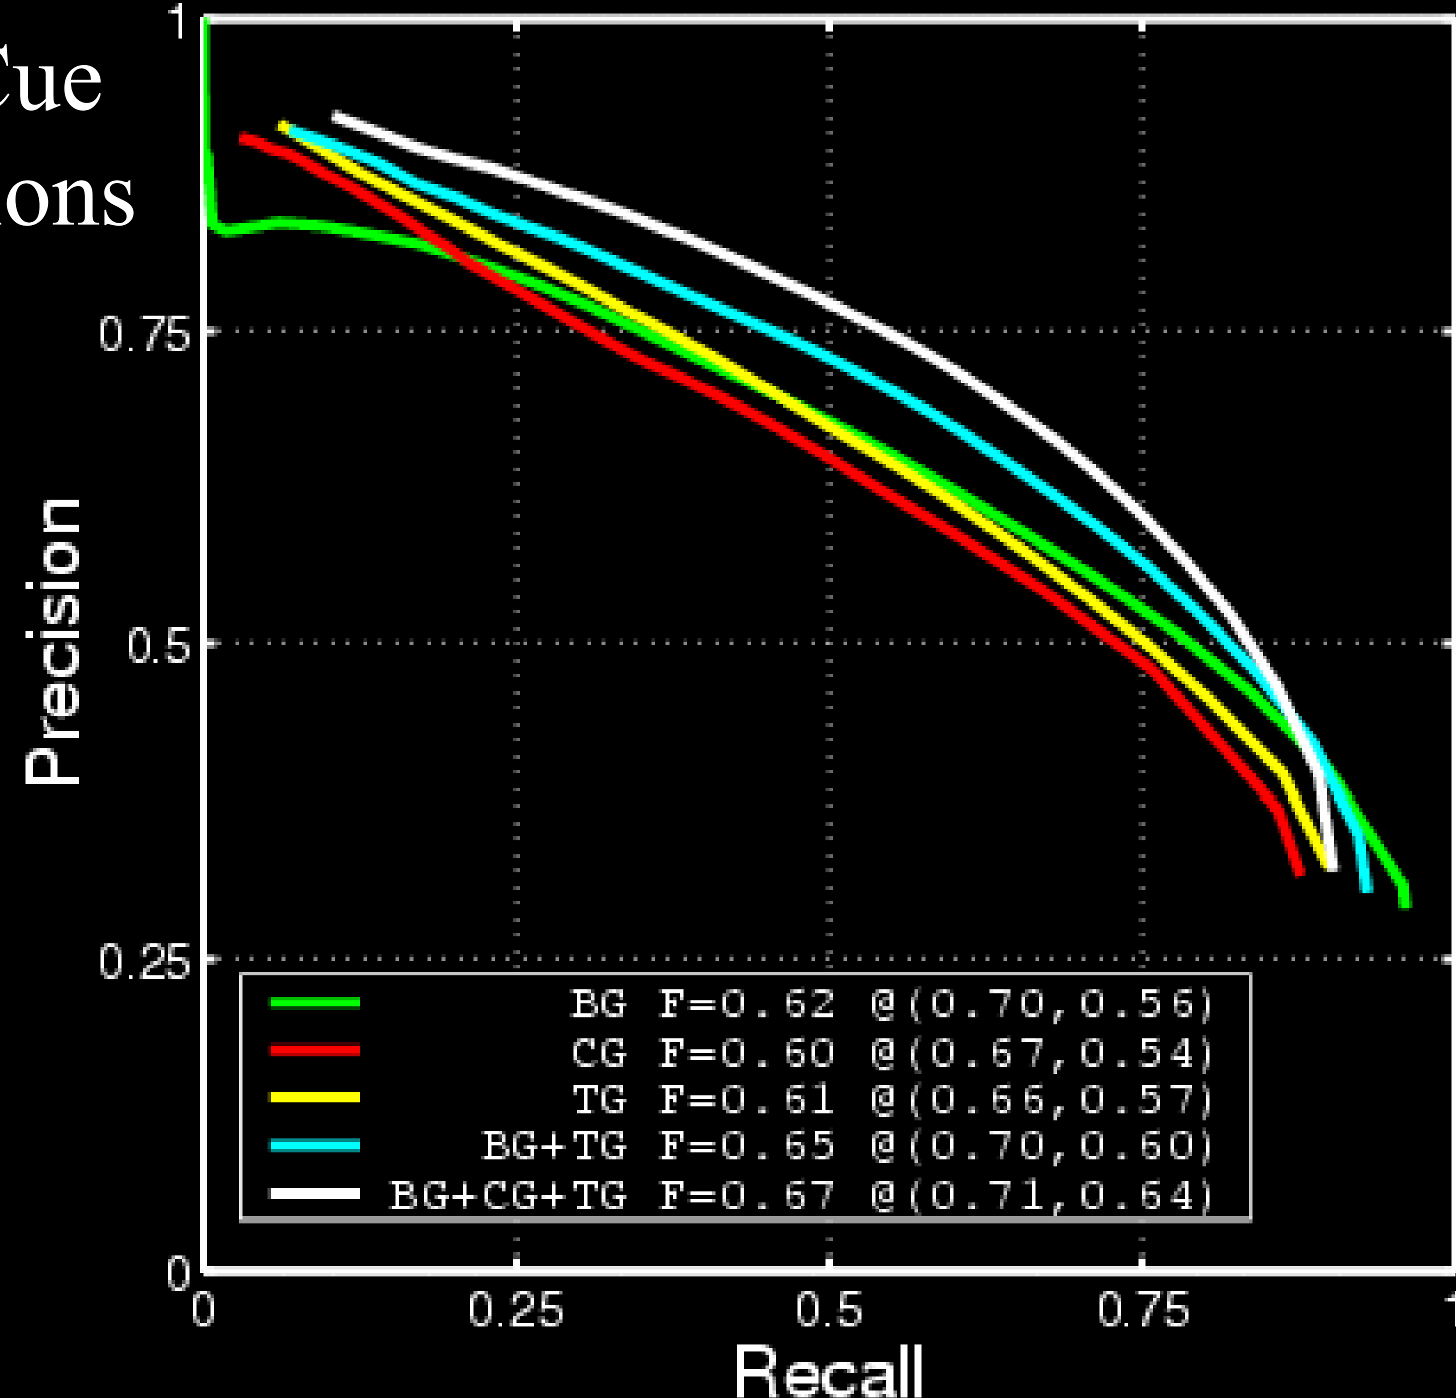

Outline

1. Collect Data Set of Human segmented images
2. Learn Local Boundary Model for combining brightness, color and texture
3. Global framework to capture closure, continuity
4. Detect and localize junctions
5. Integrate low, mid and high-level information for grouping and figure-ground segmentation

Exploiting global constraints: Image Segmentation as Graph Partitioning

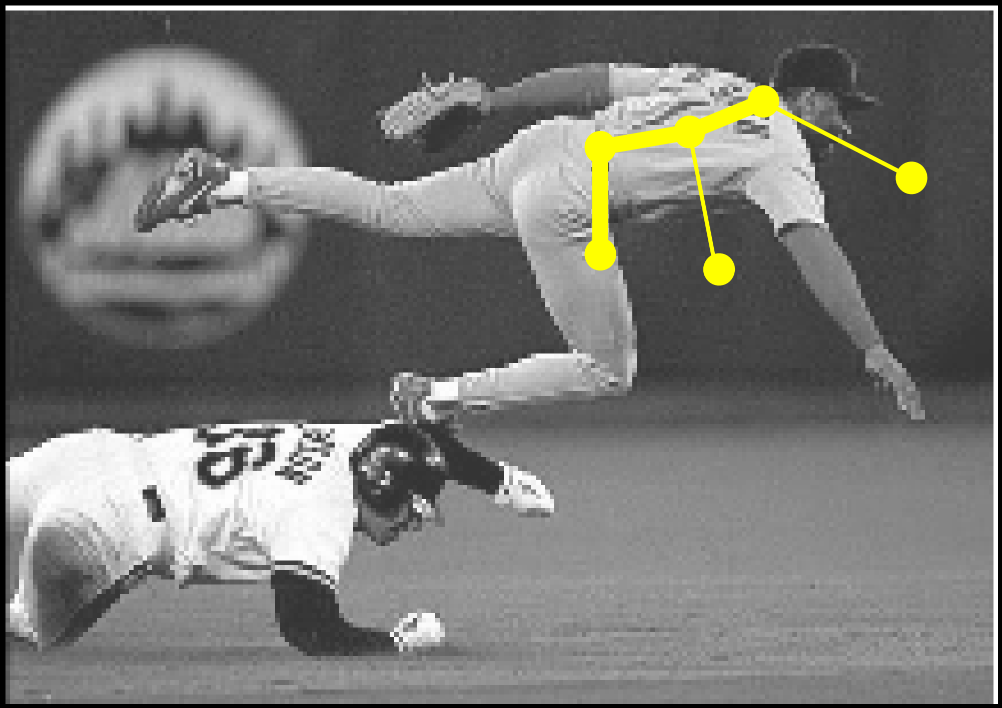

V: image pixels

E: connections between pairs of nearby pixels

Partition graph so that similarity within group is large and similarity between groups is small -- ***Normalized Cuts*** [Shi & Malik 97]

W_{ij} small when intervening contour strong, small when weak..

C_{ij} = max Pb(x,y) for (x,y) on line segment ij; $W_{ij} = \exp (- C_{ij} / \sigma$ □

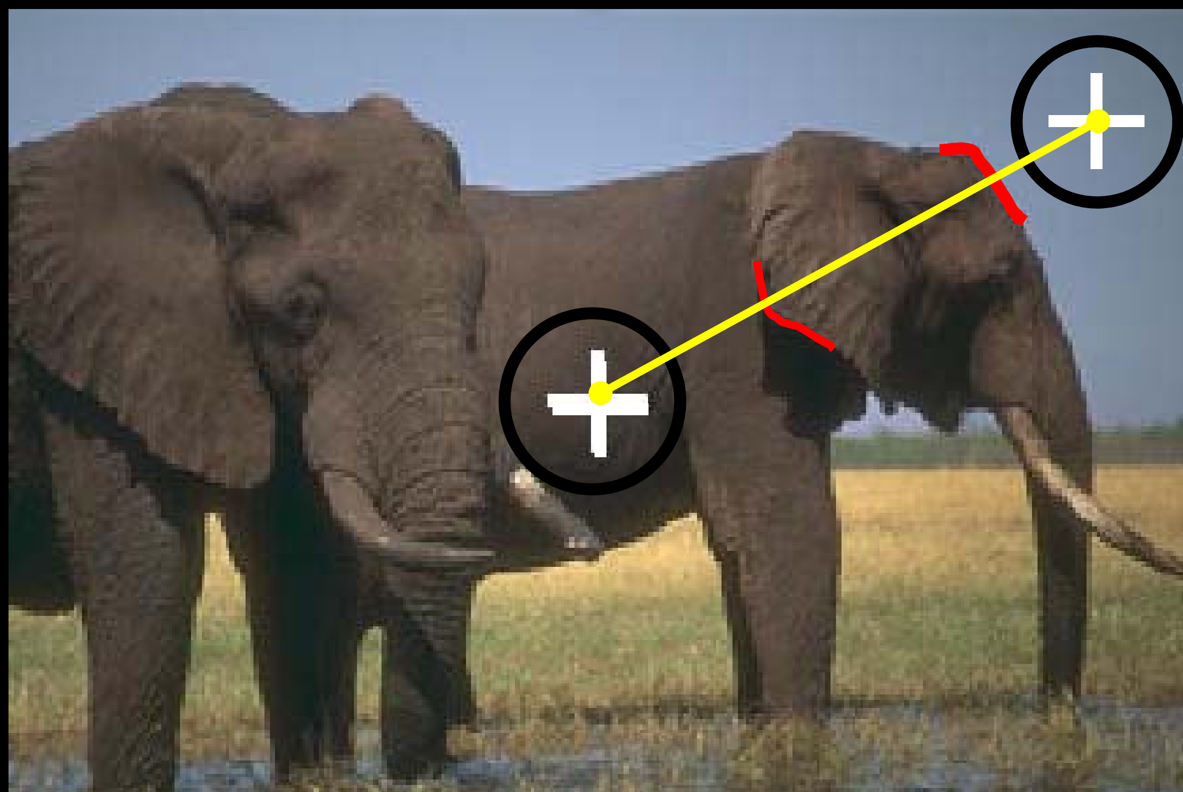

Eigenvectors carry contour information

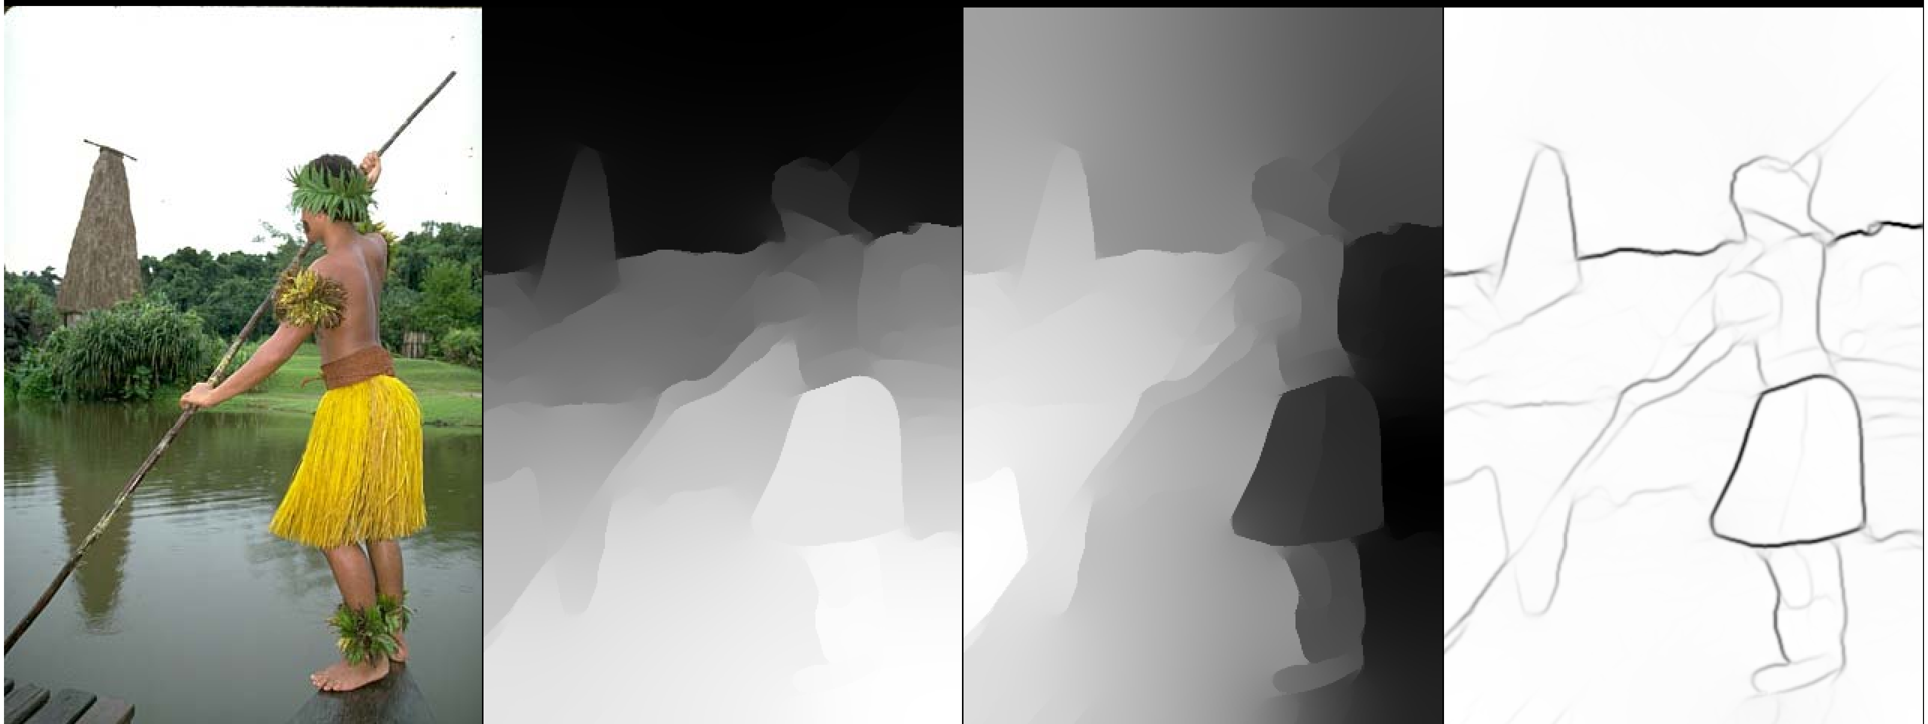

We do not try to find regions from the eigenvectors, so we avoid the “broken sky” artifacts of Ncuts …

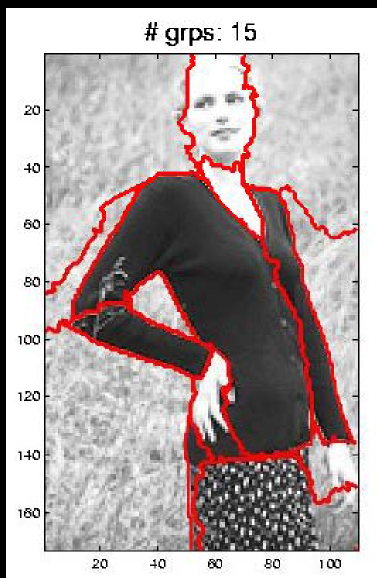

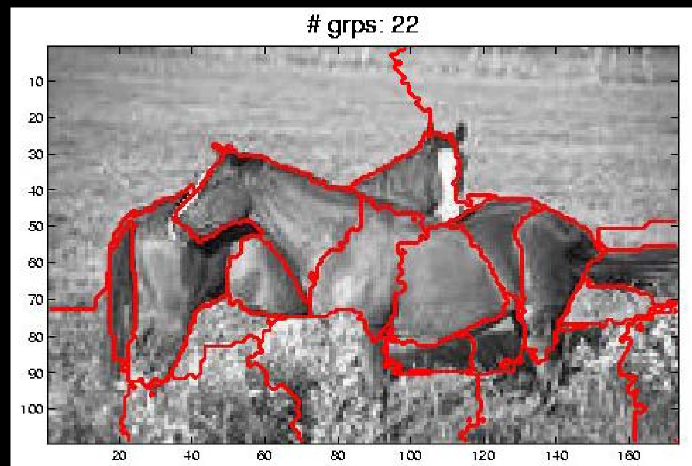

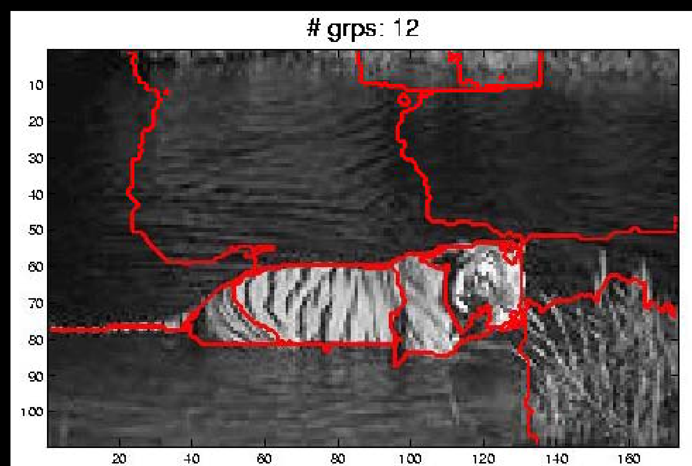

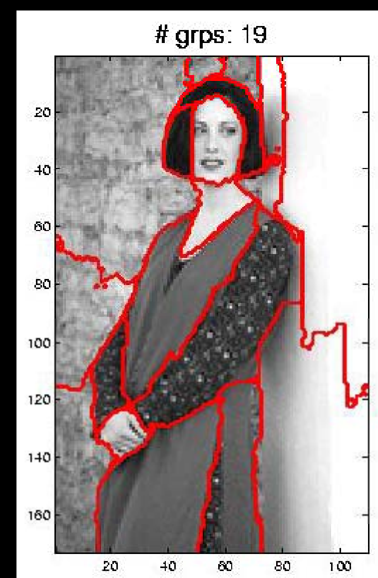

Key idea – compute edges *on ncut eigenvectors,* sum over first k:

$$sPb(x,y,\theta) = \sum_{j=1}^{k} \frac{1}{\sqrt{\lambda_j}} \cdot sPb_{\mathbf{v_j}}(x,y,\theta)$$

where $sPb_{\mathbf{v_j}}(x,y,\theta)$ is the output of a Gaussian derivative on the j-th eigenvector of $(D-W)\mathbf{v} = \lambda D\mathbf{v}$

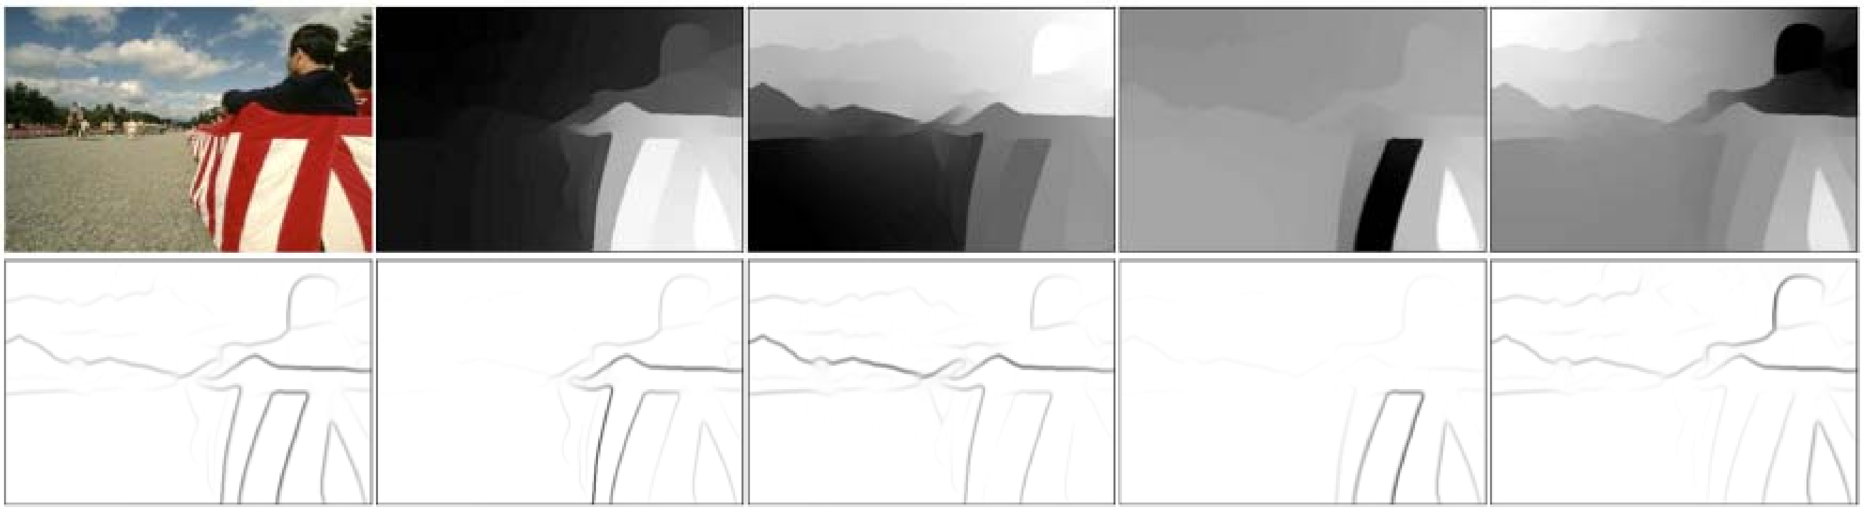

Figure 1. **Top:** Original image and first four generalized eigenvectors. **Bottom:** Maximum response over orientations θ of $sPb(x,y,\theta)$, and of $sPb_{\mathbf{v_j}}(x,y,\theta)$ for each eigenvector $\mathbf{v_j}$.

The Benefits of Globalization

Maire, Arbelaez, Fowlkes, Malik, CVPR 08

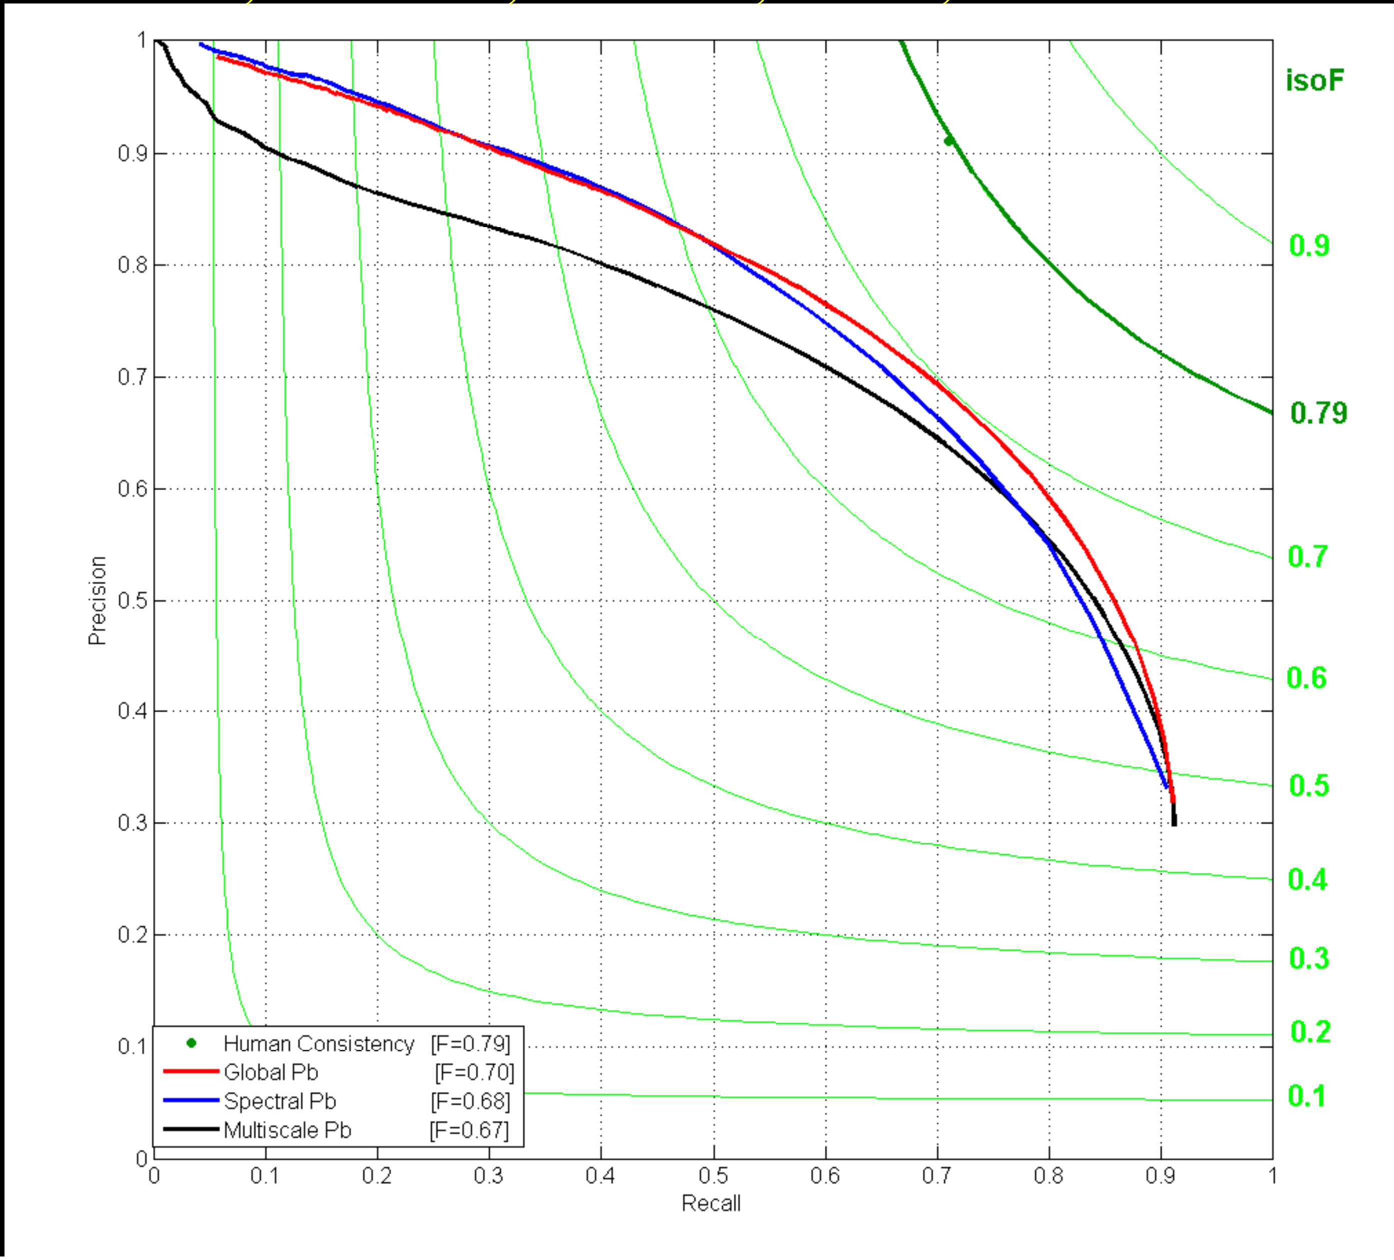

Comparison to other approaches

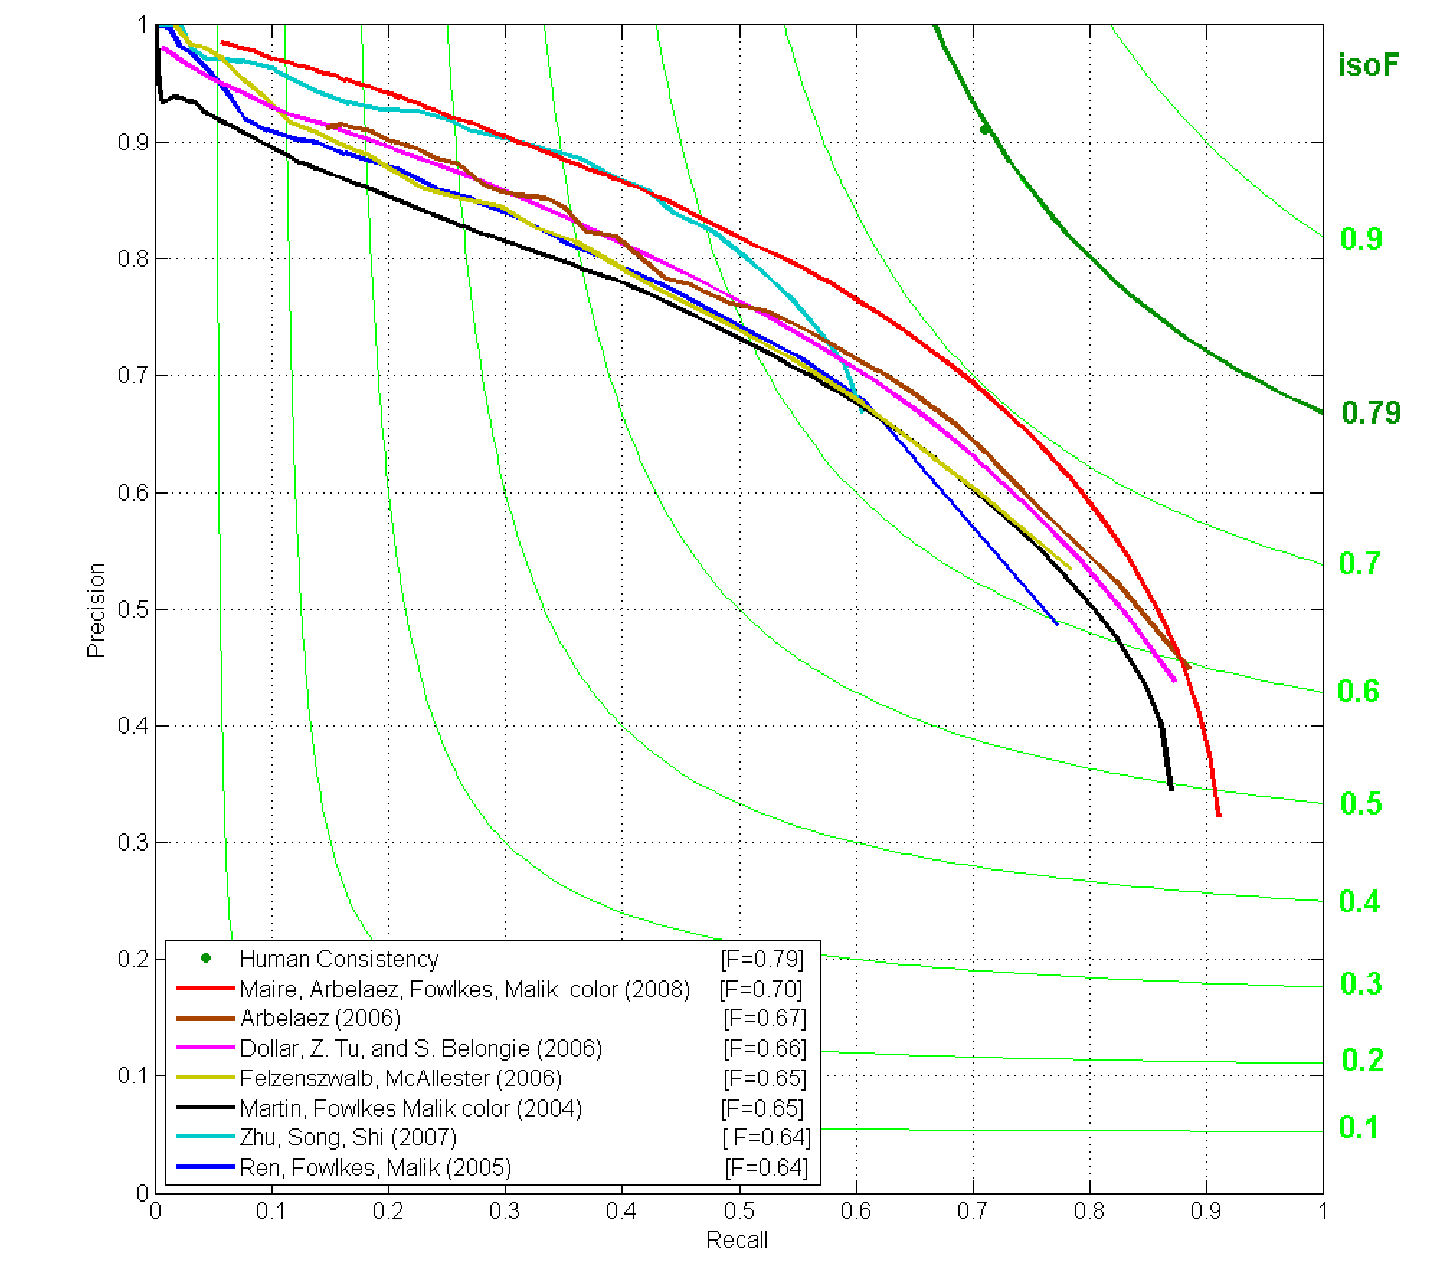

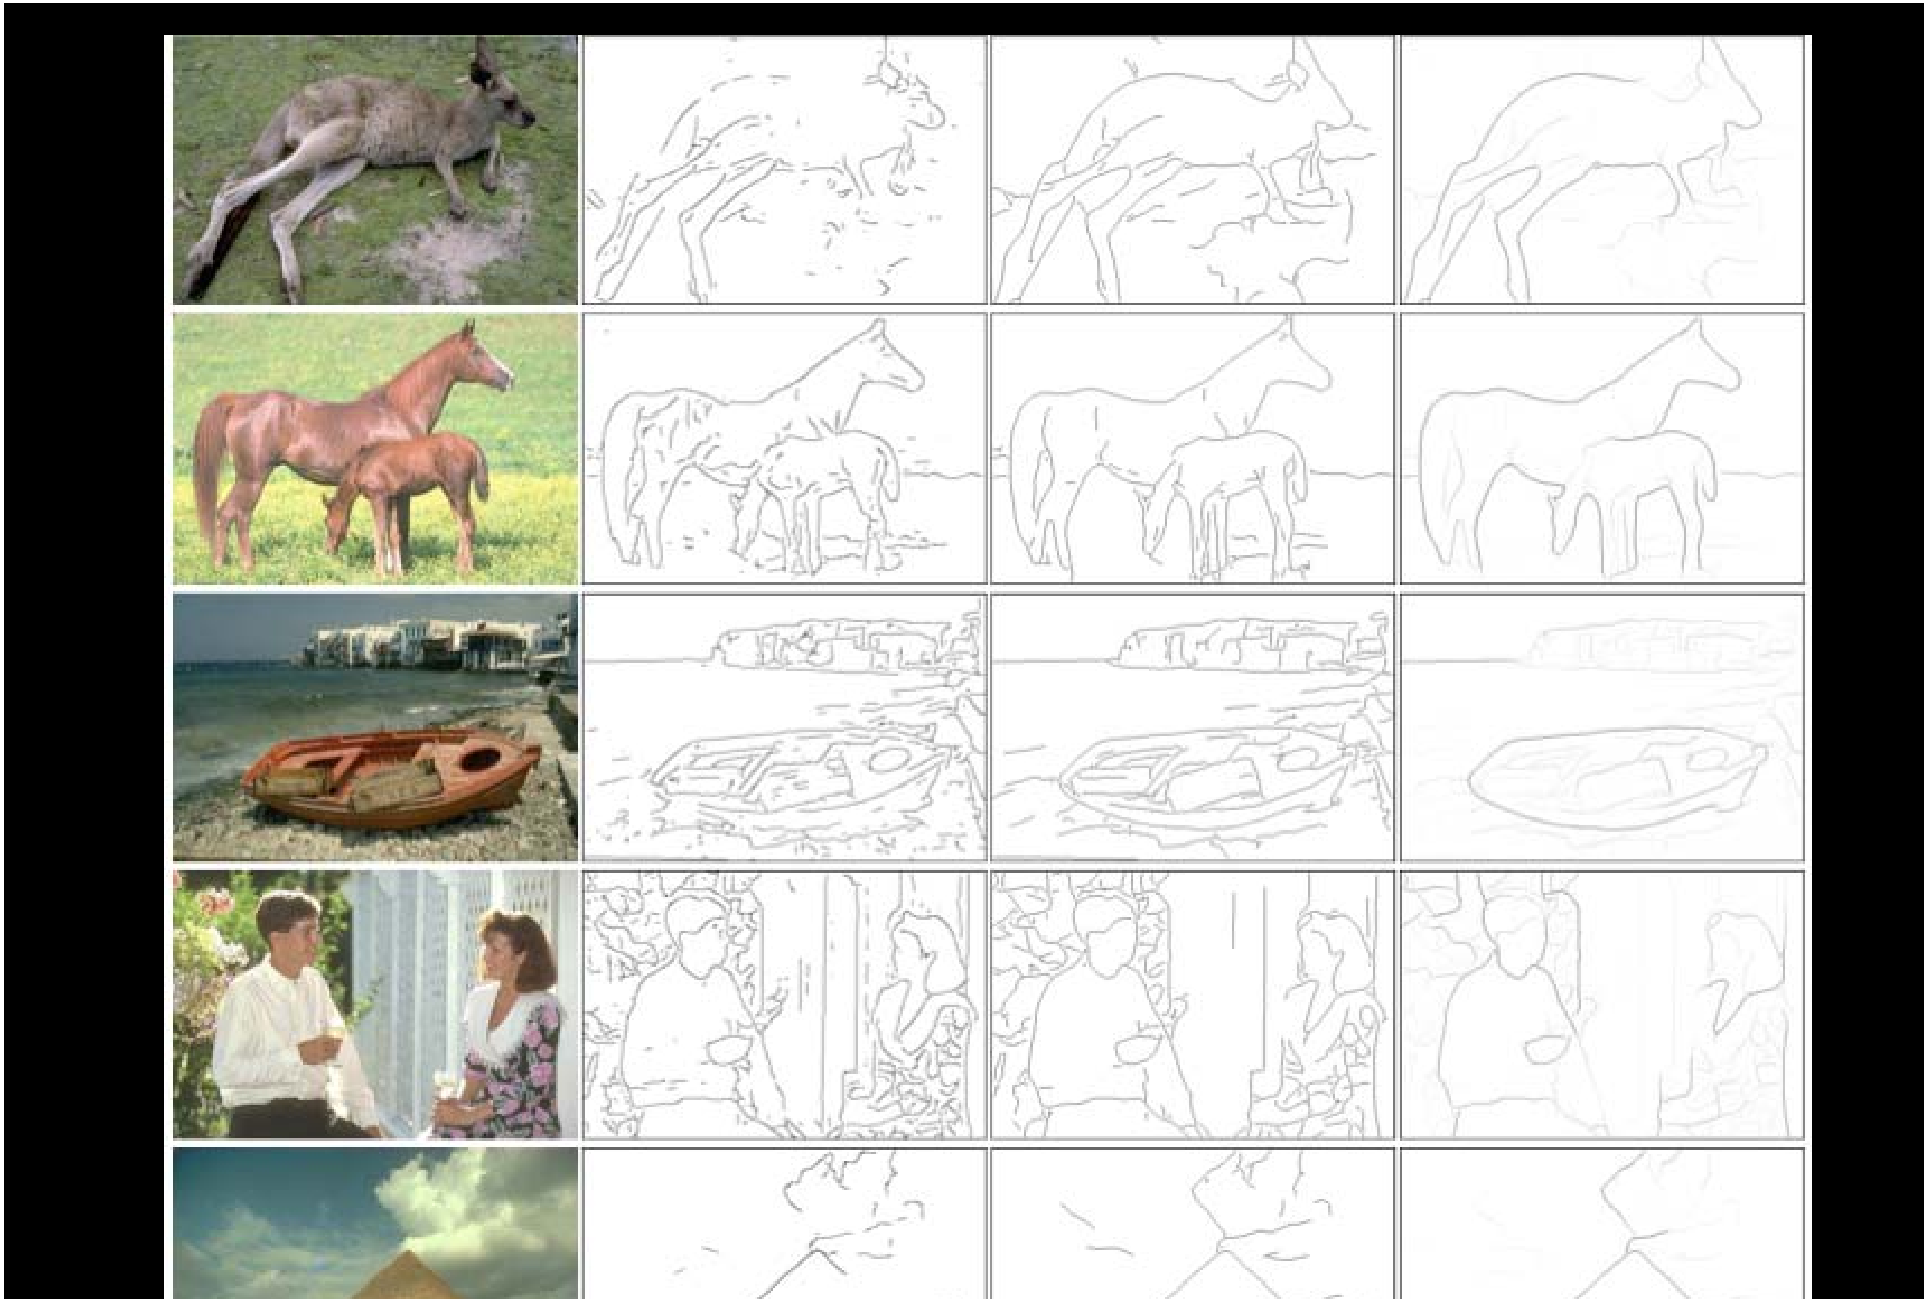

Figure 3. When compared with the local detector Pb, our detector gPb reduces clutter and completes contours. **From left to right:** Original image, thresholded Pb, thresholded gPb, and gPb. The thresholds shown correspond to the points of maximal F-measure on the curves in Figure 2.

Outline

1. Collect Data Set of Human segmented images
2. Learn Local Boundary Model for combining brightness, color and texture
3. Global framework to capture closure, continuity
4. Detect and localize junctions
5. Integrate low, mid and high-level information for grouping and figure-ground segmentation

Detecting Junctions

1. Estimate the optimal junction location $L = (x_L, y_L)$ by minimizing its weighted distance from the contours $\{C_i\} \in I_N$

$$L = argmin_{(x,y)\in I_N} \sum_i w_i d(C_i, (x, y)) \quad (7)$$

2. Update the weight w_i of each contour C_i in order to select only those contours passing close to the junction:

$$w_i = |C_i| \cdot \exp(-d(C_i, L)^2/\epsilon^2) \quad (8)$$

where $|C_i| = \sum_{(x,y)\in C_i} Pb(x, y)$ is the total contrast

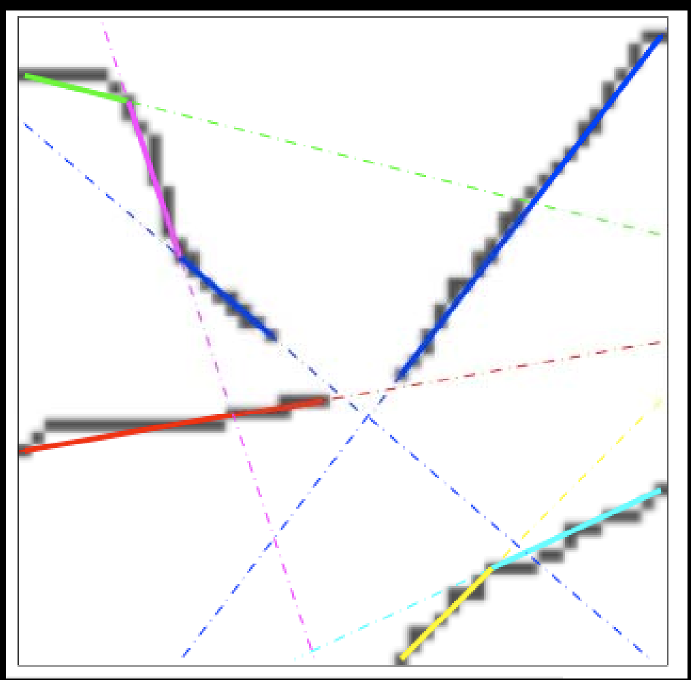

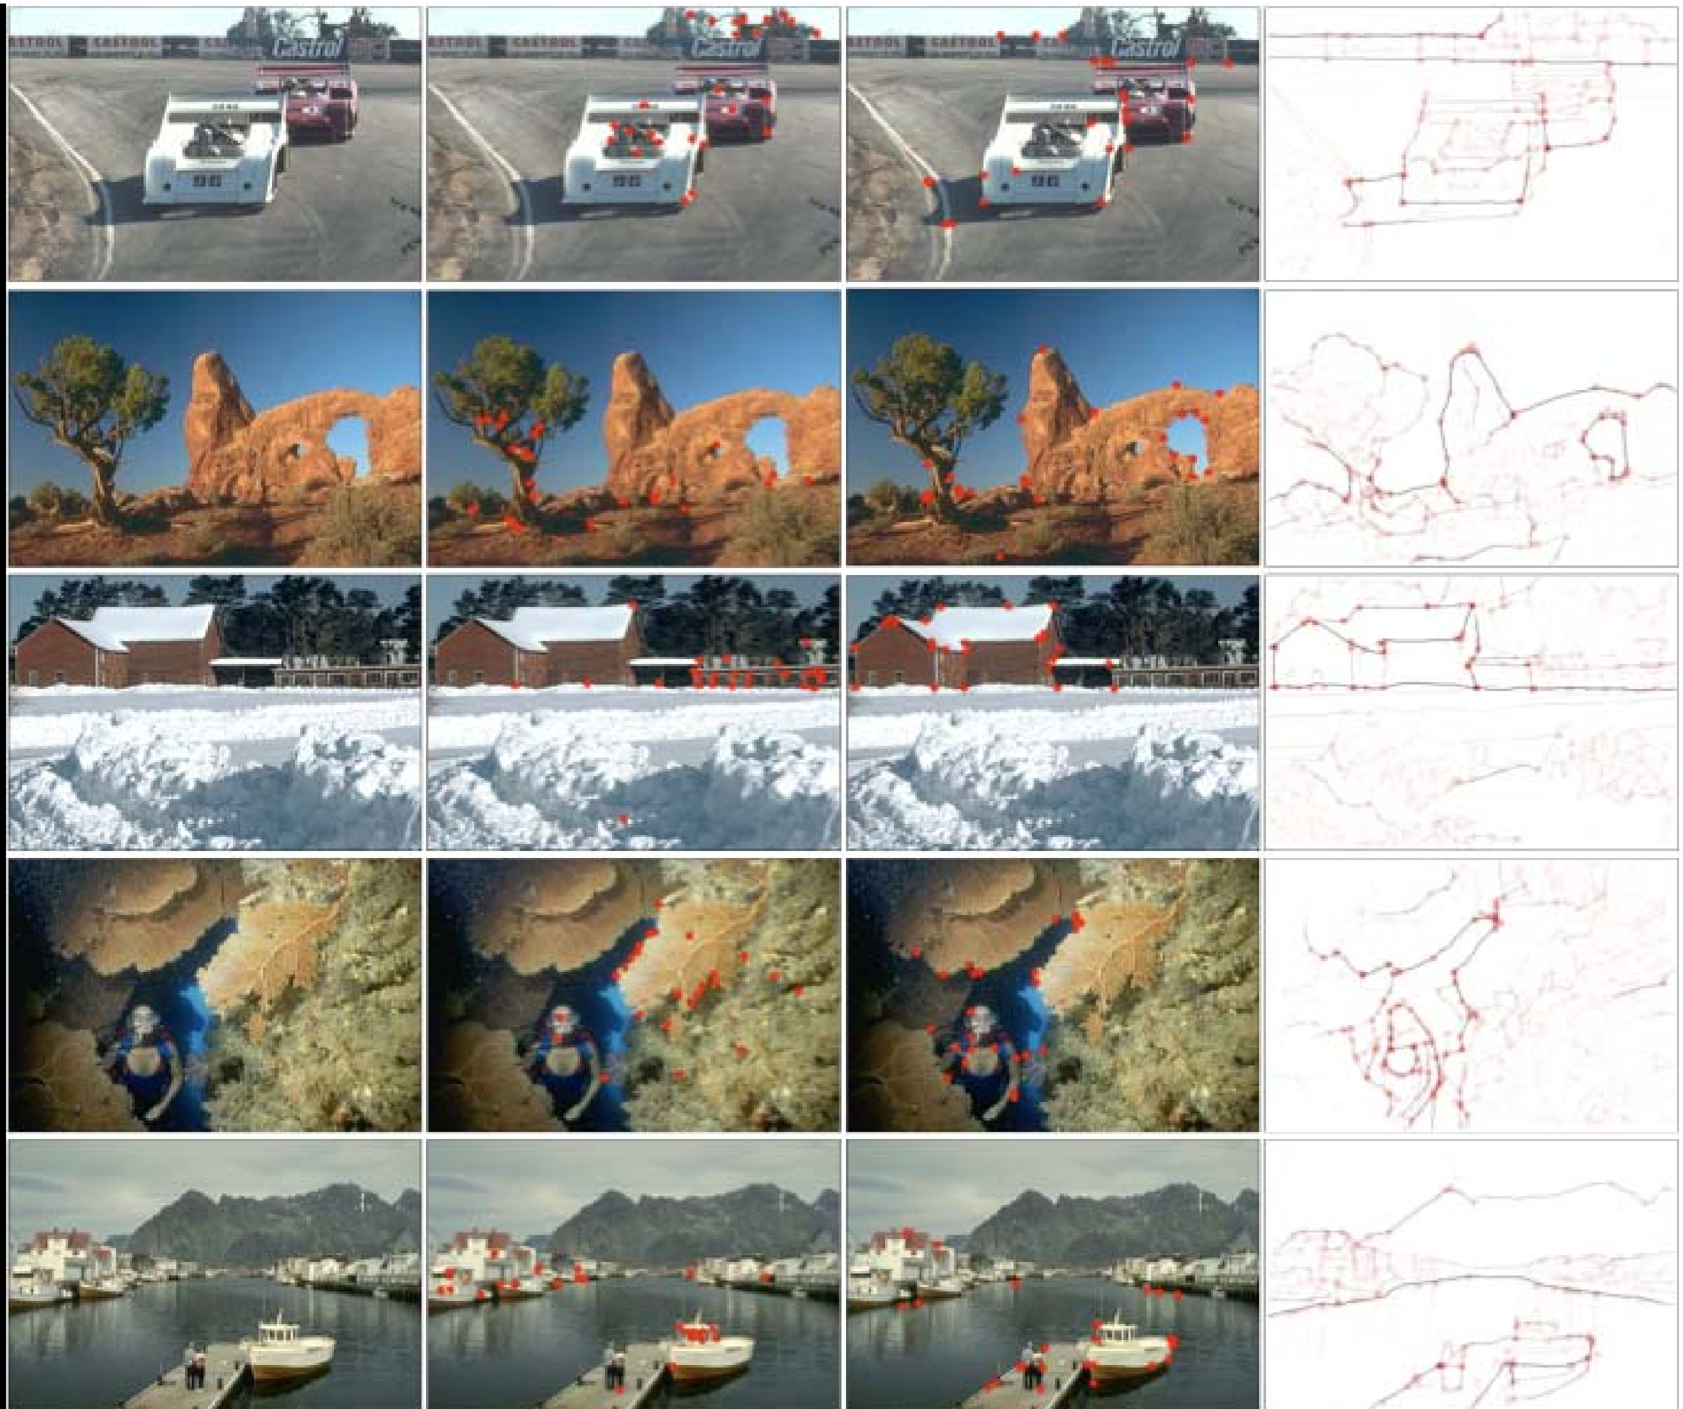
Castrol

Benchmarking corner detection

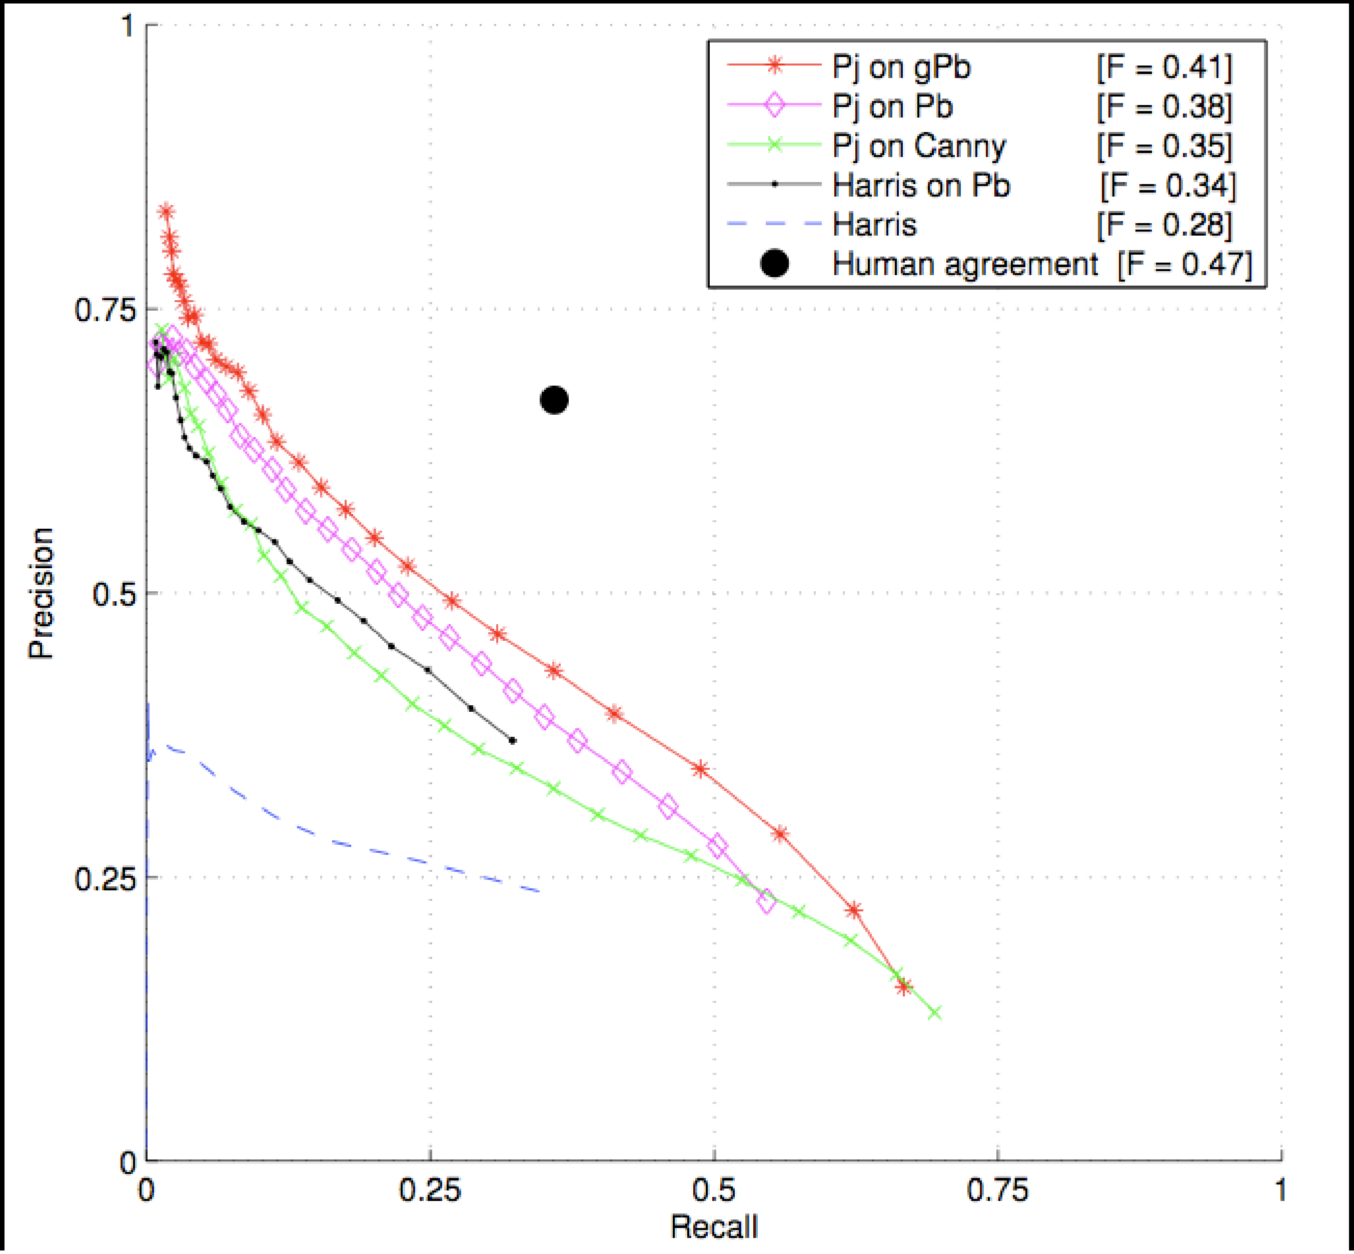

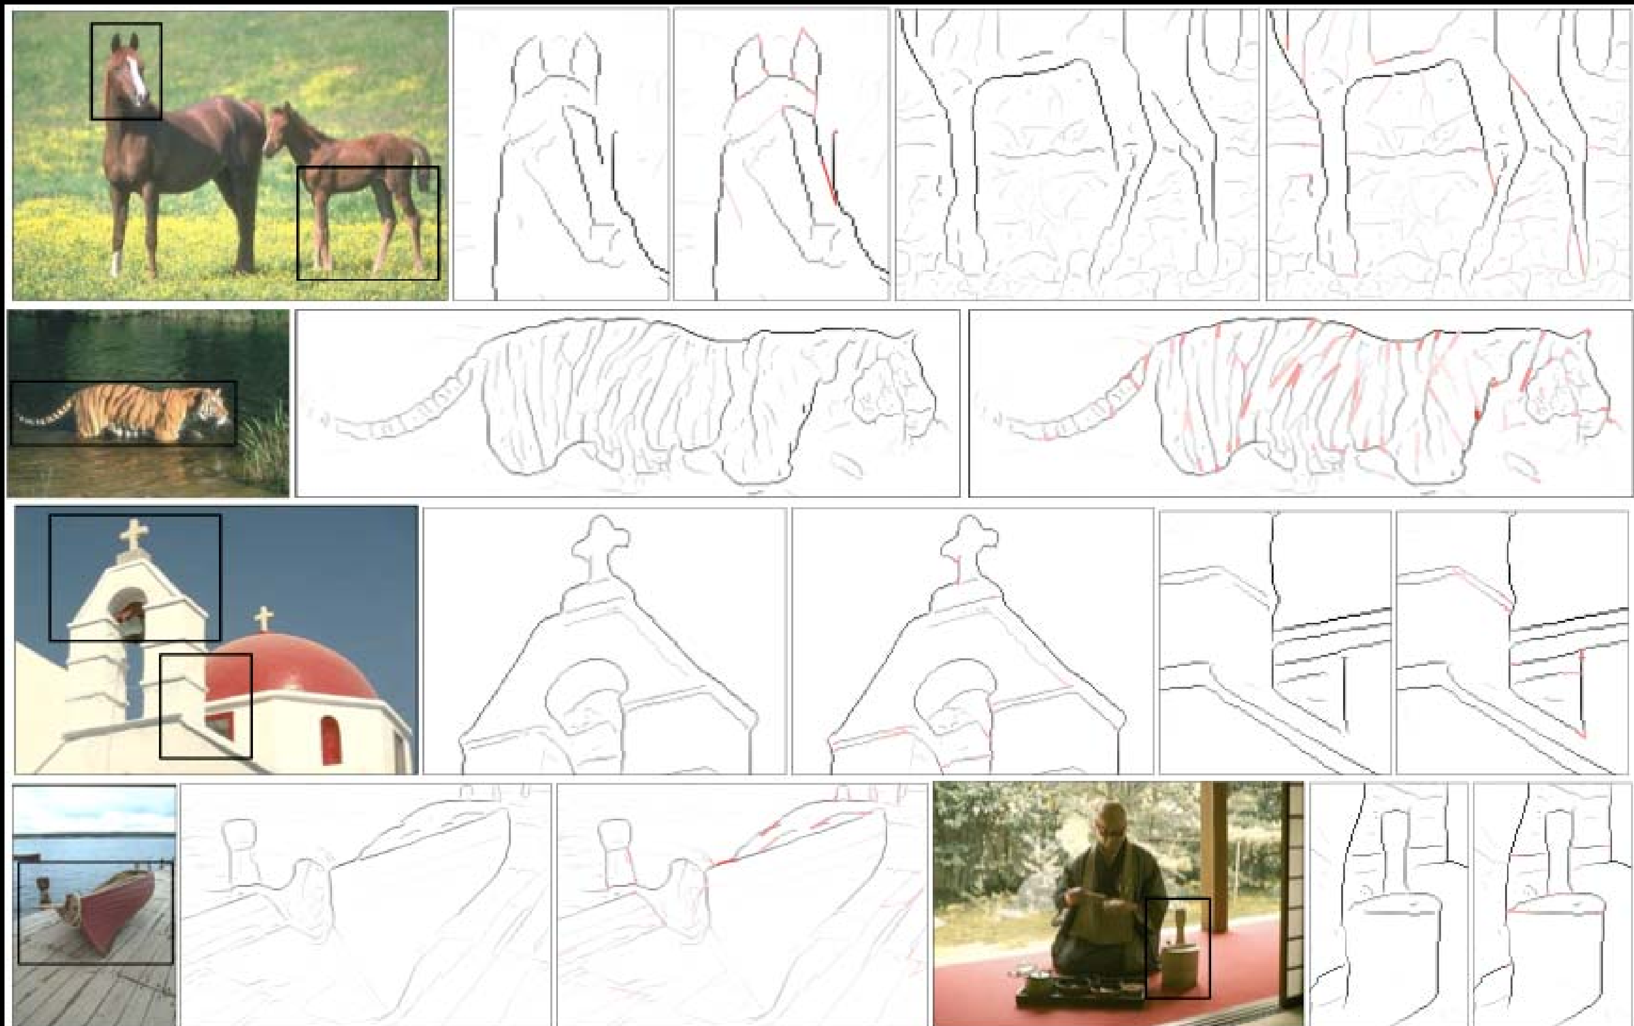

Figure 9. Junction restoration by extension of existing contours to detected junction locations. The magnified view of each boxed area shows contours before and after junction restoration. Extended contours are shown in red, with brightness corresponding to their estimated strength.

Better object recognition using previous version of Pb

- Ferrari, Fevrier, Jurie and Schmid (PAMI 08)

The main reason for preferring the Berkeley detector over the traditional Canny detector, is the inclusion of texture and color segmentation cues, in addition to brightness. Moreover, it treats edge detection as a pixel classification problem and trains a classifier from natural images with human-annotated boundaries. This results in less clutter edgels inside textured areas, and longer, smoother boundaries around textured objects (e.g. giraffes). Using this detector instead

- Shotton, Blake and Cipolla (PAMI 08)

	Classification ROC AUC	Detection RP AUC
Canny	0.9127	0.8498
Berkeley [38]	0.9275	0.8871
BEL [18] Natural	0.9029	0.8354
BEL [18] Horse	0.9518	0.8976

The Berkeley detector performs considerably better than Canny, especially for detection. While the

Outline

1. Collect Data Set of Human segmented images
2. Learn Local Boundary Model for combining brightness, color and texture
3. Global framework to capture closure, continuity
4. Detect and localize junctions
5. Integrate low, mid and high-level cues for grouping and figure-ground segmentation
    1. Ren, Fowlkes, Malik, IJCV '08
    2. Fowlkes, Martin, Malik, JOV '07
    3. Ren, Fowlkes, Malik, ECCV '06

Power laws for contour lengths

Fig. 2 We take each boundary contour C and break it up at local curvature maxima (corners). Shown are a few examples of this decomposition. The contour segment length $|L|$ is large for large-scale features such as the back of an animal, and is small in the vicinity of fine-scale details such as the head

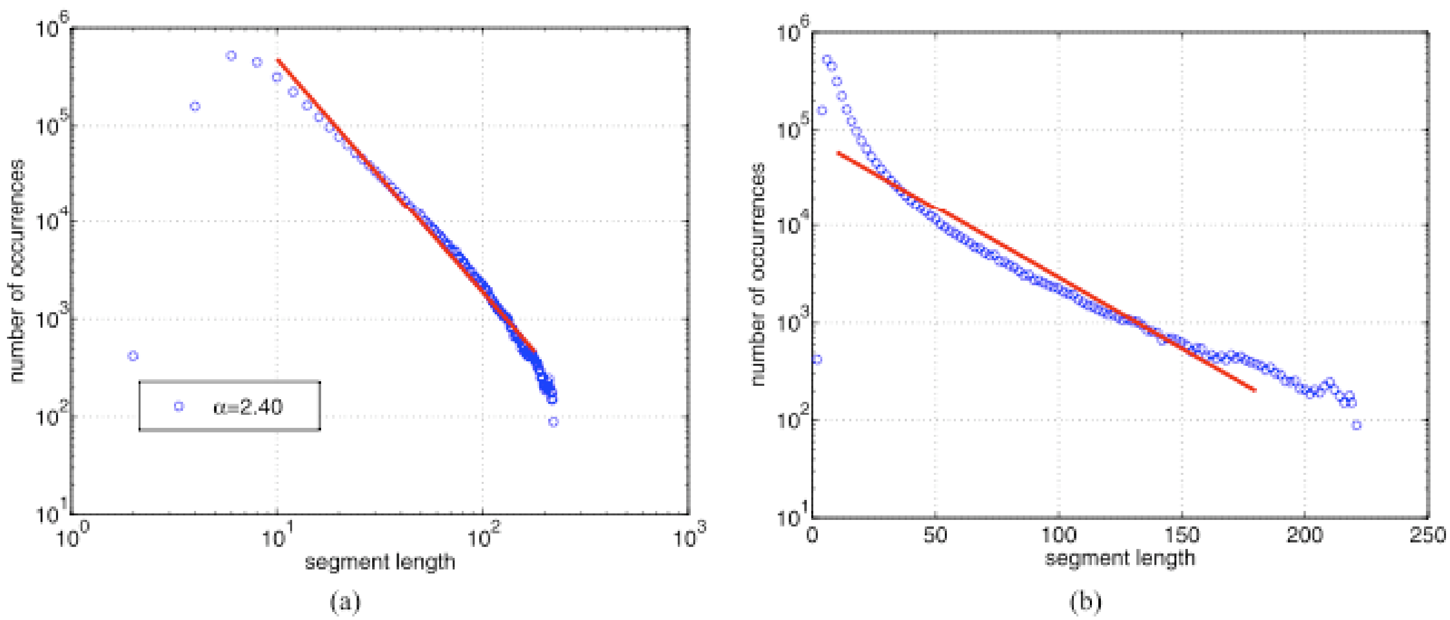

Fig. 3 Empirical distributions of approximately straight contour segment length $|L|$. (**a**) The marginal distribution of $|L|$ on a log-log scale. It closely follows a power law with $\alpha = 2.40$, in direct contradiction with the Markov model that predicts an exponential distribution. (**b**) As a comparison, the same distribution on a semi-log scale and an exponential fit. The power law fits the data much better than the exponential

Convexity

[Metzger 1953, Kanizsa and Gerbino 1976]

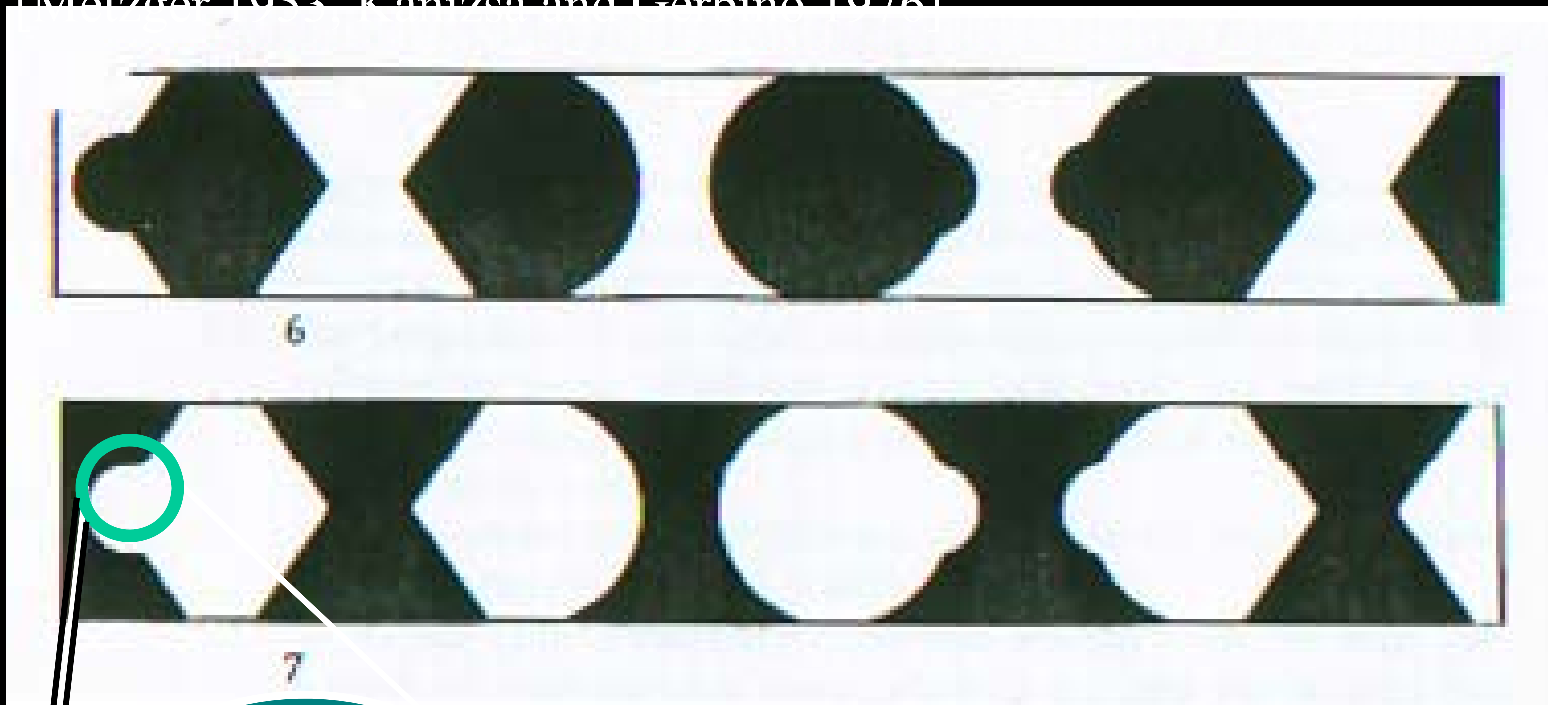

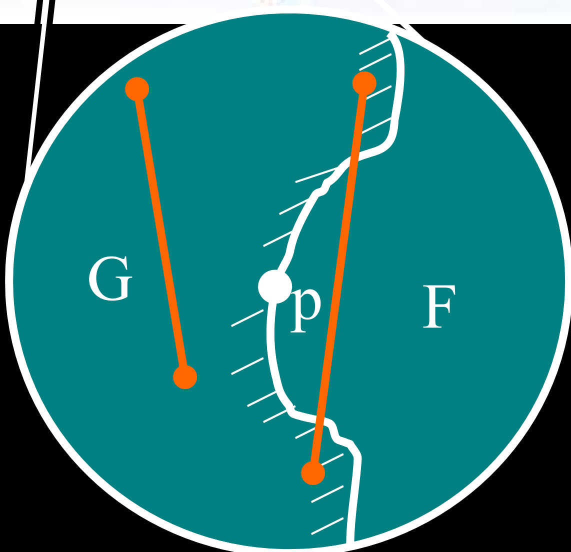

Conv_G = percentage of straight lines that lie completely within region G

$$\mathrm{Convexity}(p) = \log(\mathrm{Conv}_F \,/\, \mathrm{Conv}_G)$$

Figural regions tend to be convex

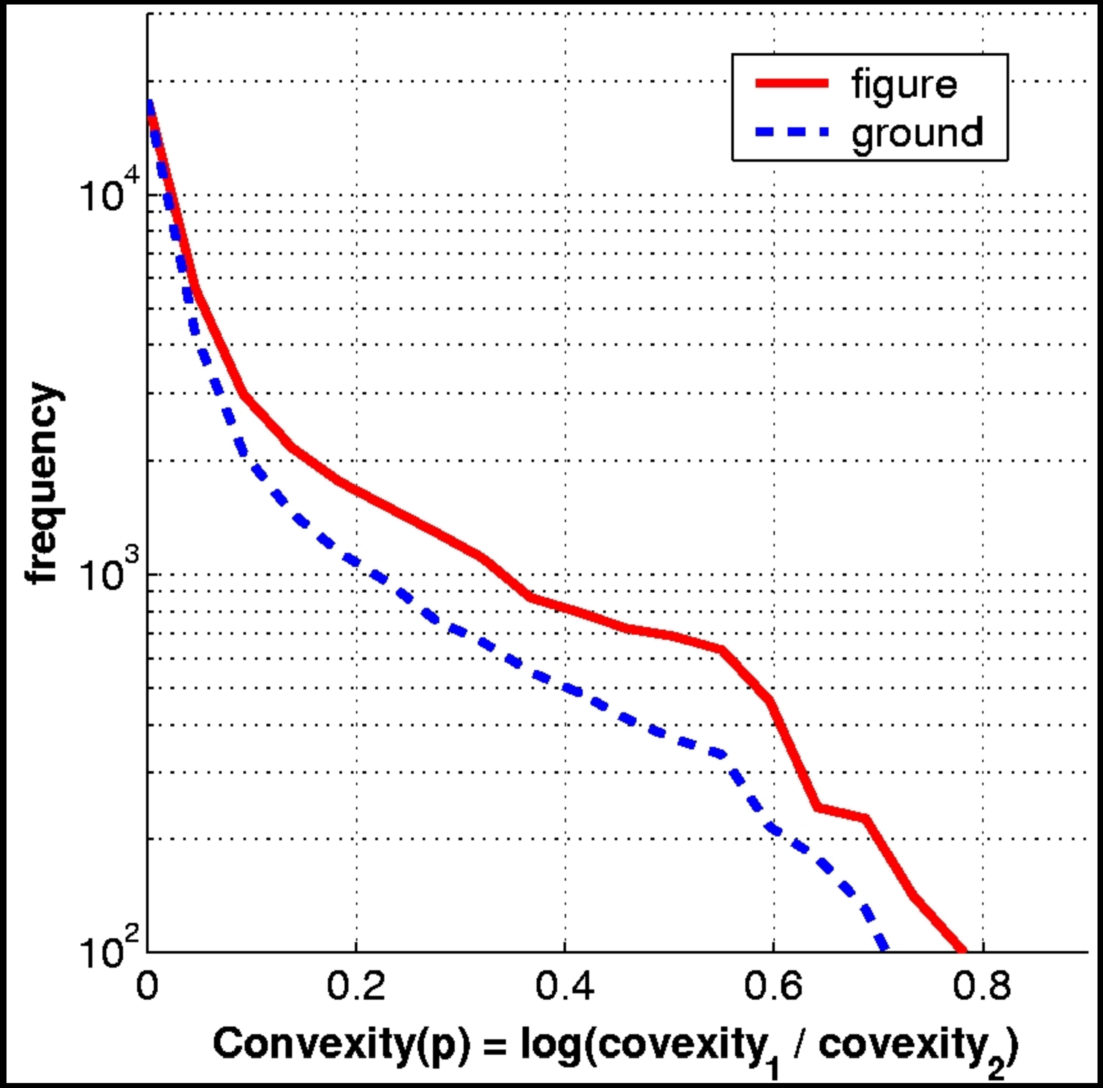

Lower Region

[Vecera, Vogel & Woodman 2002]

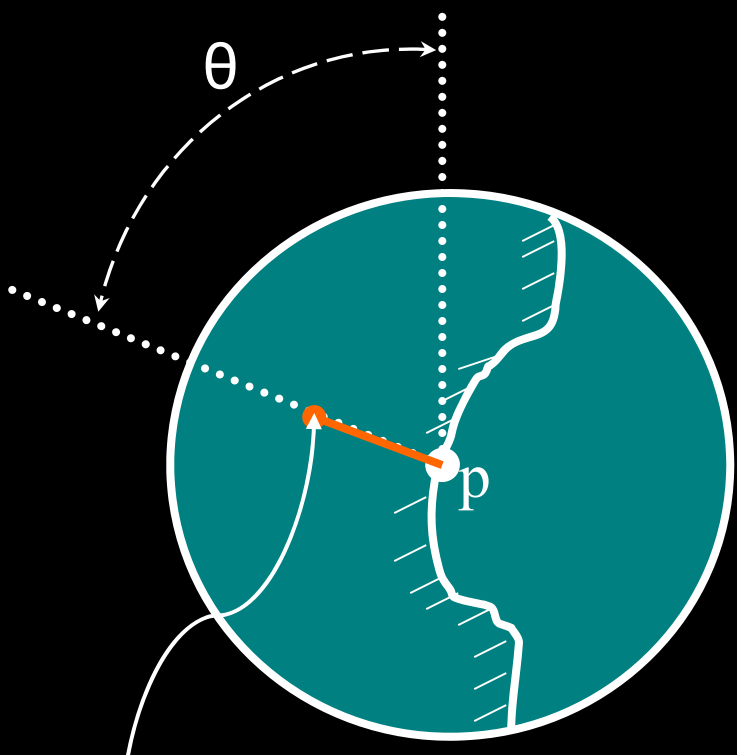

center of mass

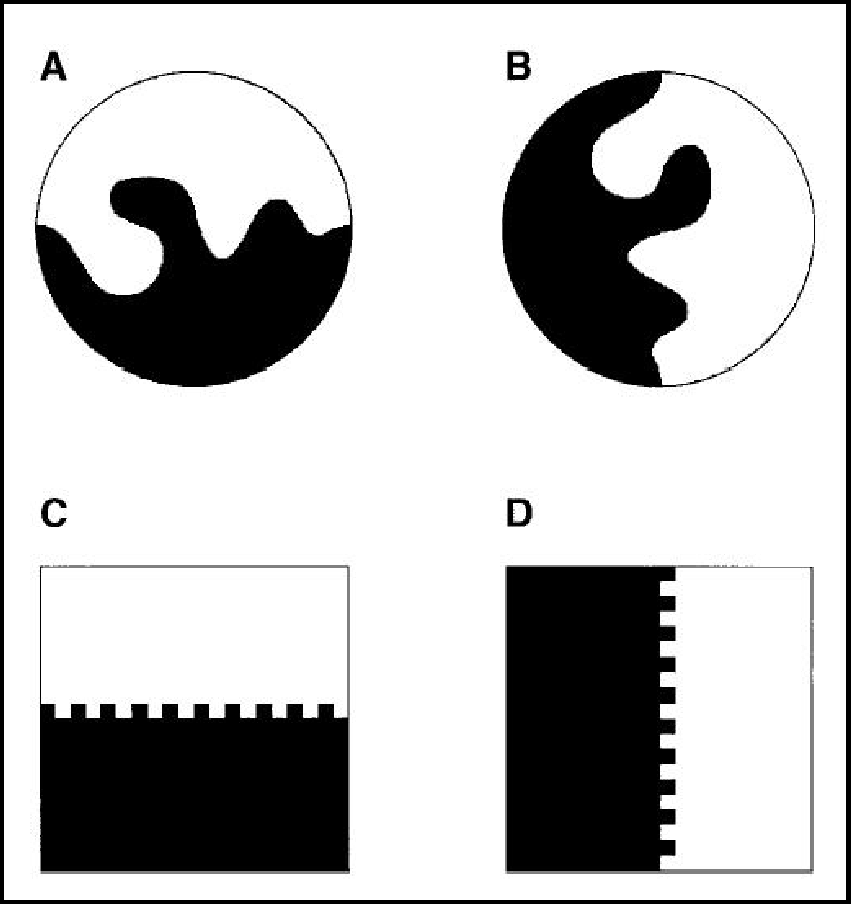

$$\mathrm{LowerRegion}(p) = \theta_G$$

Figural regions tend to lie below ground regions

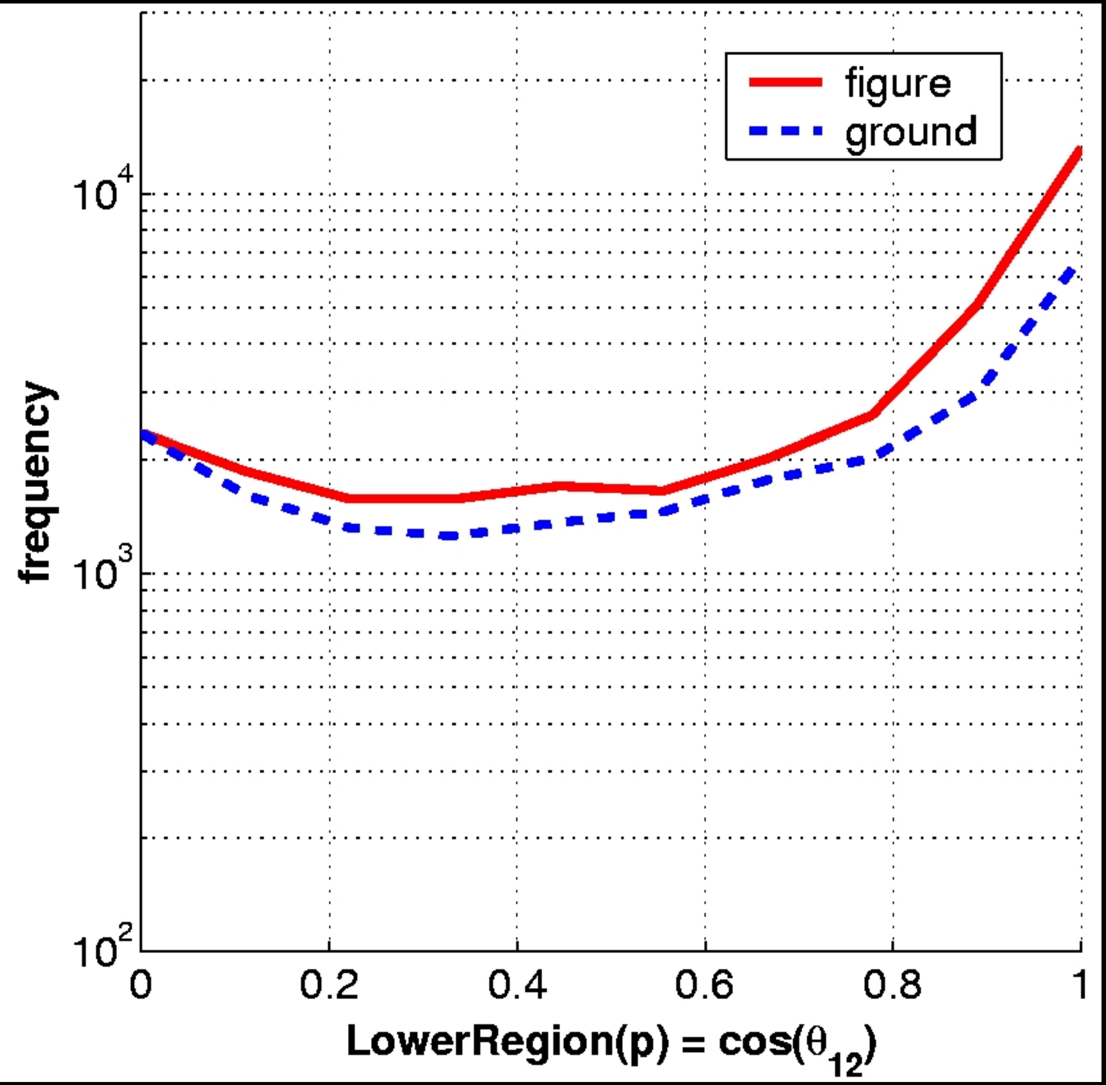

Ren, Fowlkes, Malik ECCV '06

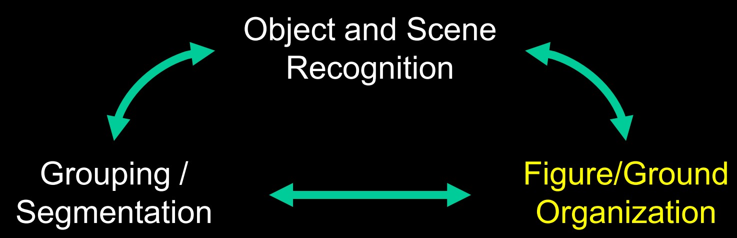

- Human subjects label groundtruth figure/ground assignments in natural images.
- Shapemes encode high-level knowledge in a generic way, capturing local figure/ground cues.
- A conditional random field incorporates junction cues and enforces global consistency.

Forty years of contour detection

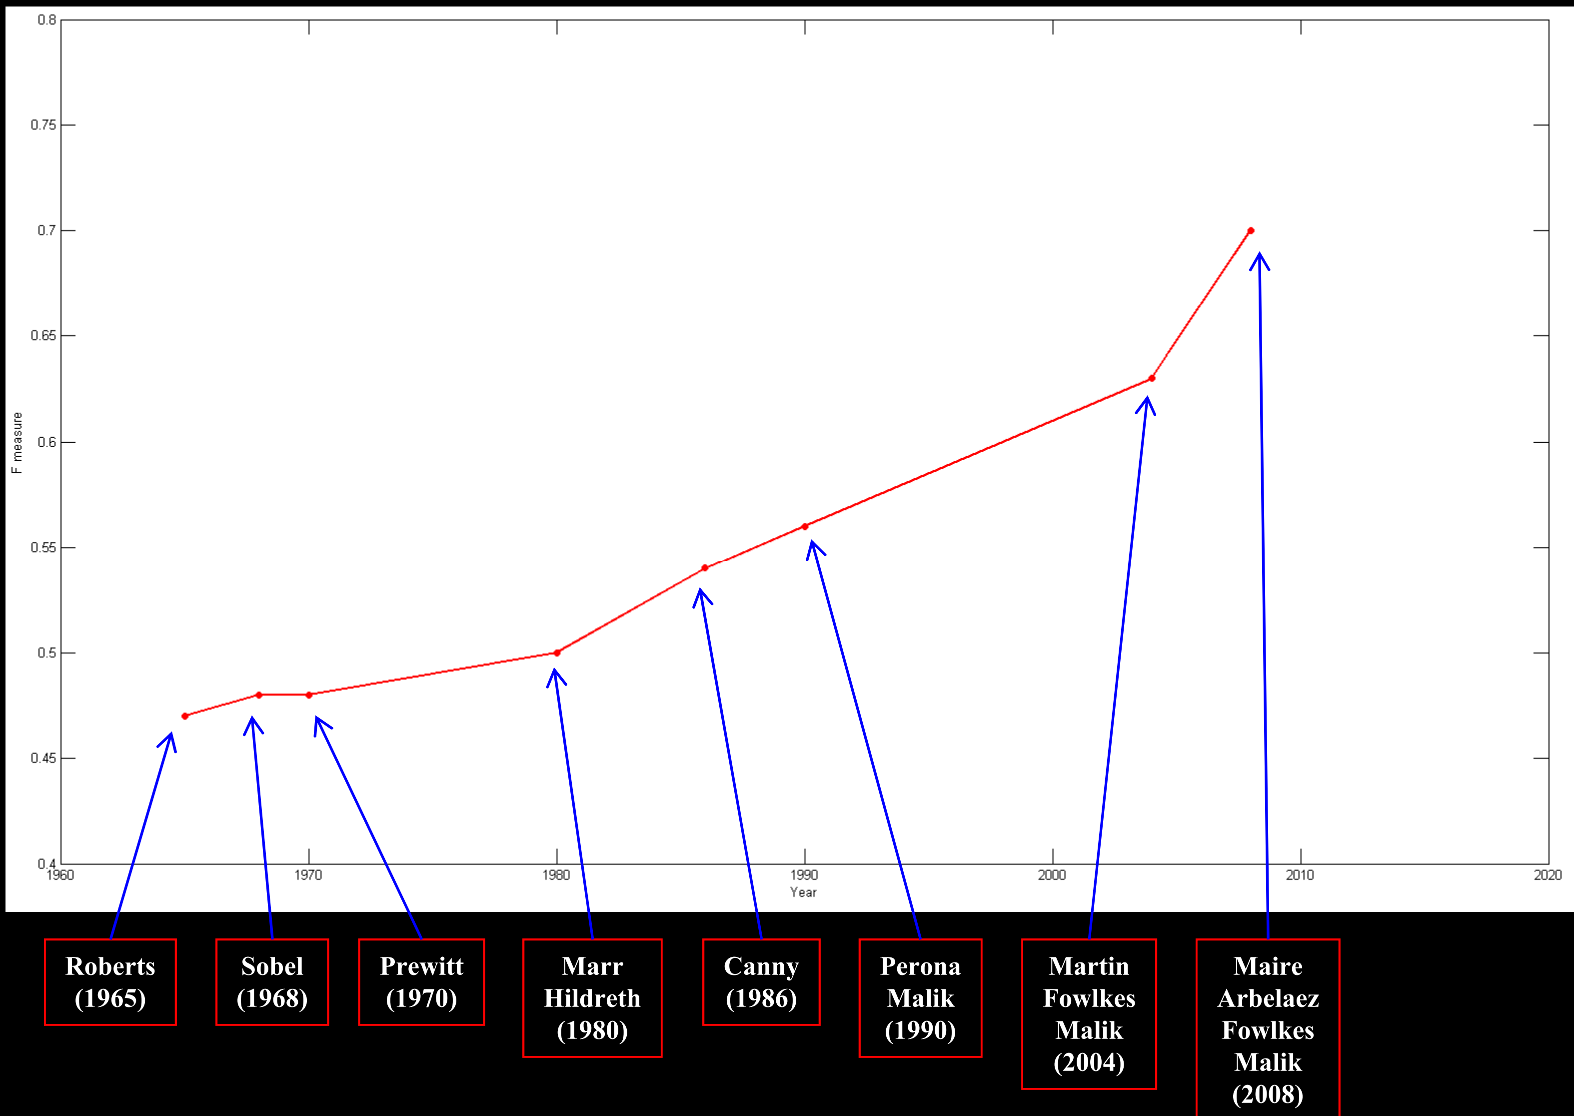

Forty years of contour detection

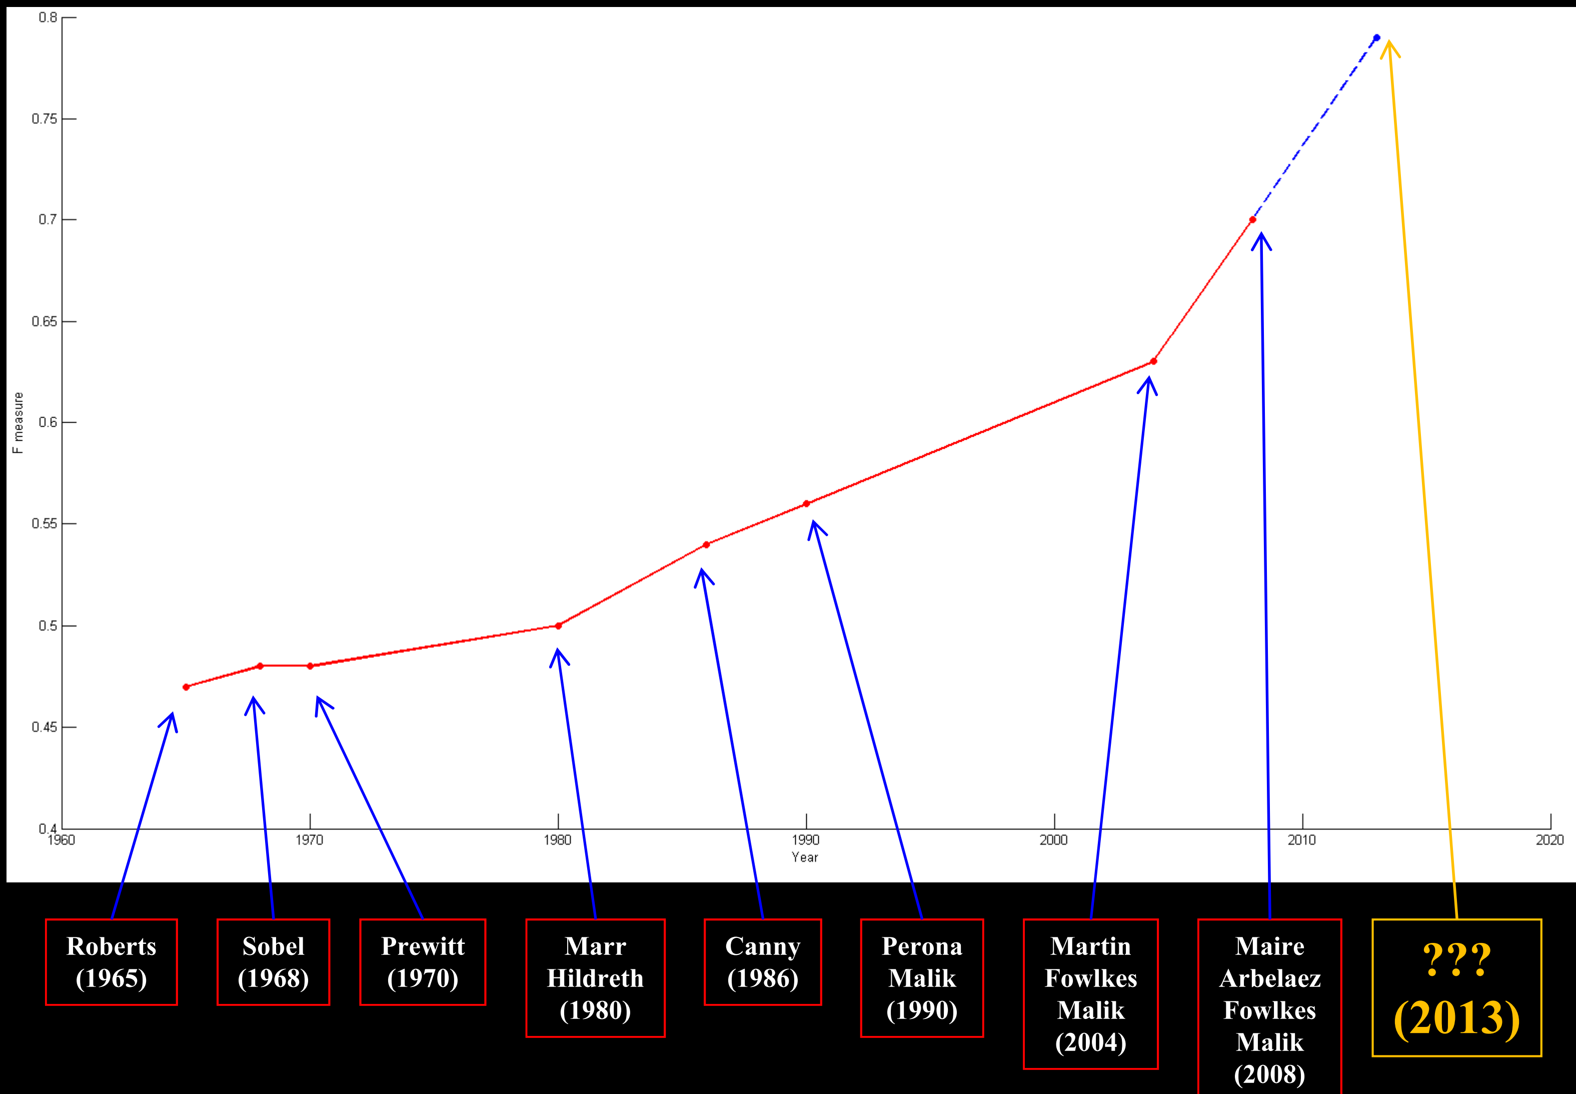

Curvilinear Grouping

- Boundaries are smooth in nature!
- A number of associated visual phenomena

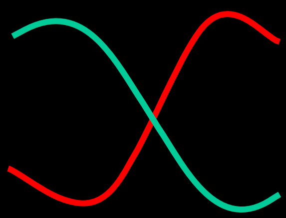
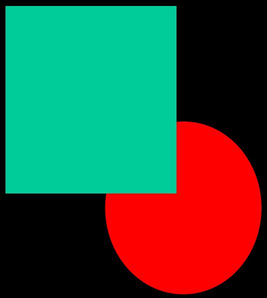
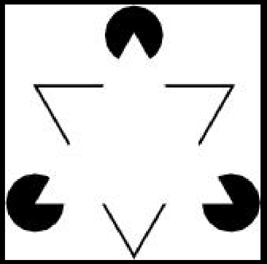

Computational Photography

Computer Vision

CSE 576, Spring 2008

Richard Szeliski
Microsoft Research

Computational Photography

- photometric camera calibration
- high-dynamic range imaging & tone mapping
- flash photography

Readings

- Debevec and Malik, Recovering High Dynamic Range Radiance Maps from Photographs. In *SIGGRAPH 97*.
- S. B. Kang et al. High dynamic range video. *SIGGRAPH 2003*.
- D. Lischinski. Interactive local adjustment of tonal values. *SIGGRAPH 2006.*
- G. Petschnigg *et al*. Digital photography with flash and no-flash image pairs. *SIGGRAPH 2004.*
- P. Pérez et al. Poisson image editing. *SIGGRAPH 2003*

Sources

Some of my slides are from:

6.098 Digital and Computational Photography
6.882 Advanced Computational Photography

Spring 2006



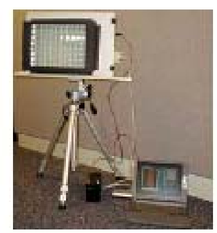

home | syllabus | problem sets and solutions | handouts | links

Bill Freeman and Frédo Durand

http://groups.csail.mit.edu/graphics/classes/CompPhoto06/

Sources

Some of my slides are from:

15-463 (15-862): Computational Photography
Computer Science Department
Carnegie Mellon University

INSTRUCTOR: Alexei (Alyosha) Efros (Office hours: Thursdays 2:30-3:30, NSH 4207)
TA: Jim McCann (Office hours: Tuesdays 5-6, NSH 4228)
UNIVERSITY UNITS: 12
SEMESTER: Fall 2007
NEWSGROUP: cmu.cs.class.cs463 (read this for important information!)
WEB PAGE: http://graphics.cs.cmu.edu/courses/15-463/
LOCATION: WeH 5312
TIME: T R 12:00--1:20 PM

COURSE OVERVIEW:
Computational Photography is an emerging new field created by the convergence of computer graphics, computer vision and photography. Its role is to overcome the limitations of the traditional camera by using computational techniques to produce a richer, more vivid, perhaps more perceptually meaningful representation of our visual world.

Alexei (Alyosha) Efros

http://graphics.cs.cmu.edu/courses/15-463/

But first, …

… for something (a little) different …

Panography - http://www.flickr.com/search/?q=panography

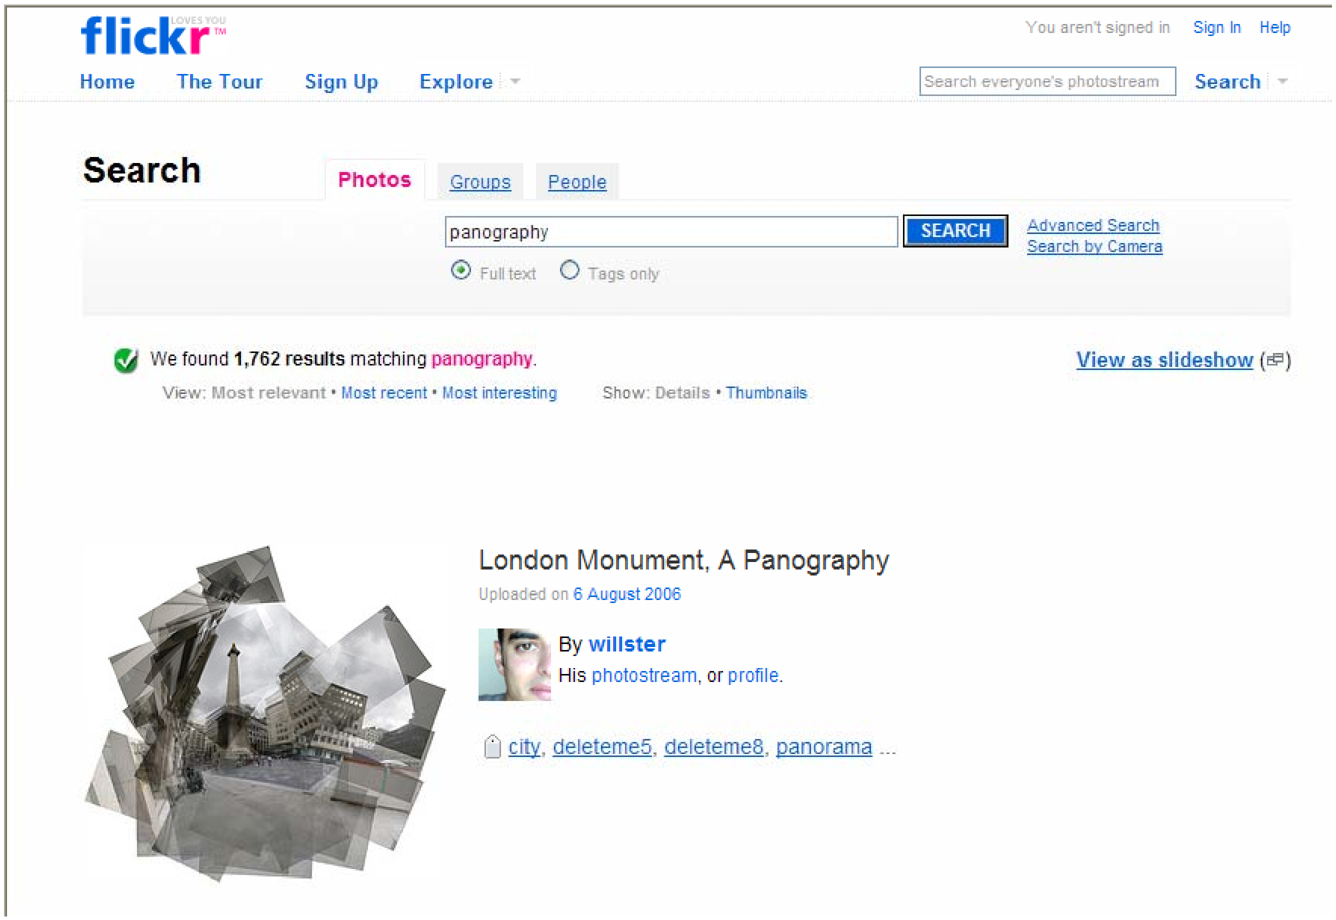

Panography - http://www.flickr.com/search/?q=panograph

Tokyo Skyline Panograph

Uploaded on 30 July 2006

By **Chalky Lives**

Chalky Lives' photostream, or profile.

tower, skyline, tokyo, photo ...

Times Square Panograph

Uploaded on 2 August 2006

By **Chalky Lives**

Chalky Lives' photostream, or profile.

ny, newyork, advertising, construction ...

Sleeping Beauty Castle (Panograph #7)

Uploaded on 28 December 2006

By **targeteer2k**

His photostream, or profile.

christmas, xmas, sleeping, people ...

Panography

What kind of motion model?

What kind of compositing?

Can you do “global alignment”?

High Dynamic Range Imaging (HDR)

slides borrowed from

15-463: Computational Photography

Alexei Efros, CMU, Fall 2007,

Paul Debevec, and my talks

Problem: Dynamic Range

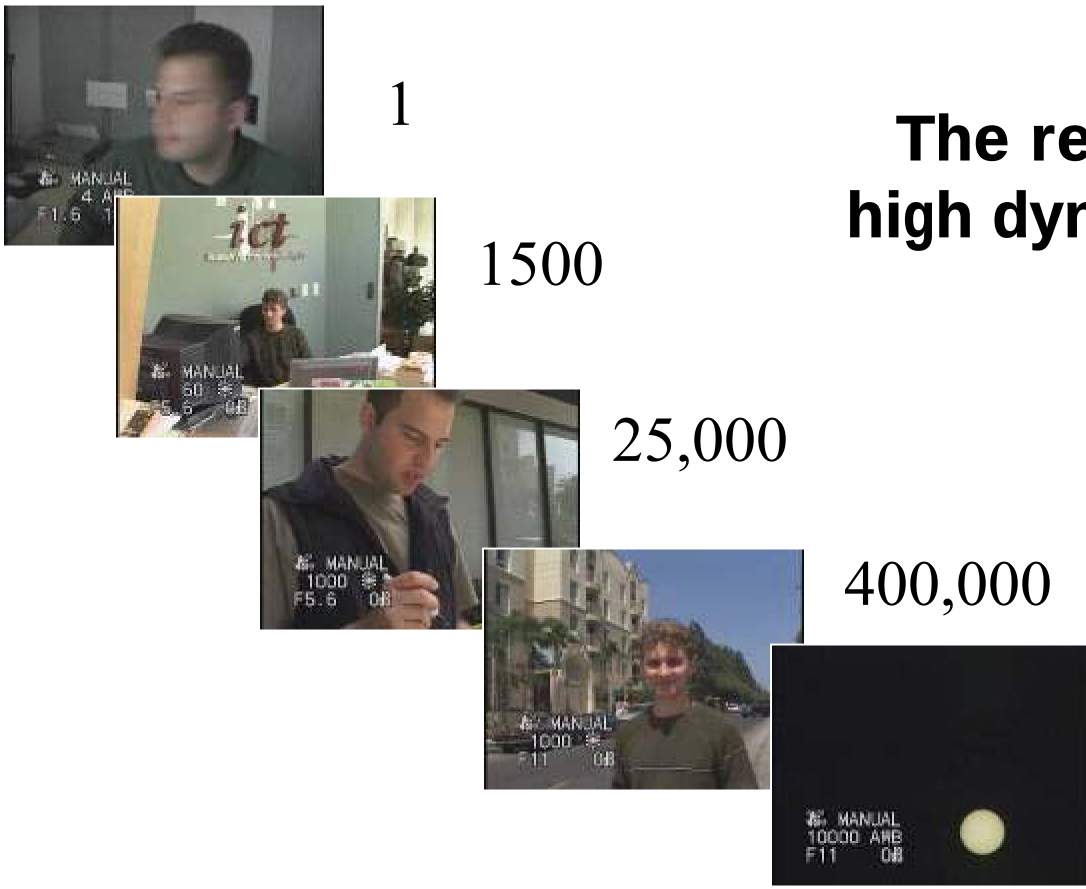

1

1500

25,000

400,000

2,000,000,000

The real world is high dynamic range.

Problem: Dynamic Range

Typical cameras have limited dynamic range

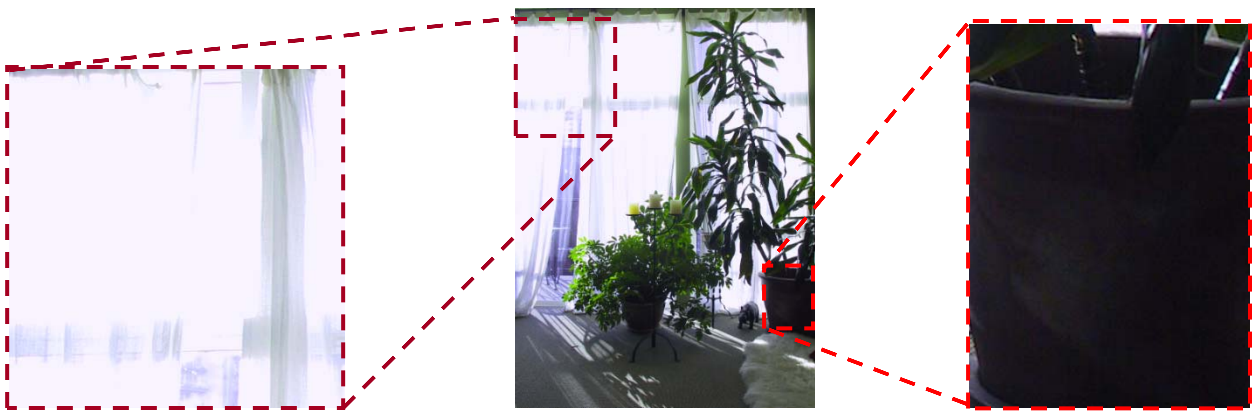

What can we do?

Solution: merge multiple exposures

Varying Exposure

HDR images — multiple inputs

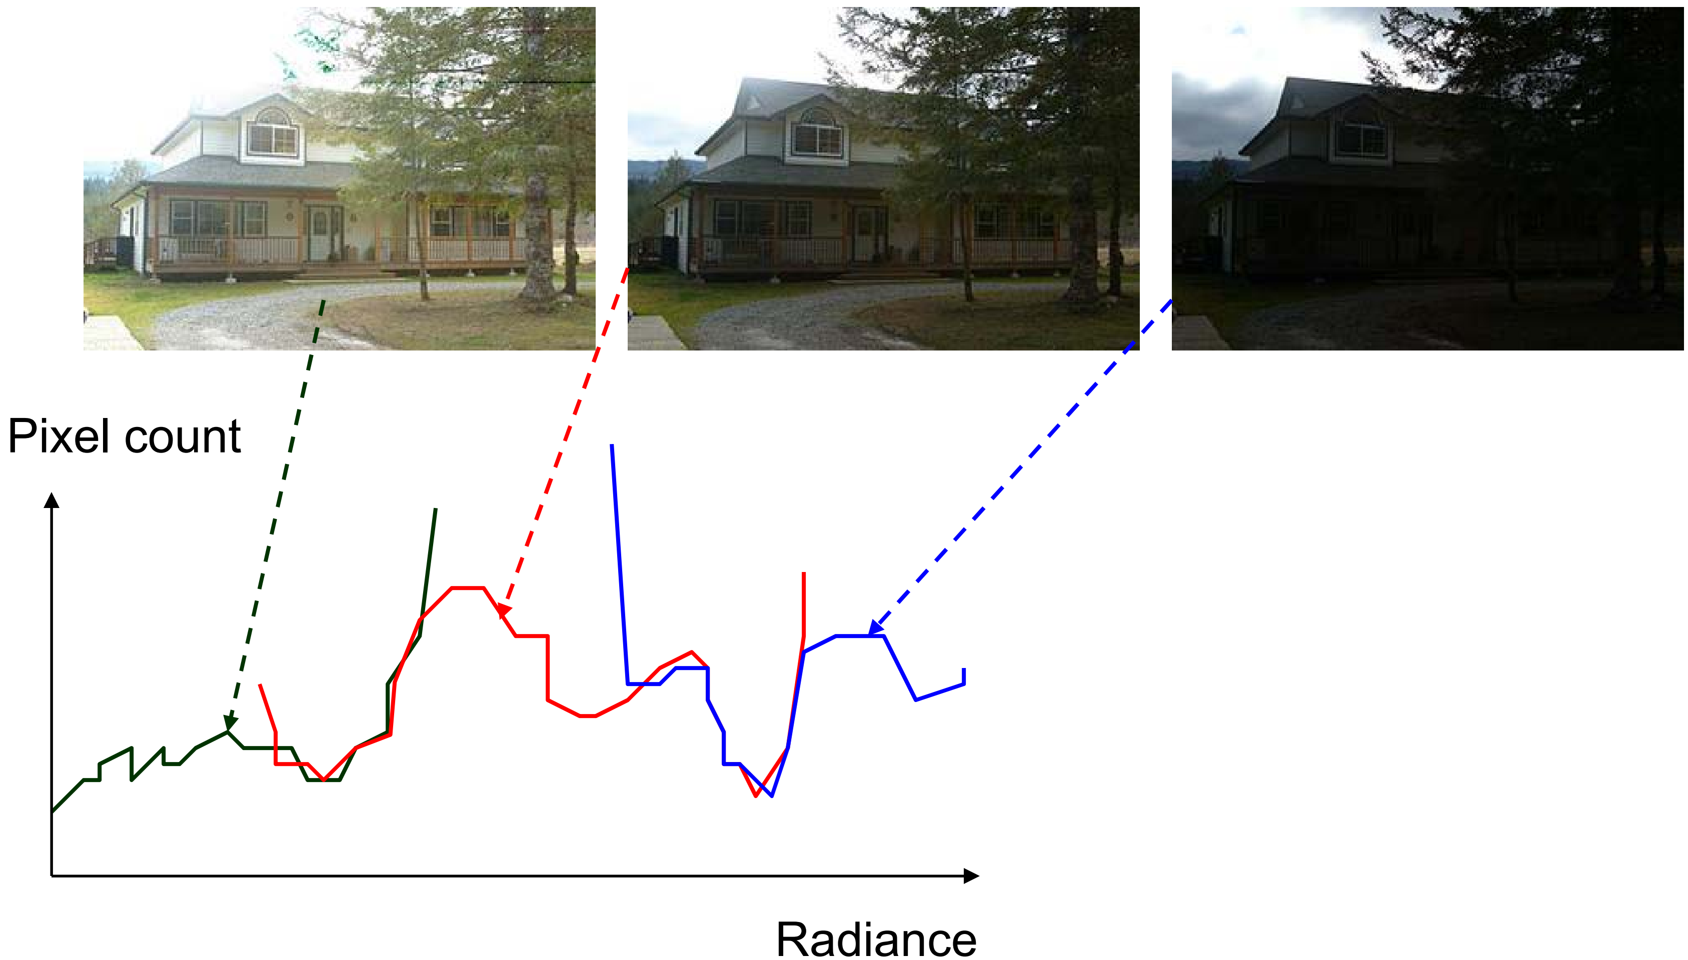

HDR images — merged

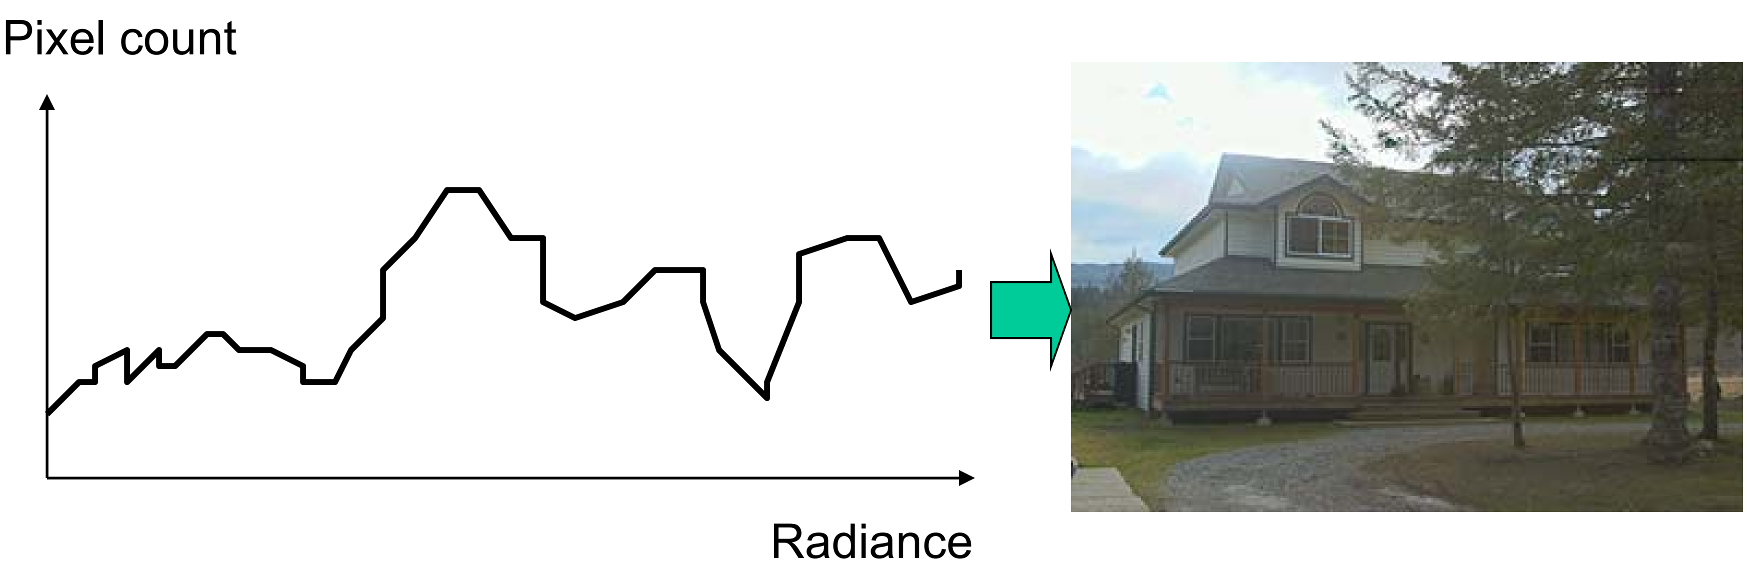

Camera is not a photometer!

Limited dynamic range

$\Rightarrow$Use multiple exposures?

Unknown, nonlinear response

$\Rightarrow$ Not possible to convert pixel values to radiance

Solution:

- Recover response curve from multiple exposures, then reconstruct the ***radiance map***

Imaging system response function

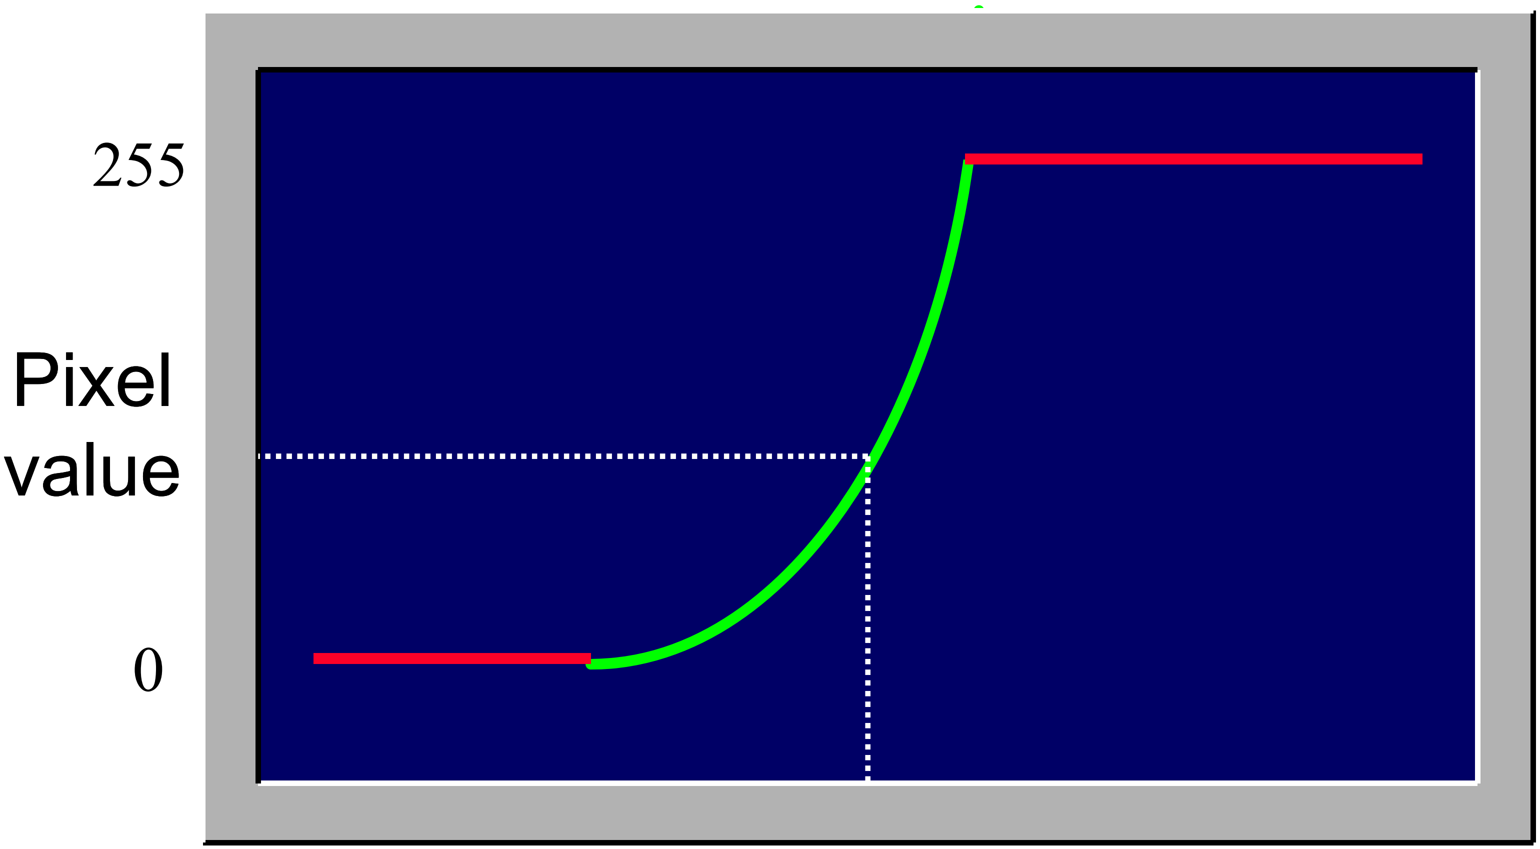

log Exposure = log (Radiance * Δt)

(CCD photon count)

Camera Calibration

Geometric

- How pixel **coordinates** relate to **directions** in the world

Photometric

- How pixel **values** relate to **radiance** amounts in the world
- Per-pixel **transfer** and **blur**

Camera sensing pipeline

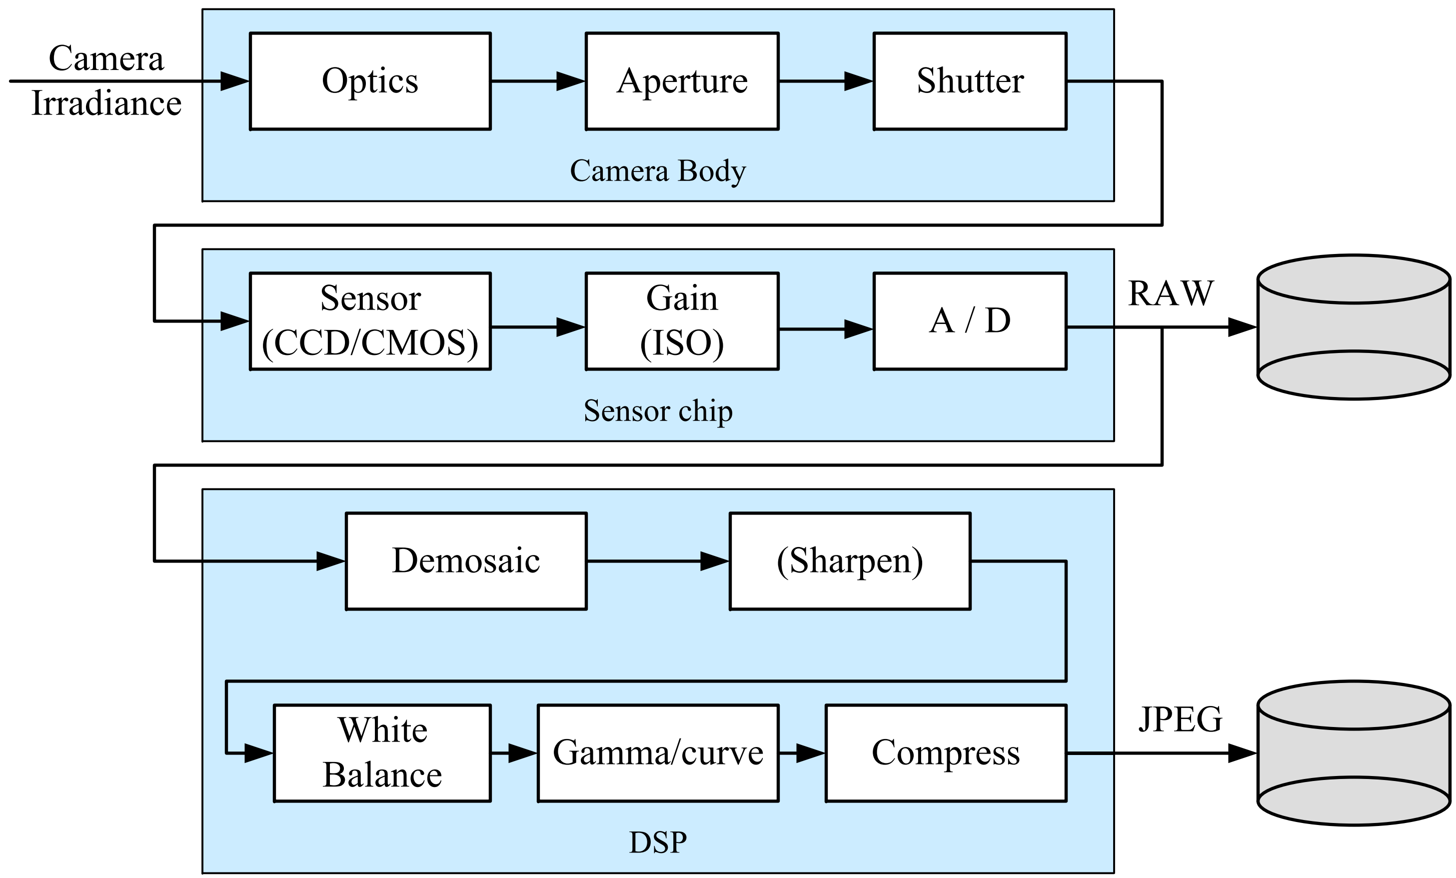

Camera sensing pipeline

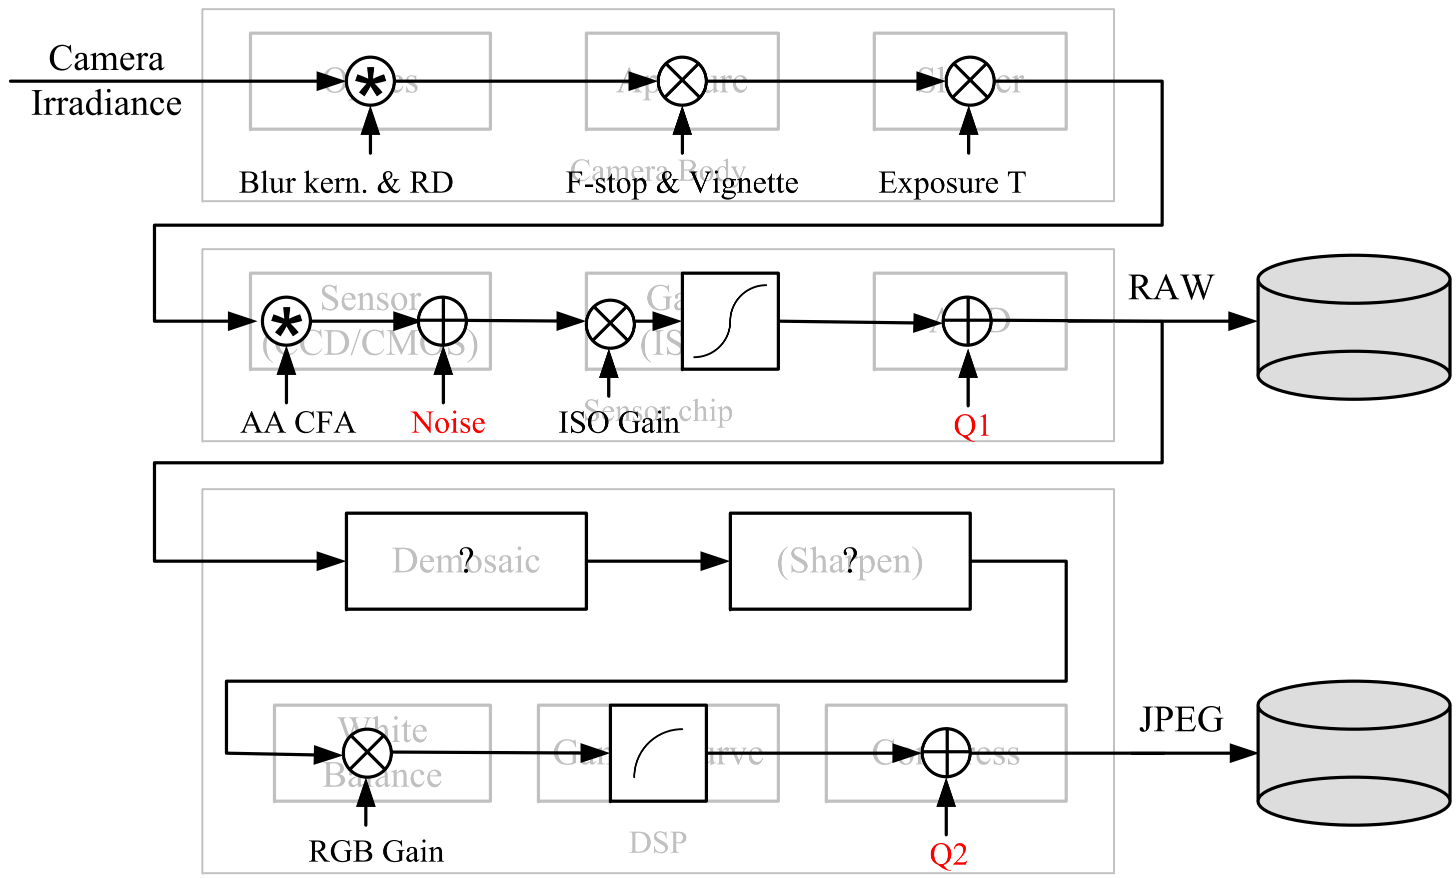

Recovering High Dynamic Range Radiance Maps from Photographs

Paul Debevec
Jitendra Malik

Computer Science Division
University of California at Berkeley

SIGGRAPH'97, August 1997

Ways to vary exposure

- Shutter Speed (*)

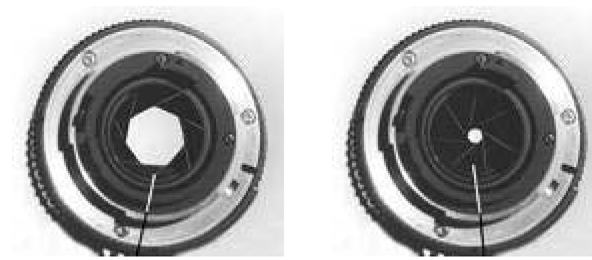

- F/stop (aperture, iris)
- Neutral Density (ND) Filters

Shutter Speed

Ranges: Canon D30: 30 to 1/4,000 sec.

(1997) Sony VX2000: ¼ to 1/10,000 sec.

Pros:

- Directly varies the exposure
- Usually accurate and repeatable

Issues:

- Noise in long exposures

Shutter Speed

Note: shutter times usually obey a power series – each "stop" is a factor of 2

¼, 1/8, 1/15, 1/30, 1/60, 1/125, 1/250, 1/500, 1/1000 sec

Usually really is:

¼, 1/8, 1/16, 1/32, 1/64, 1/128, 1/256, 1/512, 1/1024 sec

The Algorithm

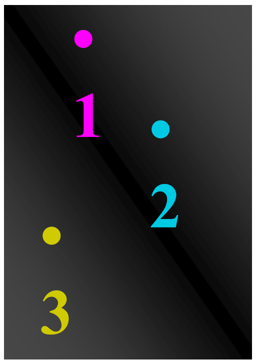

Δt = 1/64 sec

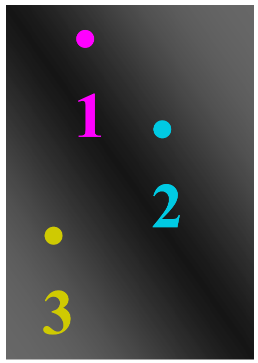

Δt = 1/16 sec

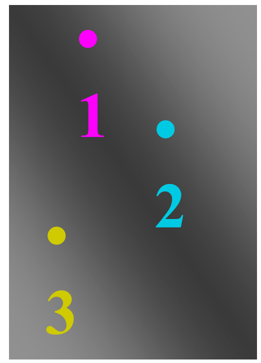

Δt = 1/4 sec

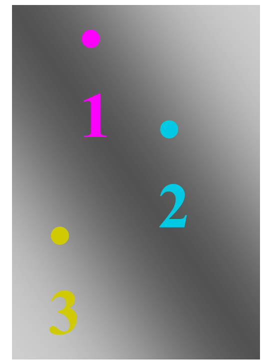

Δt = 1 sec

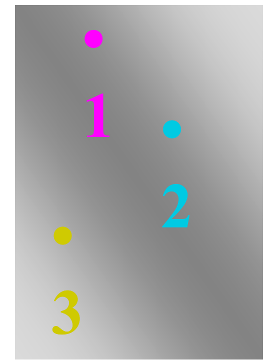

Δt = 4 sec

Pixel Value Z = f(Exposure)

Exposure = Radiance $\times$ Δt

log Exposure = log Radiance + log Δt

Response Curve

Assuming unit radiance for each pixel

After adjusting radiances to obtain a smooth response curve

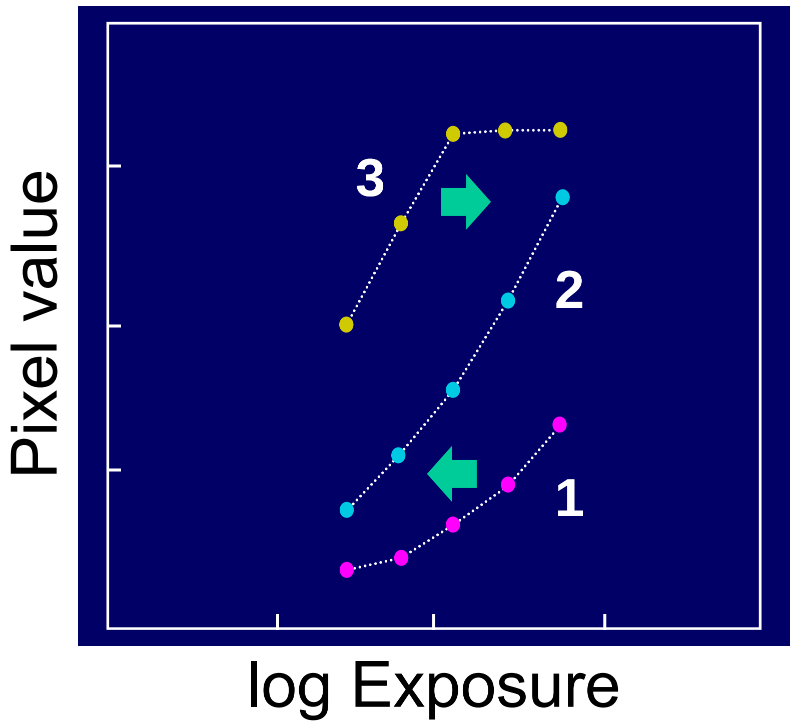

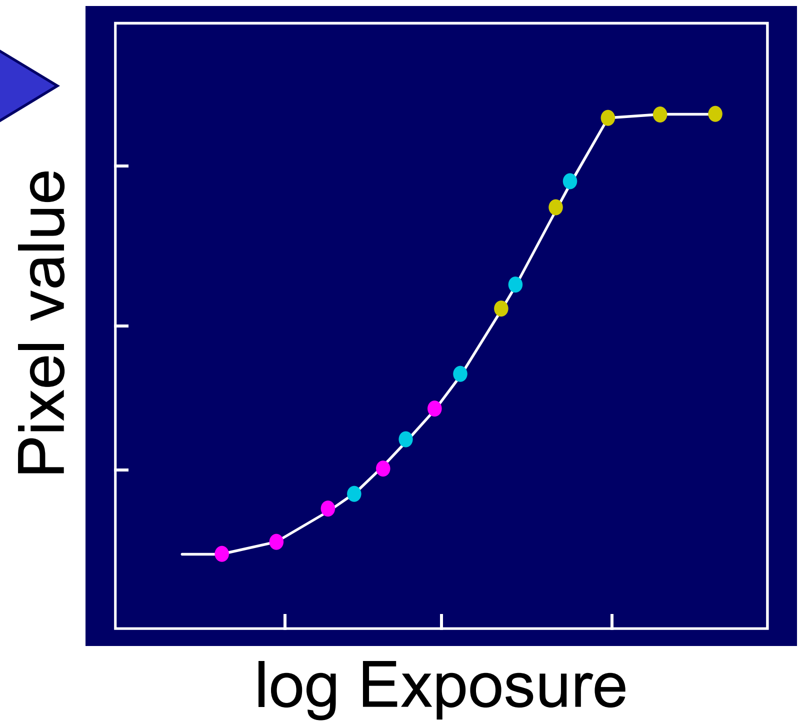

The Math

Let *g(z)* be the *discrete* inverse response function

For each pixel site *i* in each image *j*, want:

$$\ln Radiance_i + \ln \Delta t_j = g(Z_{ij})$$

Solve the over-determined linear system:

$$\underbrace{\sum_{i=1}^{N}\sum_{j=1}^{P}\left[\ln Radiance_i + \ln \Delta t_j - g(Z_{ij})\right]^2}_{\text{fitting term}} + \lambda \underbrace{\sum_{z=Z_{min}}^{Z_{max}} g''(z)^2}_{\text{smoothness term}}$$

MatLab code

```
function [g,lE]=gsolve(Z,B,l,w)

n = 256;
A = zeros(size(Z,1)*size(Z,2)+n+1,n+size(Z,1));
b = zeros(size(A,1),1);

k = 1;                %% Include the data-fitting equations
for i=1:size(Z,1)
  for j=1:size(Z,2)
    wij = w(Z(i,j)+1);
    A(k,Z(i,j)+1) = wij; A(k,n+i) = -wij; b(k,1) = wij * B(i,j);
    k=k+1;
  end
end

A(k,129) = 1;         %% Fix the curve by setting its middle value to
0
k=k+1;

for i=1:n-2           %% Include the smoothness equations
  A(k,i)=l*w(i+1); A(k,i+1)=-2*l*w(i+1); A(k,i+2)=l*w(i+1);
  k=k+1;
end

x = A\b;              %% Solve the system using SVD

g = x(1:n);
lE = x(n+1:size(x,1));
```

Results: digital camera

Kodak DCS460
1/30 to 30 sec

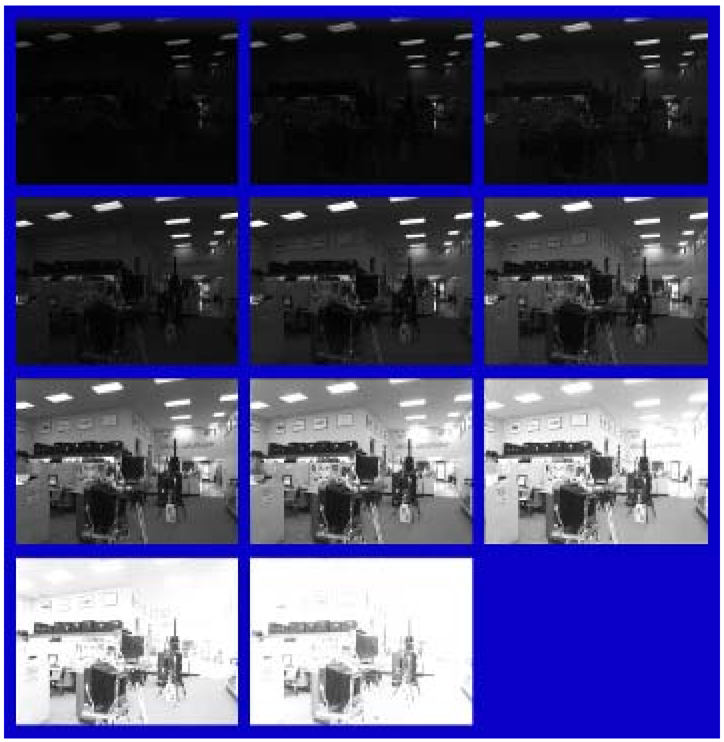

Recovered response curve

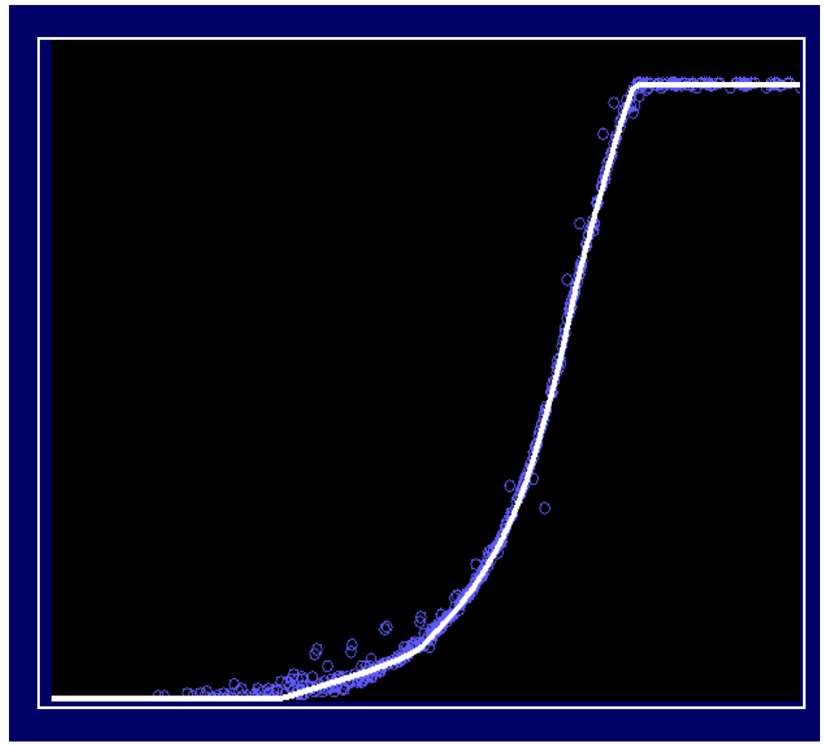

Reconstructed Radiance Map

Results: Color Film

Kodak Gold ASA 100, PhotoCD

Recovered Response Curves

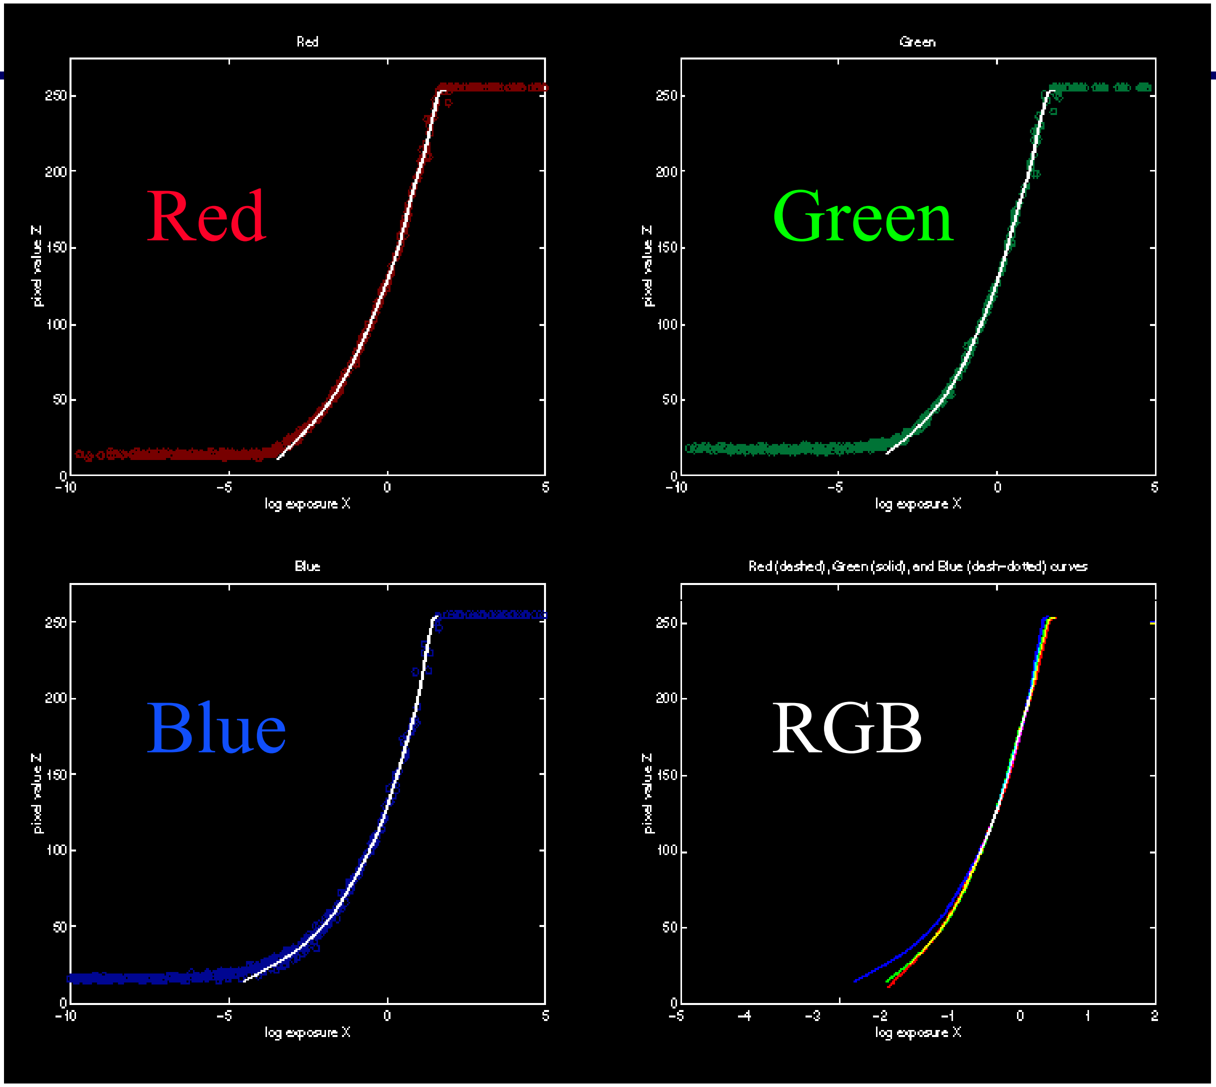

The Radiance Map

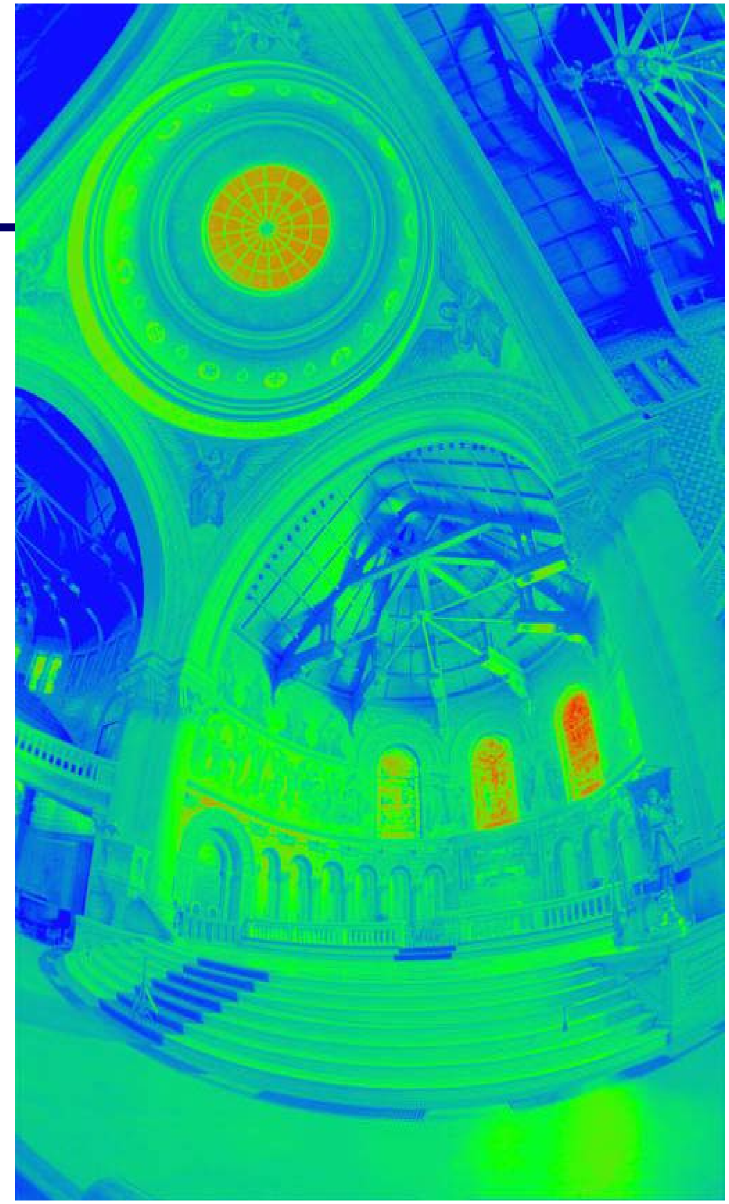

The Radiance Map

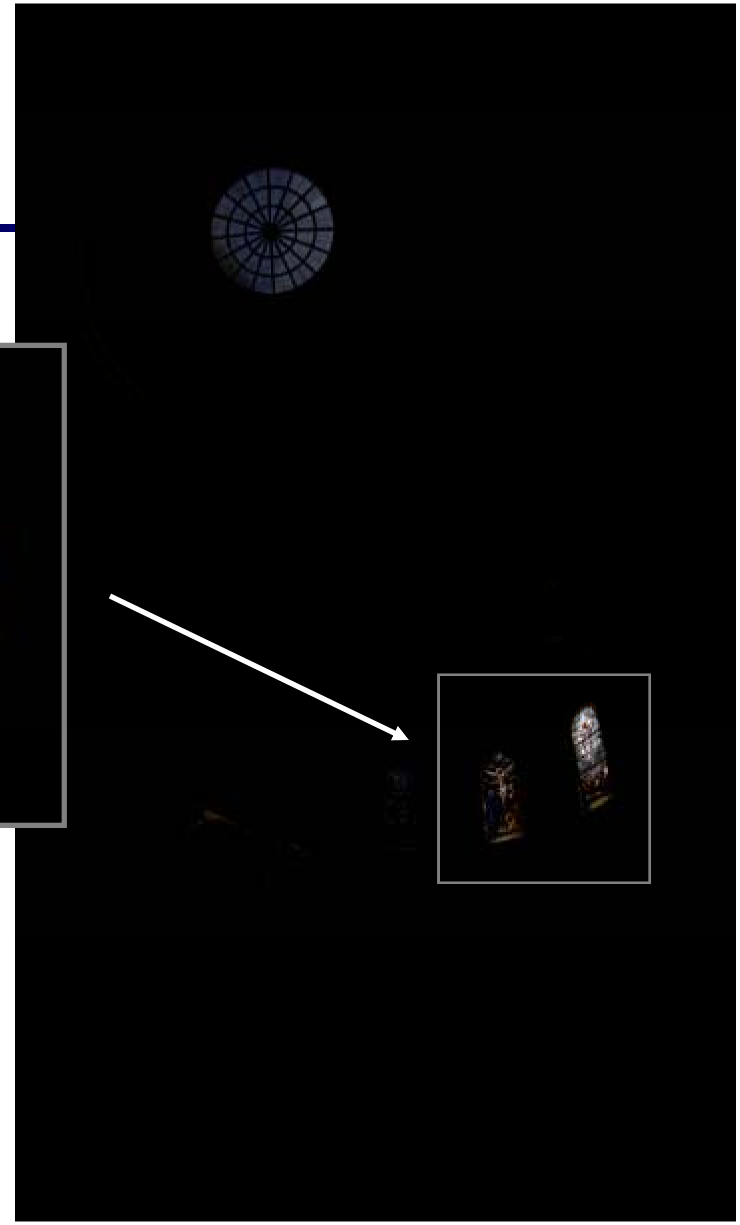

Linearly scaled to
display device

Portable FloatMap (.pfm)

12 bytes per pixel, 4 for each channel

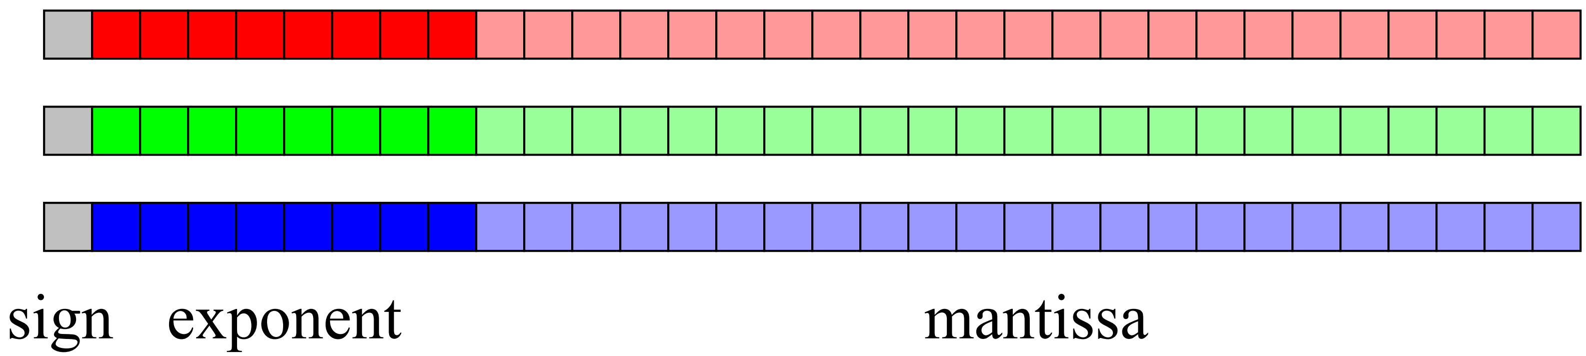

Text header similar to Jeff Poskanzer's .ppm image format:

```
PF
768 512
1
<binary image data>
```

Floating Point TIFF similar

Radiance Format (.pic, .hdr)

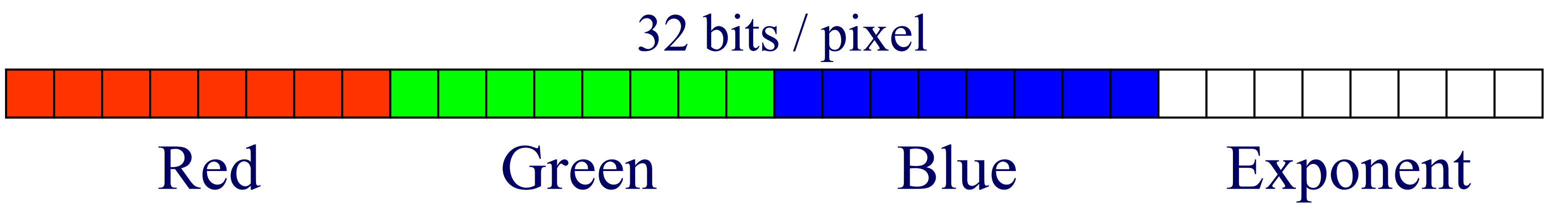

(145, 215, 87, 149) =

(145, 215, 87) * 2^(149-128) =

(1190000, 1760000, 713000)

(145, 215, 87, 103) =

(145, 215, 87) * 2^(103-128) =

(0.00000432, 0.00000641, 0.00000259)

Ward, Greg. "Real Pixels," in Graphics Gems IV, edited by James Arvo, Academic Press, 1994

ILM's OpenEXR (.exr)

6 bytes per pixel, 2 for each channel, compressed

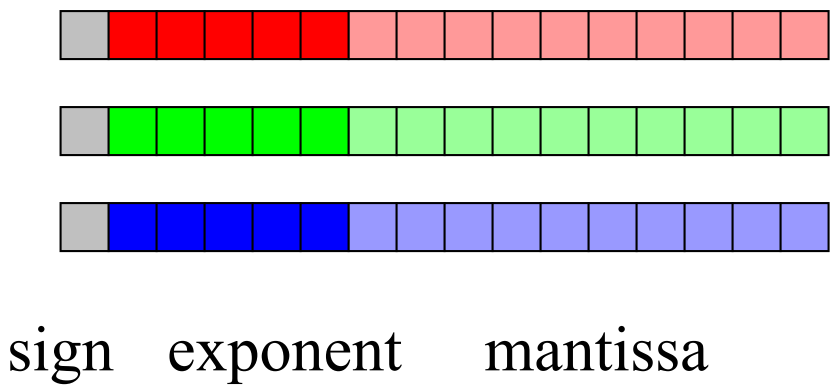

- Several lossless compression options, 2:1 typical
- Compatible with the "half" datatype in NVidia's Cg
- Supported natively on GeForce FX and Quadro FX
- Available at http://www.openexr.net/

High Dynamic Range Video

Sing Bing Kang, Matt Uyttendaele,
Simon Winder, Rick Szeliski

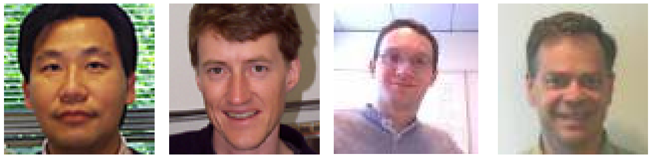

[SIGGRAPH’2003]

HDR images — merged



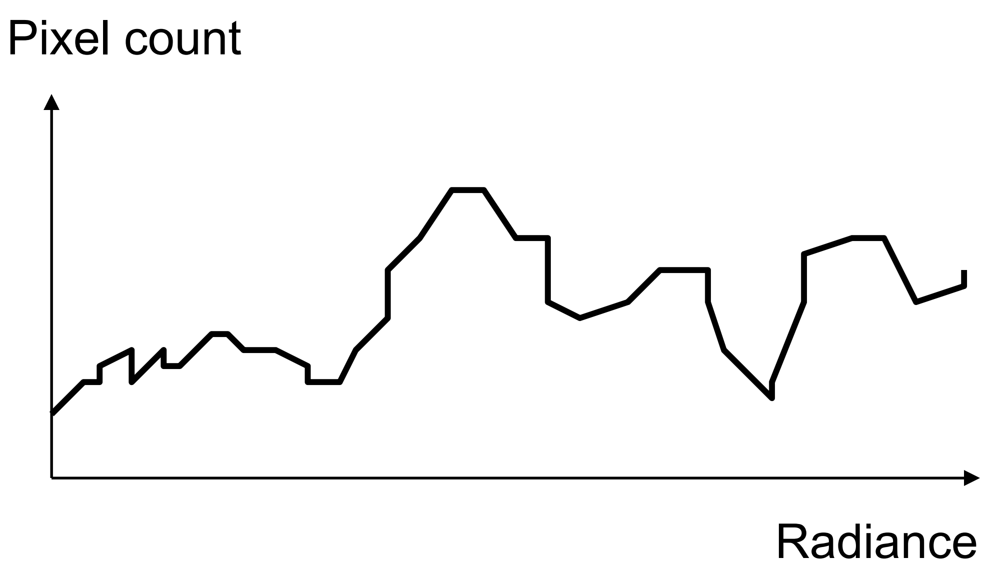

What about scene motion?

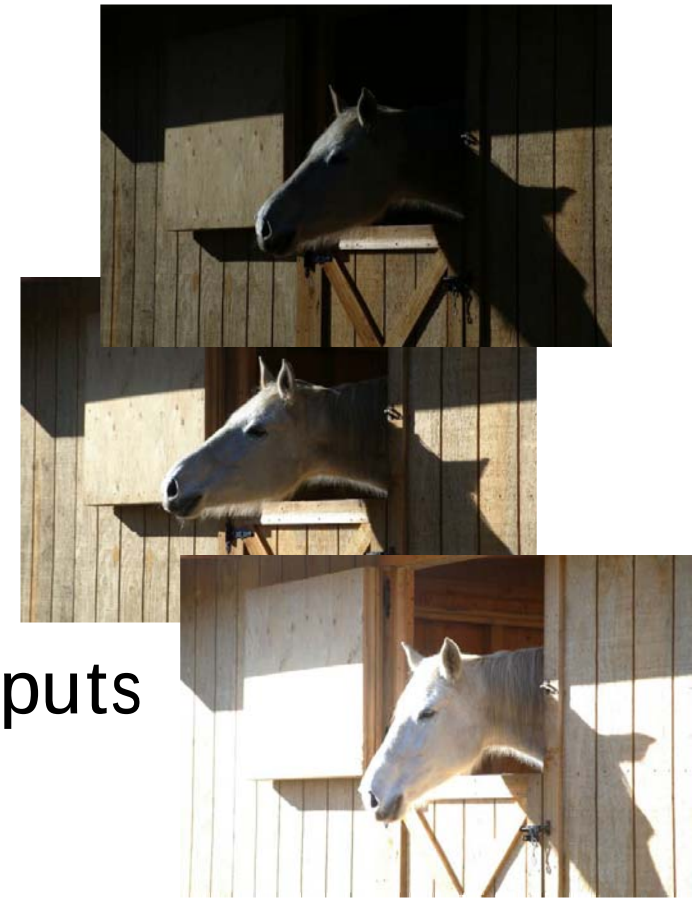

Inputs

Tonemapped output
(no compensation or
consistency check)

With motion compensation

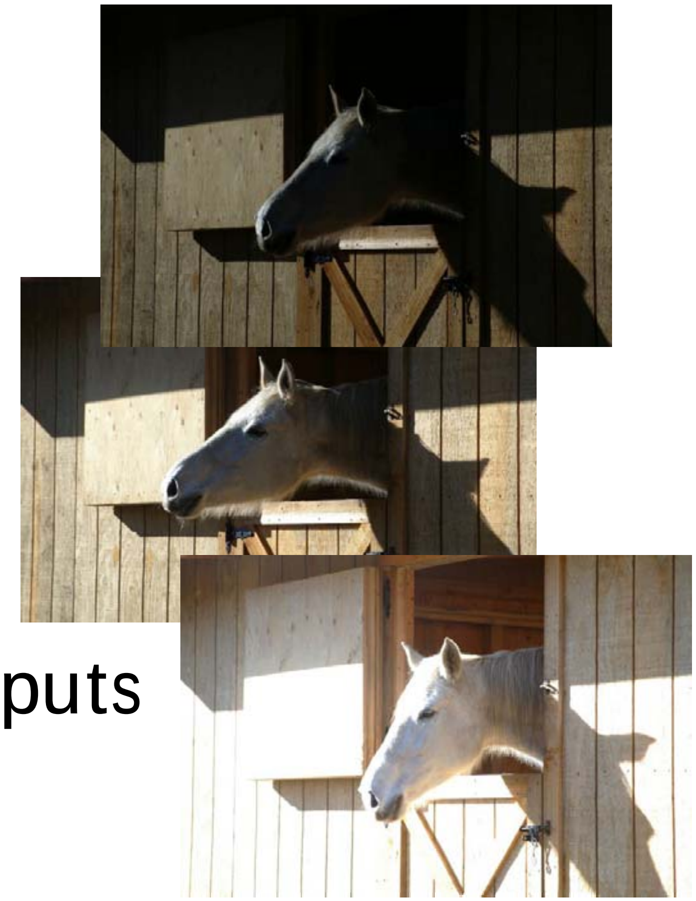

Inputs

Tonemapped output
(global+local compensation)

Registration (global)

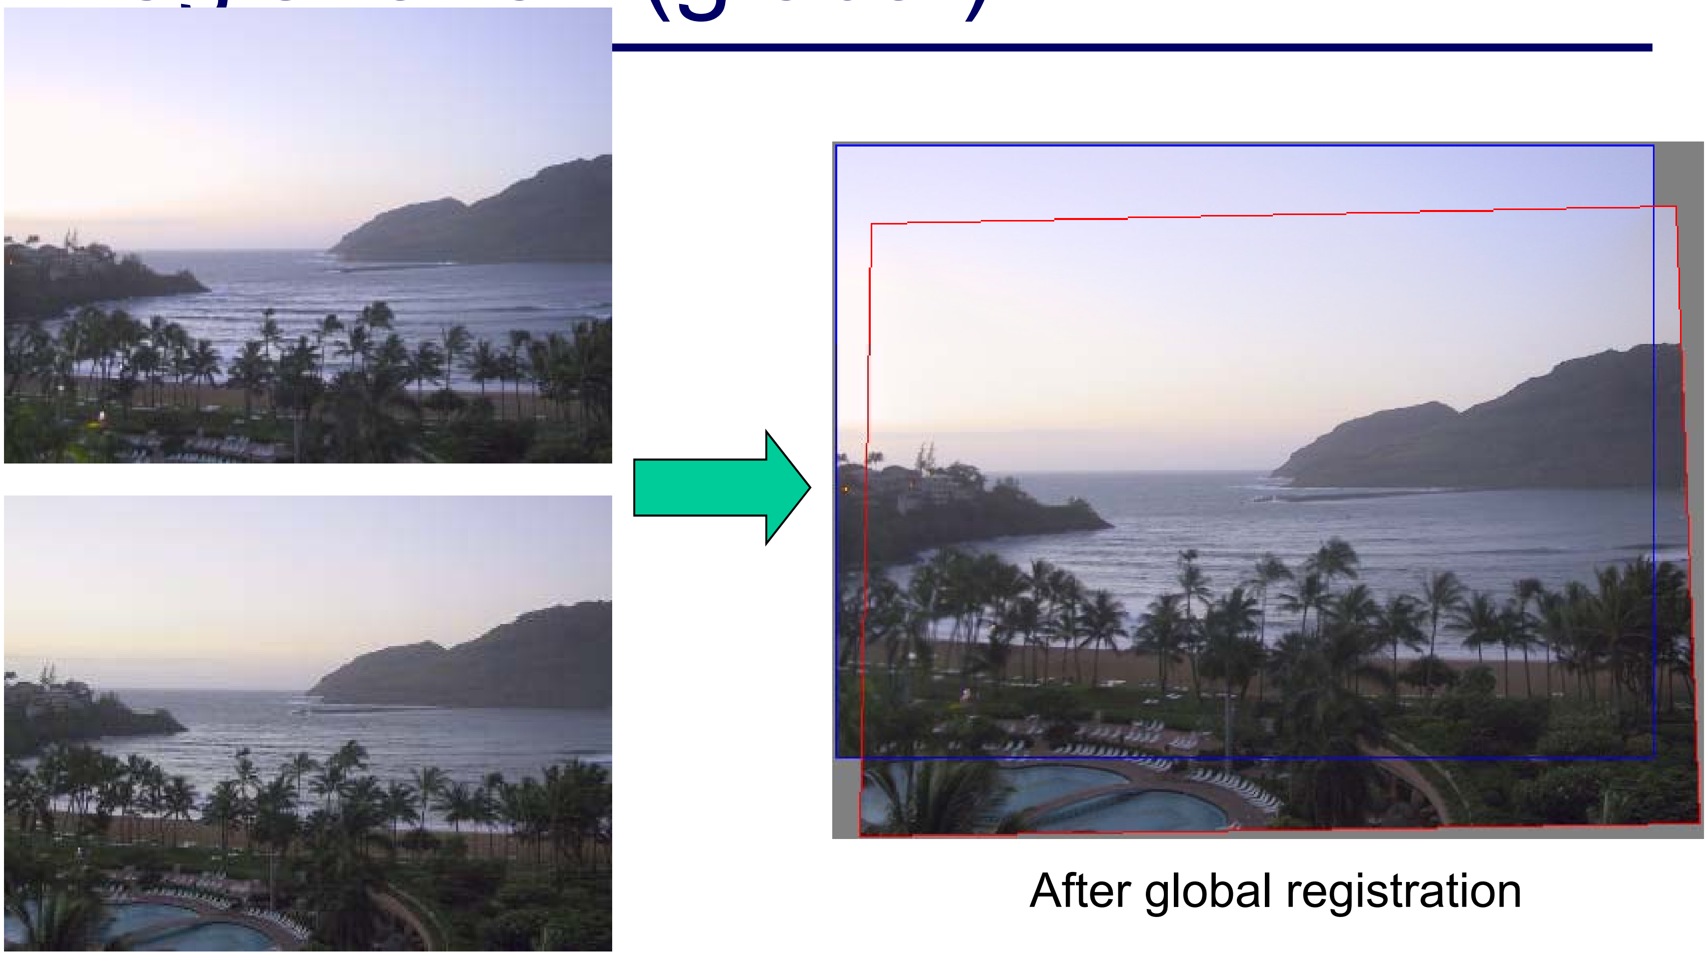

After global registration

Registration (local)

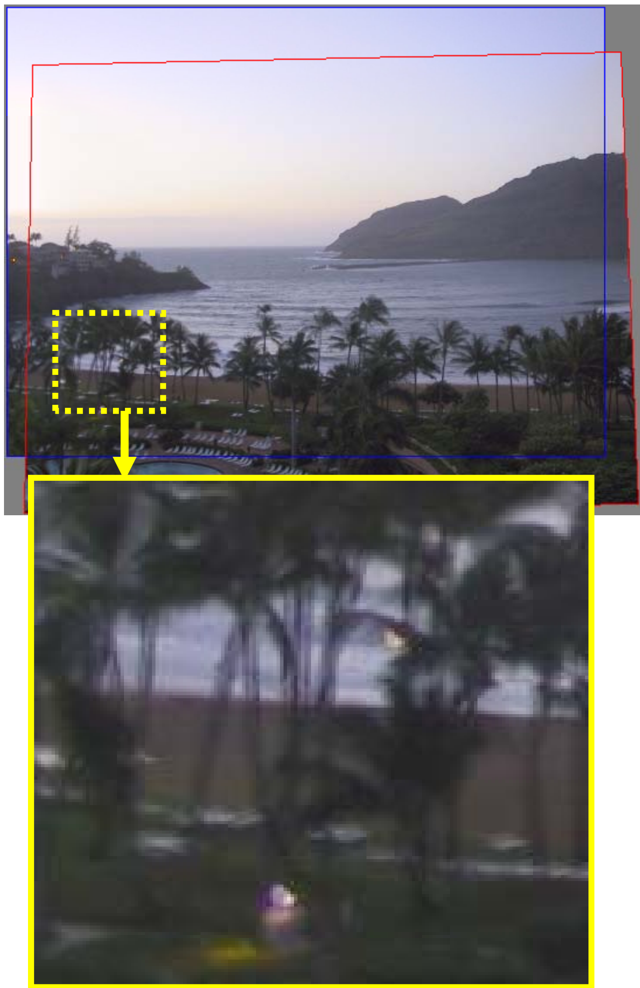

After local registration

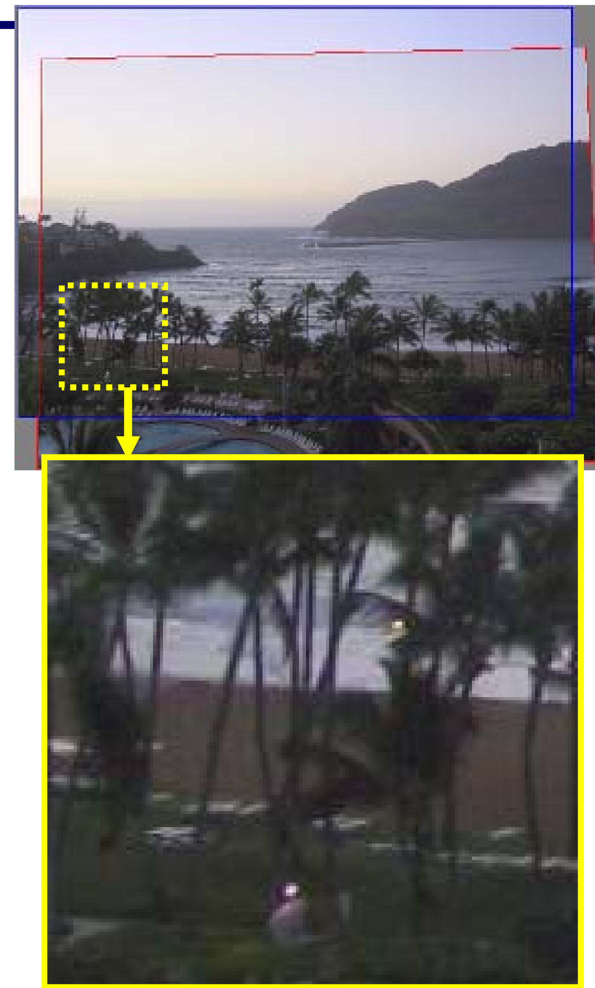

Now What?

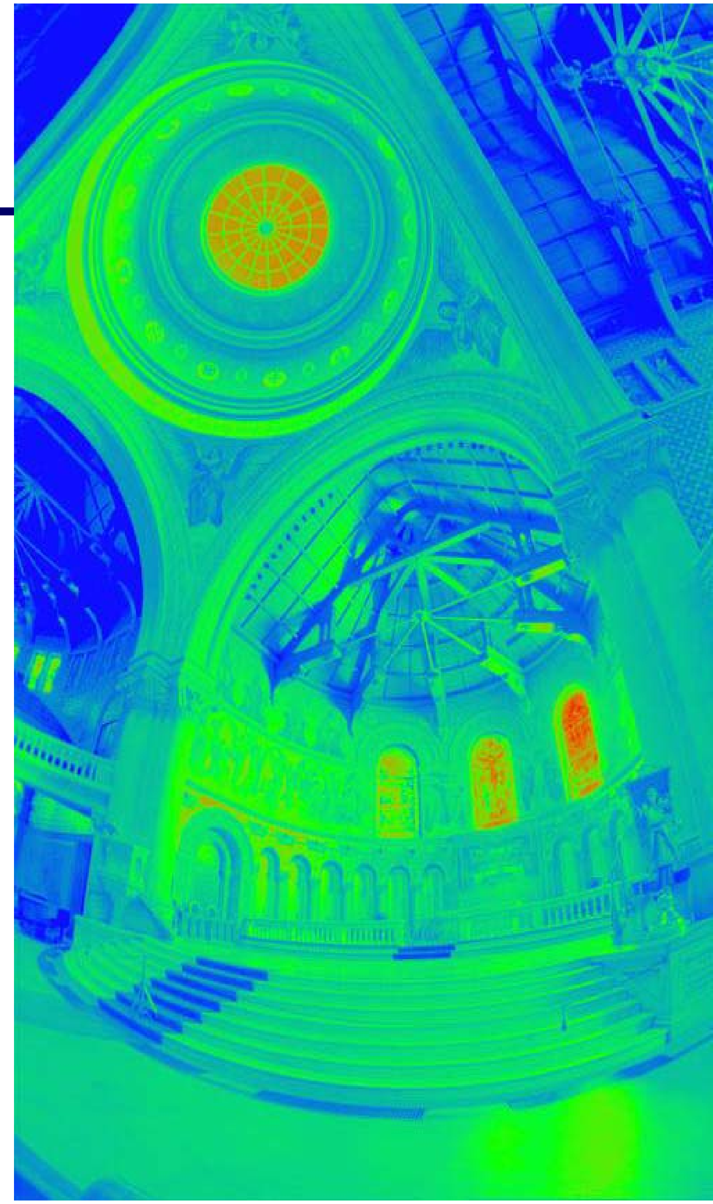

Tone Mapping

Tone Mapping

How can we do this?

Linear scaling?, thresholding? Suggestions?

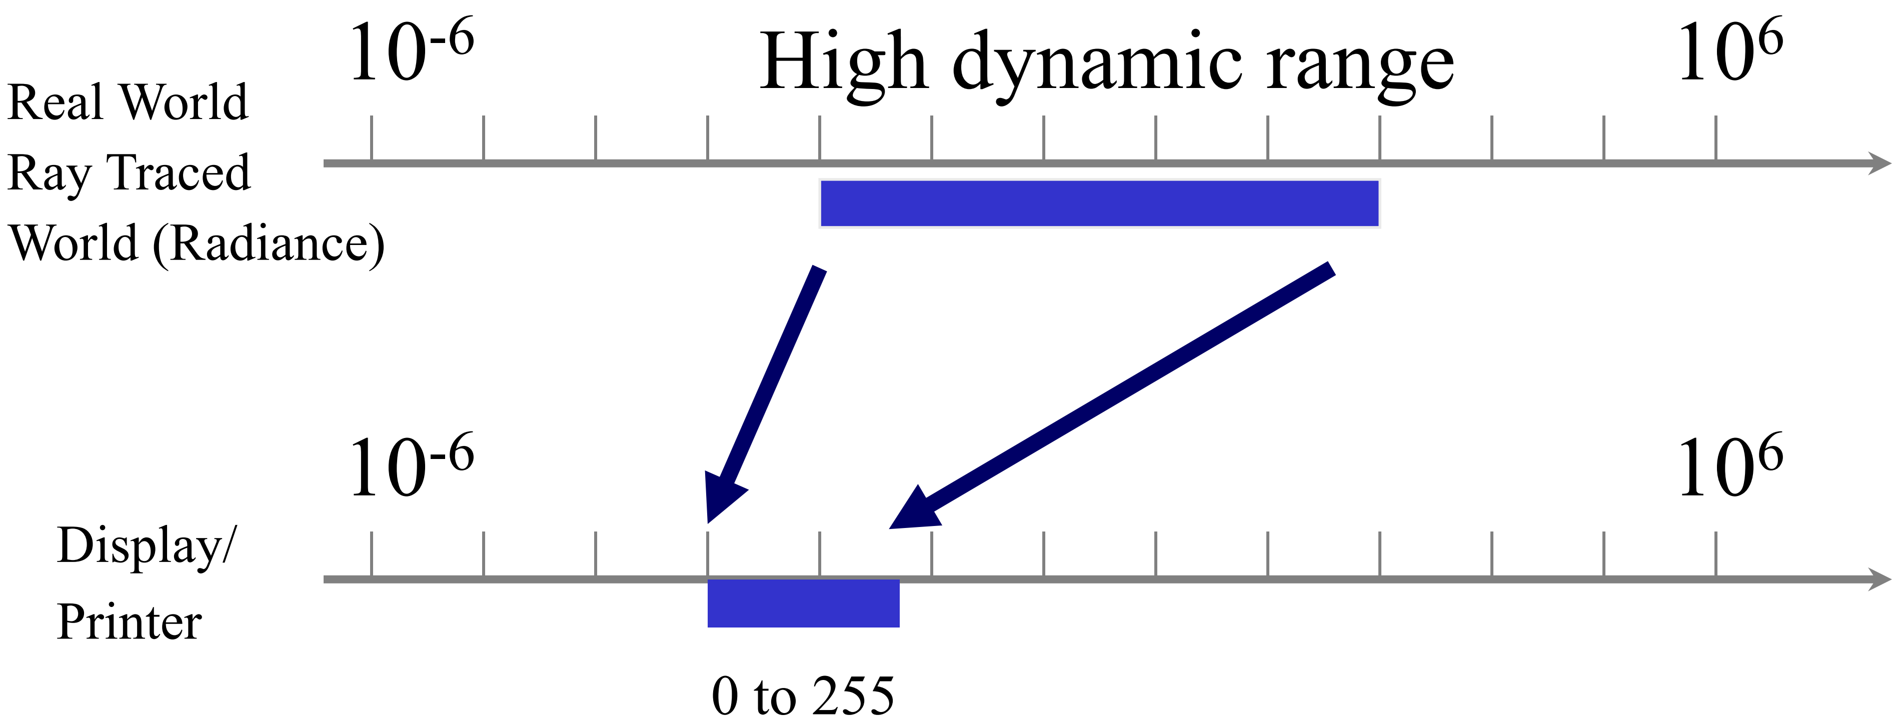

Simple Global Operator

Compression curve needs to

- Bring everything within range
- Leave dark areas alone

In other words

- Asymptote at 255
- Derivative of 1 at 0

Global Operator (Reinhart et al)

$$L_{display} = \frac{L_{world}}{1 + L_{world}}$$

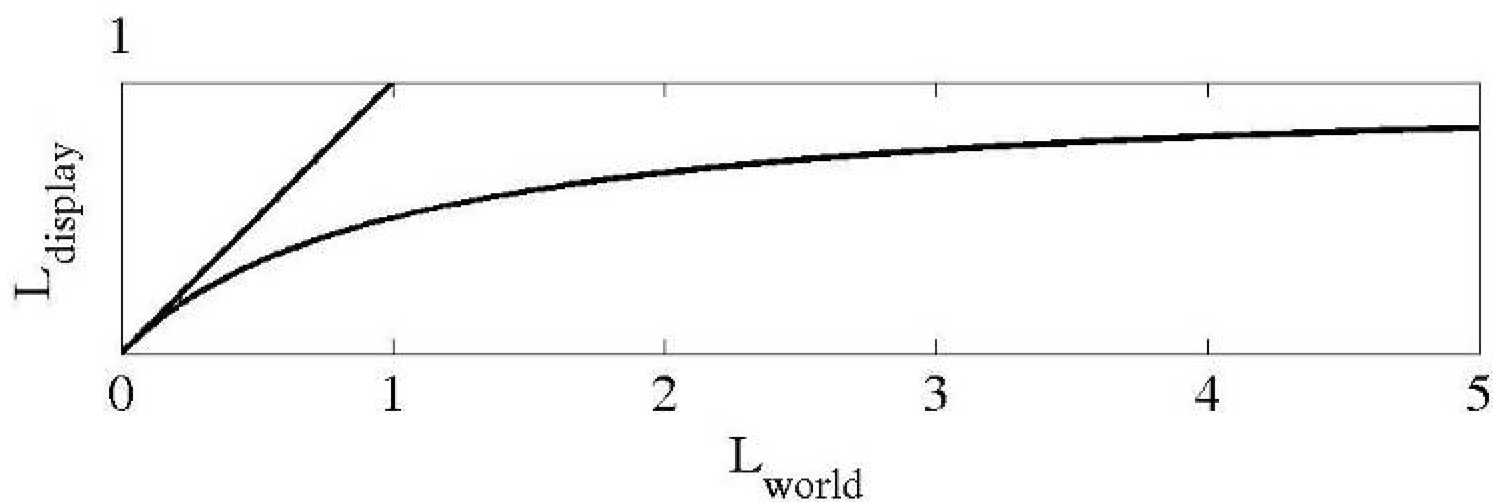

Global Operator Results

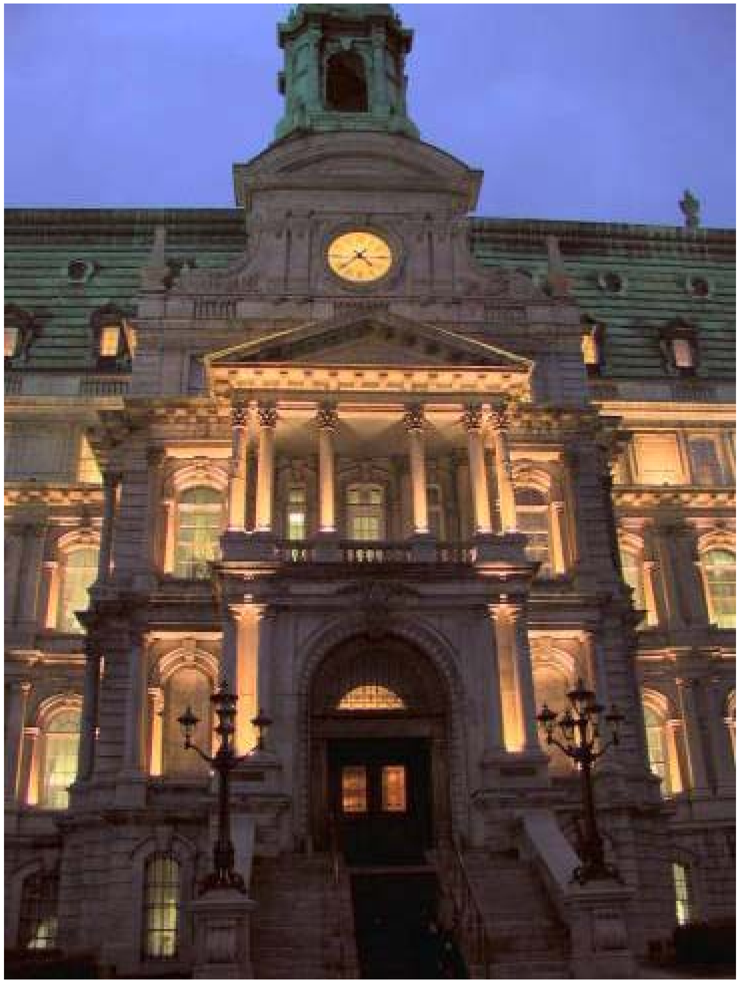

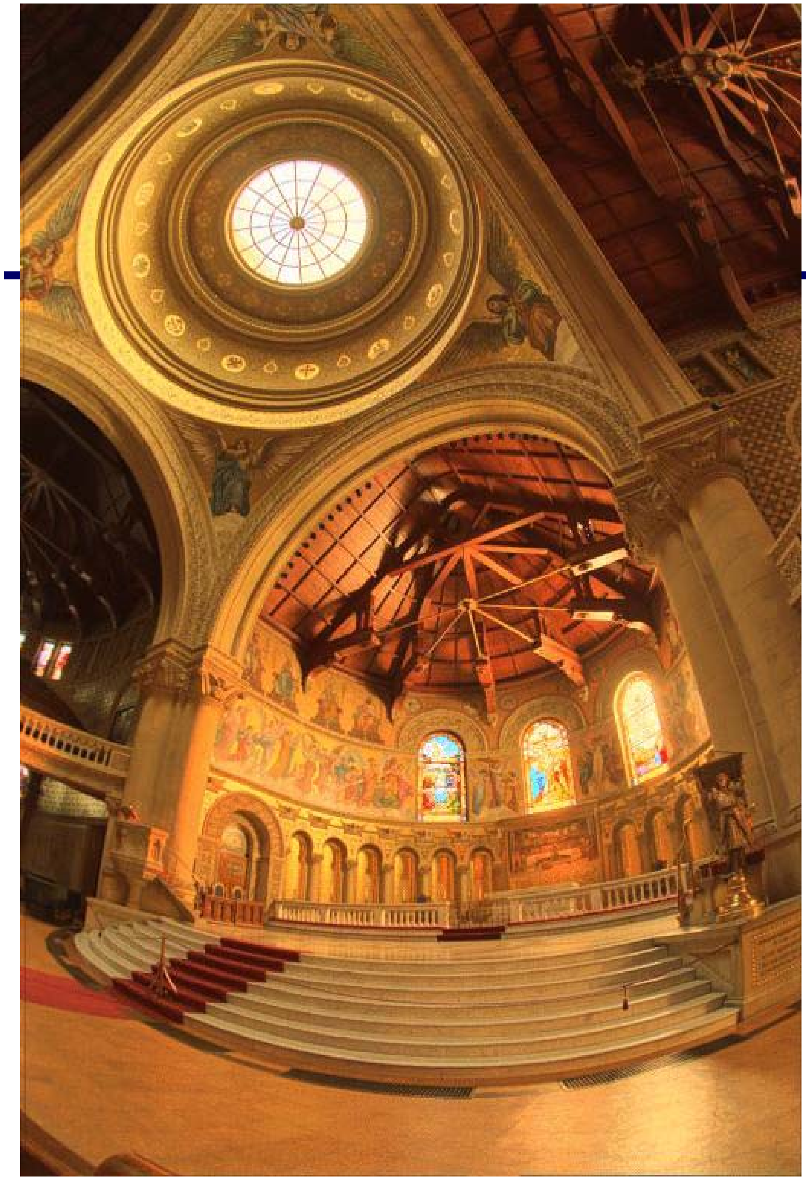

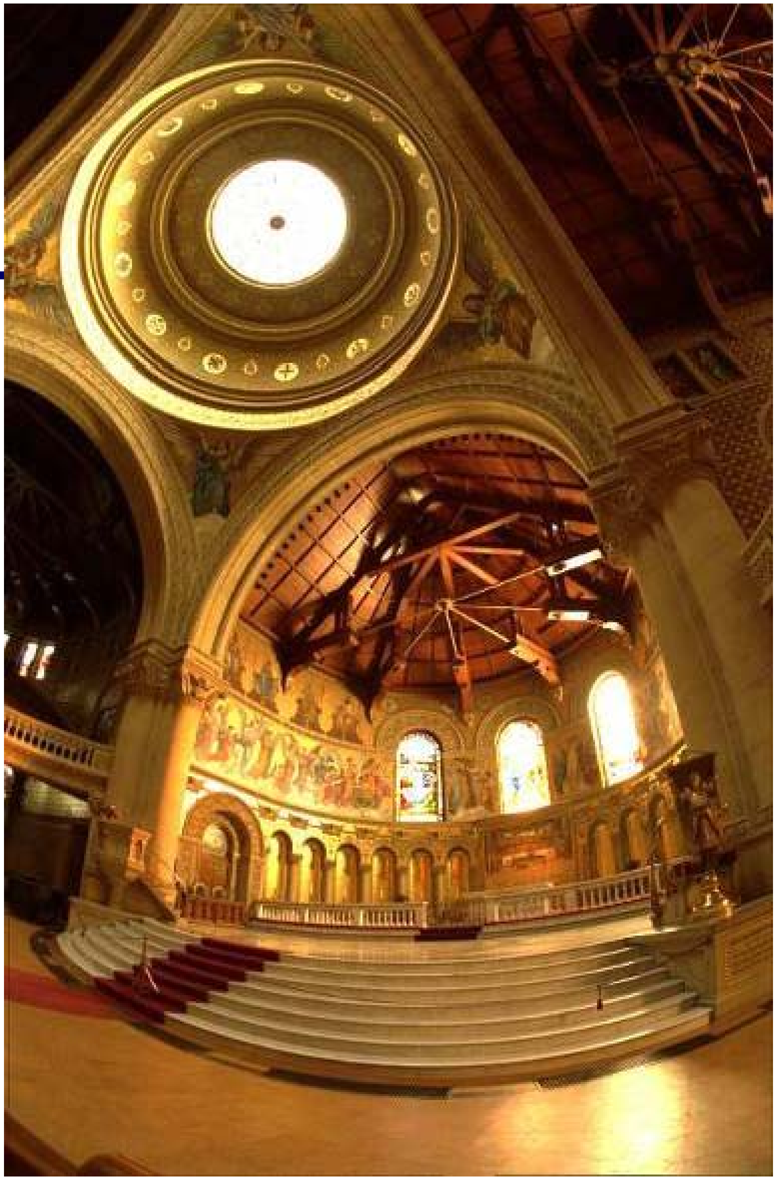

Reinhart Operator

Darkest 0.1% scaled to display device

What do *we* see?

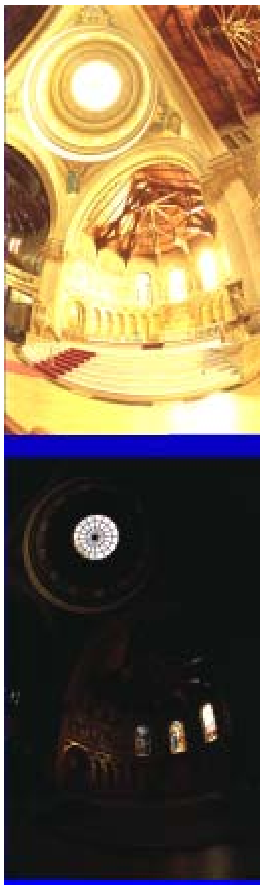

Vs.

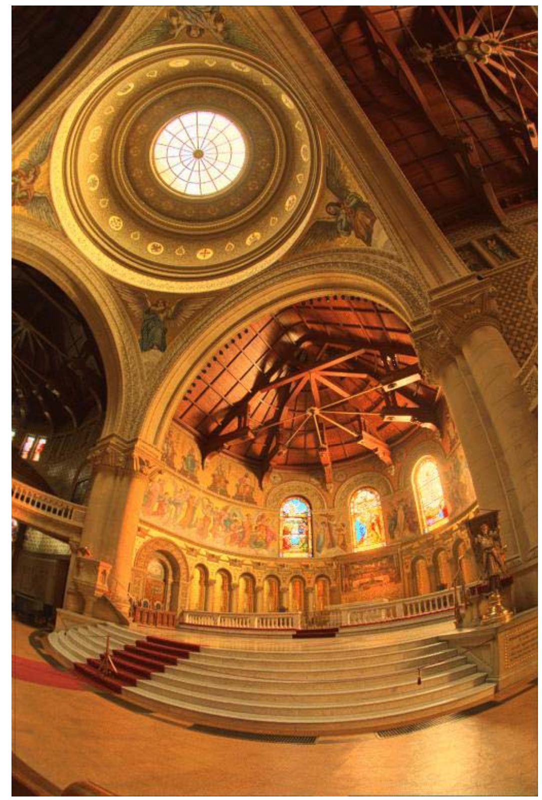

What does the eye sees?

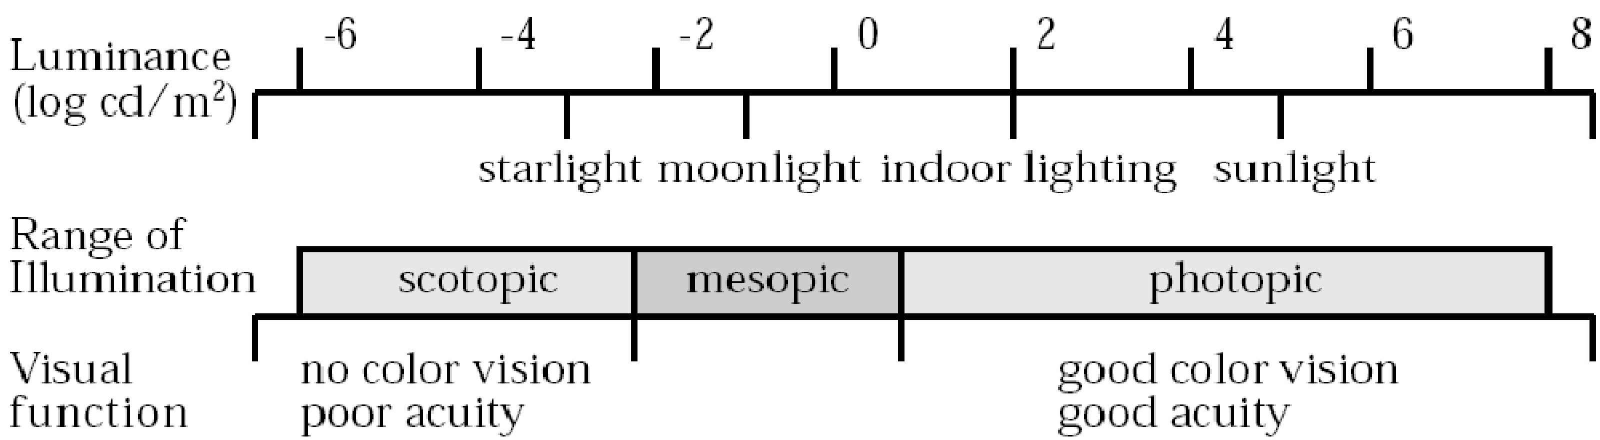

Figure 1: The range of luminances in the natural environment and associated visual parameters. After Hood (1986).

The eye has a huge dynamic range

Do we see a true radiance map?

Metamores

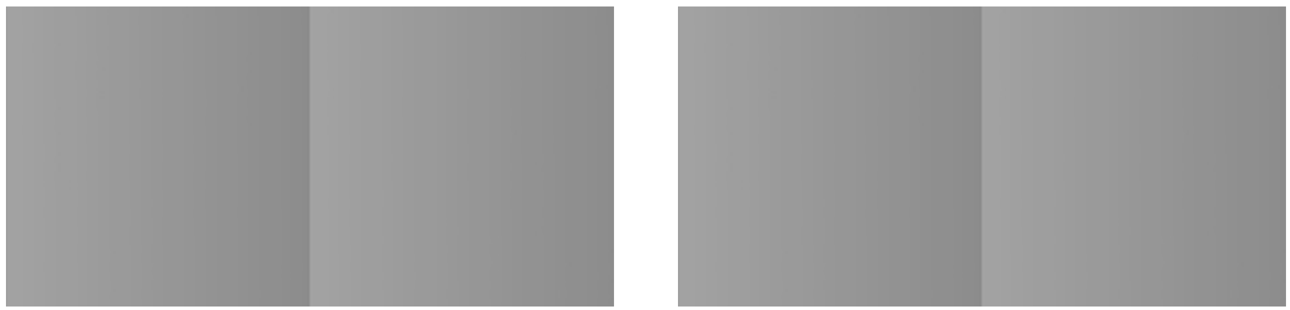

Craik-O'Brien Cornsweet Effect

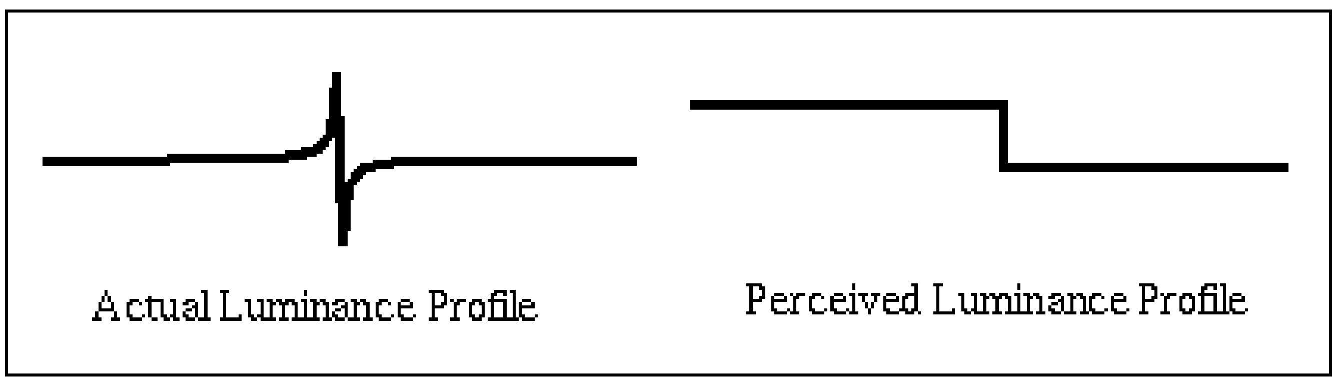

Can we use this for range compression?

Fast bilateral filtering for the display of high-dynamic-range images

Frédo Durand and Julie Dorsey
SIGGRAPH 2002.

Naïve: Gamma compression

X –> X$^{\gamma}$, colors are washed-out. Why?

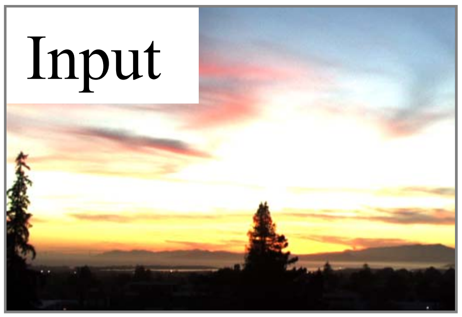

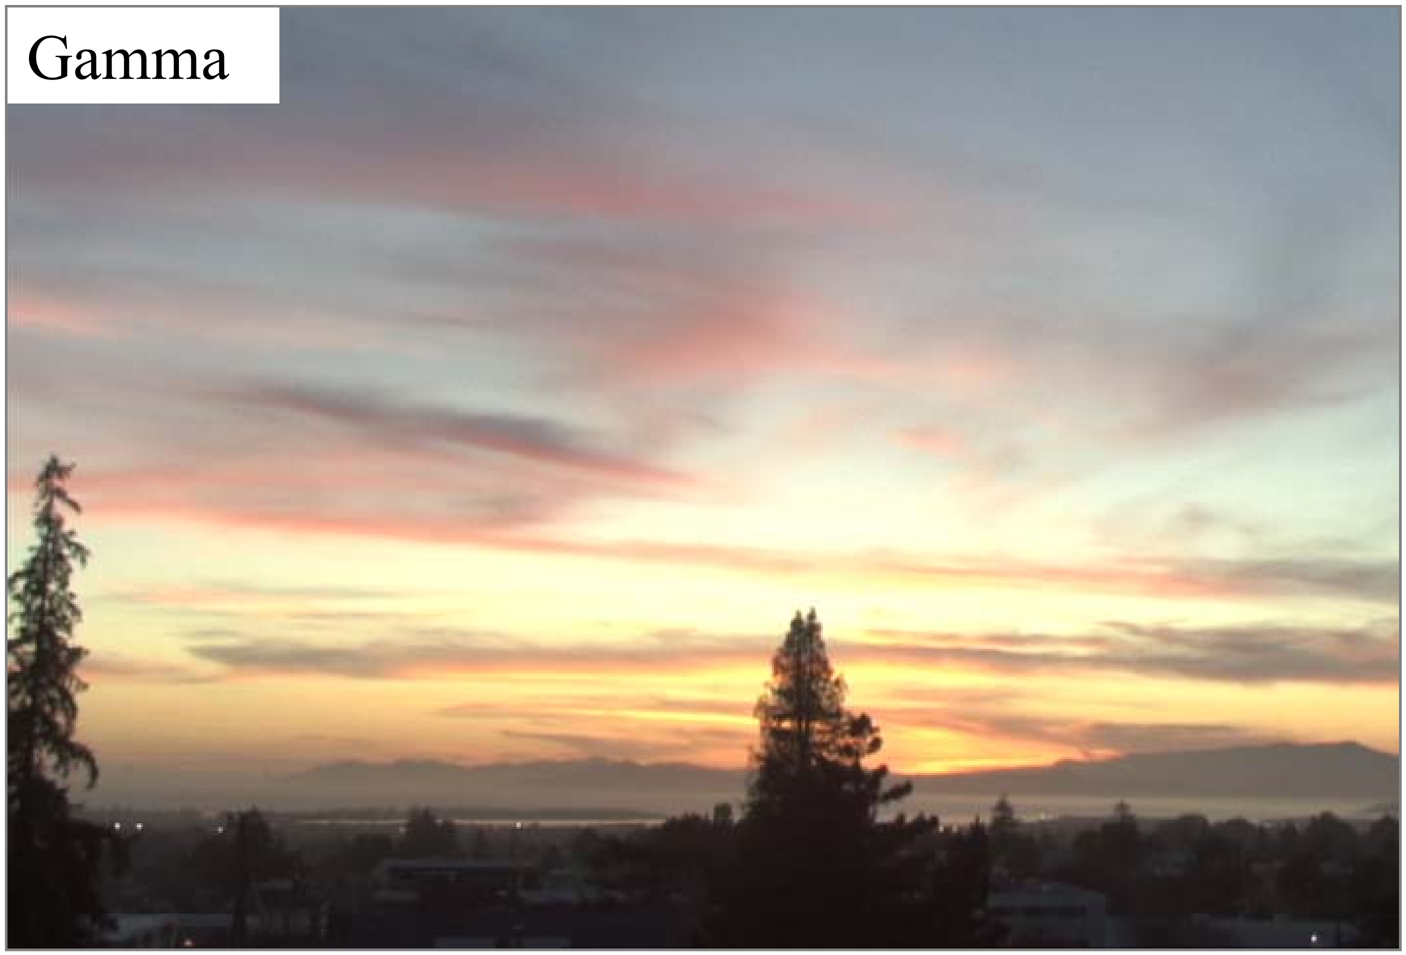

Richard Szeliski

Gamma compression on intensity

Colors are OK, details are blurred

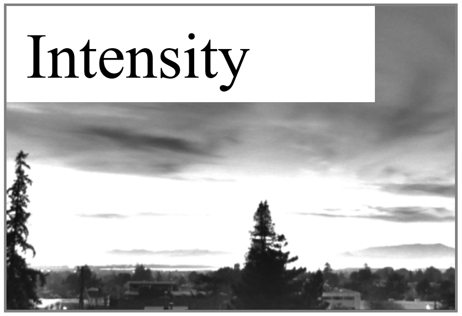

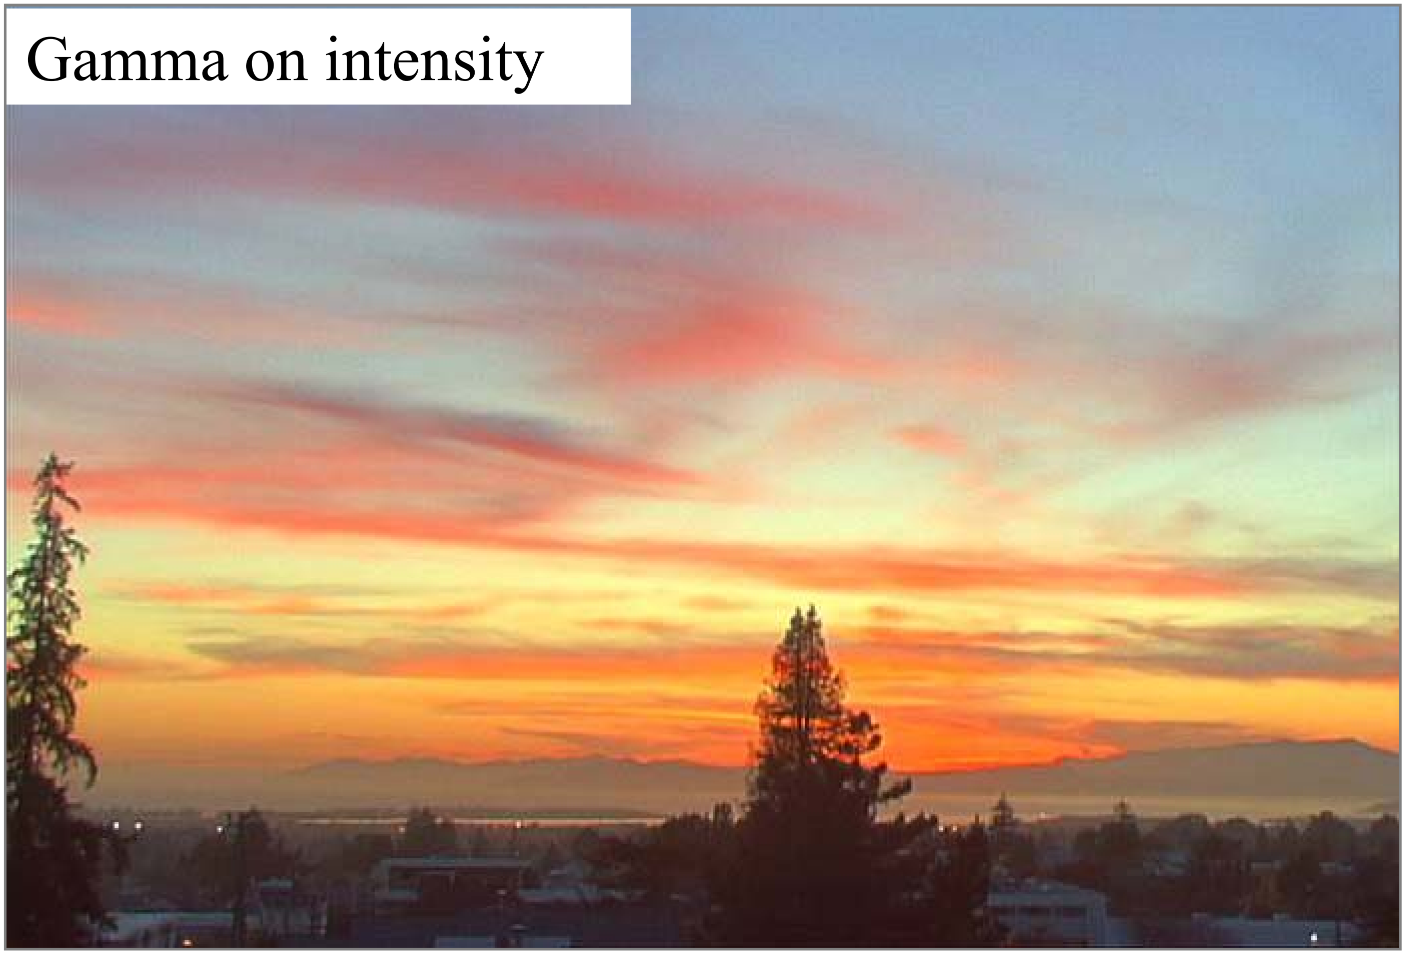

Oppenheim 1968, Chiu et al. 1993

Reduce contrast of low-frequencies, keep high

Halos

Strong edges contain high frequency

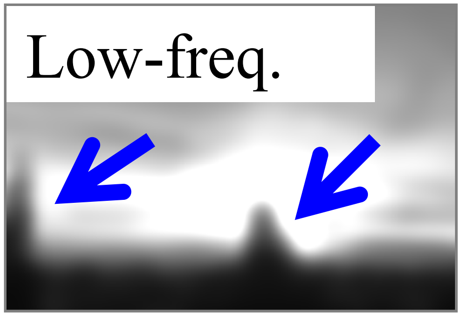

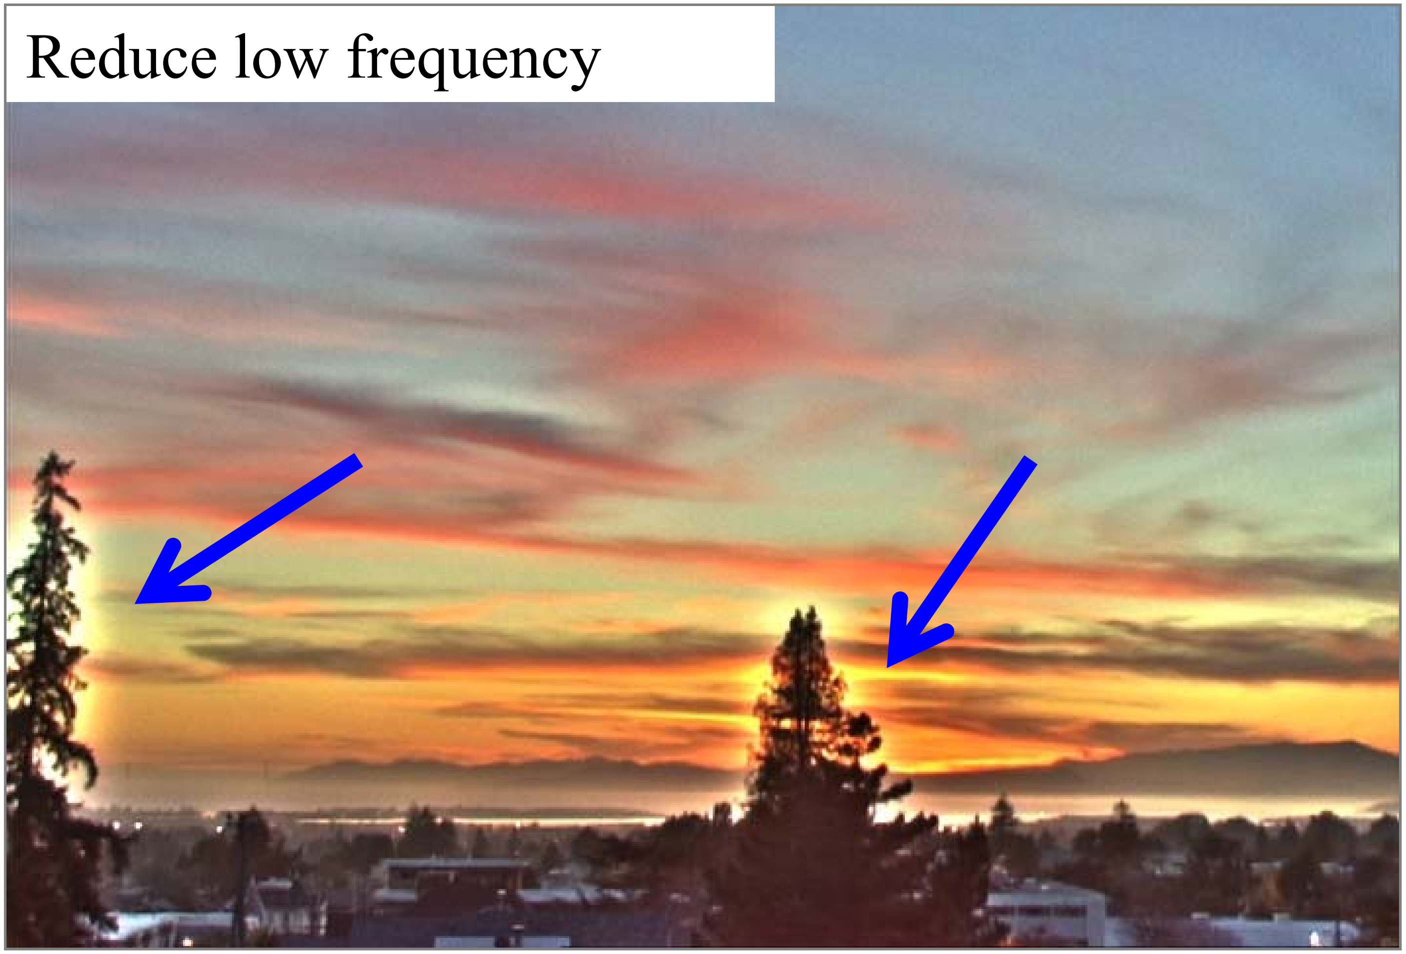

Our approach

Do not blur across edges: non-linear filtering

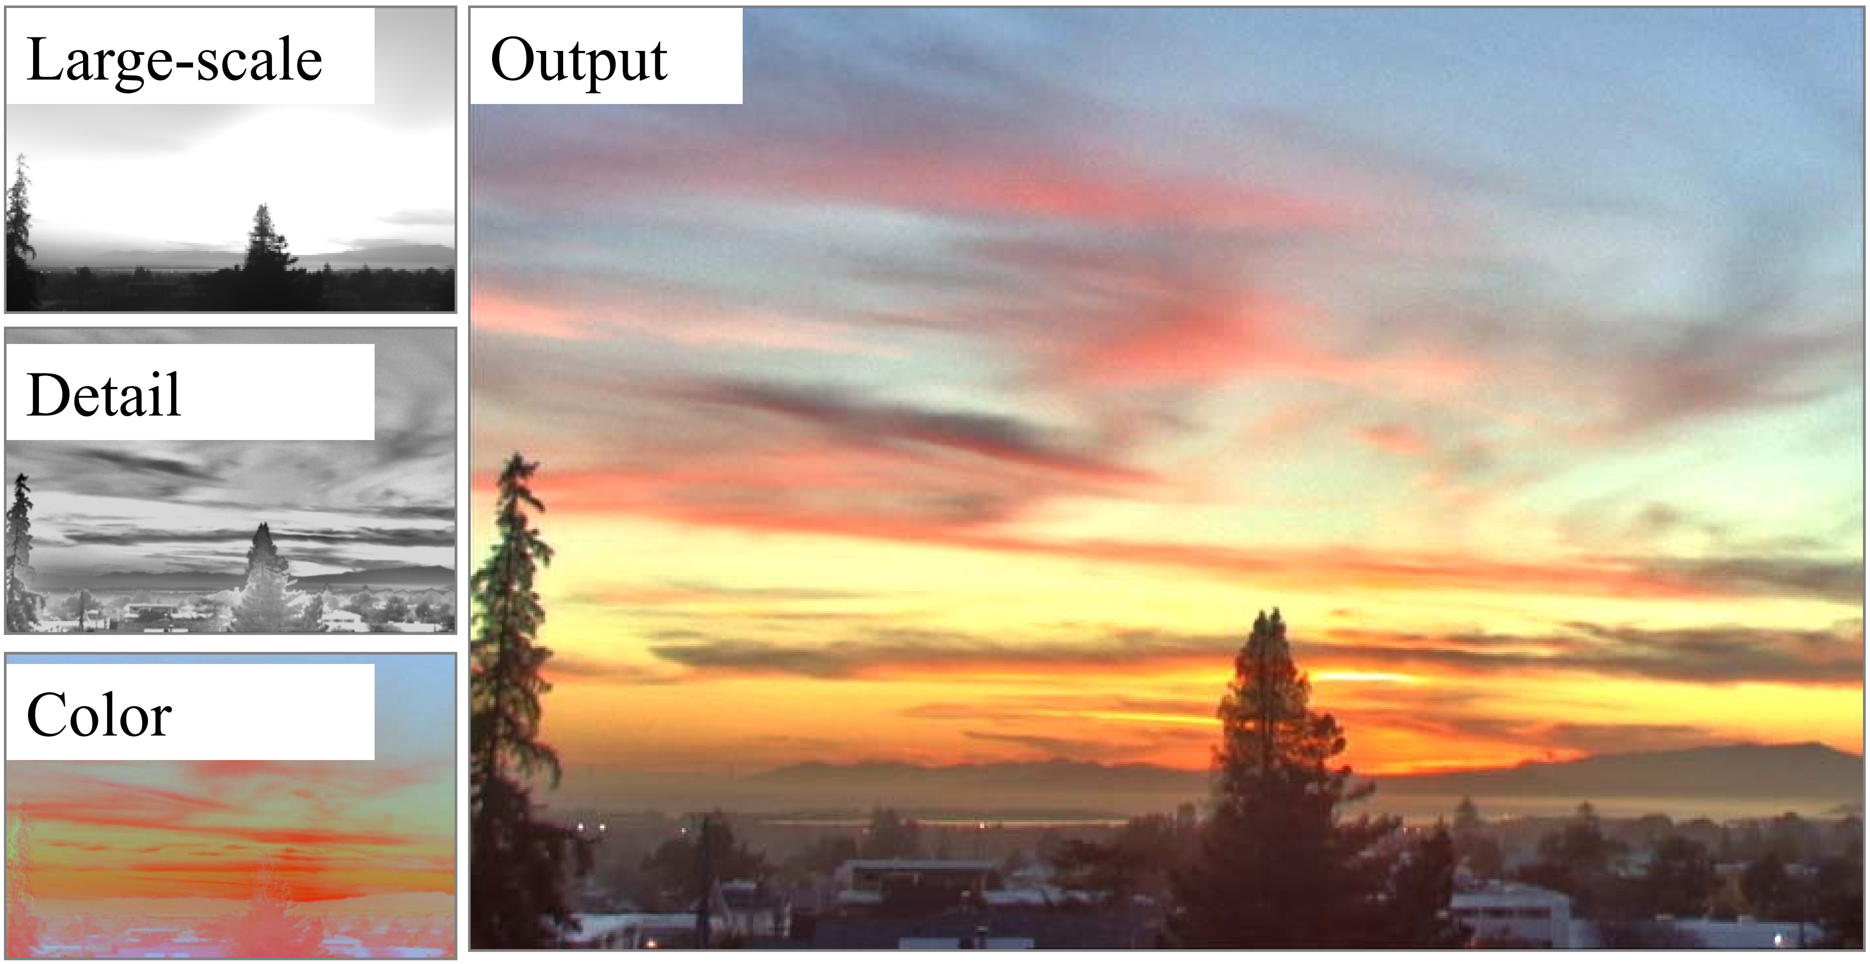

Bilateral filter

Tomasi and Manduci 1998

http://www.cse.ucsc.edu/~manduchi/Papers/ICCV98.pdf

Related to

- SUSAN filter [Smith and Brady 95] http://citeseer.ist.psu.edu/smith95susan.html
- Digital-TV [Chan, Osher and Chen 2001] http://citeseer.ist.psu.edu/chan01digital.html
- sigma filter http://www.geogr.ku.dk/CHIPS/Manual/f187.htm

Start with Gaussian filtering

Output is blurred

$$J = f \otimes I$$

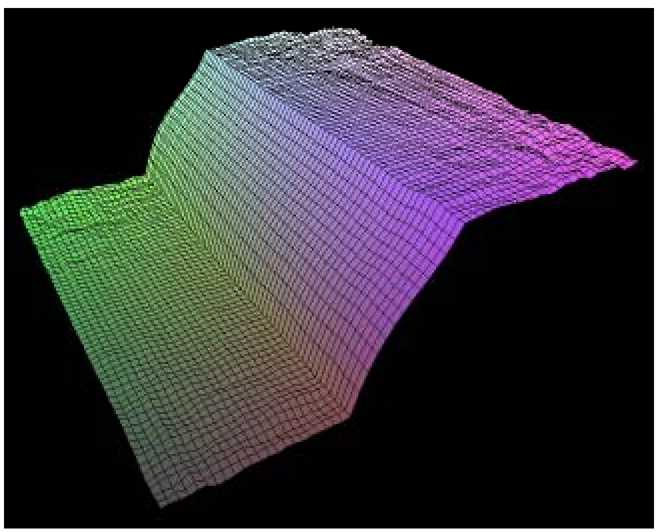

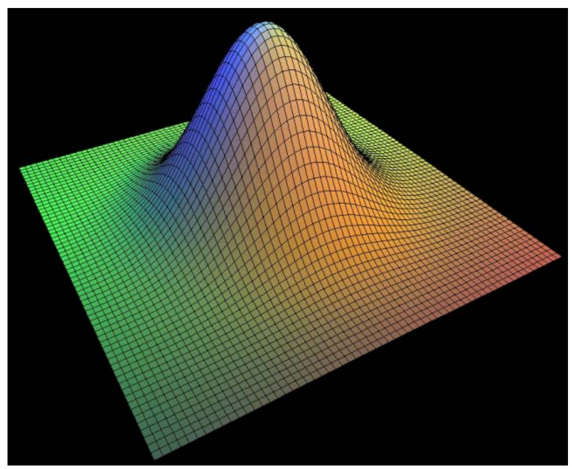

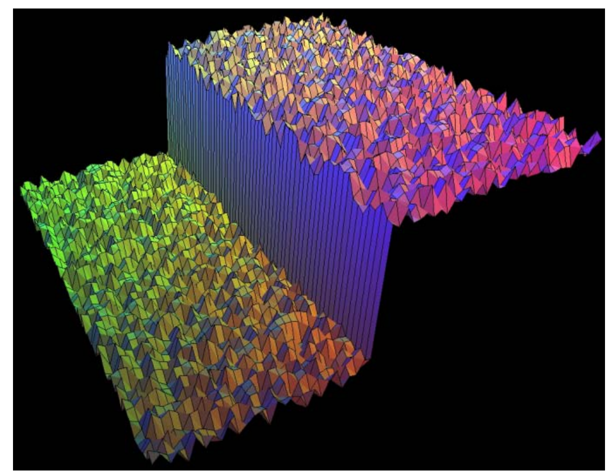

output ← input

Bilateral filtering is non-linear

The weights are different for each output pixel

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi) \quad g(I(\xi)-I(x)) \quad I(\xi)$$

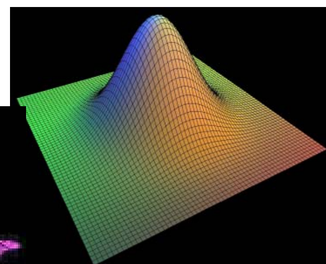

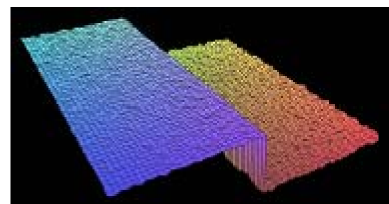

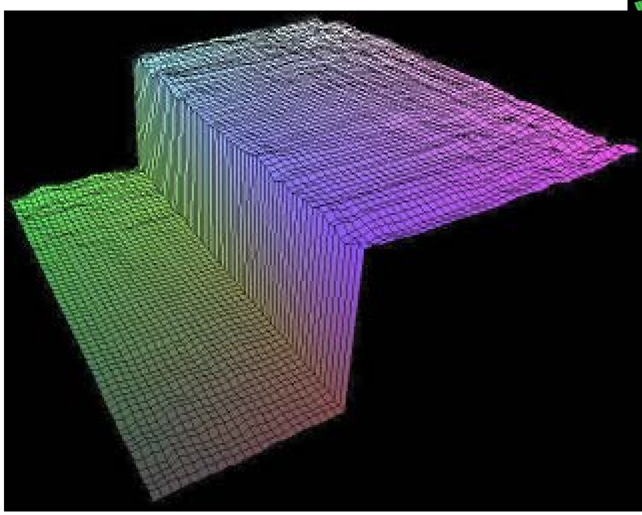

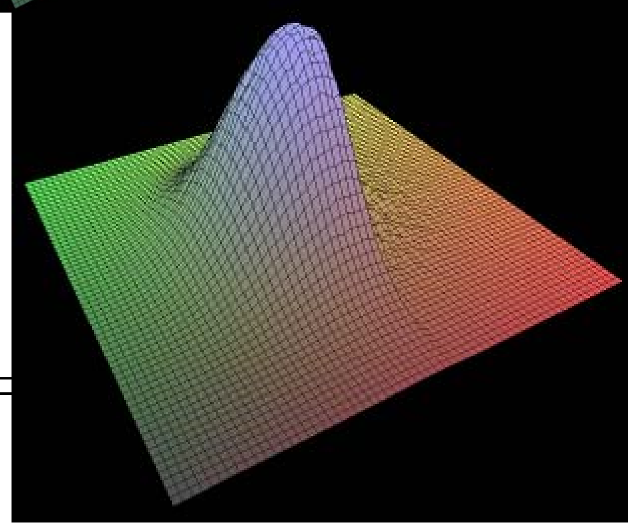

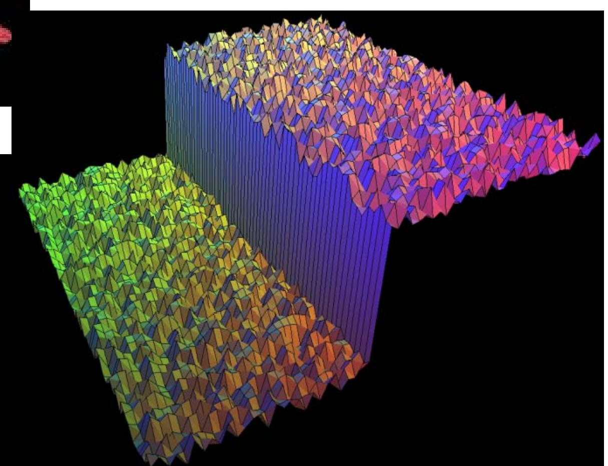

output

input

Other view

The bilateral filter uses the 3D distance

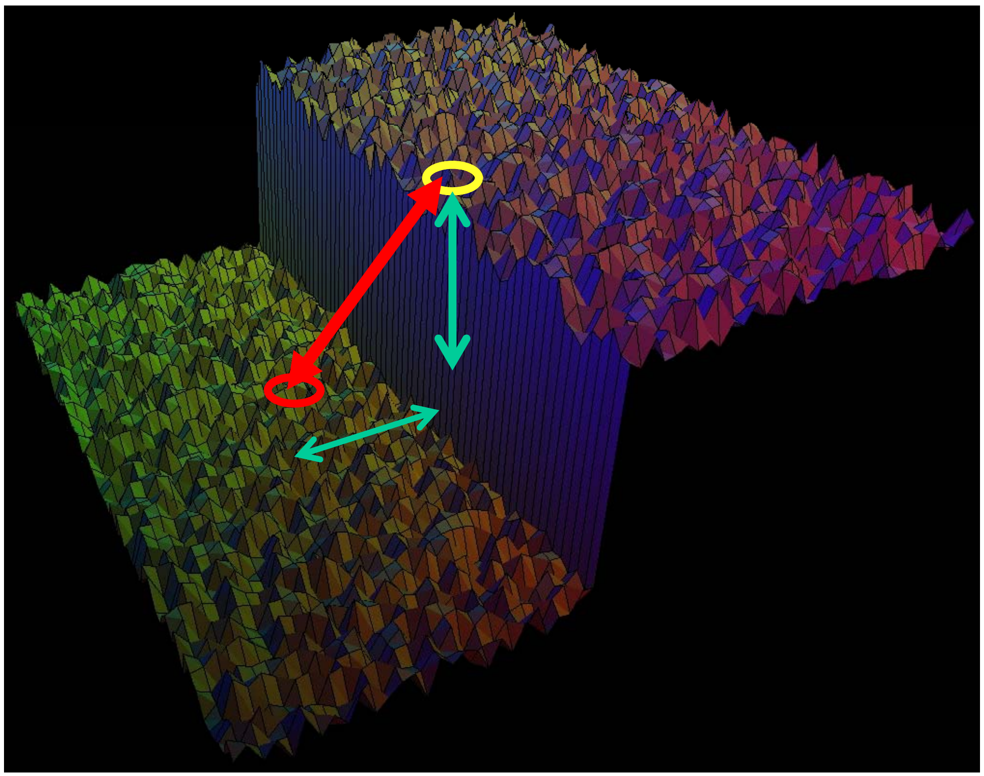

Contrast reduction

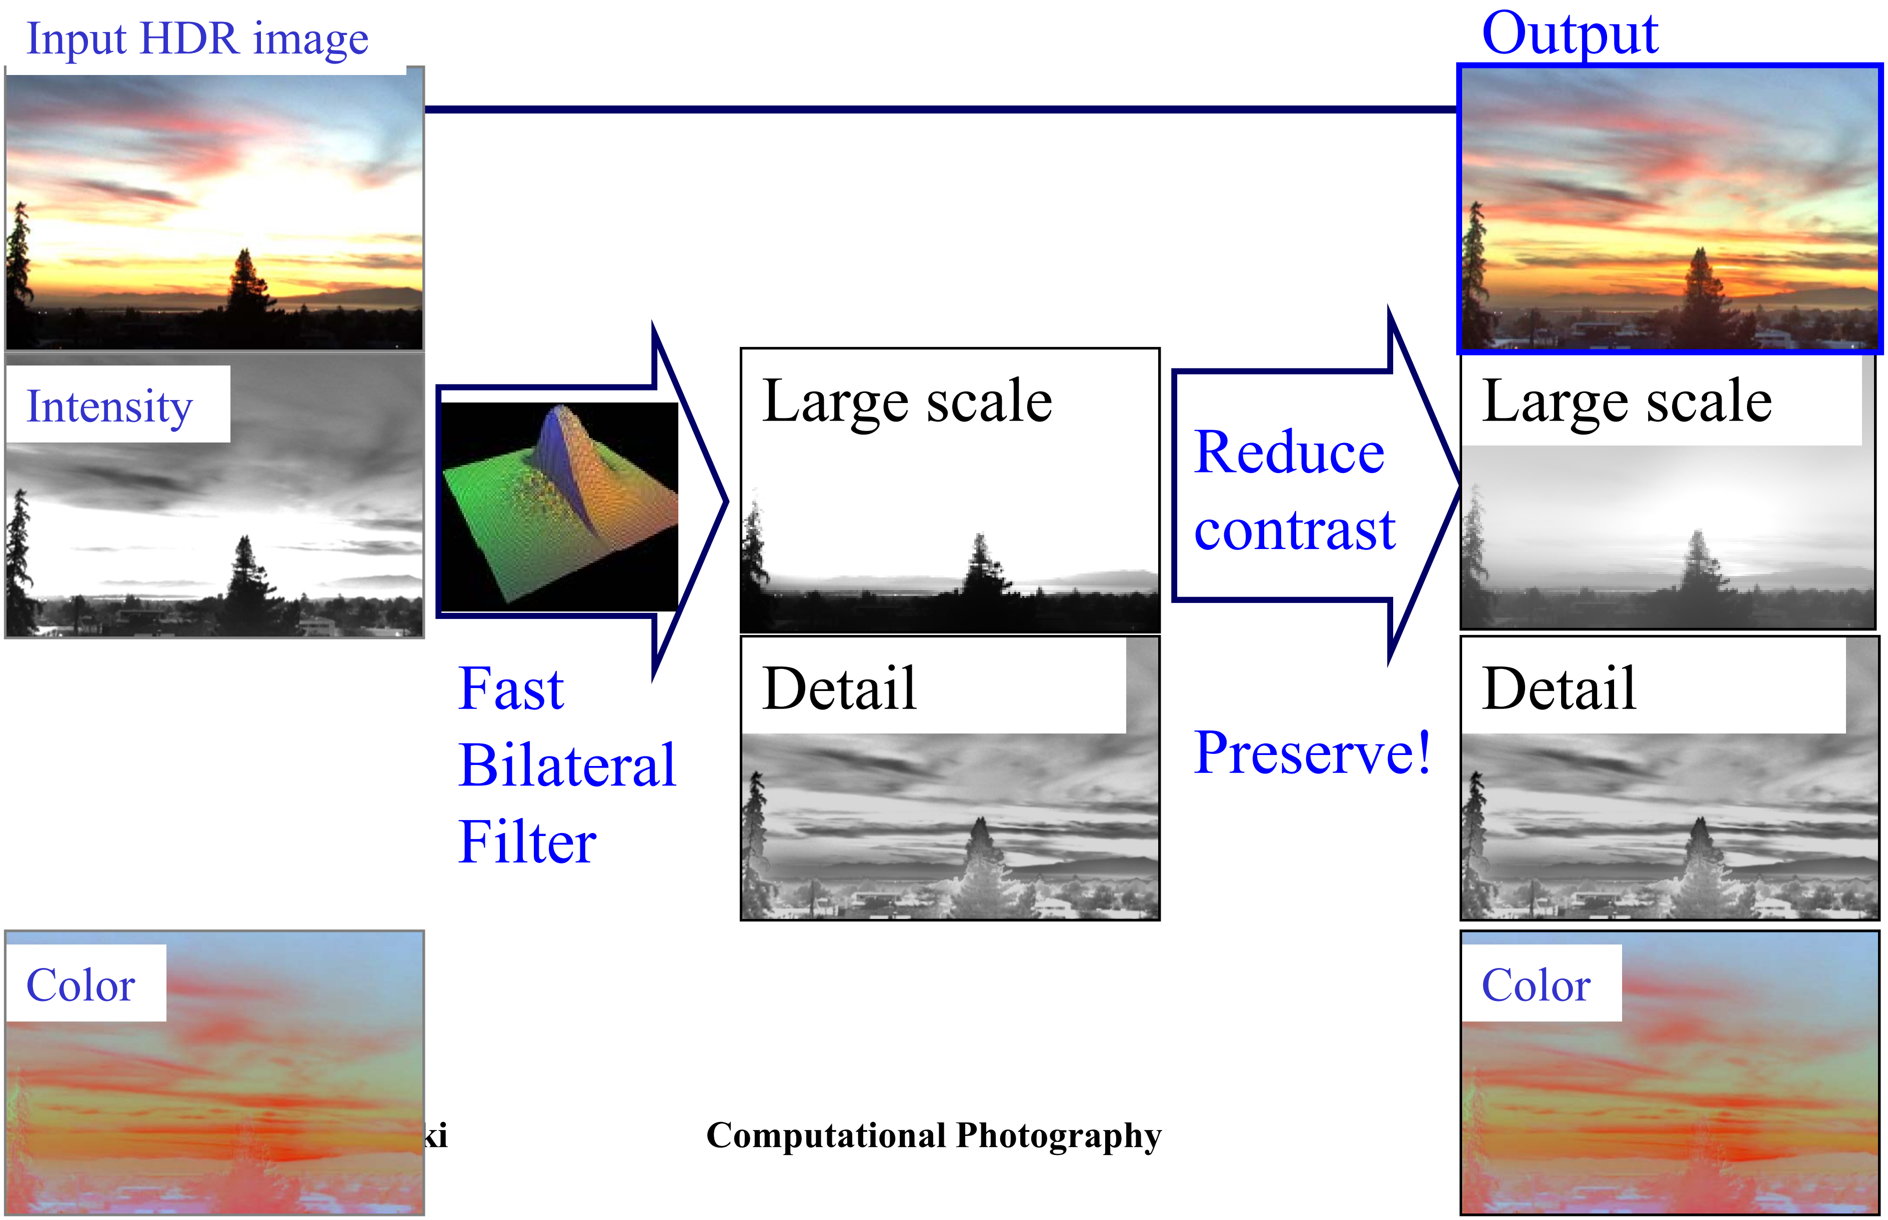

Dynamic range reduction

- To reduce contrast of base layer
  - scale in the log domain ➔ γ exponent in linear
- Set a target range: $\log_{10}$ (5)
- Compute range in the log layer: (max-min)
- Deduce γ using *division*
- Normalize so that the biggest value in the (linear) base is 1 (0 in log):
  - offset the compressed based by its max

Summary of approach

Do not blur base/gain layer: non-linear filtering

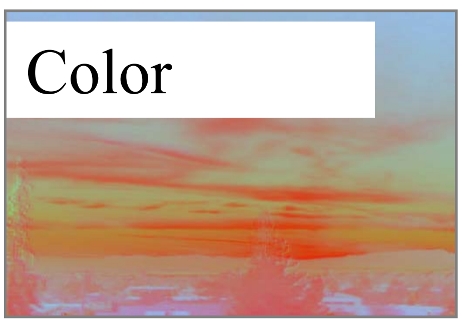

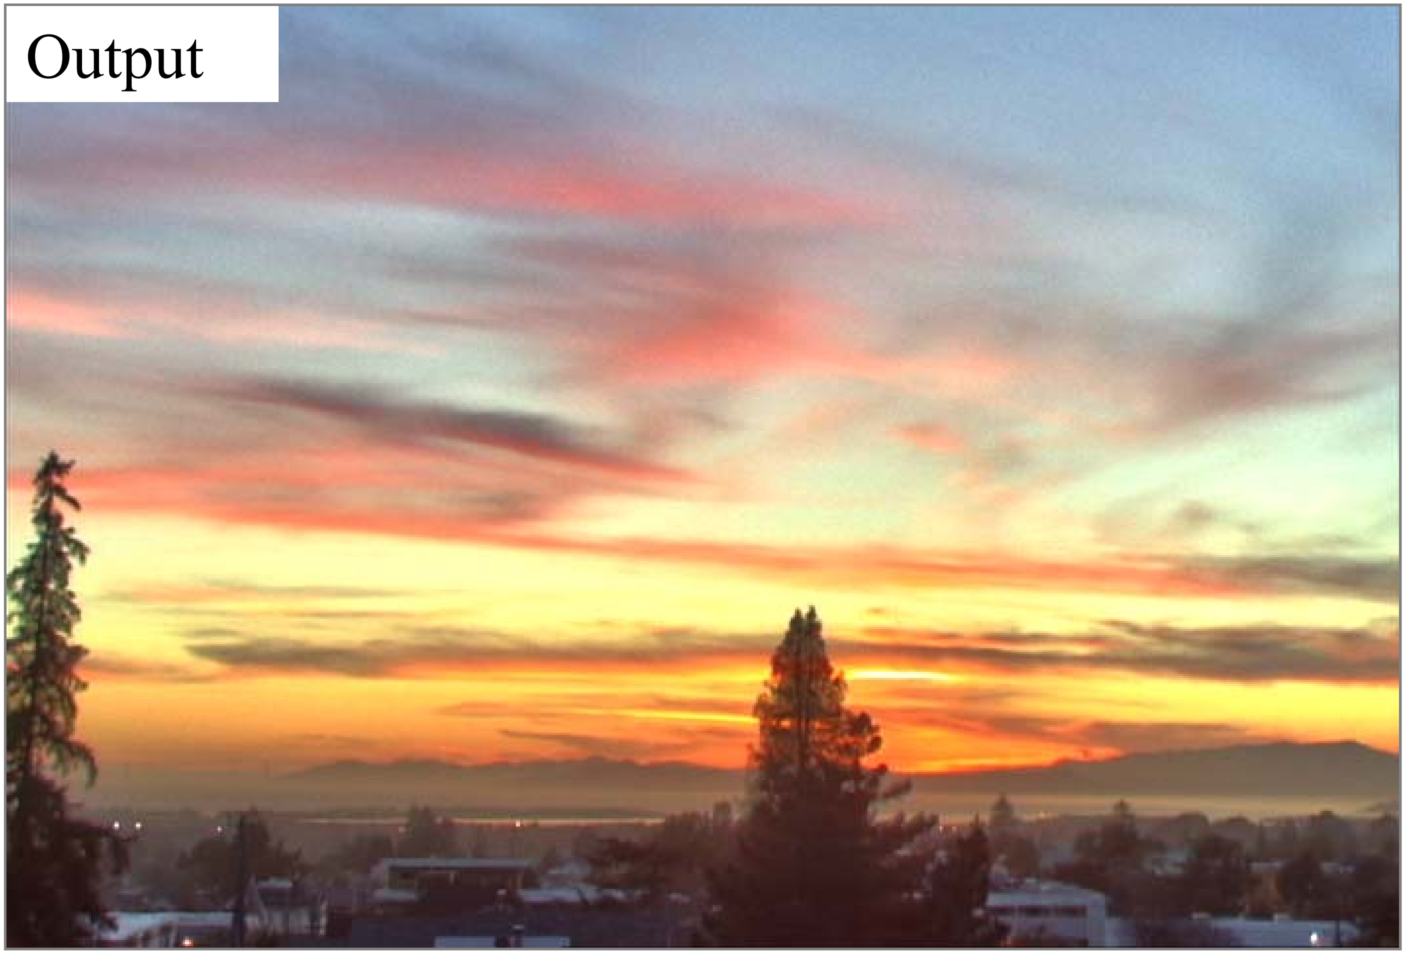

Gradient domain high dynamic range compression

Raanan Fattal, Dani Lischinski, and
Michael Werman
SIGGRAPH 2002.

Gradient Tone Mapping

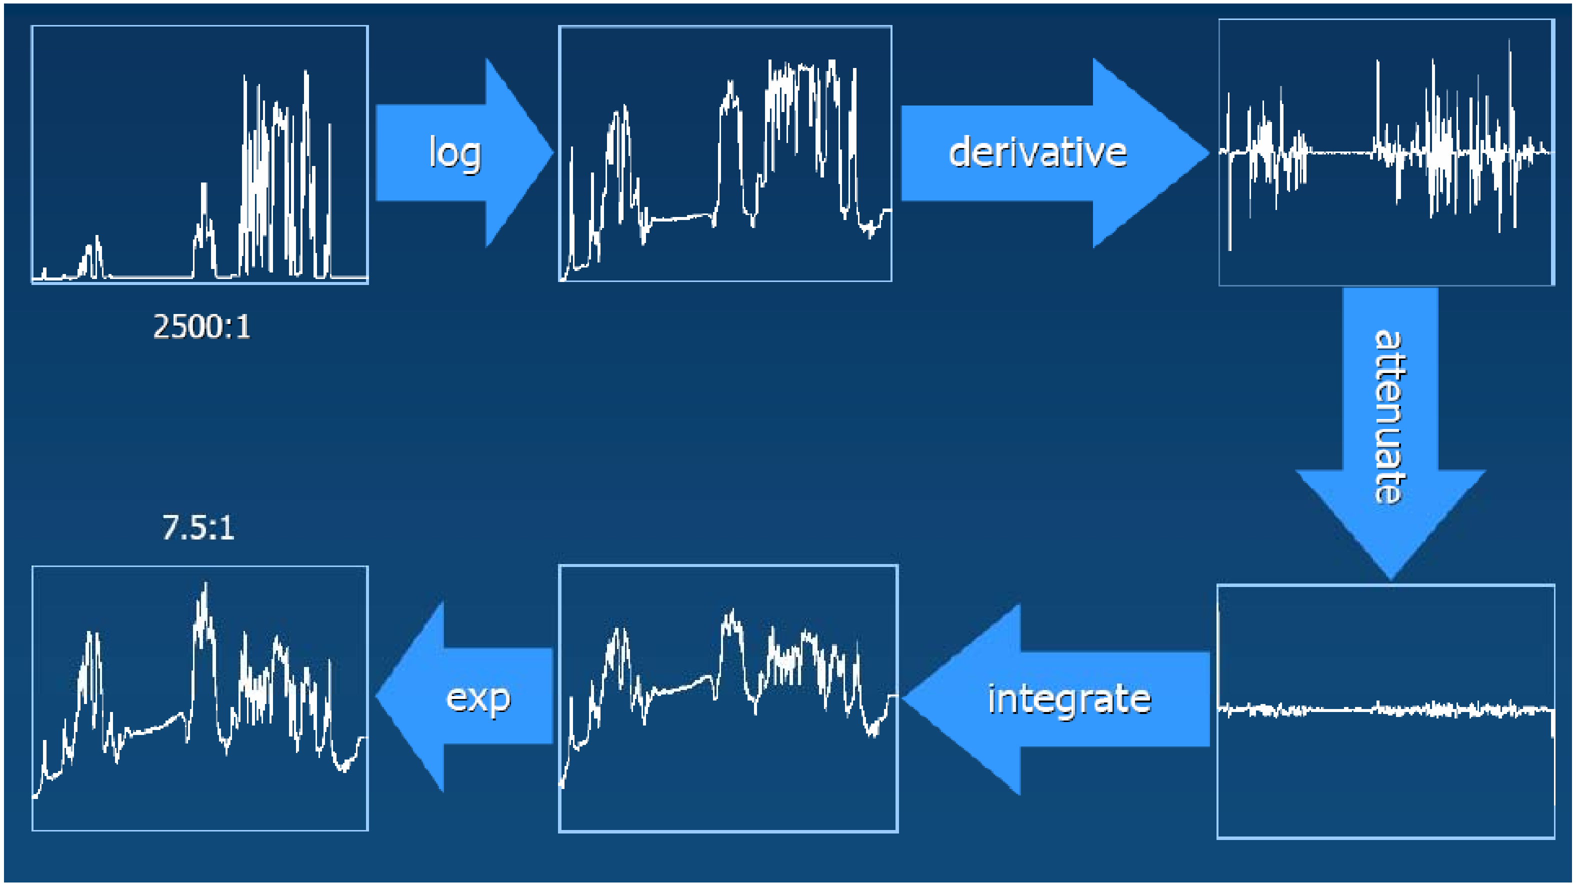

Slide from Siggraph 2005 by Raskar (Graphs by Fattal et al.)

Gradient attenuation

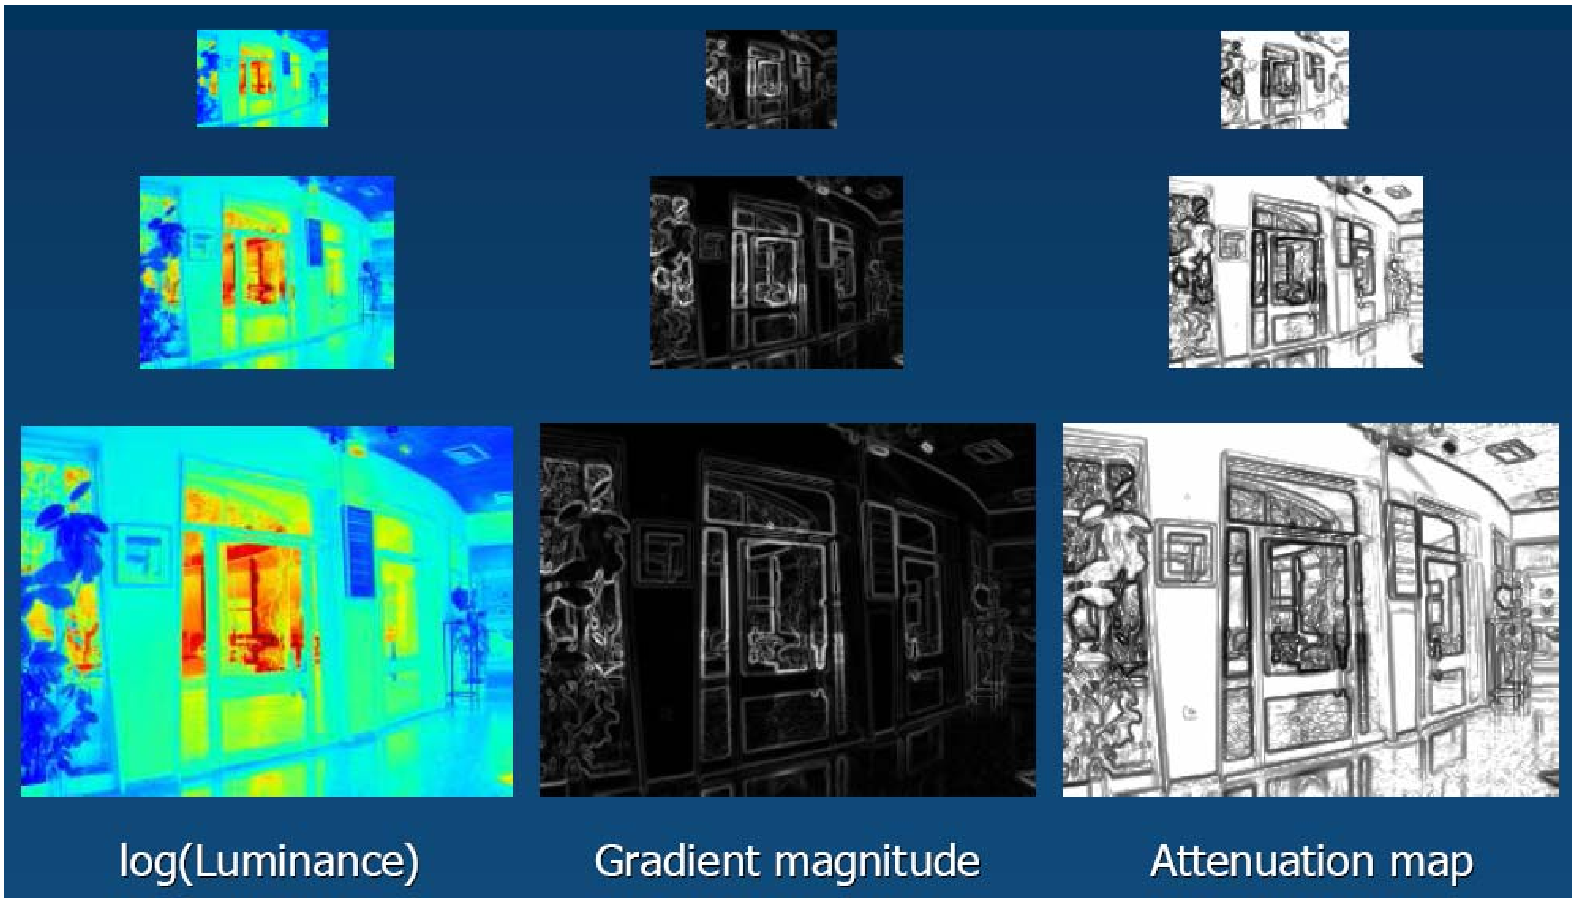

From Fattal et al.

Interactive Local Adjustment of Tonal Values

Dani Lischinski
Zeev Farbman
The Hebrew University

Matt Uyttendaele
Rick Szeliski
Microsoft Research

SIGGRAPH 2006

Tonal Manipulation

- brightness
- exposure
- contrast
- saturation
- color temperature
- …

Interpretation 1:

Interpretation 2:

Interpretation 3:

This Work is About:

New tool for interactive tonal manipulation: developing negatives in the digital darkroom.

Target material:

- HDR images: the ultimate digital negative.
- Camera RAW images: the most common digital negative.
- Ordinary snapshots.

Existing Tools

Automatic tone mapping algorithms

- Why do we need yet another tone mapping approach?
- Why interactive rather than automatic?

Image manipulation and editing packages, e.g., Adobe Photoshop.

Tone Reproduction Operators

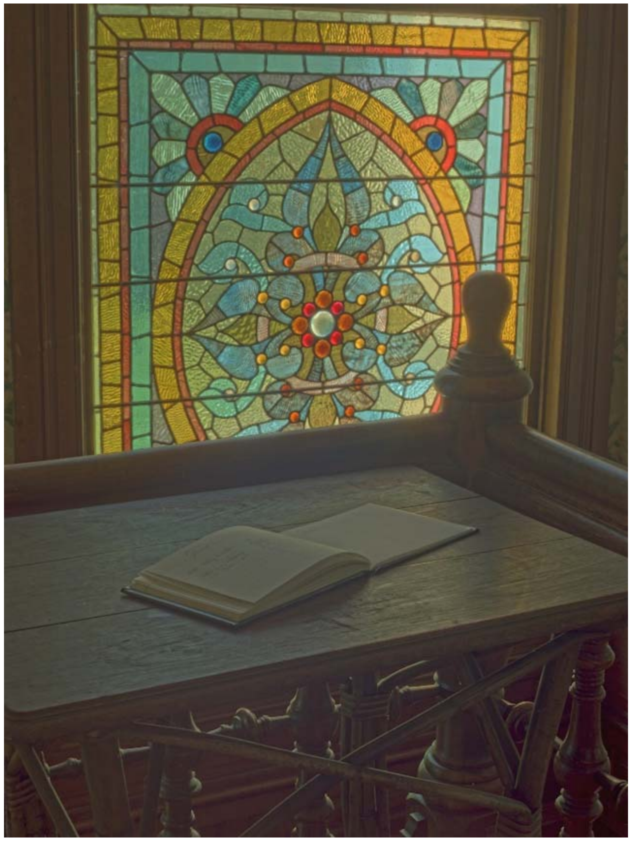

Bilateral Filtering

Durand & Dorsey 2002

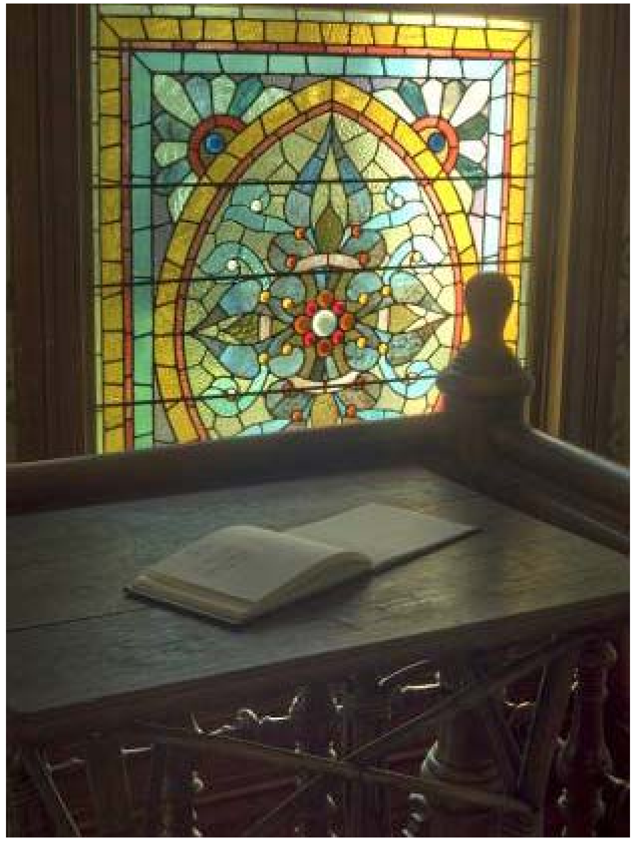

Gradient Domain

Fattal et al. 2002

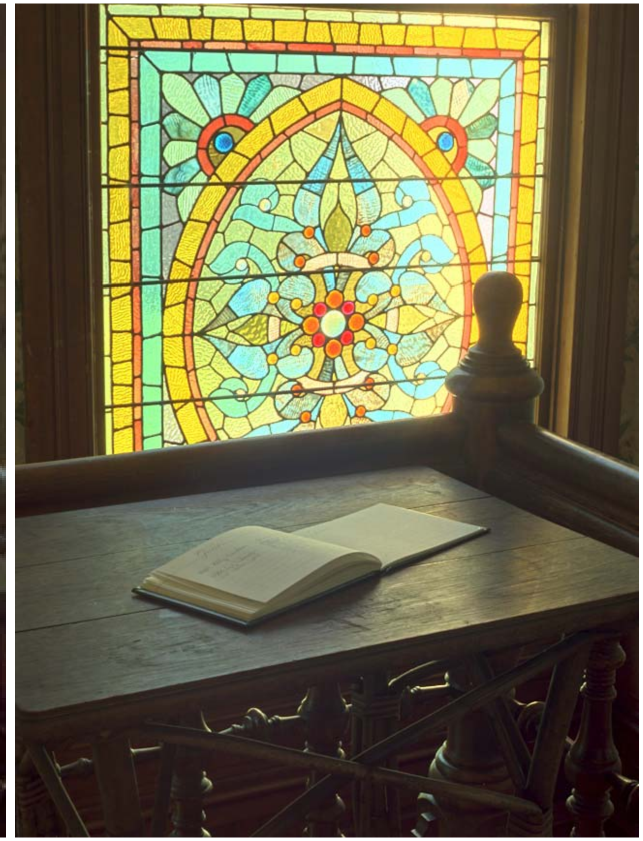

Photographic

Reinhard et al. 2002

Automatic vs. Interactive

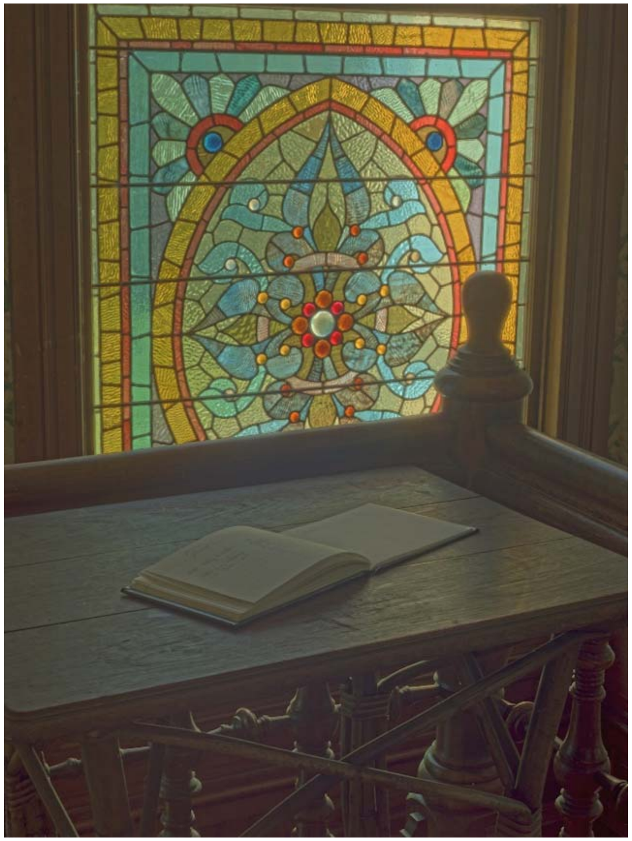

Bilateral Filtering

Durand & Dorsey 2002

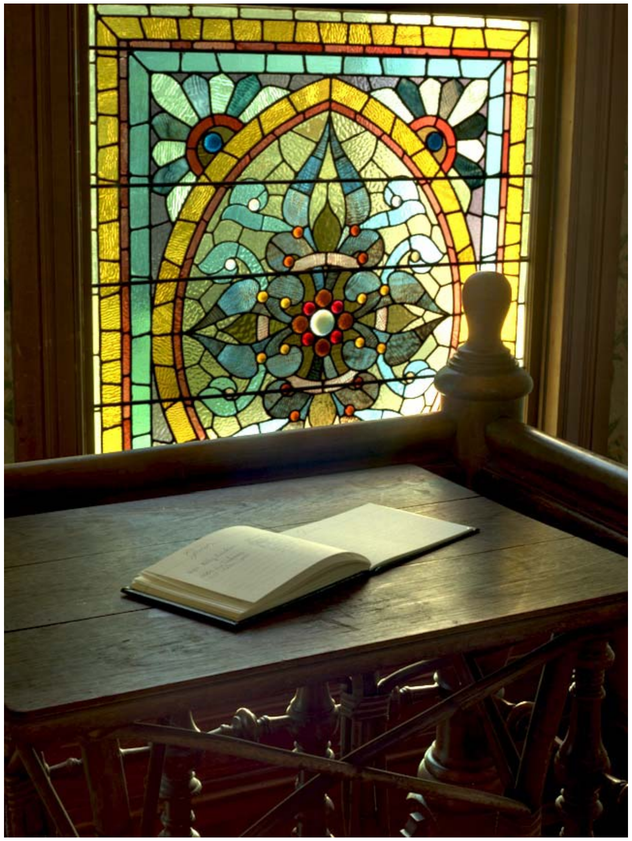

Interactive Tone Mapping

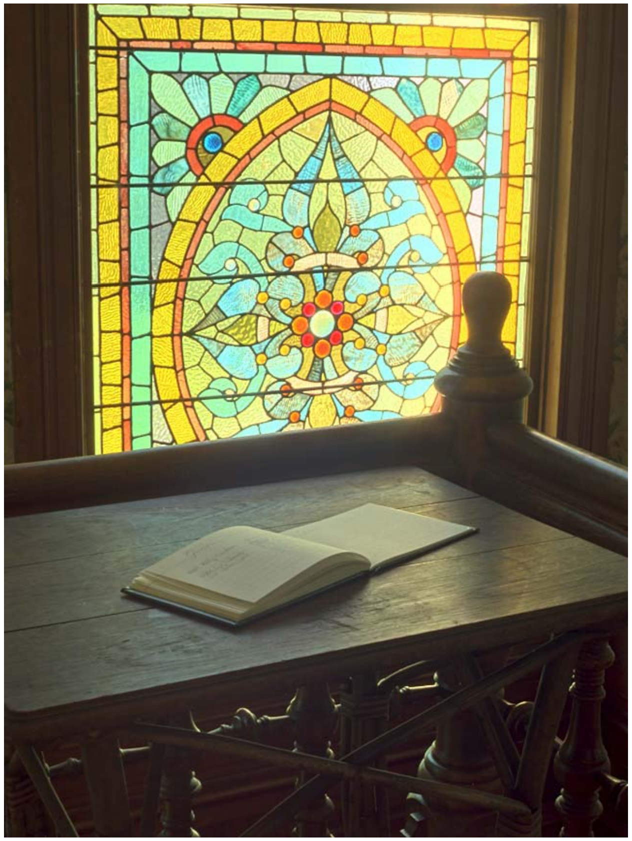

Photographic

Reinhard et al. 2002

Automatic vs. Interactive

Existing automatic TM operators are “black boxes”

- No direct control over the outcome
- No local adjustment
- Not suitable for creative/artistic work
- Results do not always look “photographic”
- Most operators not *really* automatic

But What About Photoshop?

You can do just about everything …

Adjustment Layers

Layer Masks

- Select regions
- Paint blending weights

… but you need a lot of experience, patience, and time!

Example

15 minutes in Photoshop: 3 minutes:

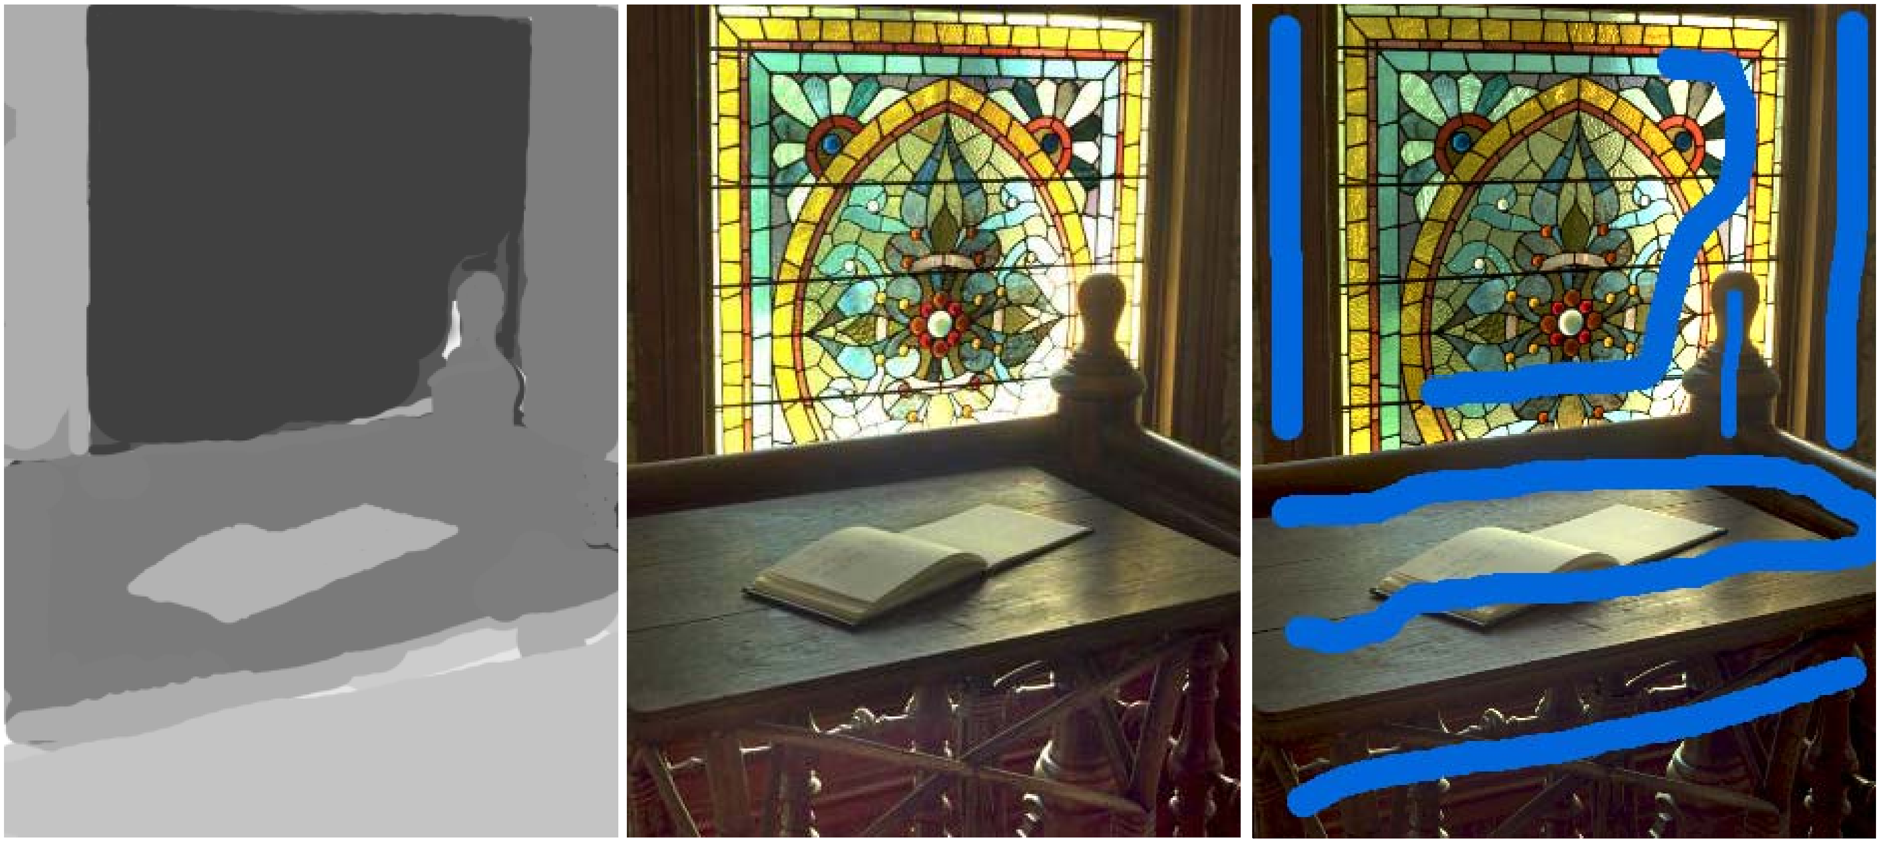

Approach

User indicates regions using scribbles.

User adjusts tonal values using sliders.

Scribbles + tonal values are interpreted as soft constraints.

Optimization framework “propagates” the constraints to the entire image.

User interface

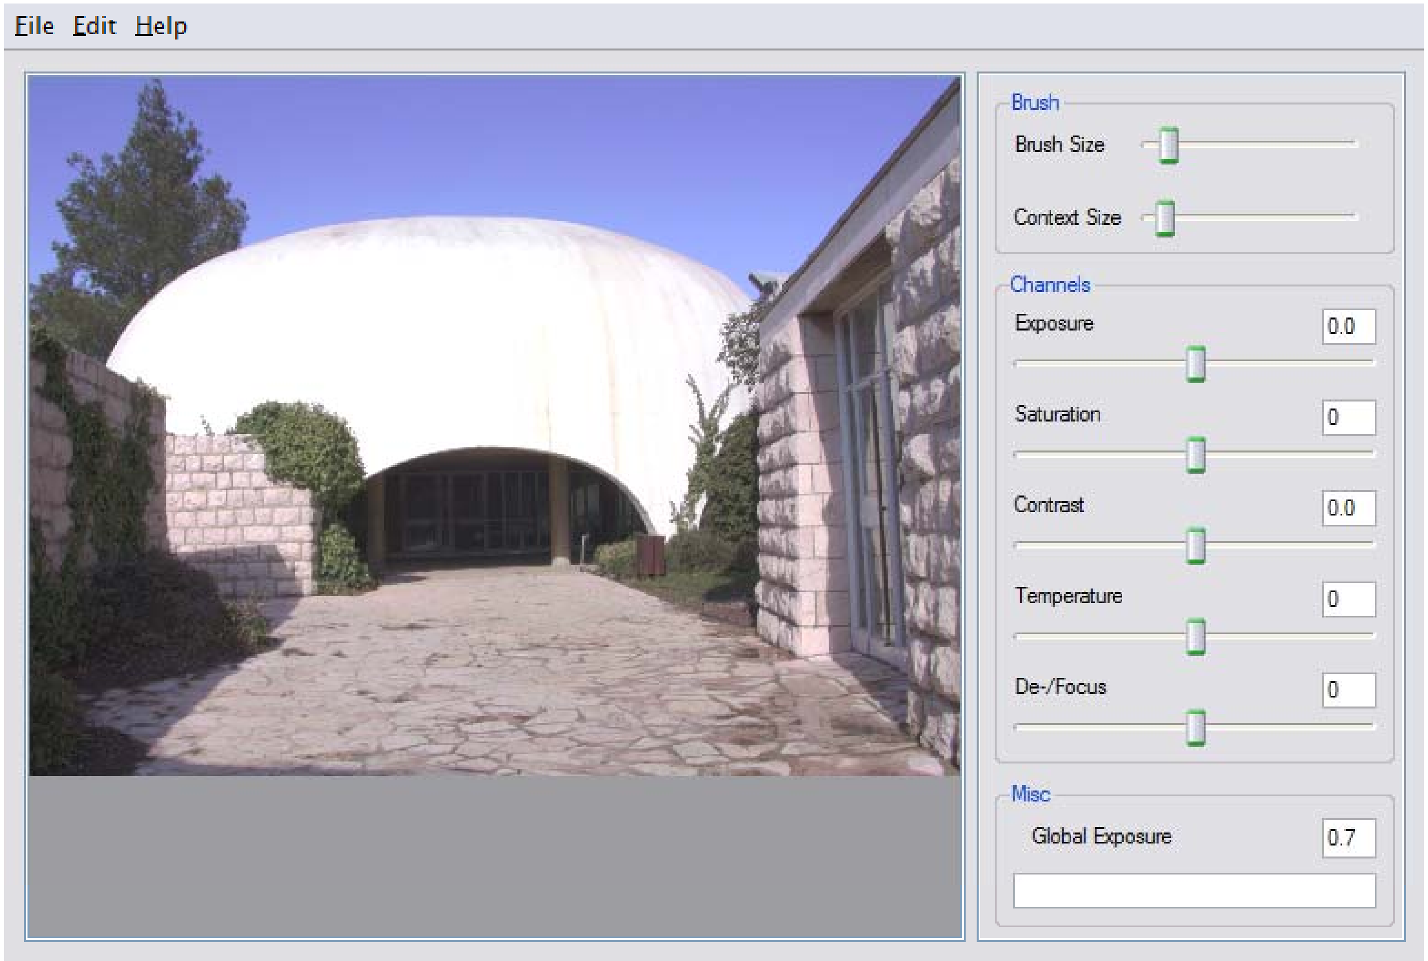

Input: constraints

Result: adjustment map

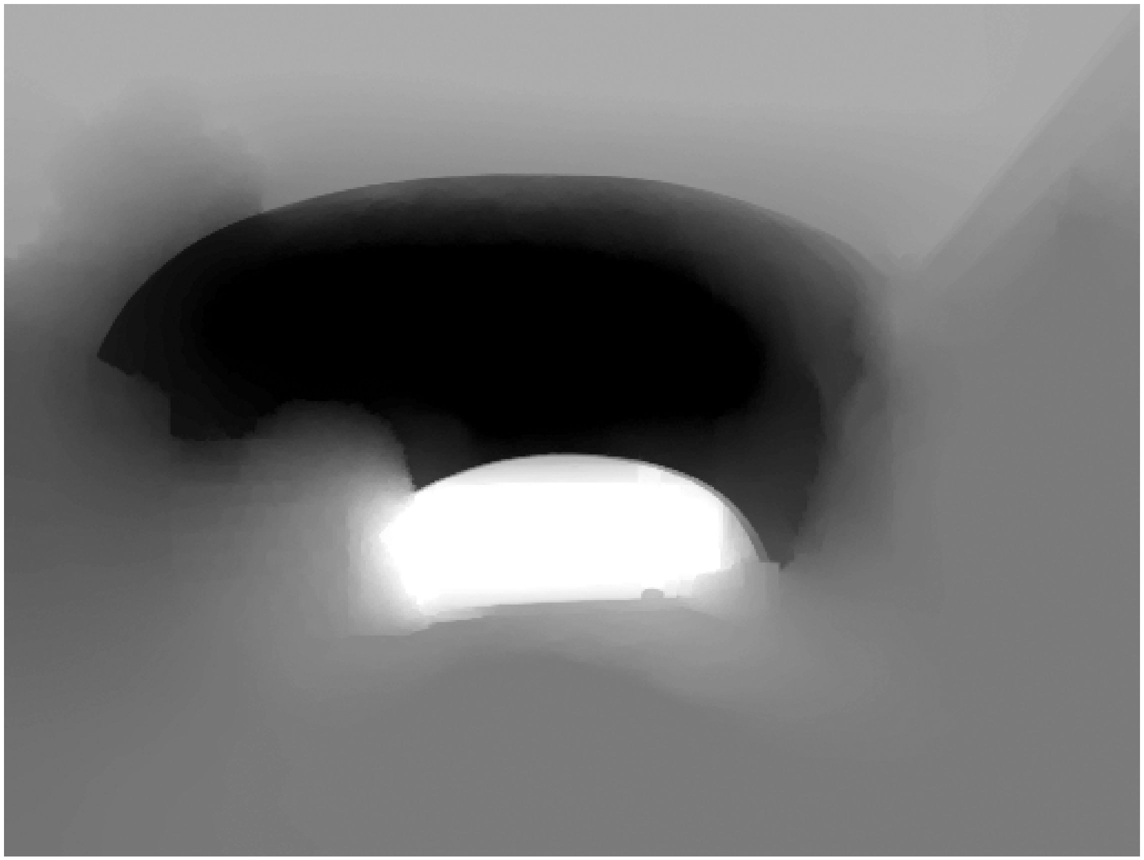

Constraint Propagation

Approximate constraints with a function whose smoothness is determined by underlying image:

$$f = \arg\min_f \left\{ \underbrace{\sum_{\mathbf{x}} w(\mathbf{x})\,(f(\mathbf{x}) - g(\mathbf{x}))^2}_{\text{data term}} + \lambda \underbrace{\sum_{\mathbf{x}} h(\nabla f, \nabla L)}_{\text{smoothness term}} \right\}$$

Our smoothness term:

$$h(\nabla f, \nabla L) = \frac{|f_x|^2}{|L_x|^\alpha + \varepsilon} + \frac{|f_y|^2}{|L_y|^\alpha + \varepsilon}$$

Influence Functions

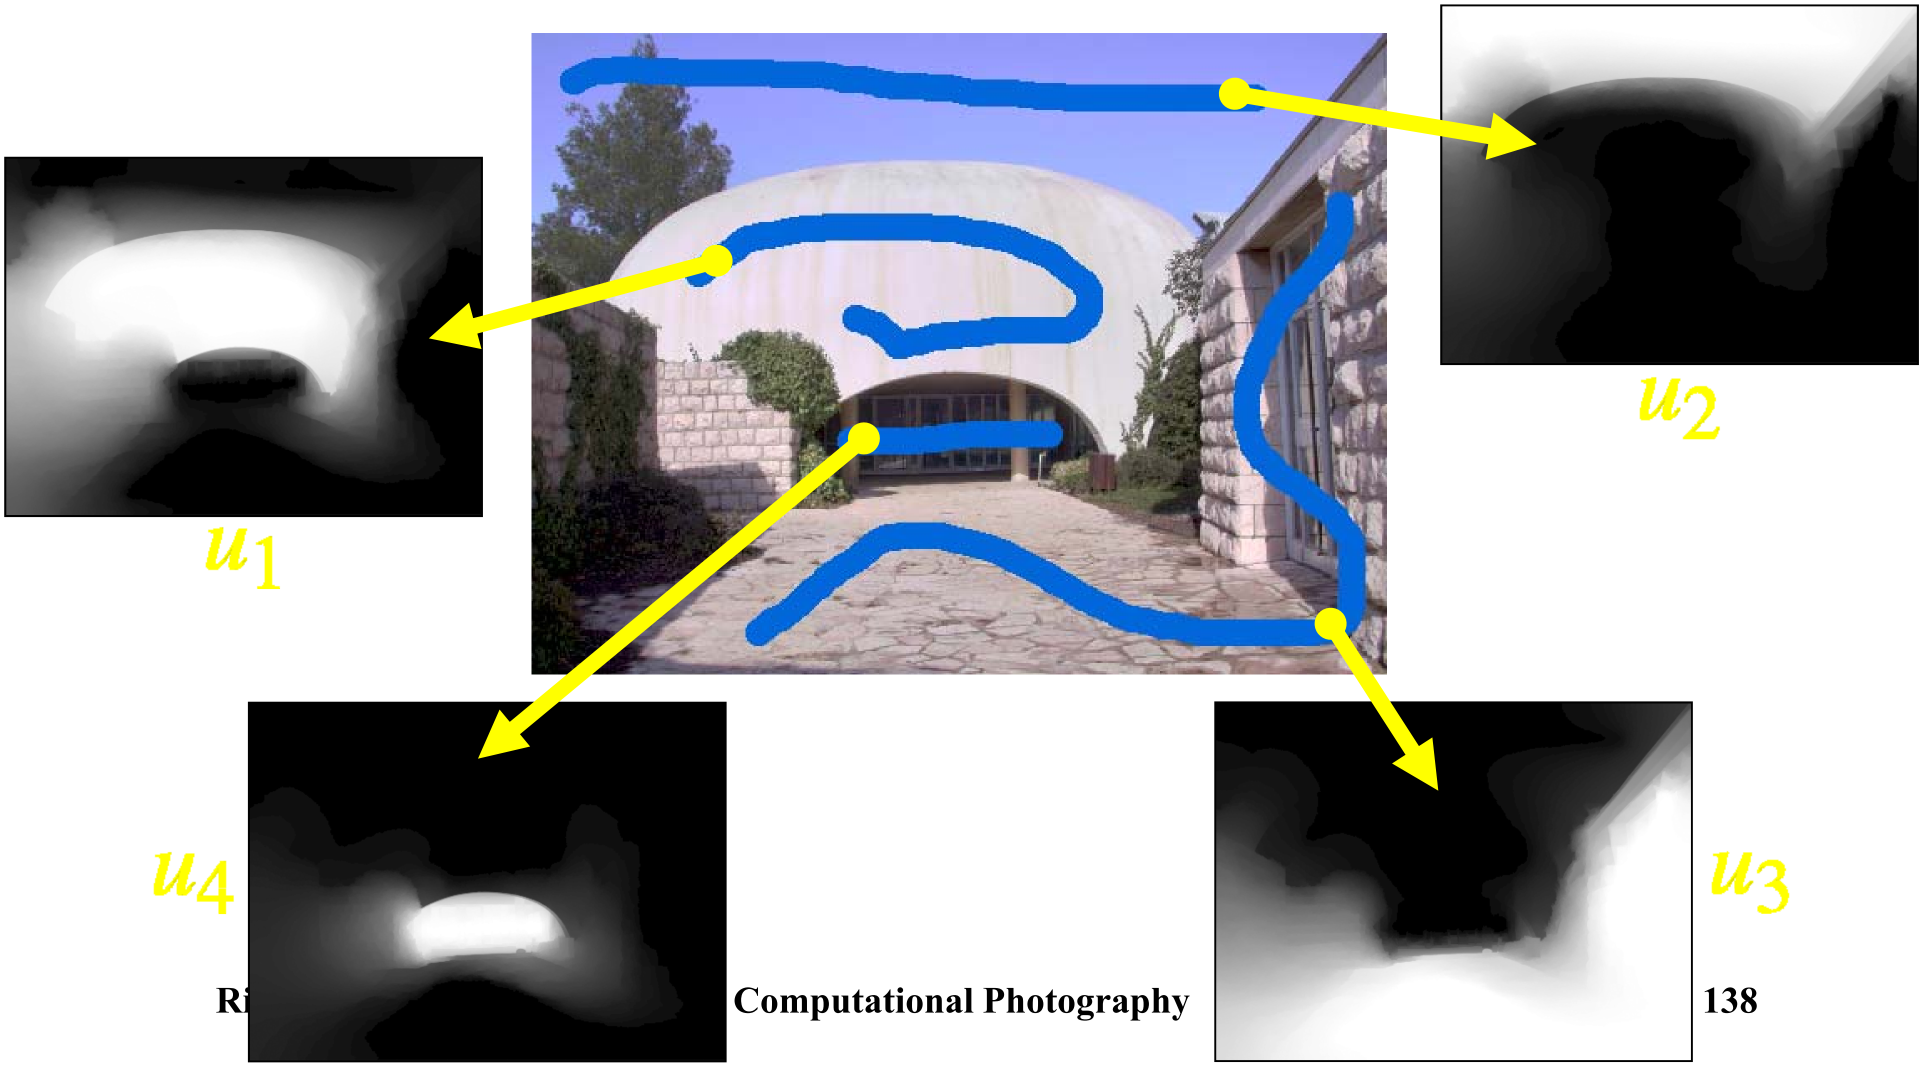

Influence Functions

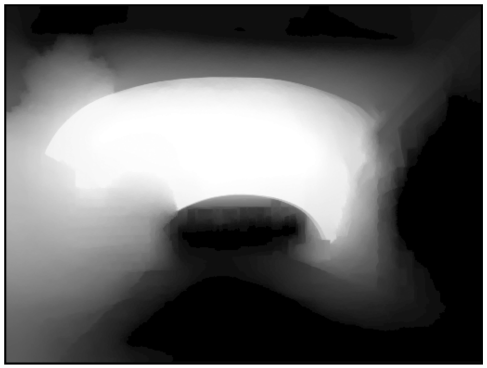

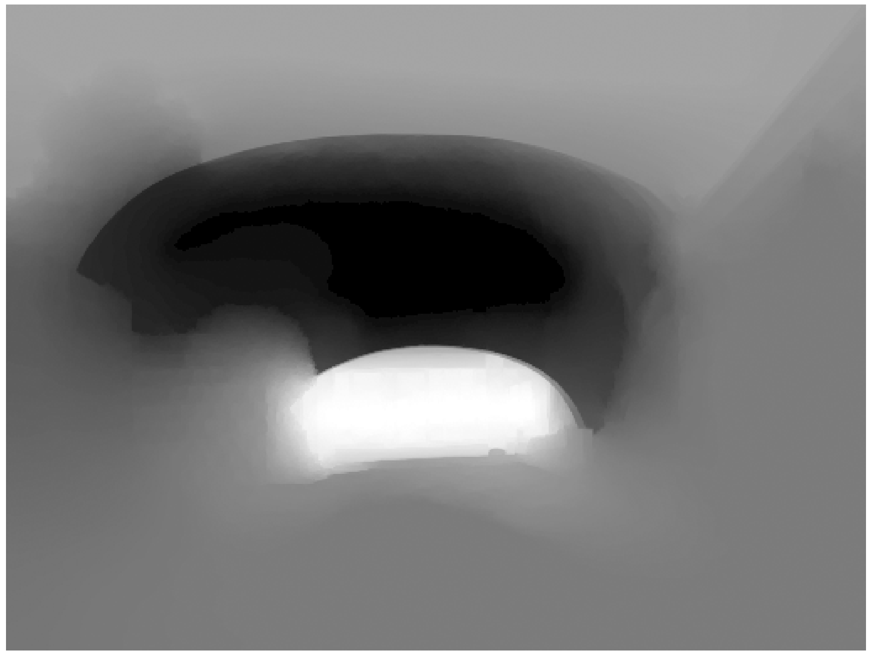

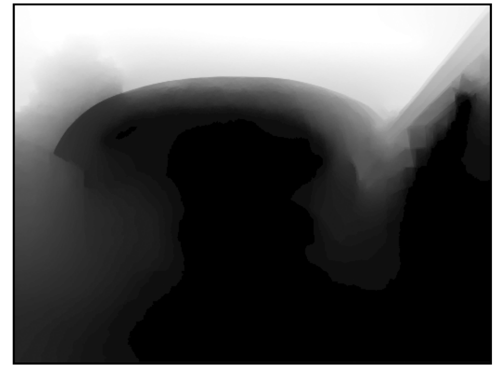

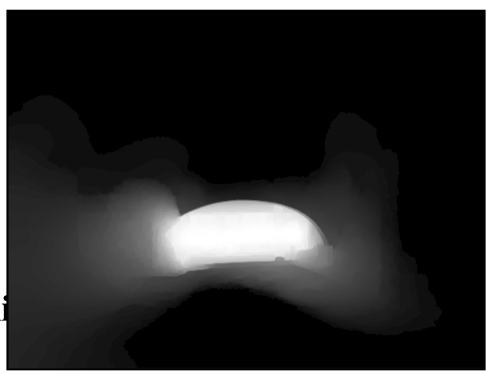

$$f = \sum_c g_c \, u_c$$

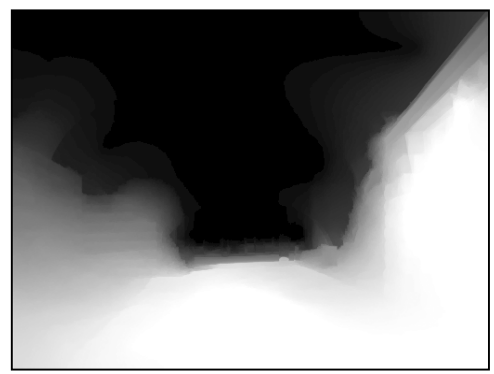

Automatic Initialization

Inspired by Ansel Adams' "Zone System".

- Segment image (very crudely) into brightness "zones"
- Determine the desired exposure for each zone
- Let the image-guided optimization produce a piecewise smooth exposure map

Results – Automatic mode

Results – Automatic mode

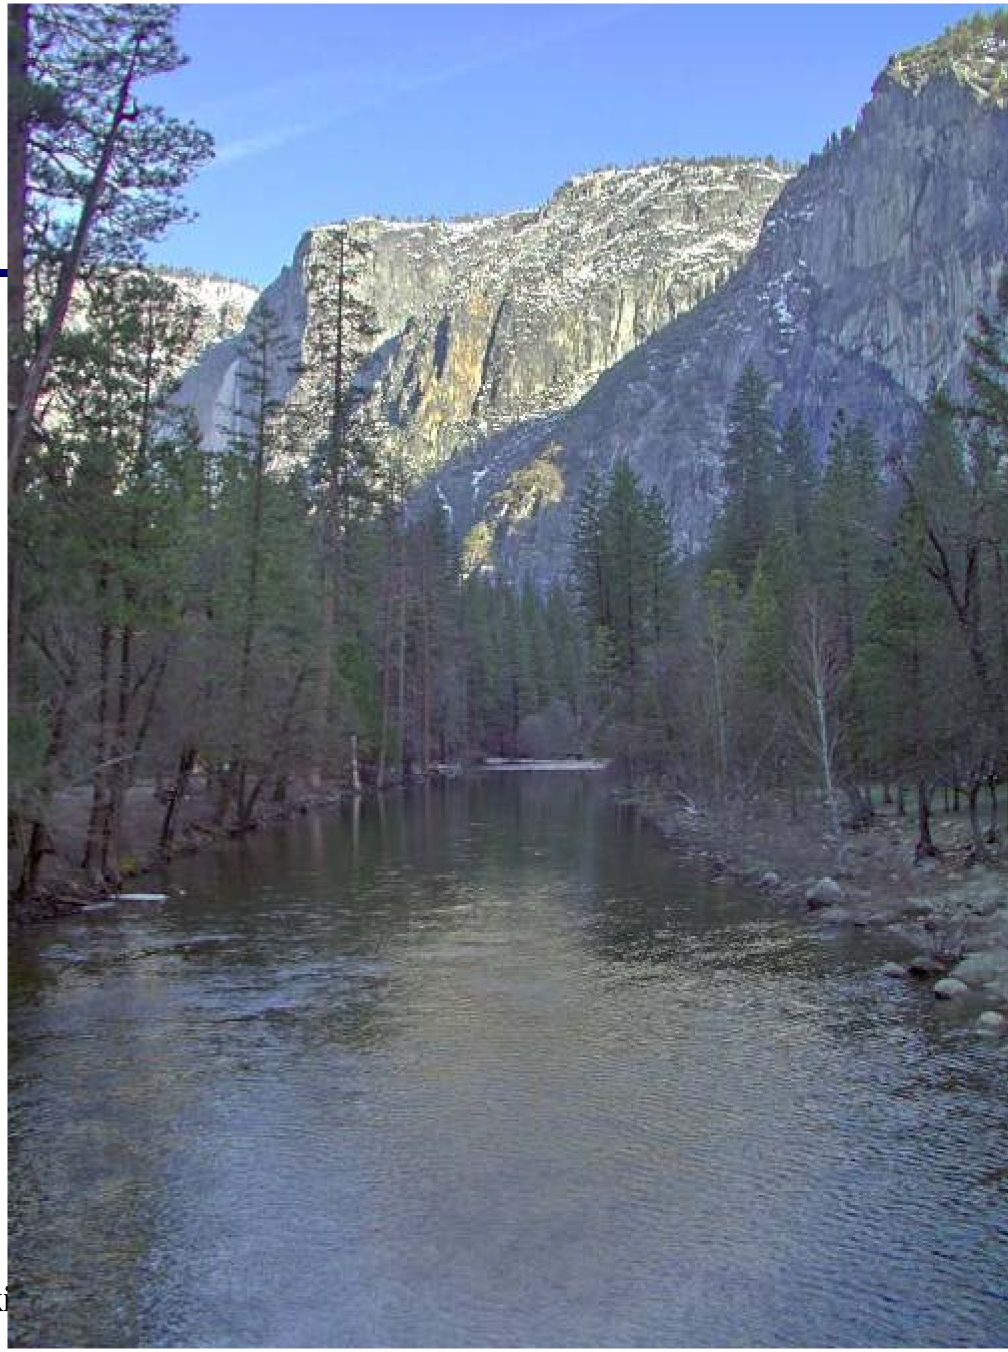

Snapshot Enhancement

Snapshot Enhancement

Spatially Variant White Balance

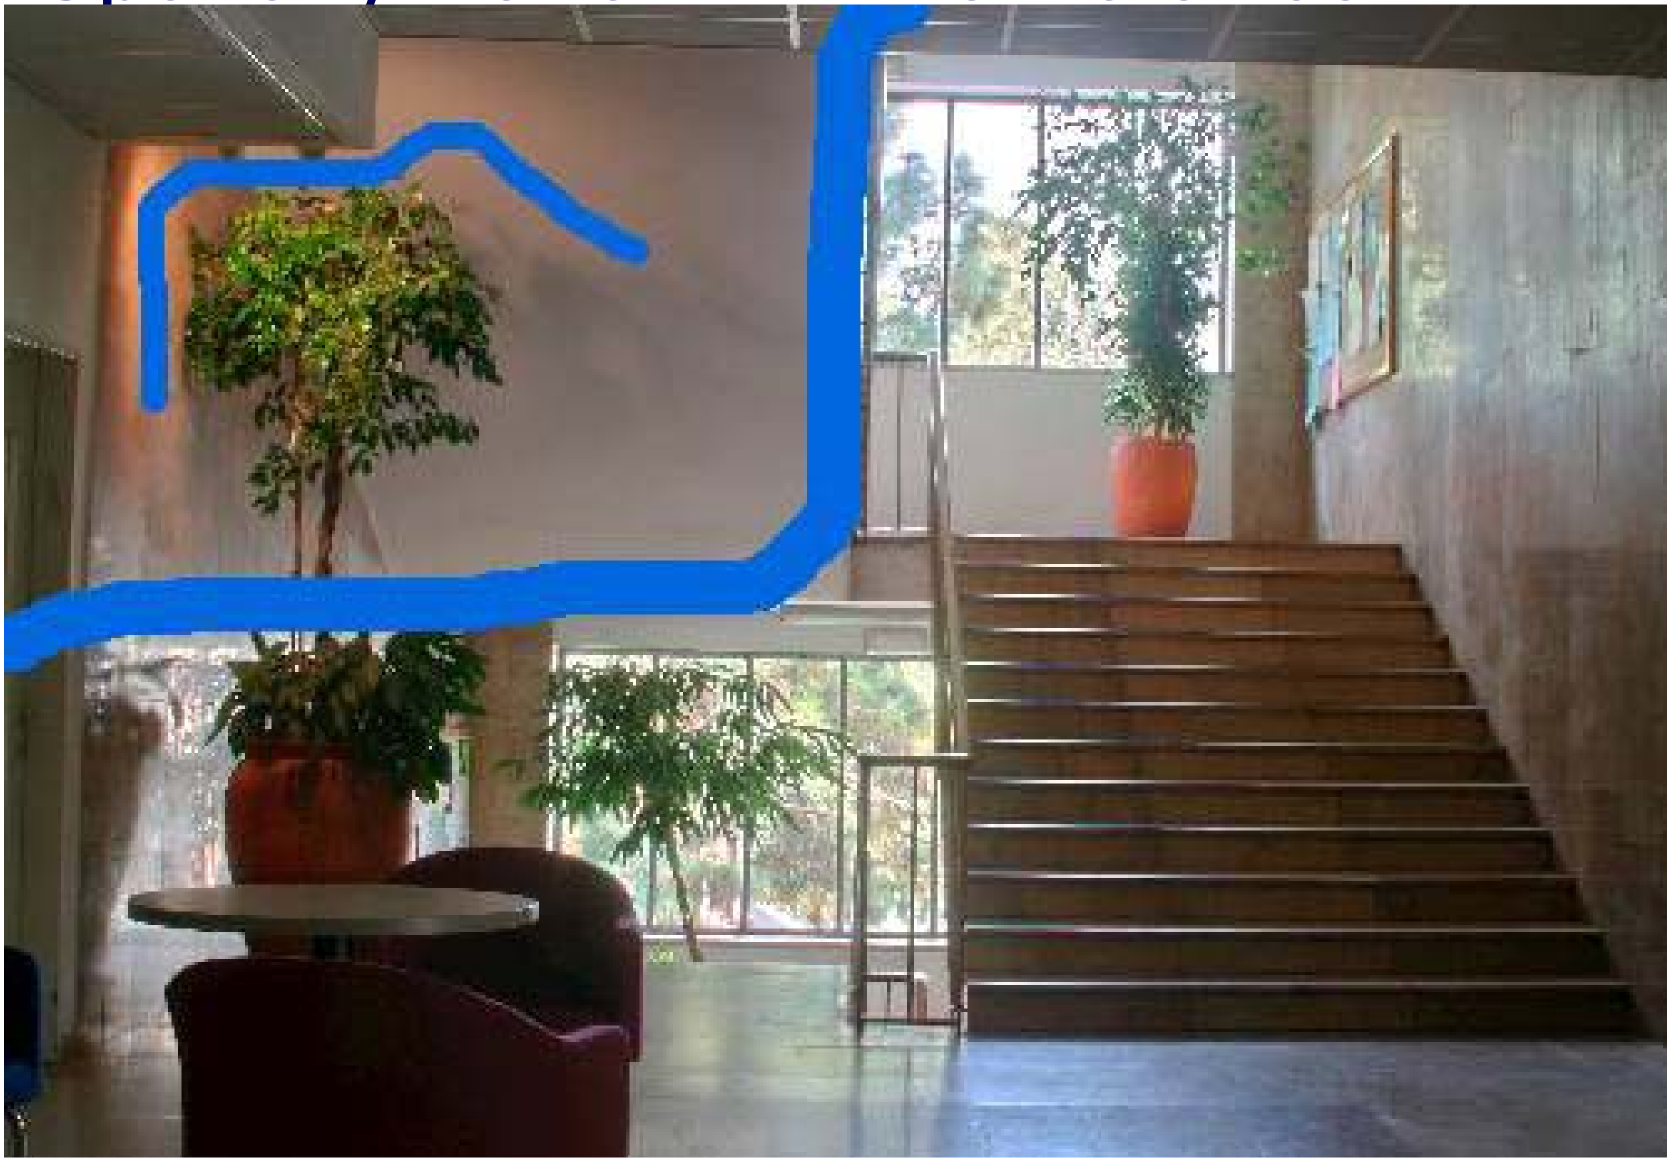

Comparison of tone mappers

Durand and Dorsey. *Fast bilateral filtering for the display of high-dynamic-range images*. SIGGRAPH 2002.

Fattal, Lischinski, and Werman. *Gradient domain high dynamic range compression*. SIGGRAPH 2002.

Li, Sharan, and Adelson. *Compressing and Companding High Dynamic Range Images with Subband Architectures*. SIGGRAPH 2005.

Li et al. 2005

Fattal et al. 2002

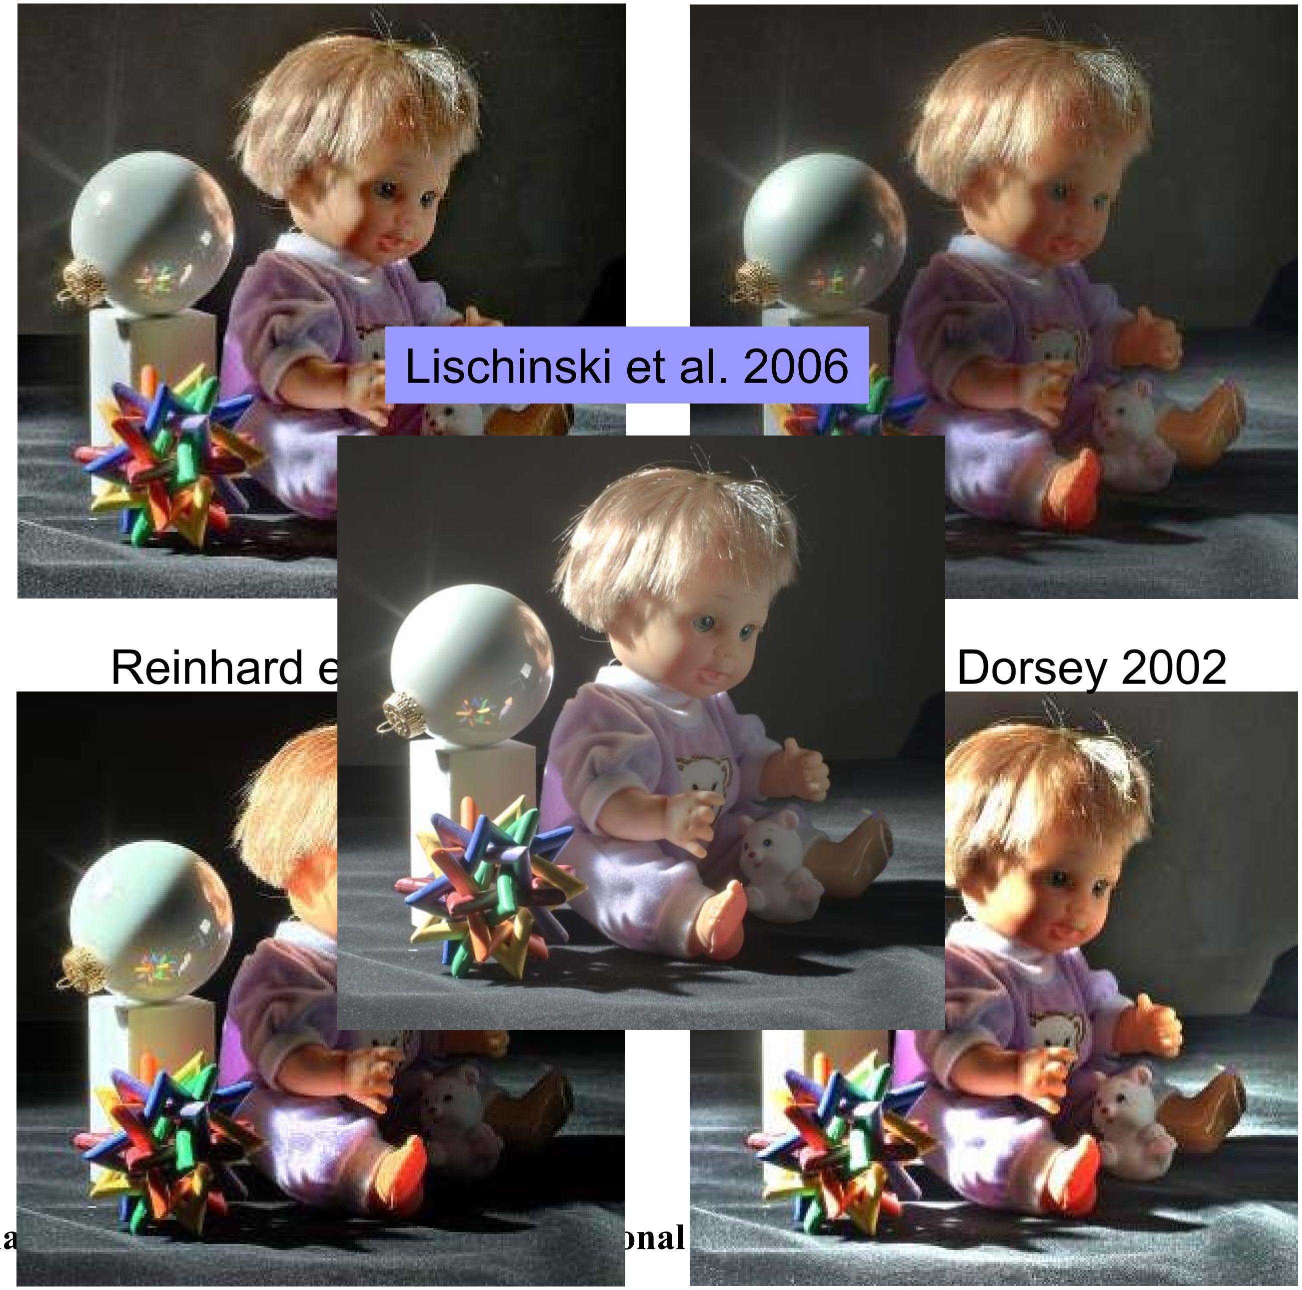

Reinhard e

Dorsey 2002

Merging flash and non-flash images

Georg Petschnigg, Maneesh Agrawala,
Hugues Hoppe, Rick Szeliski,
Michael Cohen, Kentaro Toyama

[SIGGRAPH'2004]

Flash + non-flash images

Flash photos have less noise, more detail

Non-flash photos have better color

Idea: merge them together

- But how?

non-flash

\+

flash

merged

Flash + non-flash images

Smooth non-flash photo using flash photo's edge information

Add high-frequency details from flash image

non-flash

+

flash

merged

Joint bilateral filter

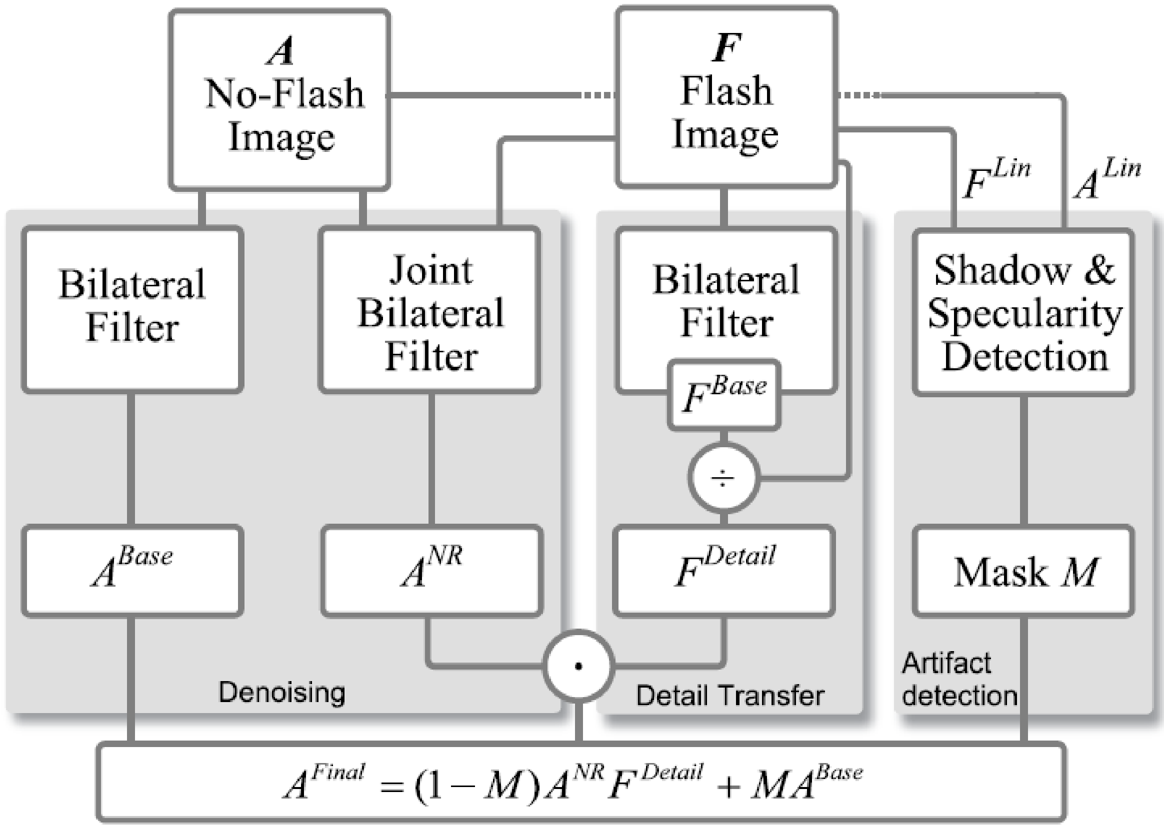

Figure 3: Overview of our algorithms for denoising, detail transfer, and flash artifact detection.

Bilateral detail filter

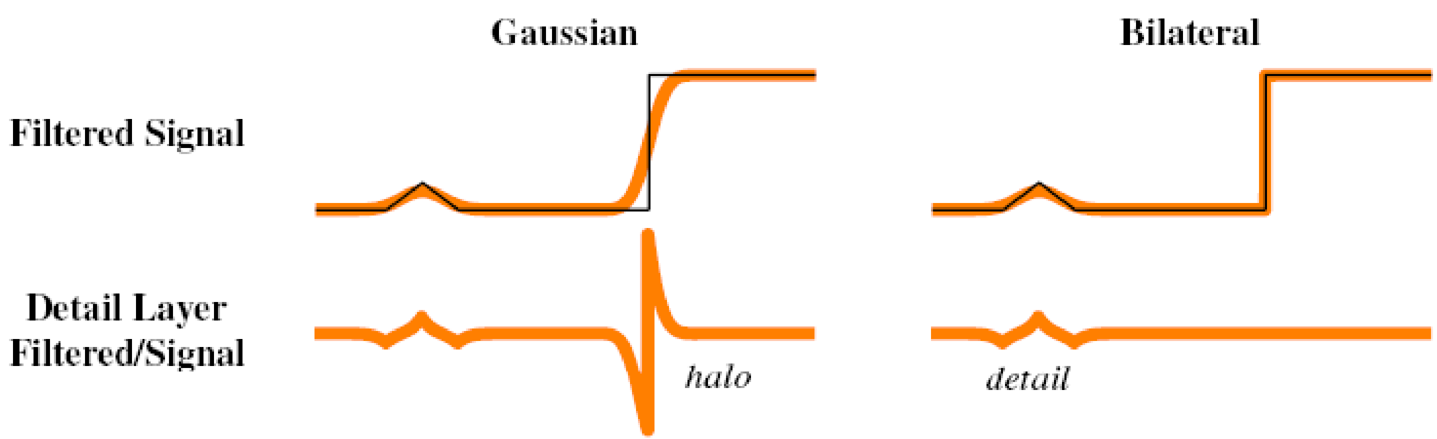

Figure 5: (left) A Gaussian low-pass filter blurs across all edges and will therefore create strong peaks and valleys in the detail image that cause halos. (right) The bilateral filter does not smooth across strong edges and thereby reduces halos, while still capturing detail.

Final result

Coda

Image Formation

Color

Filters

Pyramids

Local Features

Texture

Alignment

Flow

Stereo

SFM

Recognition Intro.

Topic Models

Recognition Kernels

Voting and Parts

Context

Articulated Recognition

Photometric Stereo

Tracking

MRFS

Segmentation

Comp. Photography