Color Analysis

Oct. 23. 2013

Rei Kawakami
(rei@cvl.iis.u-tokyo.ac.jp)

Color in computer vision

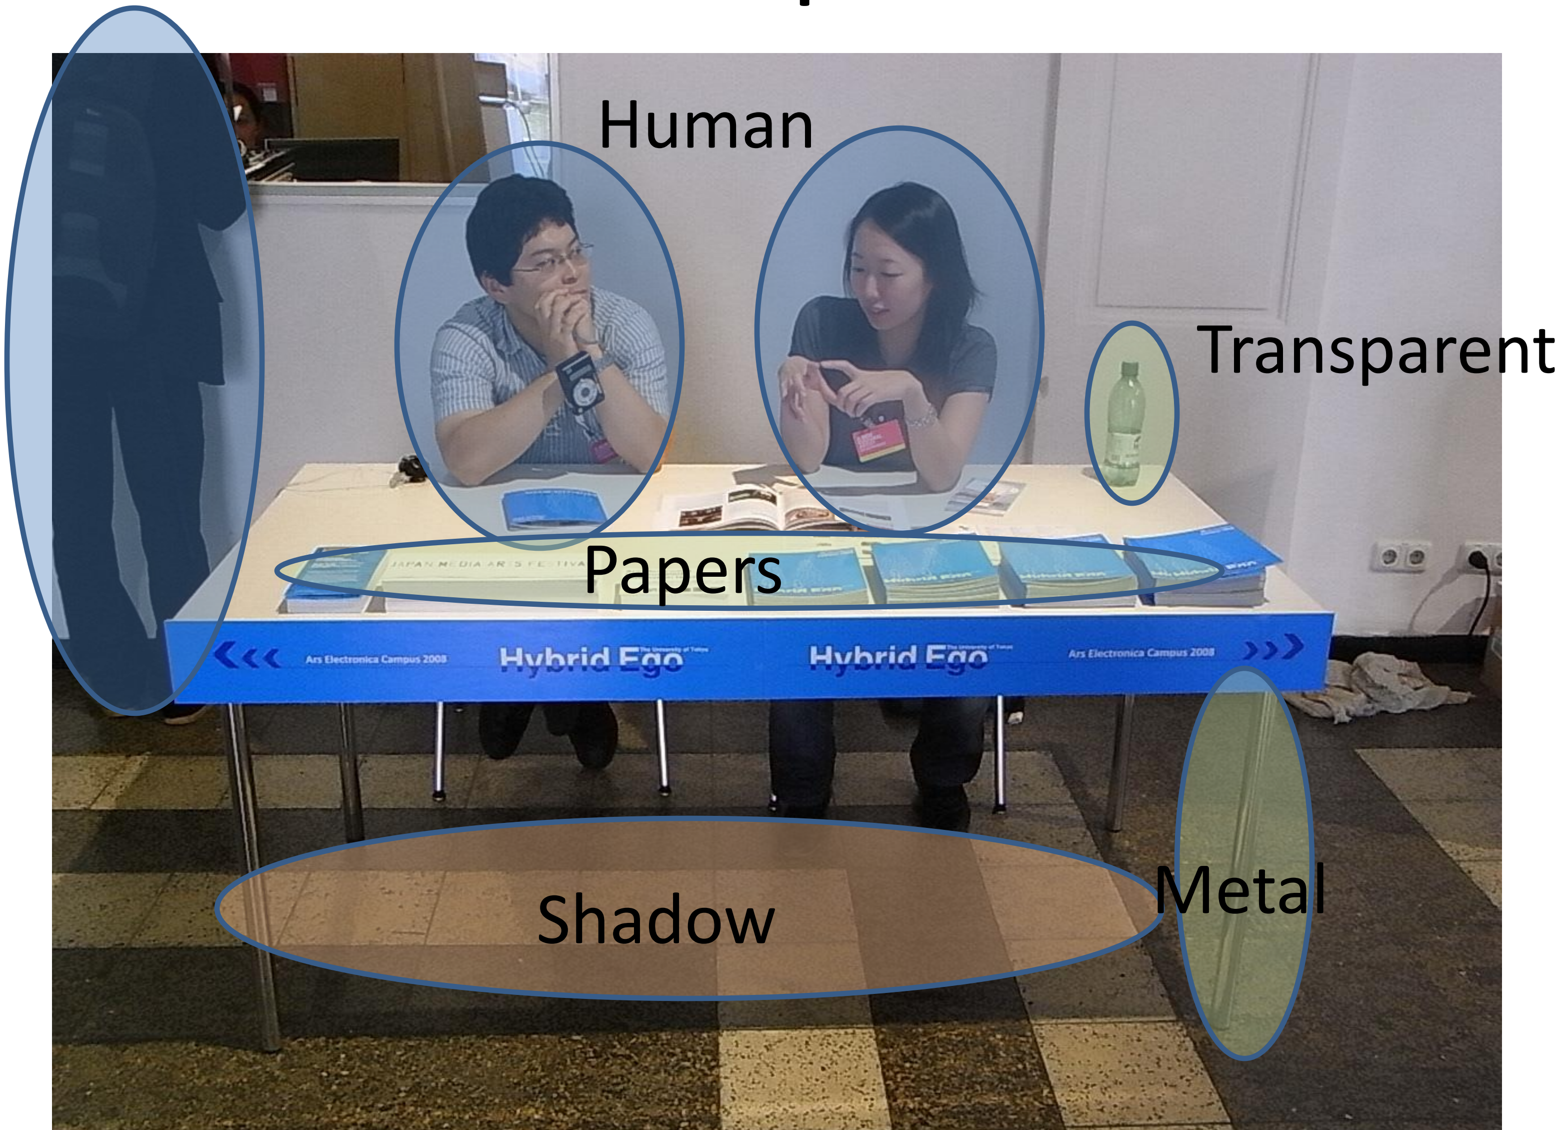

Major topics related to color analysis

- Image segmentation
- BRDF acquisition
- Radiometric camera calibration

Today's topic

- Intrinsic image decomposition
- Image similarity
- Color constancy

Next week's topic

- Photometric stereo, Multiplexed illumination, Image matting

Search them for further information

Oral: 4.8% (60), Overall: 28.0% (352)
Image segmentation papers: 23

Sylvain Paris, Fredo Durand (MIT)

International Conference on Computer Vision and Pattern Recognition (CVPR) 2007, Poster

A TOPOLOGICAL APPROACH TO HIERARCHICAL SEGMENTATION USING MEAN SHIFT

Image segmentation

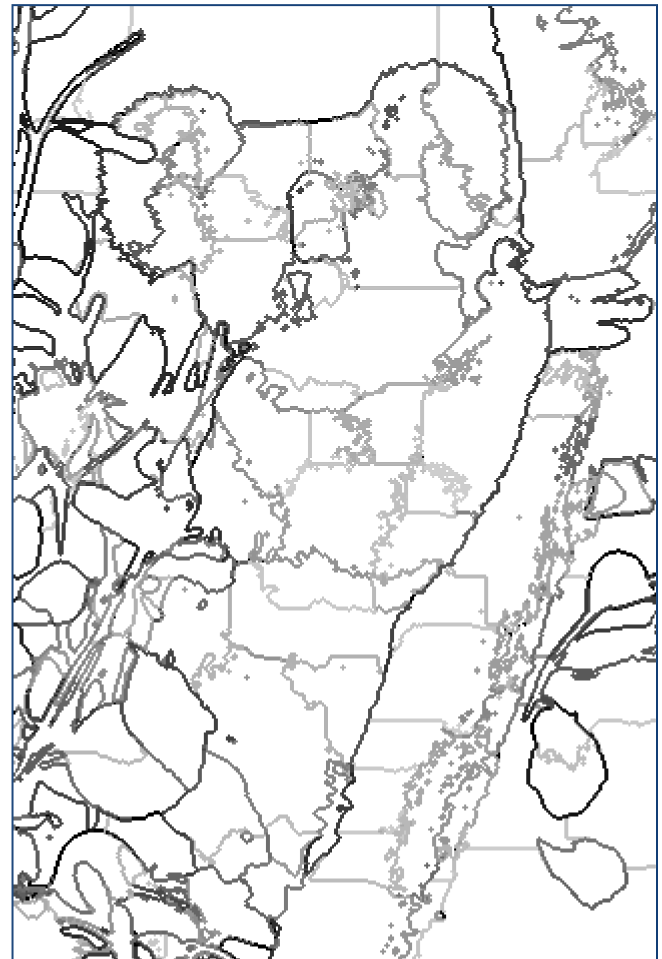

Related work

Minimum distance classifier

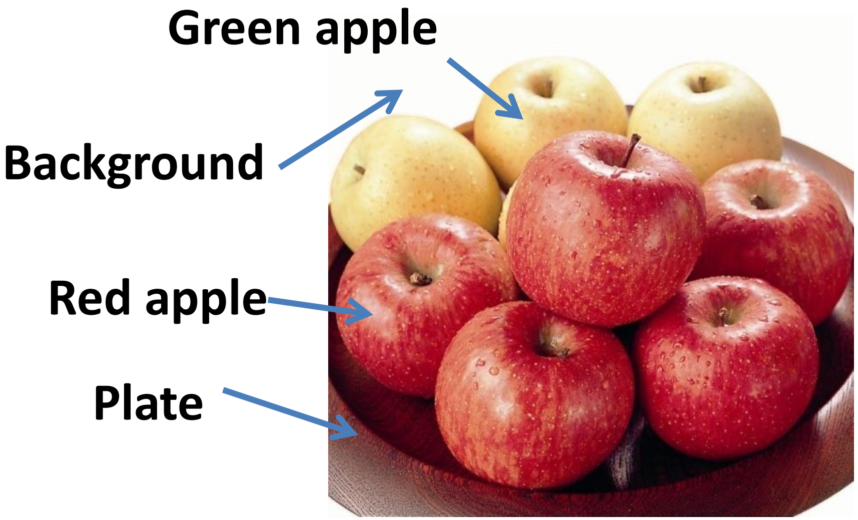

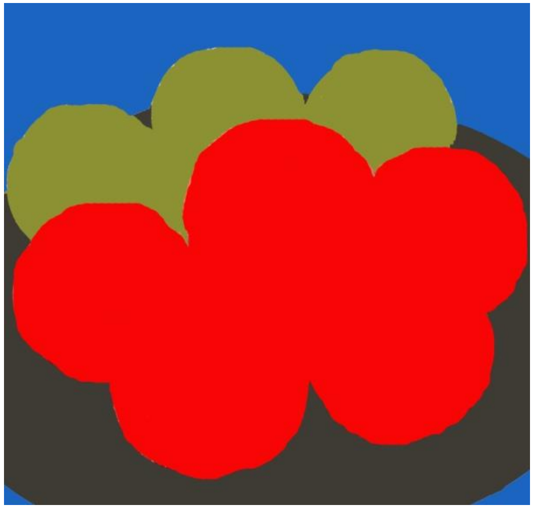

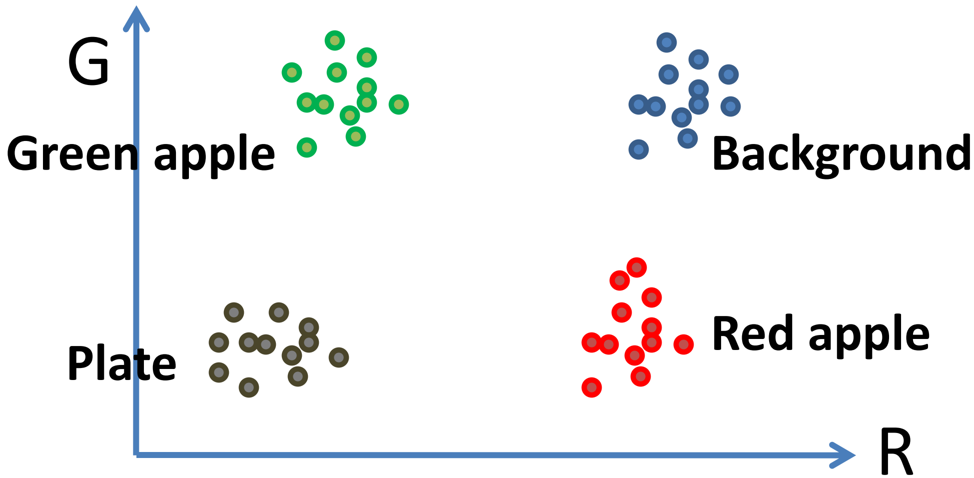

Minimum distance classifier

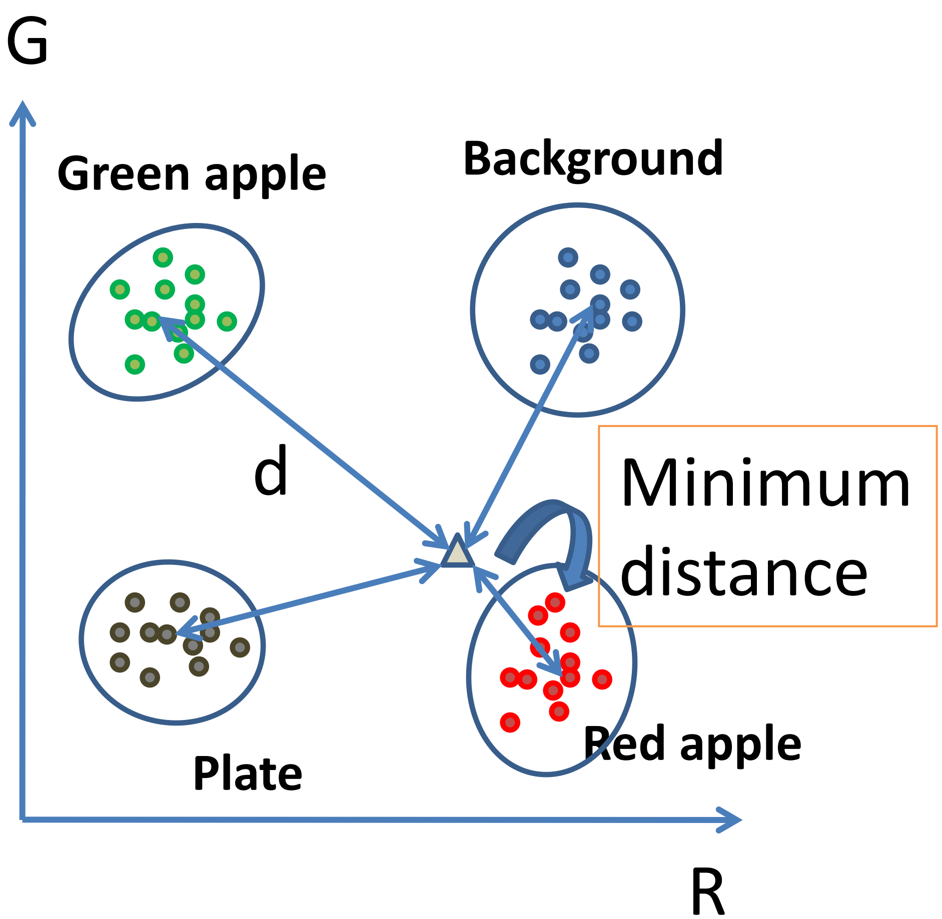

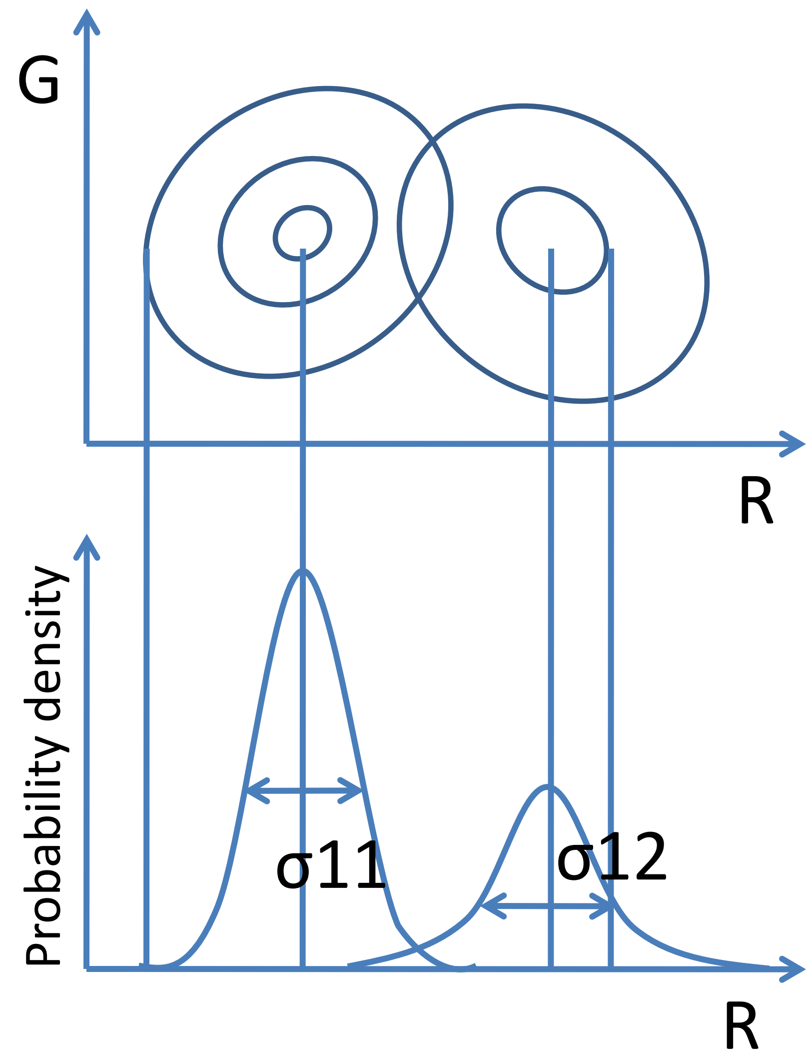

K-means clustering

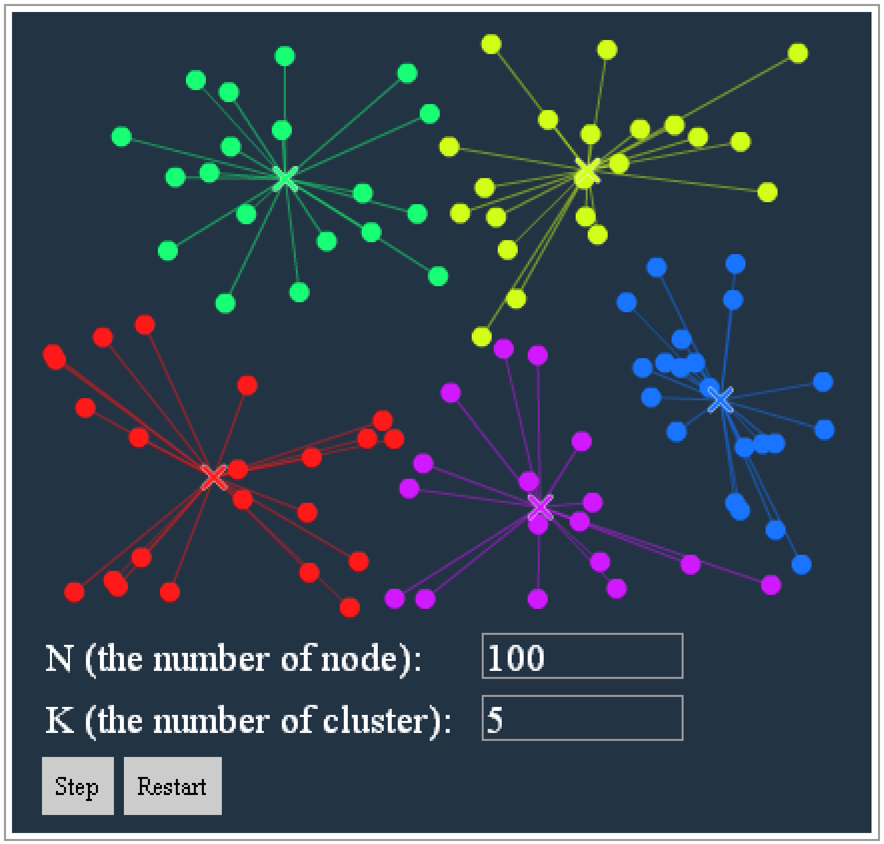

- Label a class randomly for each point
- Calculate the center
- Re-label the class to the nearest one for each pixel

http://d.hatena.ne.jp/nitoyon/20090409/kmeans_visualise

Region growing

adjacent pixel: similar feature vector → same region

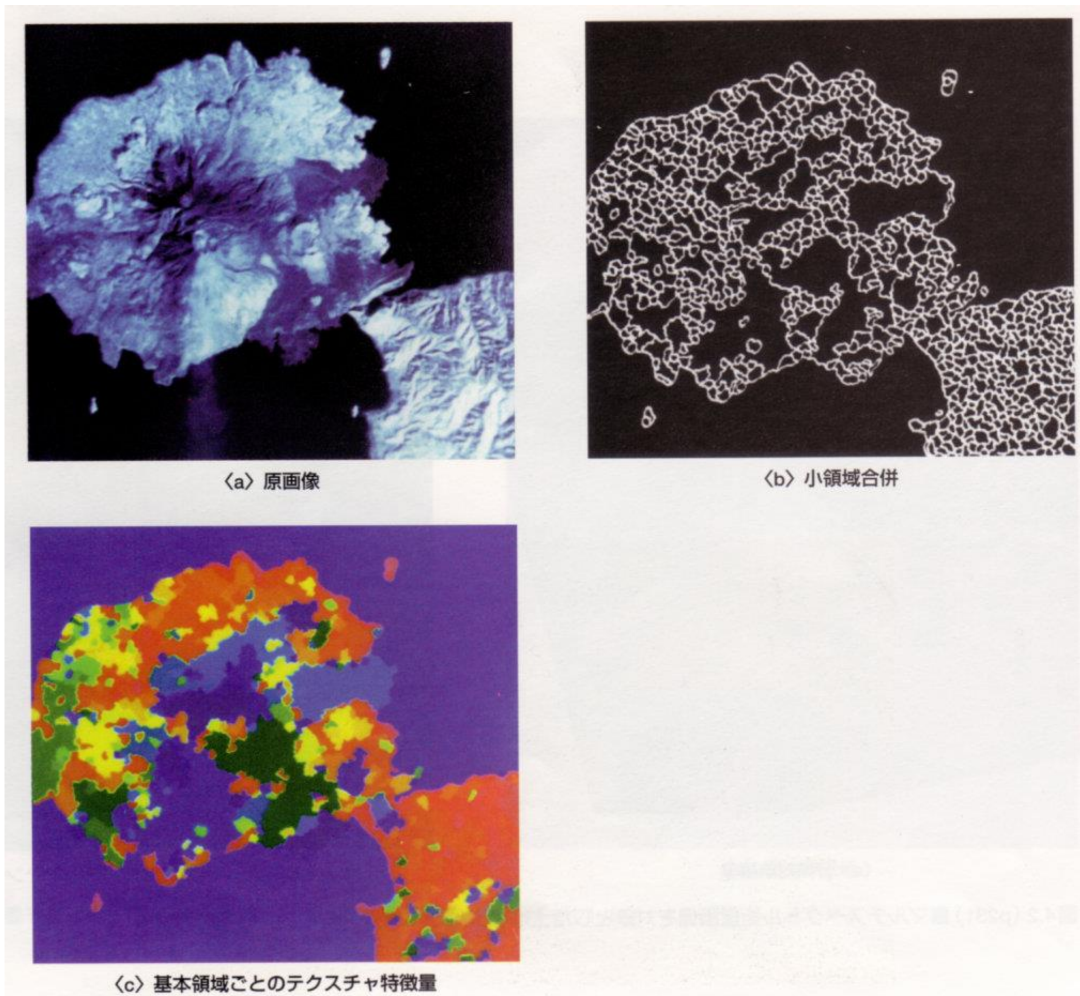

Lazy snapping (Graph-cut)

(a) Girl (4/2/12) (b) Ballet (4/7/14) (c) Boy (6/2/13)

(c) Grandpa (4/2/11) (d) Twins (4/4/12)

Included in Microsoft Expression

Graph cut (Min-Cut/Max-Flow): Concept

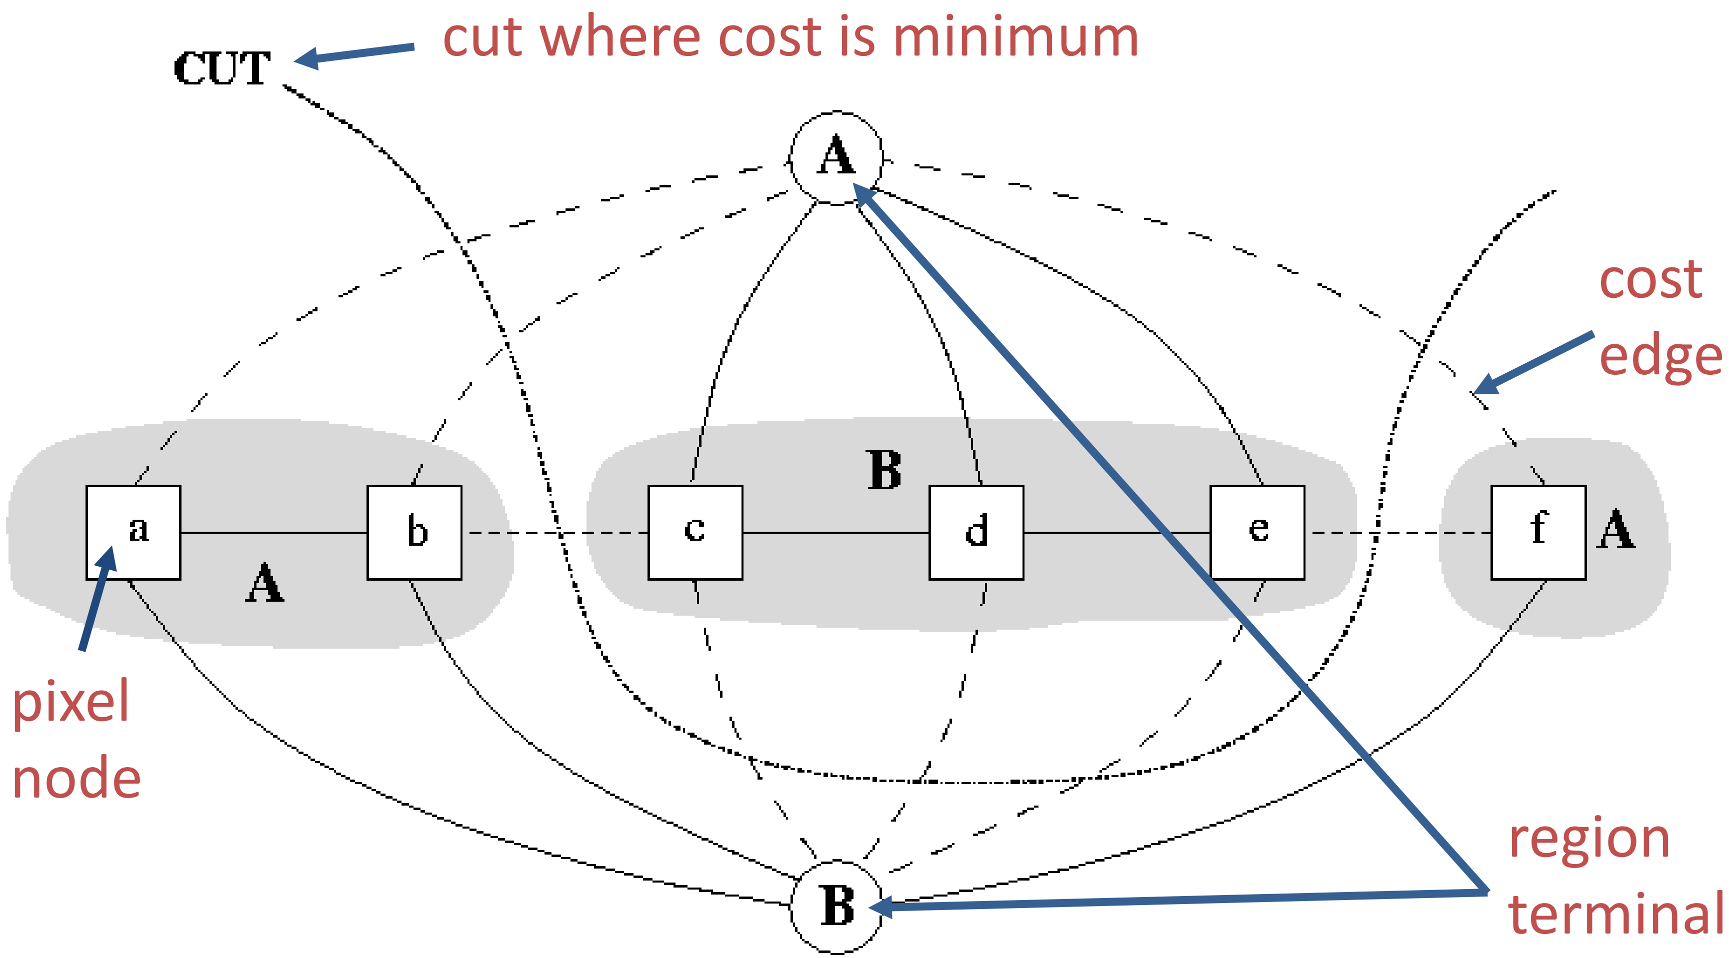

Graph cut (Min-Cut/Max-Flow): Cost function

$$E(X) = \sum_{i \in v} E_1(x_i) + \lambda \sum_{(i,j) \in \varepsilon} E_2(x_i, x_j)$$

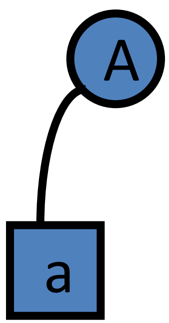

color similarity
between
pixel and region

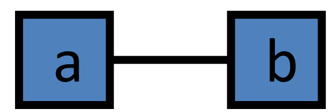

color similarity
between
adjacent pixel

Probabilistic (top-down) approach

- Use of priors (combined with recognition)
  - TextonBoost (Texture cue)

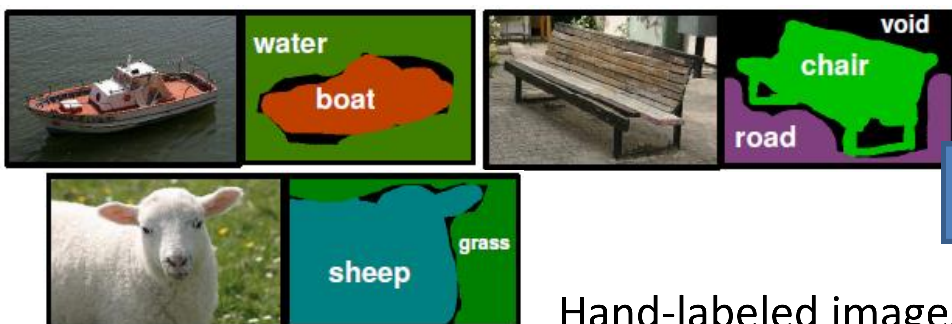

Hand-labeled images

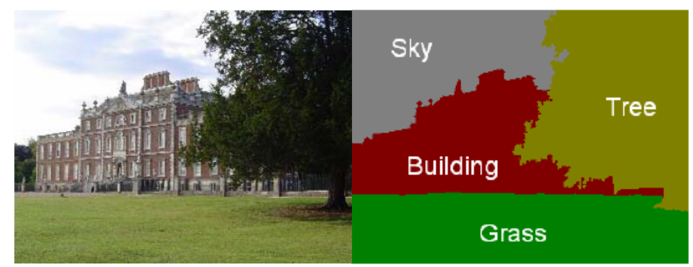

Segmentation result

Advantage of mean shift

- No priors nor human operation are required
  - Unsupervised segmentation as a pre-processing


Input image


Mean-shift A

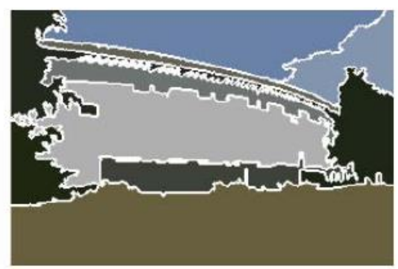
Mean-shift B

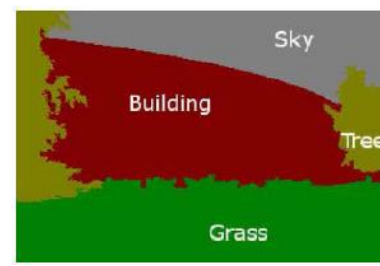

Hand labeled

Advantage of the proposed method

- Fast computation
  - Gaussian mean-shift (time-consuming)
  - Do not sacrifice accuracy for speed

- Hierarchical segmentation
  - Morse theory
  - Topological decomposition

Method

Mean shift segmentation

$$\mathbf{y}_{j+1} = \frac{\sum_{i=1}^{n} K(\mathbf{y}_j - \mathbf{x}_i)\,\mathbf{x}_i}{\sum_{i=1}^{n} K(\mathbf{y}_j - \mathbf{x}_i)}$$

K : Kernel function

$\{\mathbf{x}_i\}$: Feature points (pixels)

$\{\mathbf{y}_j\}$: Series

$\mathbf{y}_0$: Seed

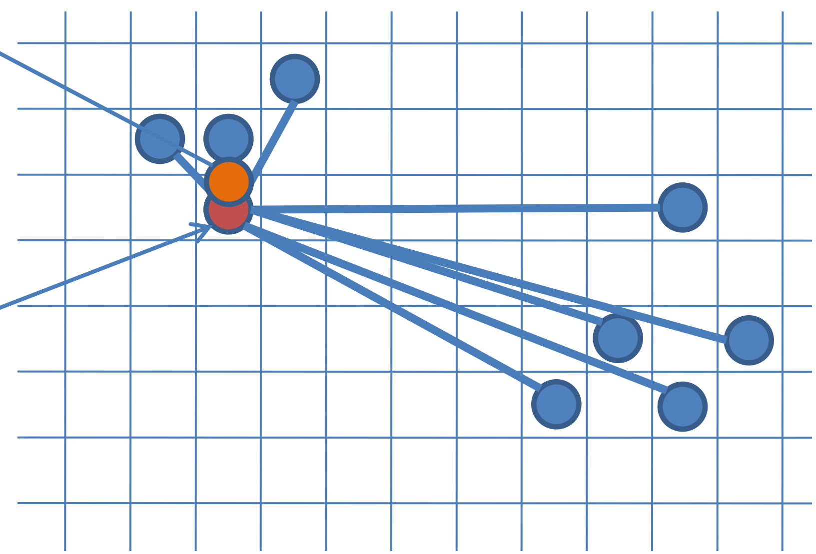

Intuitive Description

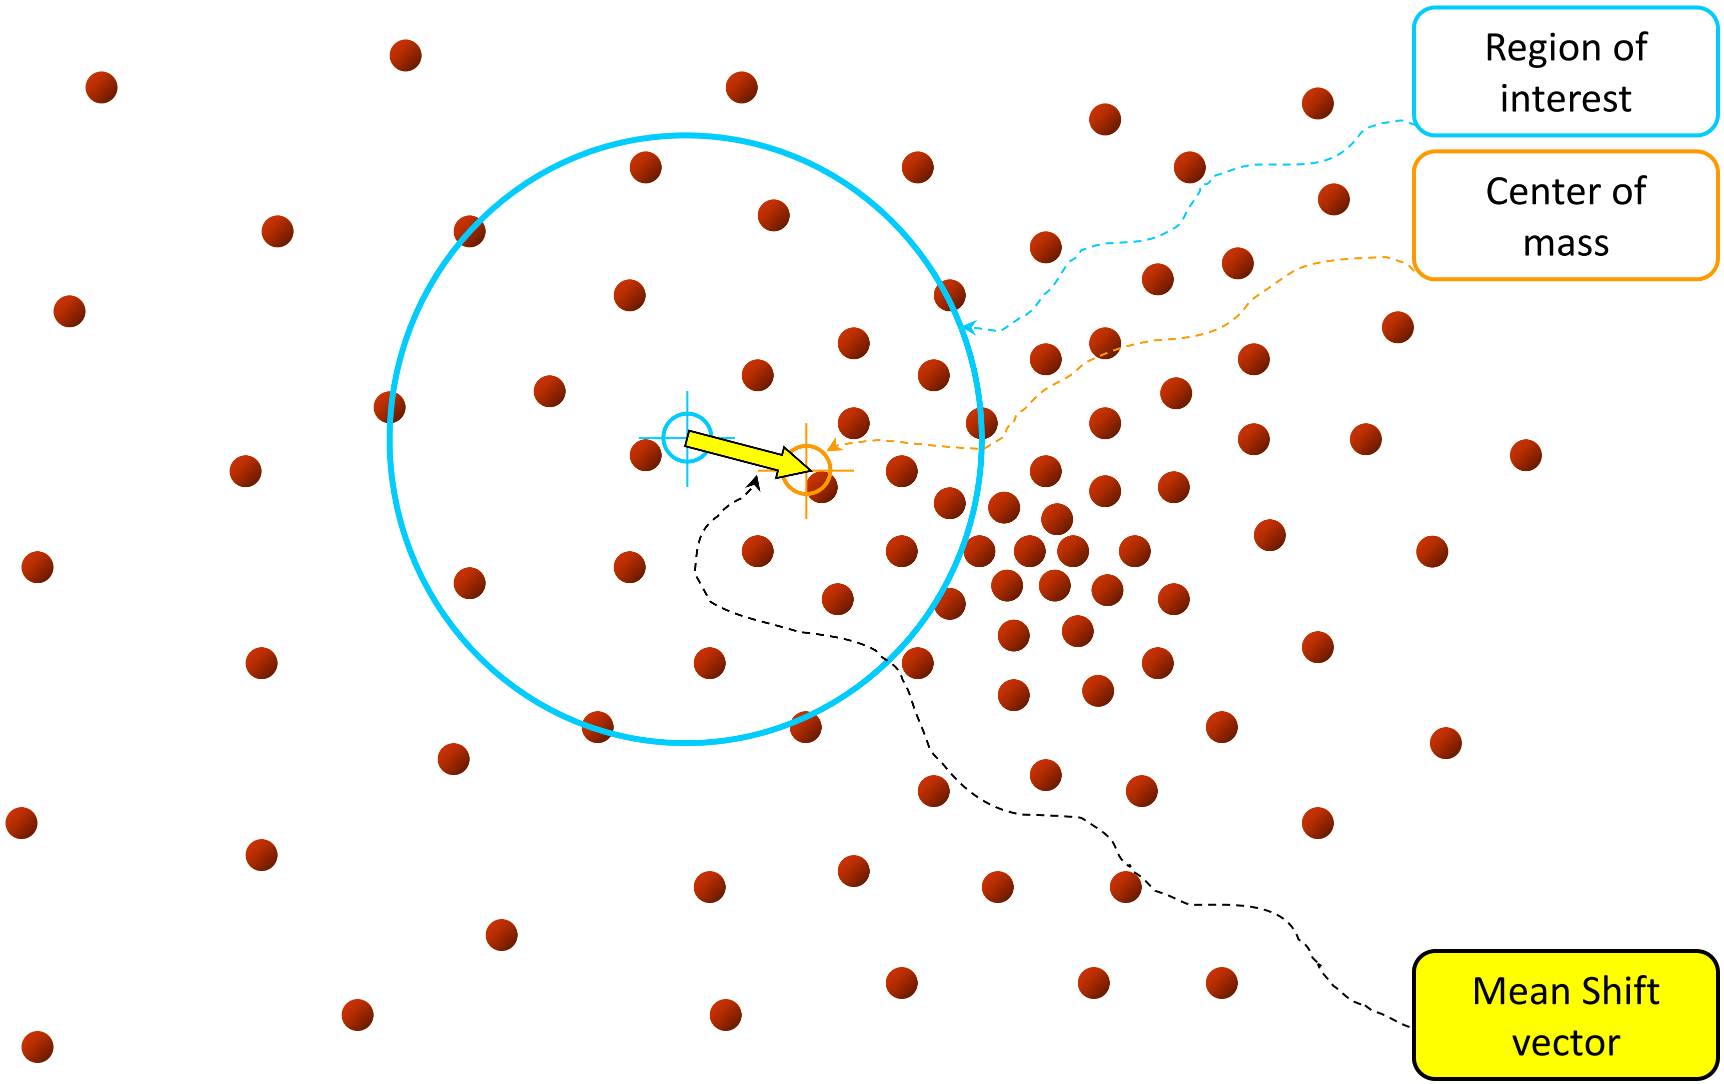

Objective : Find the densest region

Distribution of identical billiard balls

Intuitive Description

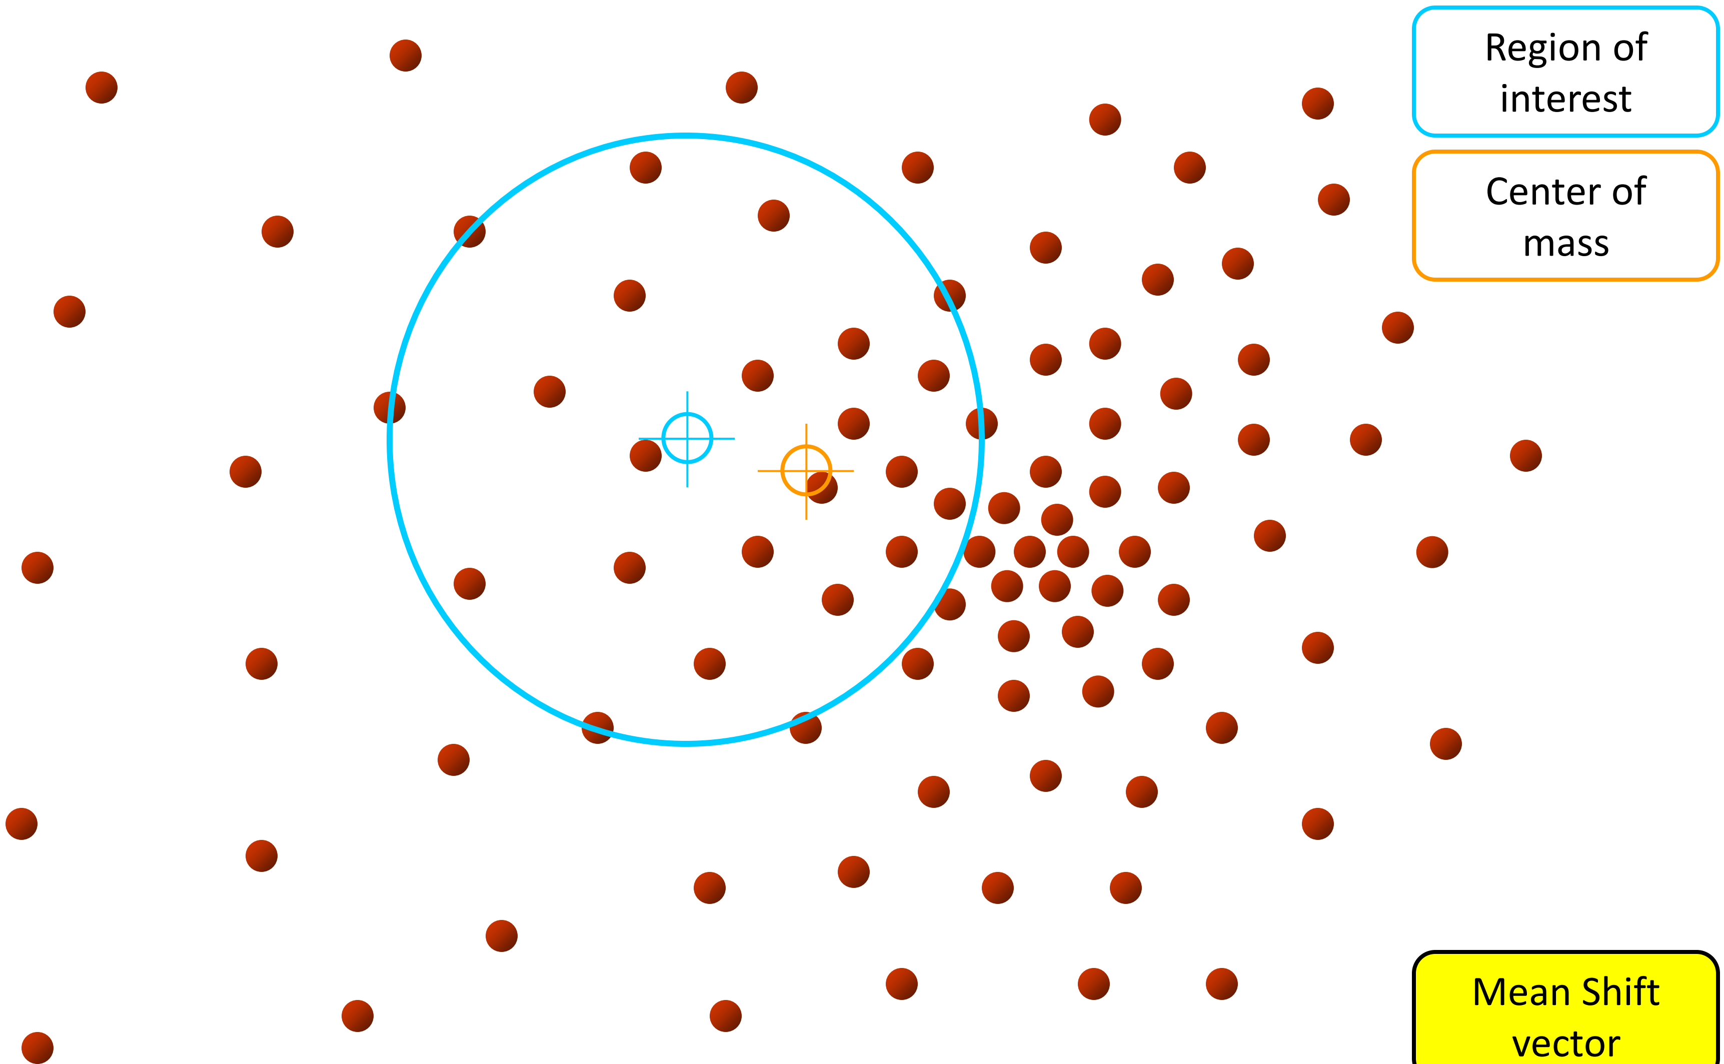

Objective : **Find the densest region**

Distribution of identical billiard balls

Intuitive Description

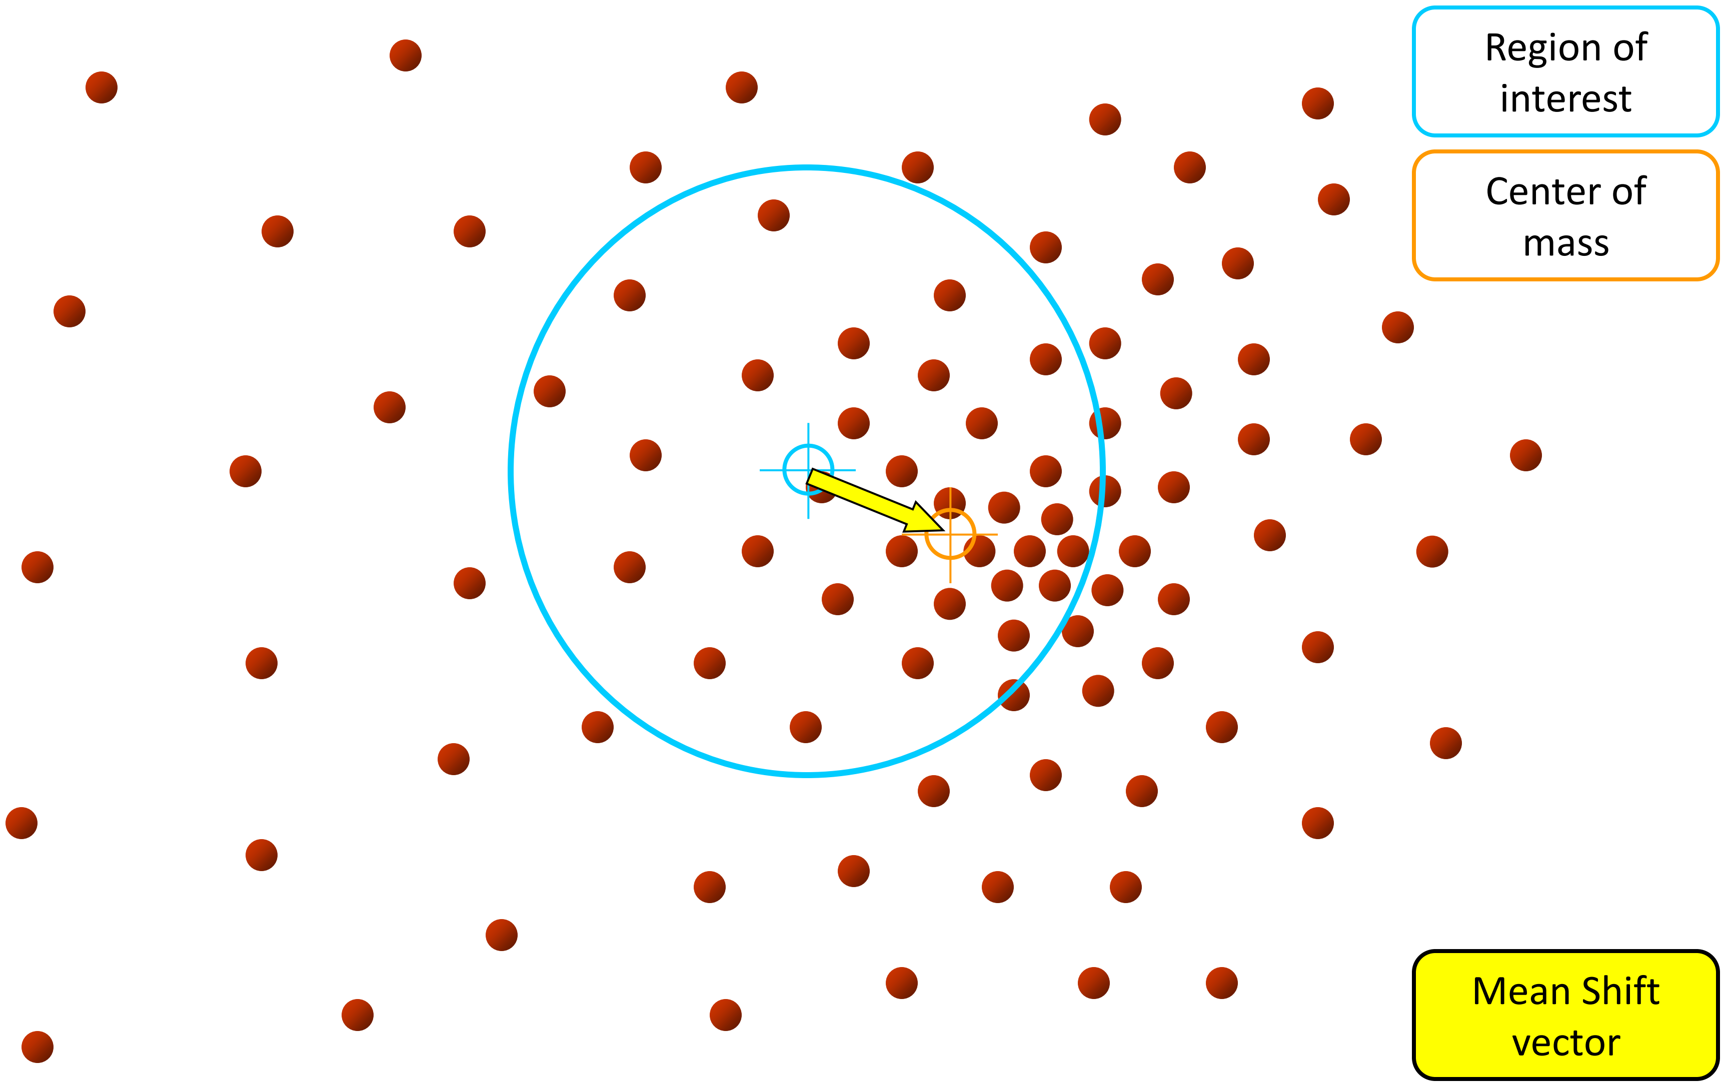

Objective : Find the densest region

Distribution of identical billiard balls

Intuitive Description

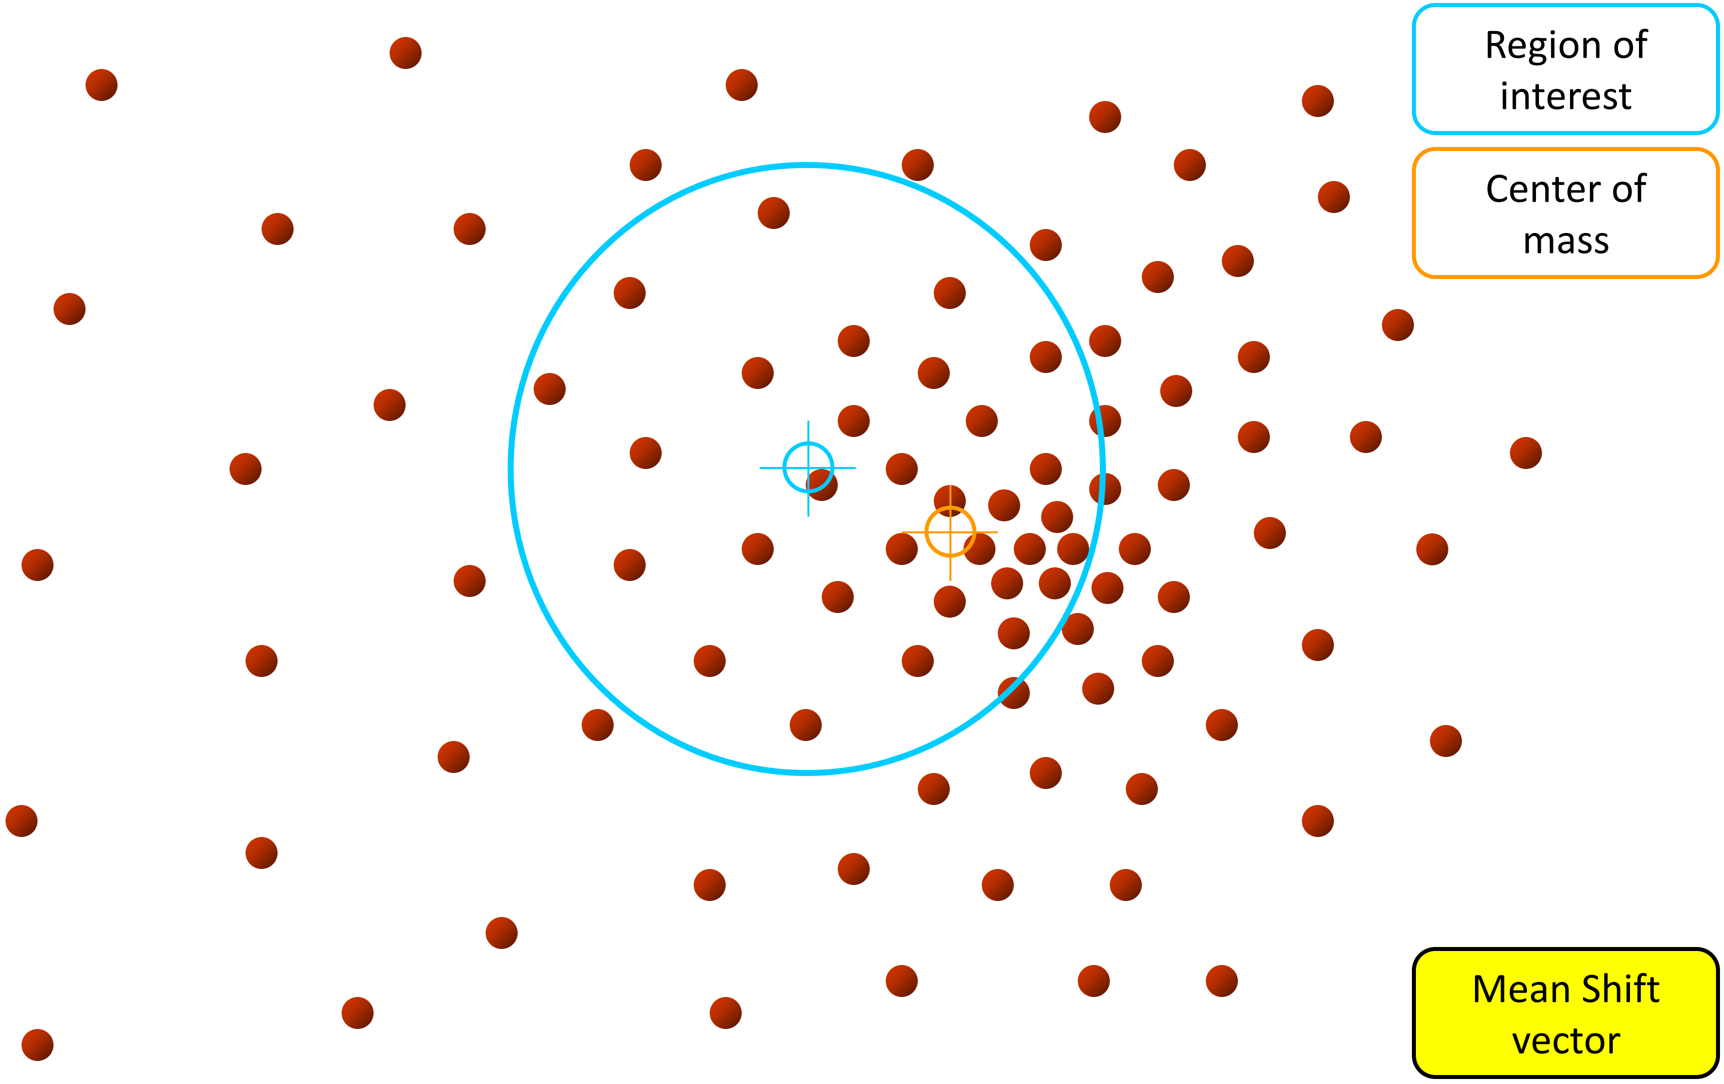

Objective : Find the densest region

Distribution of identical billiard balls

Intuitive Description

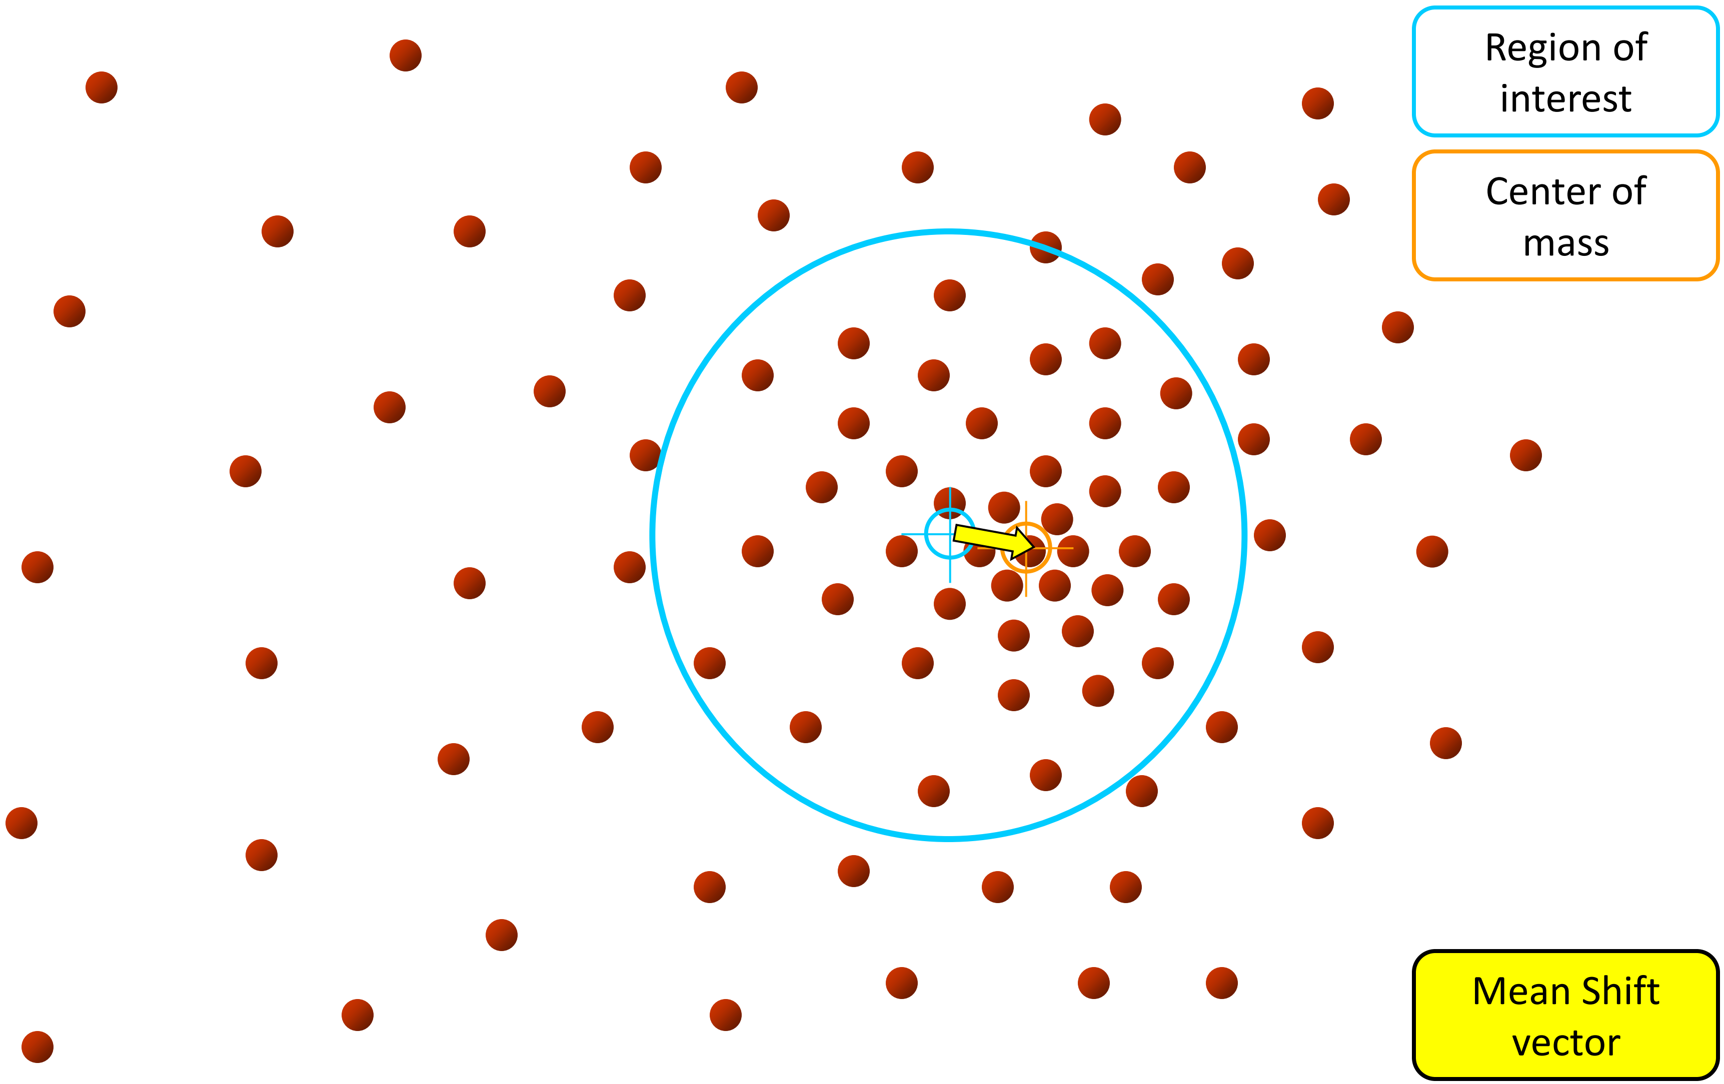

Objective : Find the densest region

Distribution of identical billiard balls

Intuitive Description

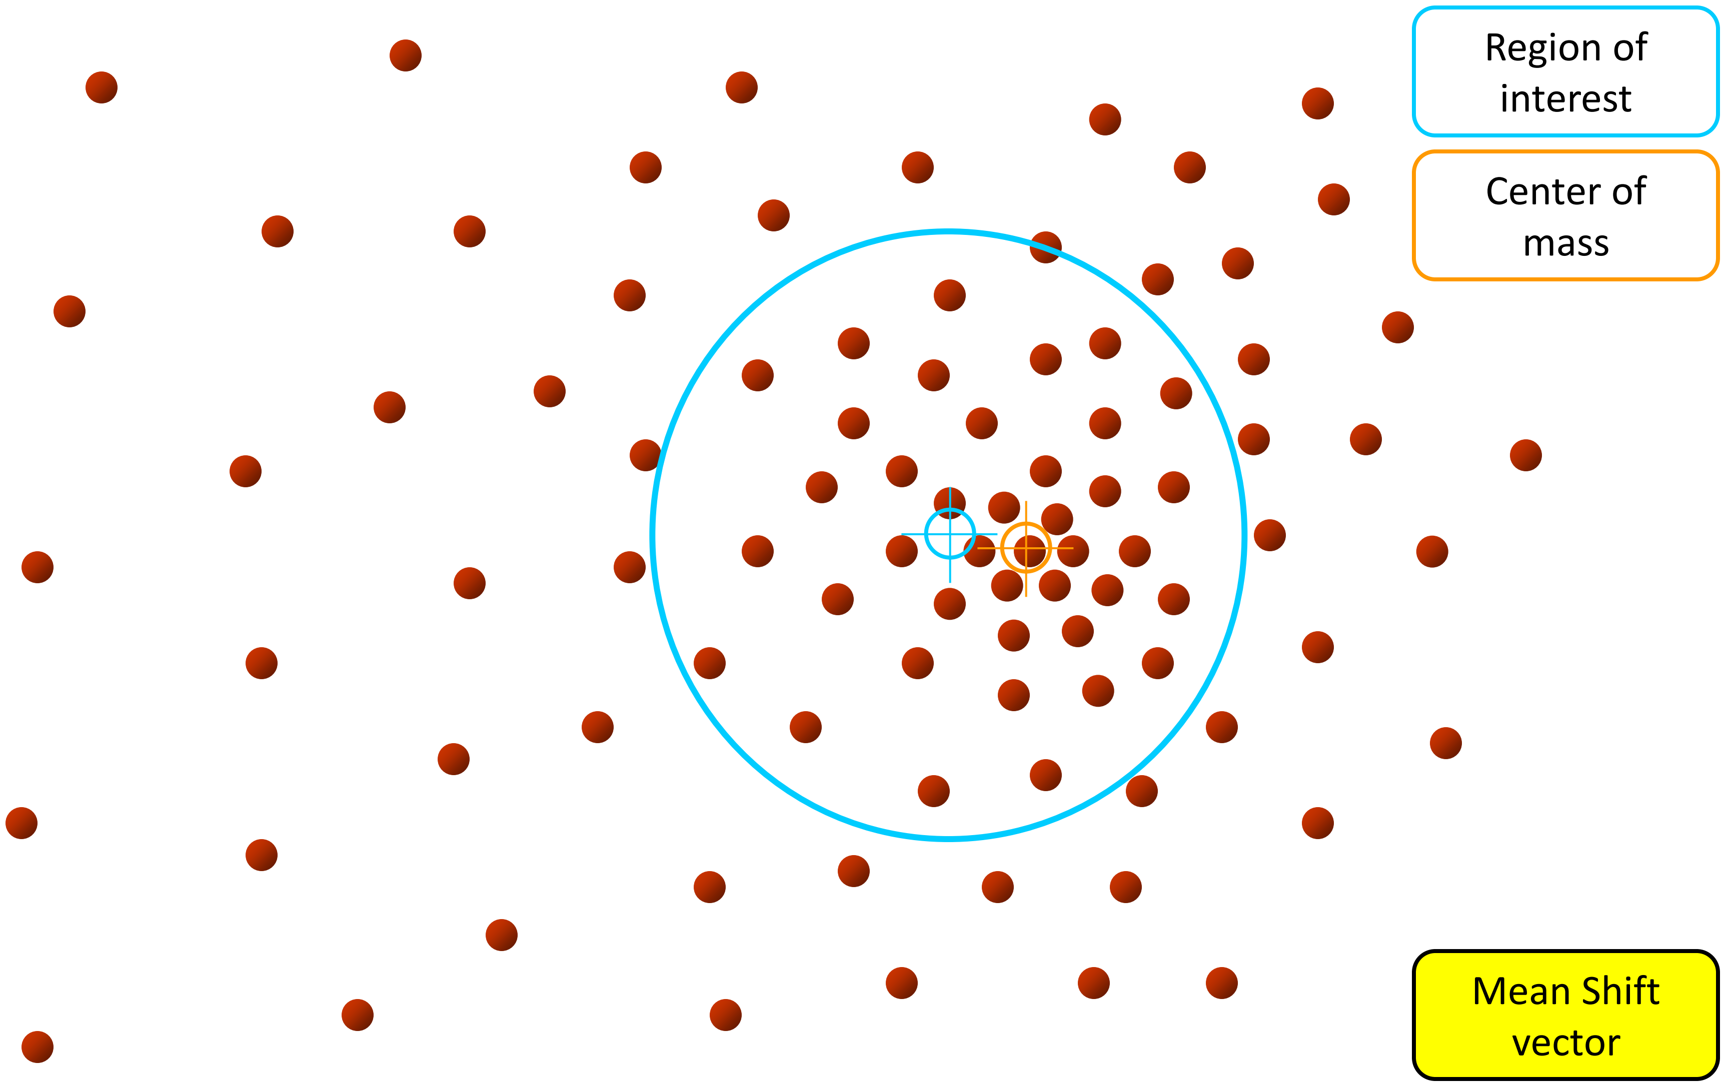

<u>Objective </u>: **Find the densest region**

Distribution of identical billiard balls

Intuitive Description

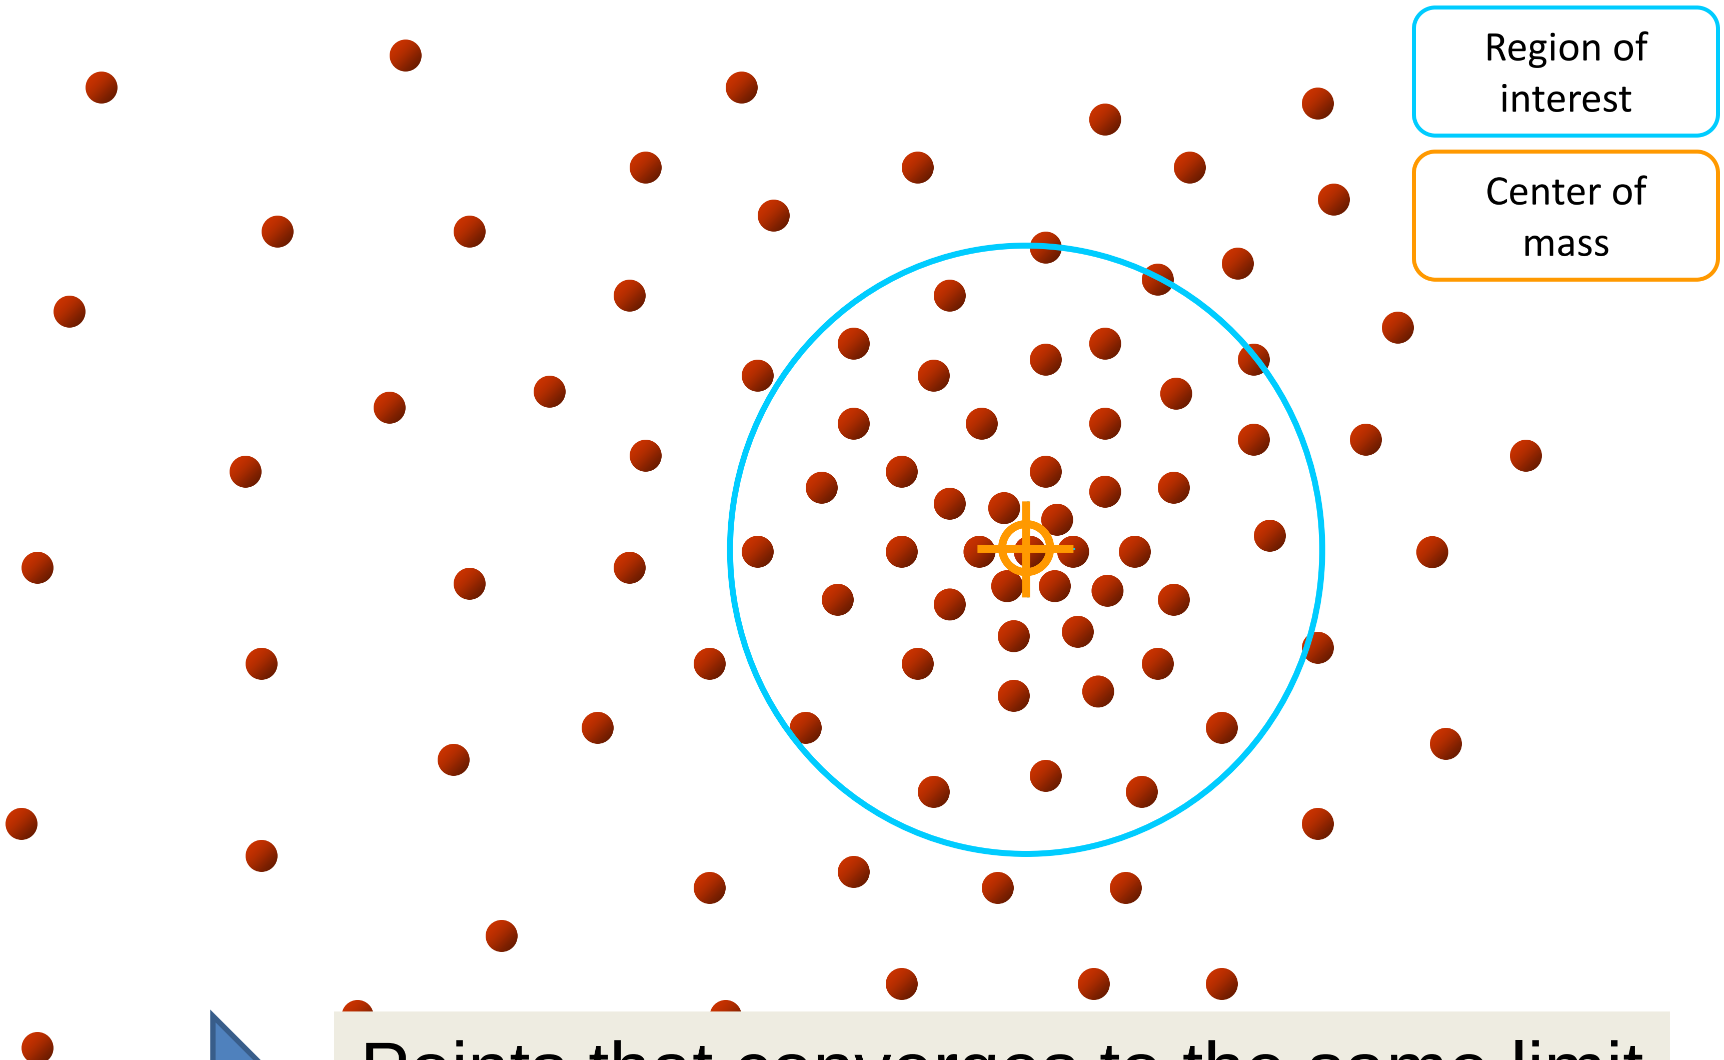

Points that converges to the same limit are grouped

Problem

- Computational time

$$\mathbf{y}_{j+1} = \frac{\sum_{i=1}^{n} K(\mathbf{y}_j - \mathbf{x}_i)\, \mathbf{x}_i}{\sum_{i=1}^{n} K(\mathbf{y}_j - \mathbf{x}_i)}$$

Underlying density function

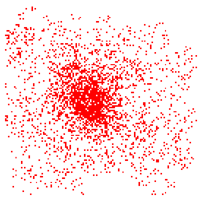
Data

Non parametric
Density gradient estimation
of a shadow kernel
(Mean Shift)

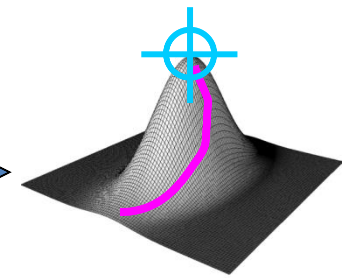
PDF

Density function: $D(\mathbf{p}) = \sum_{i=1}^{n} \tilde{K}(\mathbf{p} - \mathbf{x}_i)$

Shadow kernel

When a Gaussian kernel: $\tilde{K} = K$

Density function is computable

No iteration

Local maxima and saddles

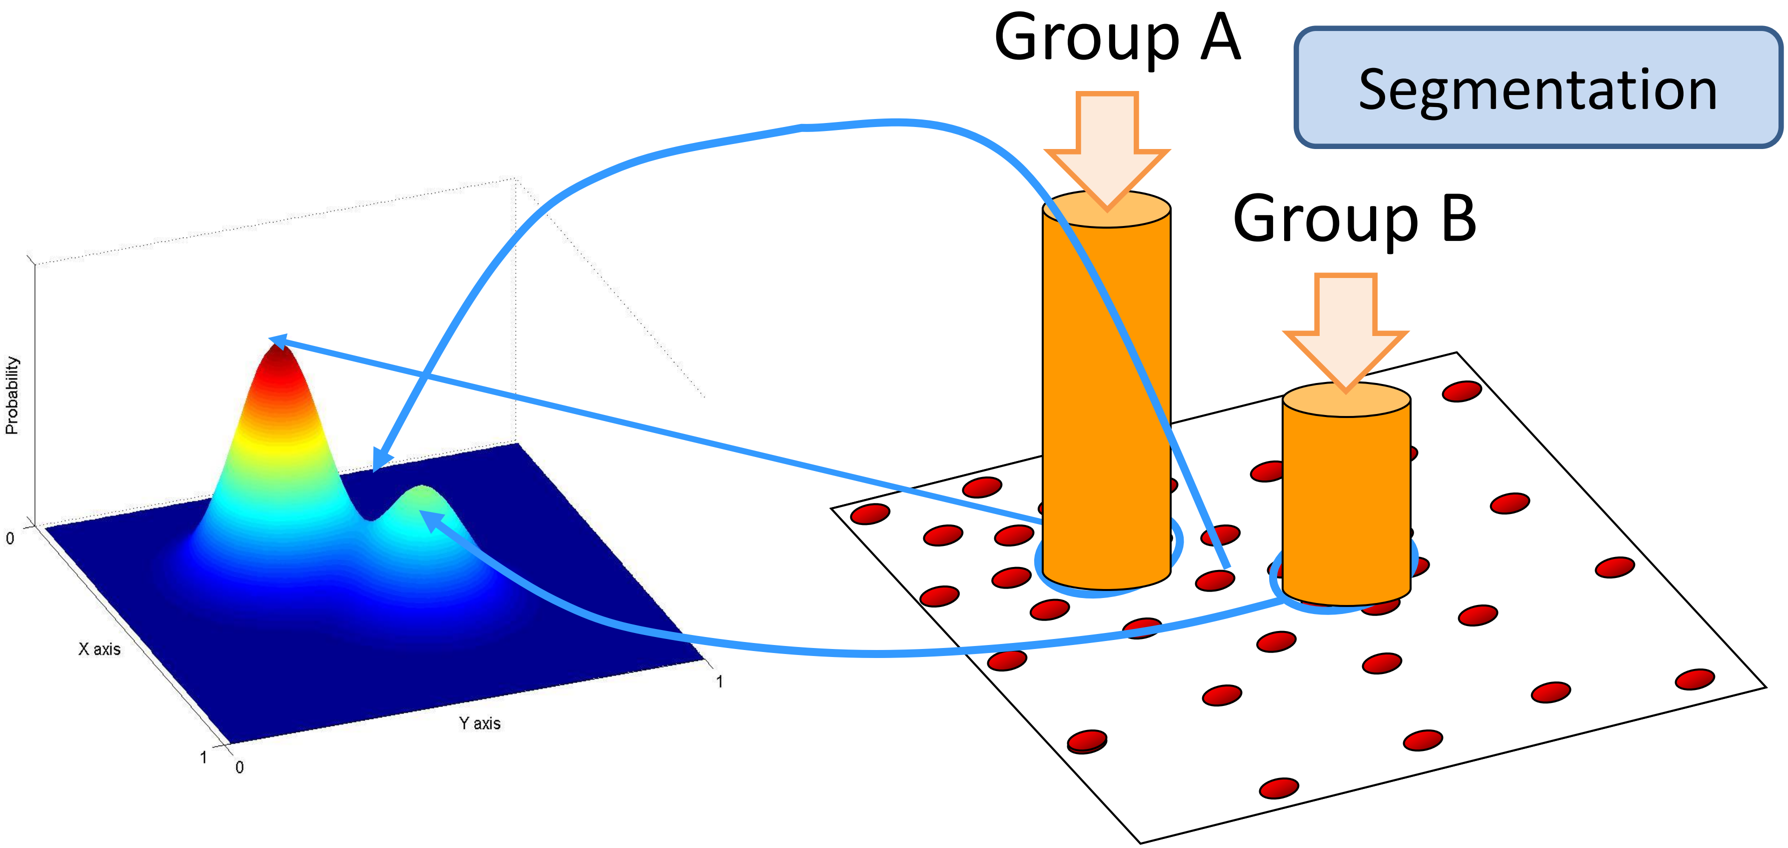

Underlying Density Function **Real Data Samples**

Problem

- Cluster hierarchy

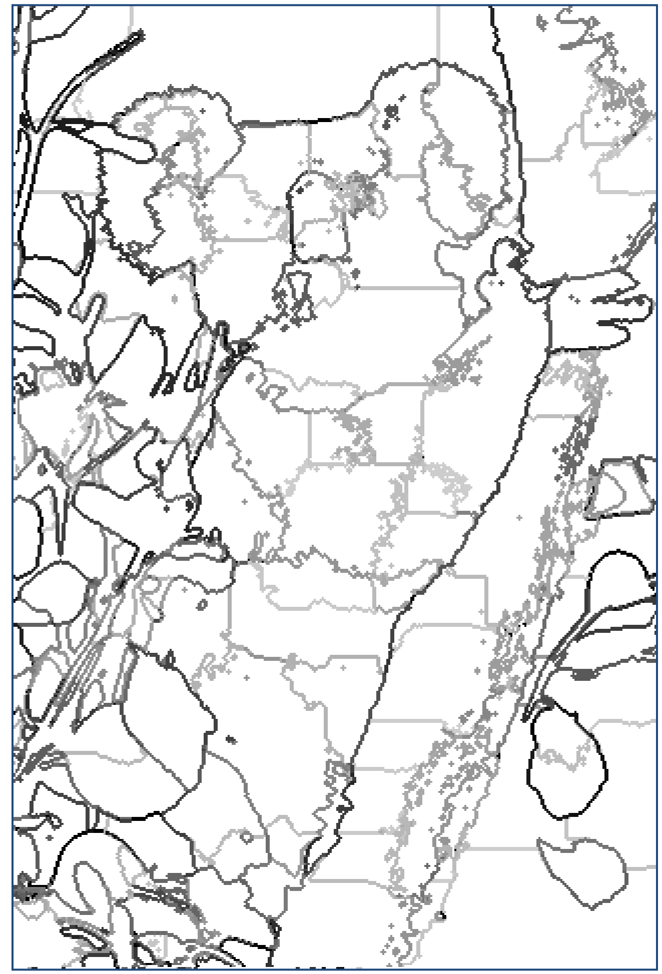

Morse theory

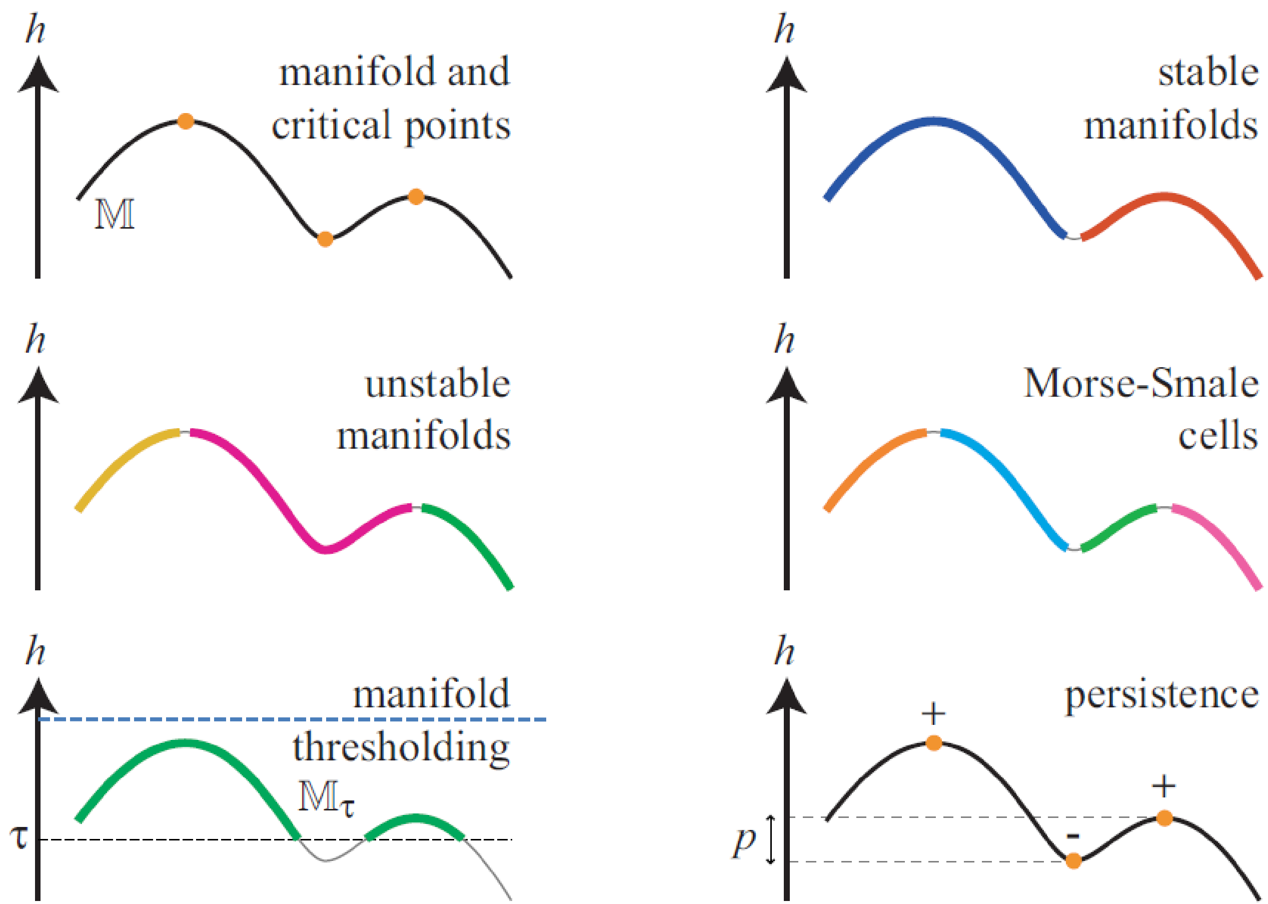

Change in p creates a topological feature → Critical point = **Positive**

Change in p removes a feature → Critical point = **Negative**

Hierarchy construction

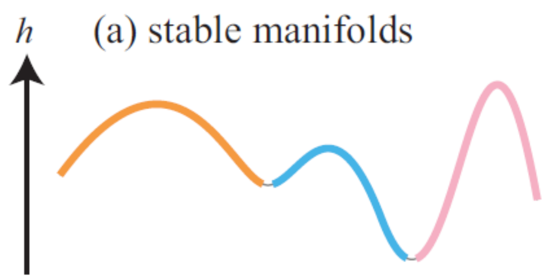

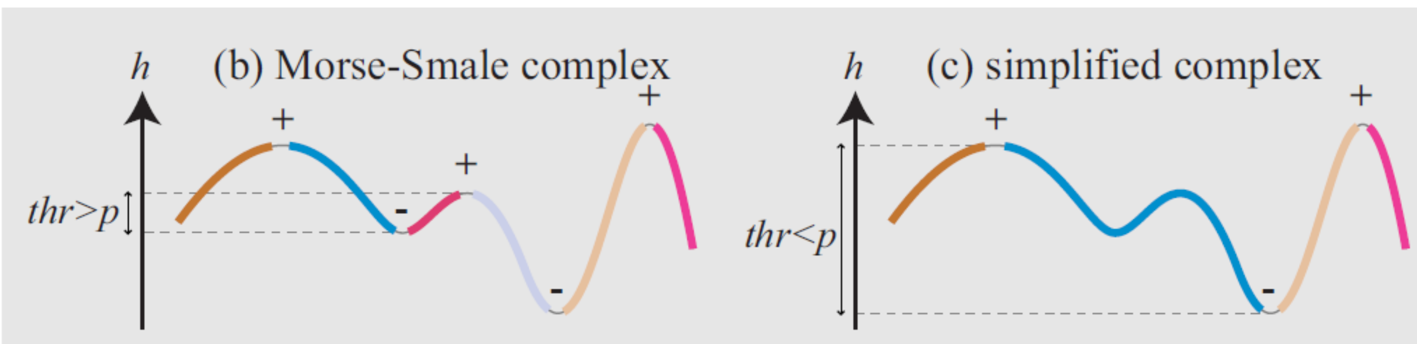

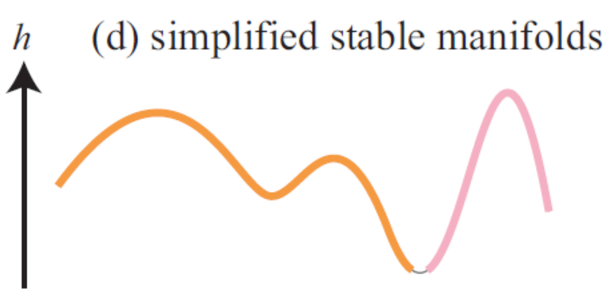

Change thr from 0 to ∞ to construct a hierarchy

Computation of density function

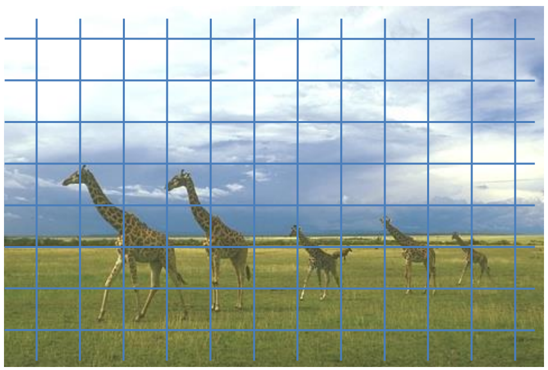

Take histogram of
5 dimensional values
(x, y, r, g, b)

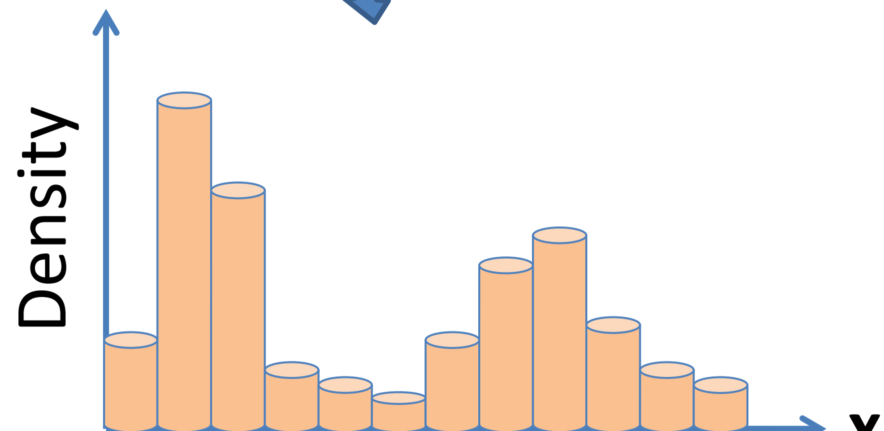

$$D(\mathbf{p}) = \sum_{i=1}^{n} \tilde{K}(\mathbf{p} - \mathbf{x}_i)$$

Calculate the convolution of each Gaussian kernels

Separability of Gaussian kernels in dimensions

Mode extraction

1. Sort g(k) by the values D(g(k))

   Position of the grid cell ➡ Computation: g1→ g4 → g2 → g3

2. When compute g(k),
   - Zero label → g(k) = local maxima
   - One label **m**(l) → g(k) is labeled with **m**(l)
   - Two or more labels → g(k) = boundary, label **b**

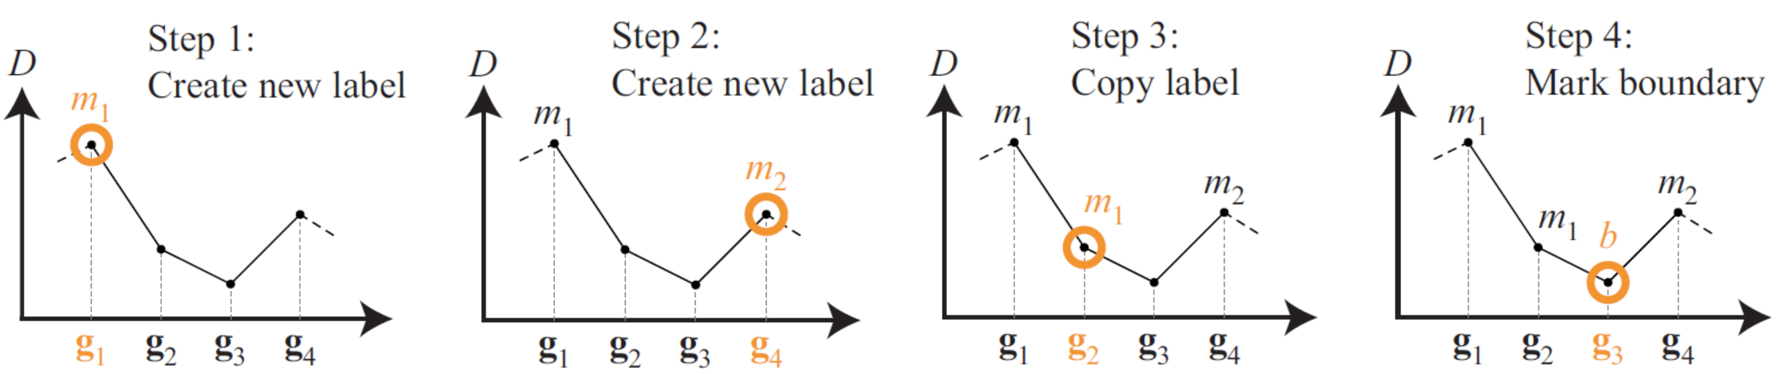

Result

Video

Quick time video

References

1. S. Paris et al., “A topological approach to hierarchical segmentation using mean shift,” CVPR 2007.
2. Y. Li et al., “Lazy snapping,” SIGGRAPH 2004.
3. J. Shotton et al., “TextonBoost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation,” ECCV 2006.
4. P. Kohli et al., “Robust higher order potentials for enforcing label consistency,” CVPR 2008.
5. http://www.wisdom.weizmann.ac.il/~deniss/vision_spring04/files/mean_shift/mean_shift.ppt
6. http://people.csail.mit.edu/sparis/

Oral: 3.9% (47), Overall: 23.5% (280)
Honorary Paper Mentions

Abhijeet Ghosh, Shruthi Achutha, Wolfgang Heidrich, Matthew O’Toole (Univ. of British Columbia)

International Conference on Computer Vision (ICCV) 2007, Oral

BRDF ACQUISITION WITH BASIS ILLUMINATION

What is this paper about?

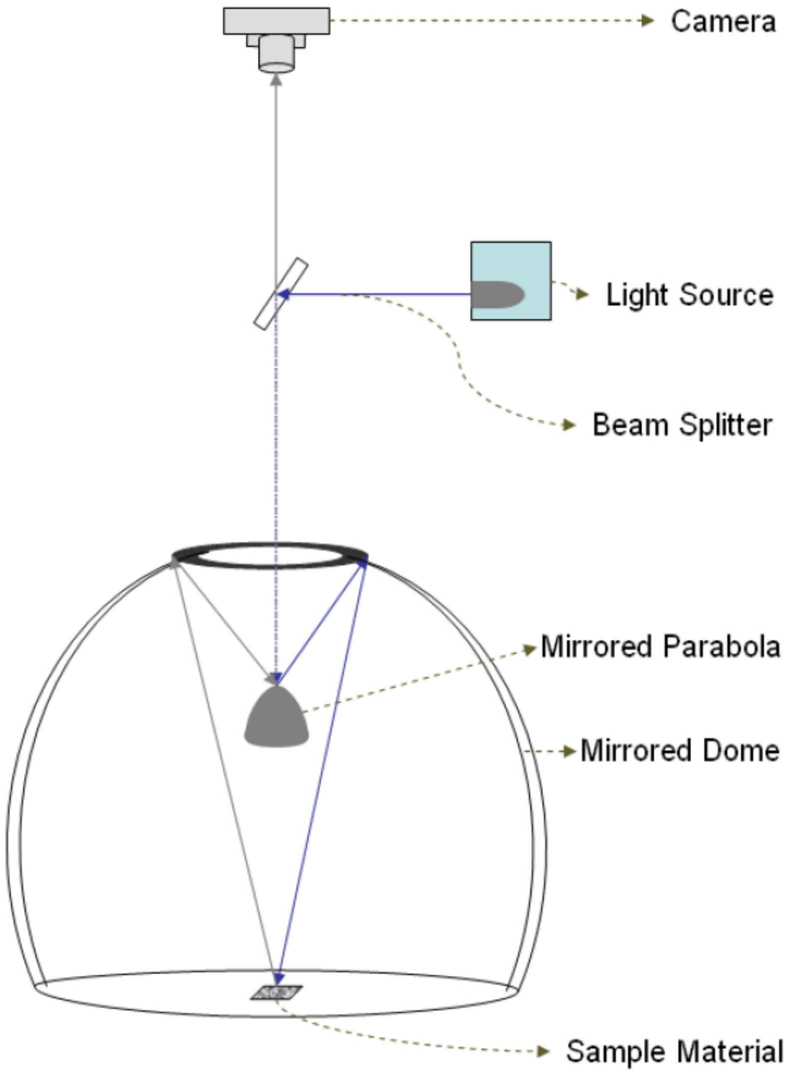

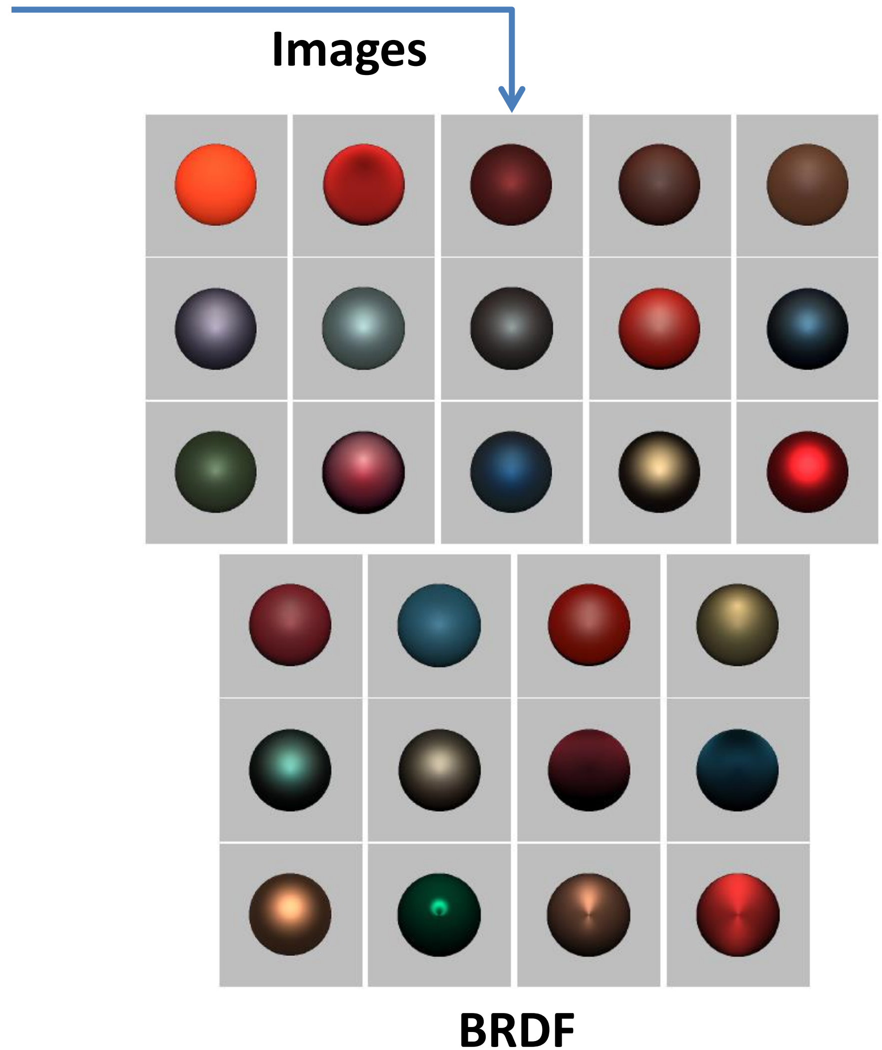

What is BRDF?

- Bidirectional Reflection Distribution Function

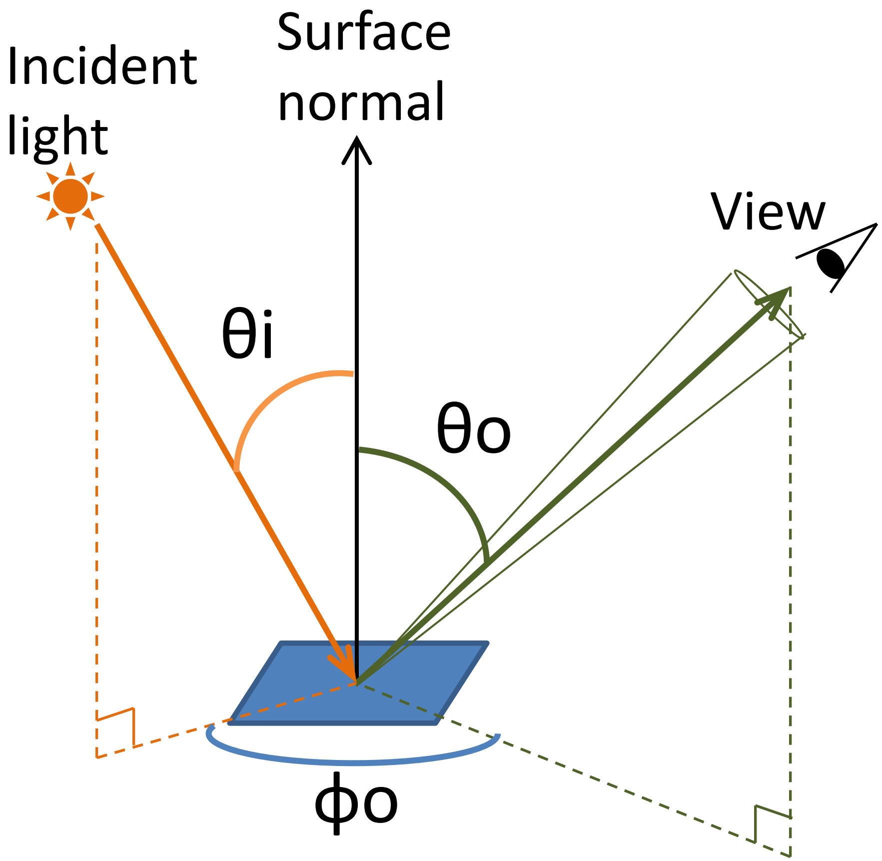

BRDF

Outgoing radiance

$$f_r(\theta_i,\phi_i,\theta_o,\phi_o)=\frac{dL(\theta_o,\phi_o)}{dE(\theta_i,\phi_i)}$$

incident irradiance

- Expresses object's reflection by 4 parameters

Reflected radiance

$$L(\omega_o) = \int \underbrace{f_r(\omega_i, \omega_o)}_{\textbf{BRDF}} \underbrace{L_i(\omega_i)\cos\theta_i d\omega_i}_{\textbf{dE (Incident irradiance/solid angle)}}$$

Outgoing radiance

$$\text{where} \quad \omega = (\theta, \phi)$$

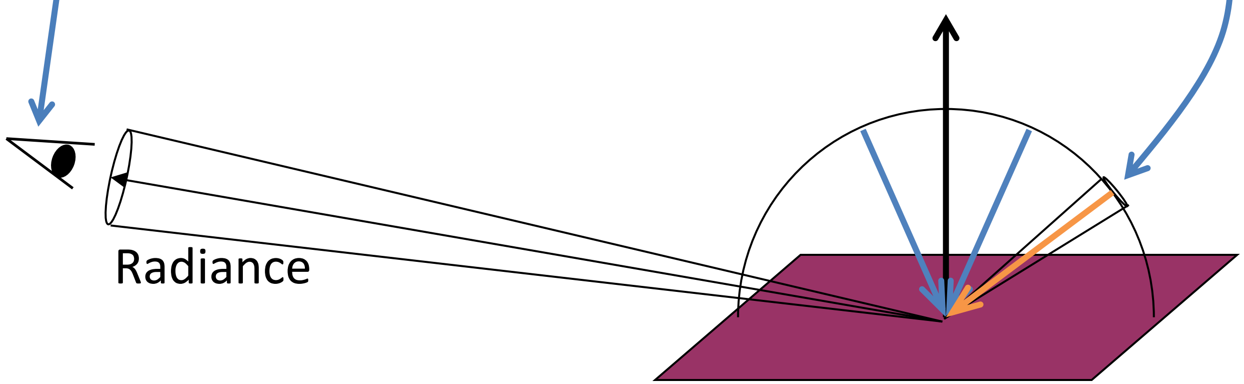

Related work

Parametric models of BRDF

- Lambertian surface
- Specular reflection
  - Phong, Oren-Nayar, Torrance-Sparrow, Blinn (simplified Torrance-Sparrow), Cook-Torrance, Beckman-Spizzichino
- Anisotropic reflection
  - Ward

Lambertian

I_d: diffuse reflection intensity
K_d: diffuse albedo
θ: angle $\rightarrow$ $\cos\theta = \mathbf{n} \cdot \mathbf{l}$

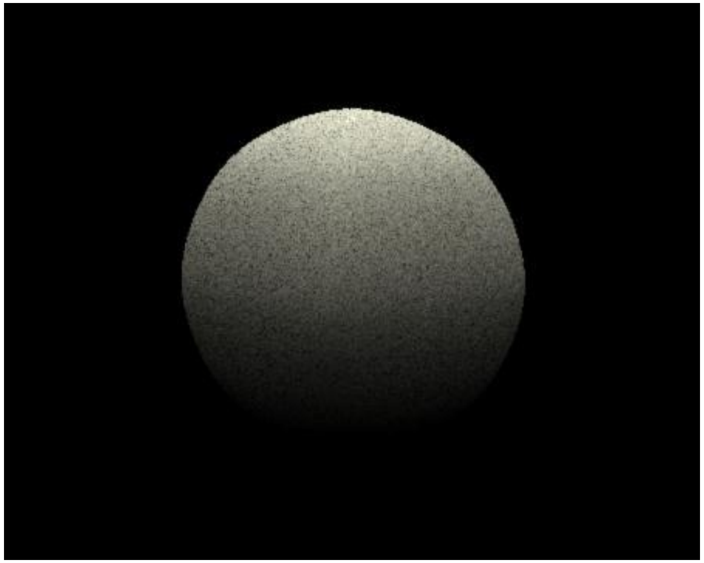

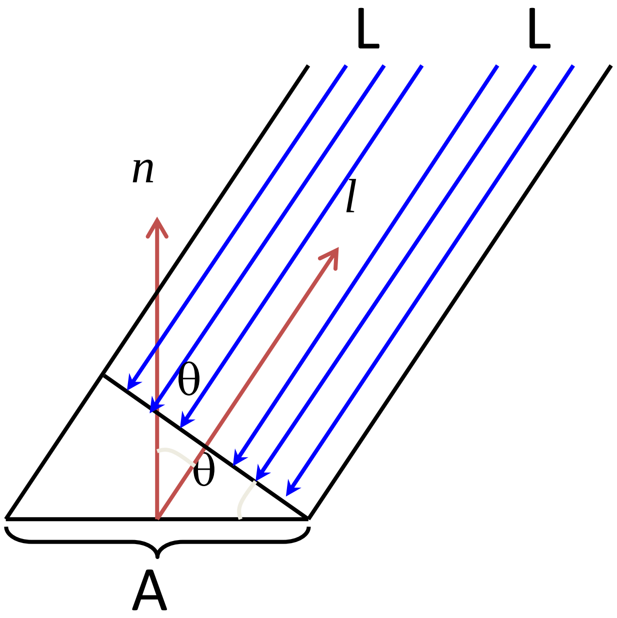

light per unit area = L
area in light direction = A cos θ
radiant flux = L × A cos θ
actual area = A
irradiance = L × A cos θ ÷ A = L cos θ

Diffuse, specular lobe, specular spike

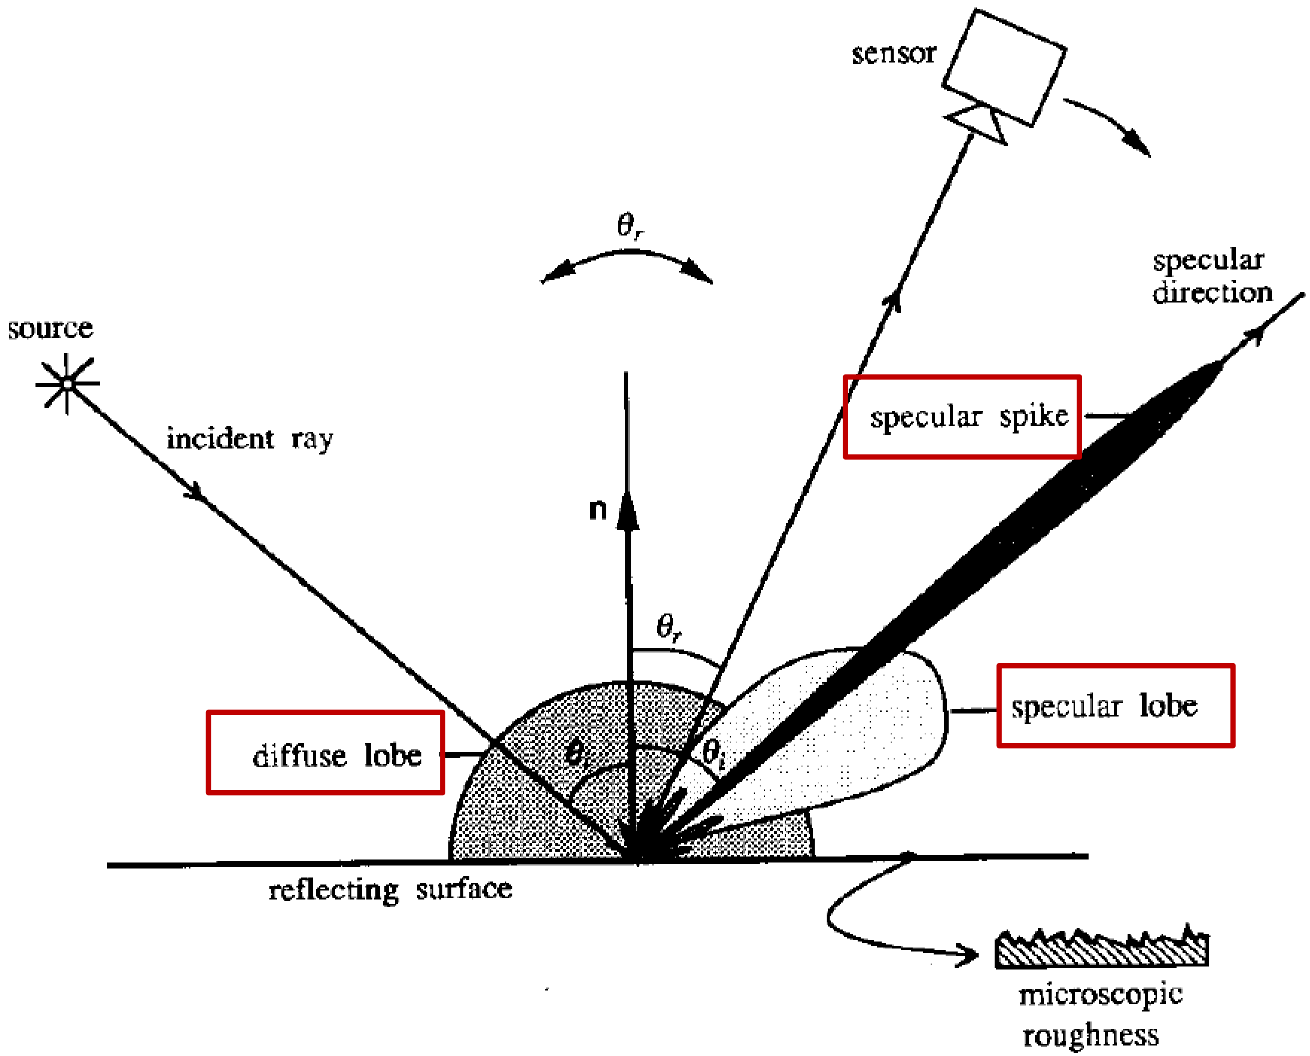

Anisotropic reflection

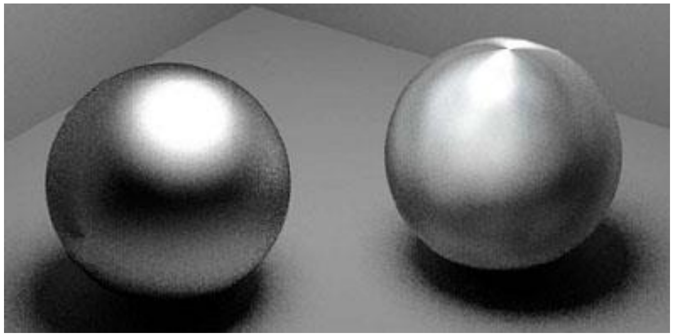




Direct measurement of BRDF

- Goniophotometers
- Light stage

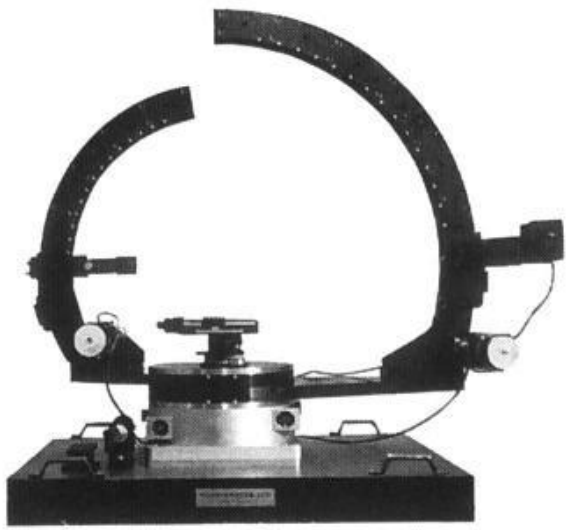

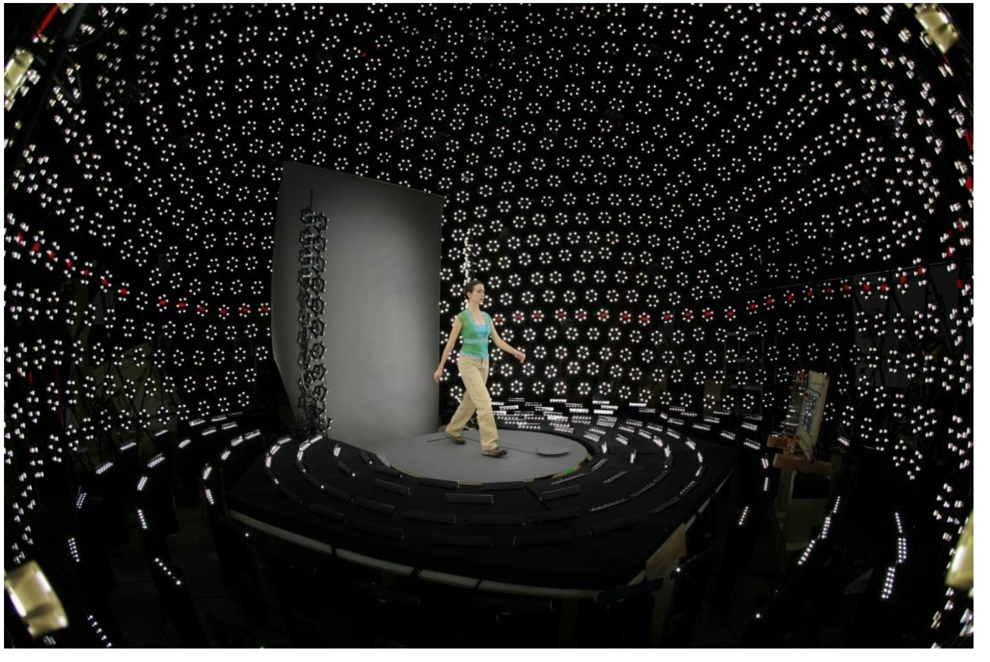

Measure impulse response using **pencils of light** ≒ **Dirac's delta function**

Efficient measurement of BRDF

- Assumption of isotoropic reflection
- Use of reflection models
- Use of a sphere for the target sample

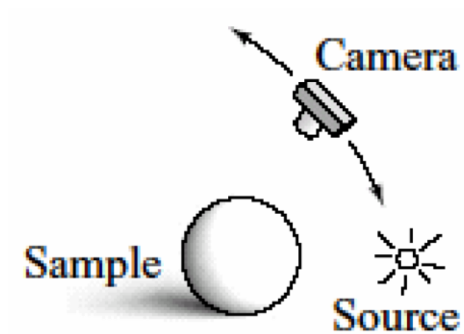

Advantage of the proposed method

- Illuminations are smooth basis functions
  - Efficient data acquisition

Method

System overview

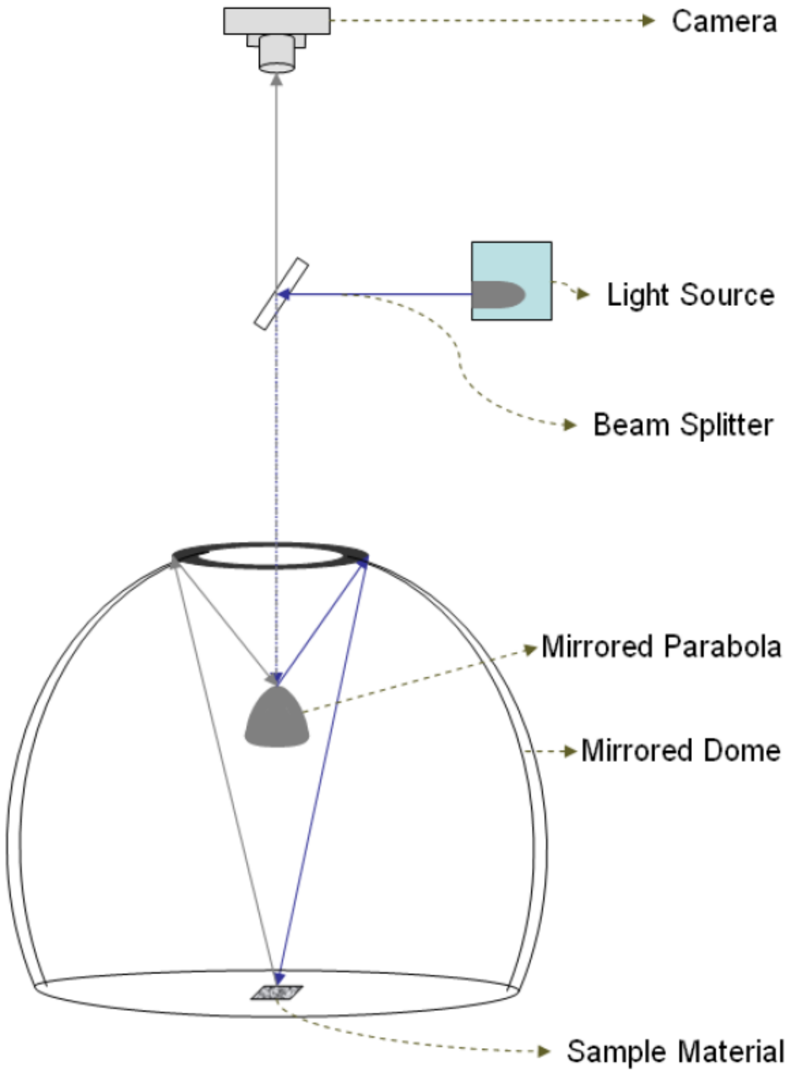

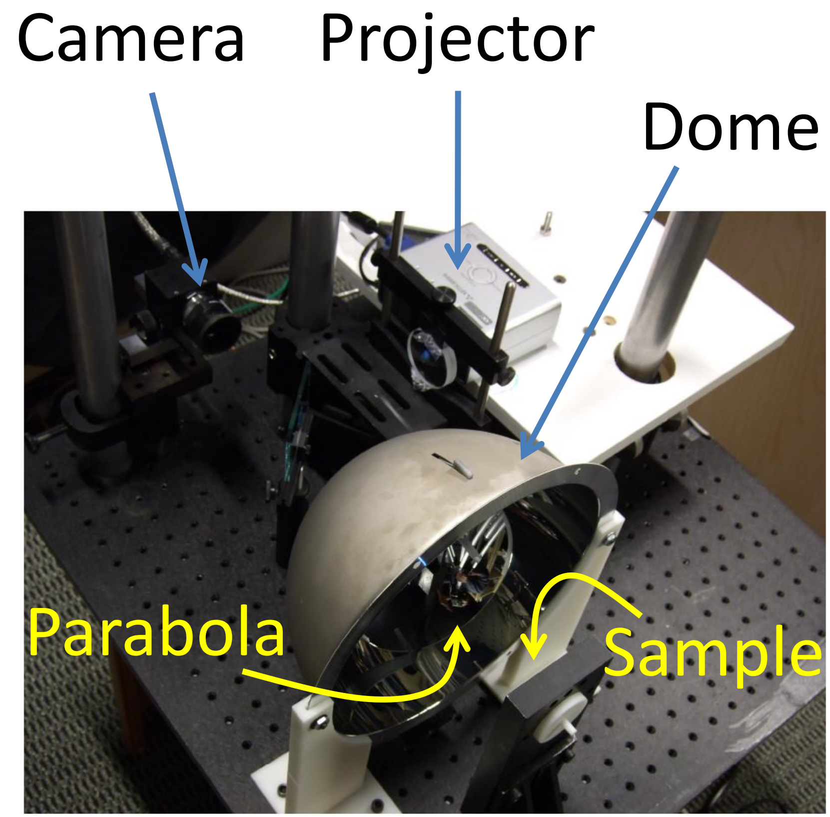

Basis functions

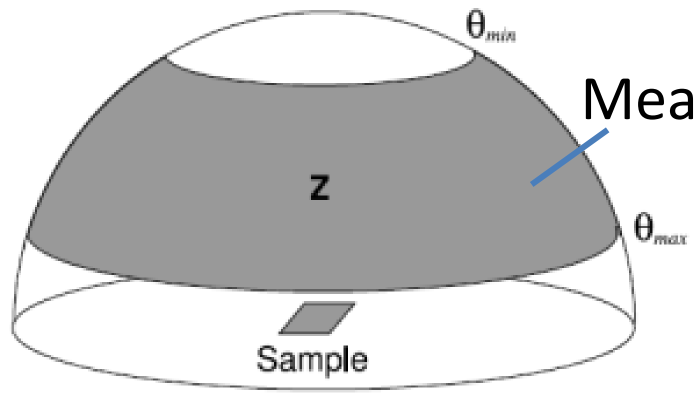

Measurement zone **Z**

$$\underbrace{\hat{f}_r(\omega_i,\omega_o)}_{\textbf{New notation of BRDF}} = \underbrace{f_r(\omega_i,\omega_o)}_{\textbf{BRDF}}\cos\theta_i \approx \sum_k \underbrace{Z_k(\omega_i)}_{\textbf{Basis functions}} z_k(\omega_o)$$

$$L_o(\omega_o) = \int_{\mathbf{Z}} f_r(\omega_i,\omega_o)L_i(\omega_i)\cos\theta_i\, d\omega_i \approx \sum_k \underbrace{z_k(\omega_o)}_{\textbf{coefficients}} \underbrace{\int_{\mathbf{Z}} Z_k(\omega_i)L_i(\omega_i)\, d\omega_i}_{\textbf{Basis functions}}$$

Measurement with basis functions

$$L(\omega_o) = \int f_r(\omega_i, \omega_o) L_i(\omega_i) \cos\theta_i d\omega_i$$

Incident radiance (illumination)

Basis function

$$\int_{\mathbf{Z}} f_r(\omega_i, \omega_o) Z_1(\omega_i) \cos\theta_i d\omega_i$$

$$= \int_{\mathbf{Z}} \left[\sum_k z_k(\omega_o) Z_k(\omega_i) \right] Z_1(\omega_i) d\omega_i$$

$$= \int_{\mathbf{Z}} \left[z_1(\omega_o) Z_1(\omega_i) + \cdots + z_K(\omega_o) Z_K(\omega_i) \right] Z_1(\omega_i) d\omega_i$$

$$= \int_{\mathbf{Z}} z_1(\omega_o) Z_1(\omega_i)^2 d\omega_i$$

$$\int_{\mathbf{Z}} Z_1(\omega_i)^2 d\omega_i = 1$$

$$= z_1(\omega_o)$$

Zonal basis functions

$$Z_l^m(\theta,\phi) = \begin{cases} \sqrt{2}\hat{K}_l^m \cos(m\phi)\hat{P}_l^m(\cos\theta) & \text{if } m > 0 \\ \sqrt{2}\hat{K}_l^m \sin(-m\phi)\hat{P}_l^{-m}(\cos\theta) & \text{if } m < 0 \\ \hat{K}_l^0 \hat{P}_l^0(\cos\theta) & \text{if } m = 0 \end{cases}$$

where $\hat{K}_l^m = \sqrt{\dfrac{(2l+1)(l-|m|)!}{2\pi \cdot (\cos\theta_{min} - \cos\theta_{max}) \cdot (l+|m|)!}}$

Results

1 minute (BRDF measurement + re-projection into spherical harmonic basis)

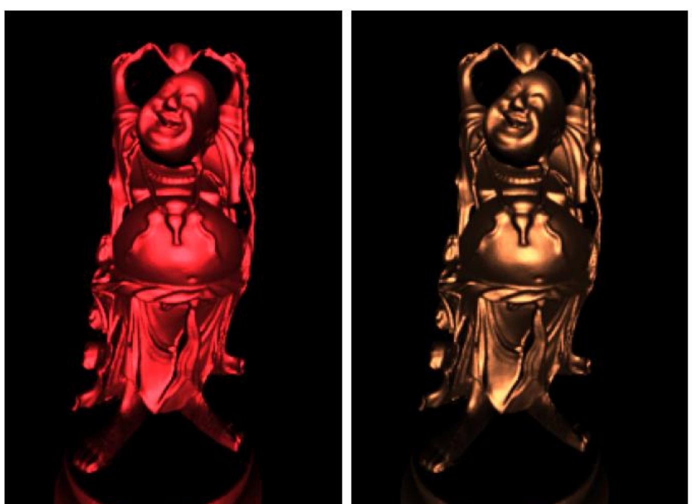

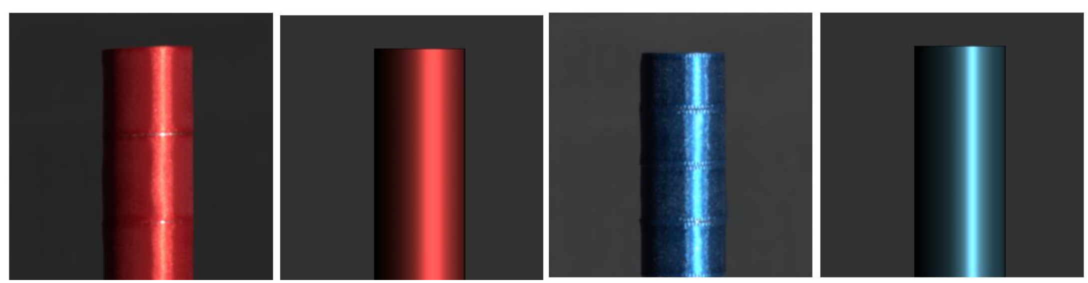

Results

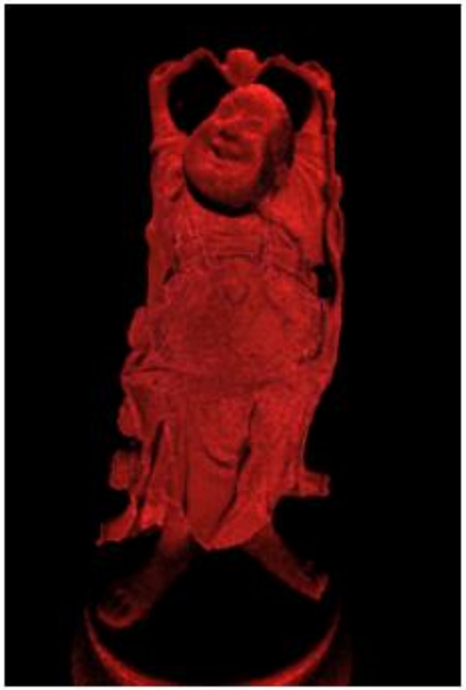

Red velvet

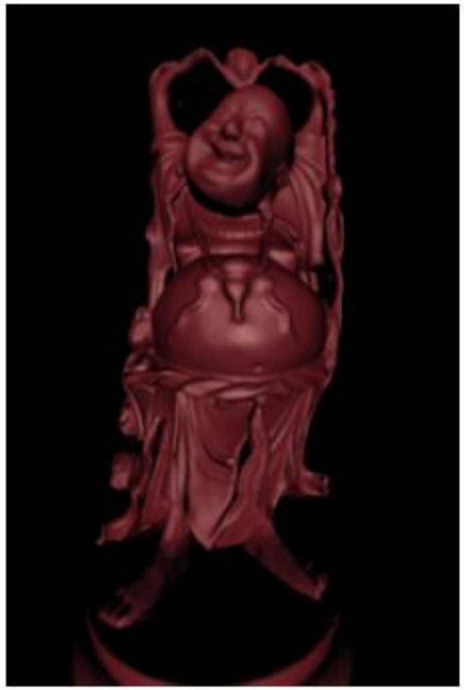

Red printer toner

Magenta plastic sheet

Chrome gold dust automotive paint

Representative set of BRDFs acquired with lower order zonal basis functions

Video

Quick time video

References

1. A. Ghosh et al., “BRDF acquisition with basis illumination,” ICCV 2007.

2. S. K. Nayar, “Surface reflection: Physical and geometrical prespectives,” TPAMI 1991.
3. Y. Sato and Y. Mukaigawa “Inverse rendering,” http://www.am.sanken.osaka-u.ac.jp/~mukaigaw/papers/CVIM-145-9.pdf, in Japanese

Oral: 4.6% (40), Overall: 27.9% (243)

Sujit Kuthirummal, Aseem Agarwala, Dan B Goldman, Shree K Nayar (Columbia University and Adobe Systems, Inc.)

European Conference on Computer Vision (ECCV) 2008, Oral

PRIORS FOR LARGE PHOTO COLLECTIONS AND WHAT THEY REVEAL ABOUT CAMERAS

Individual Photograph ← Camera (Scene, Camera, Photographer)

Scale Photographer

Credit: Snowdosker @ Flickr

Individual Photograph ⟵ (Scene, Camera, Photographer)

Internet Photo Collections

Exif Tags

Recover information about
Scenes, Cameras, and Photographers

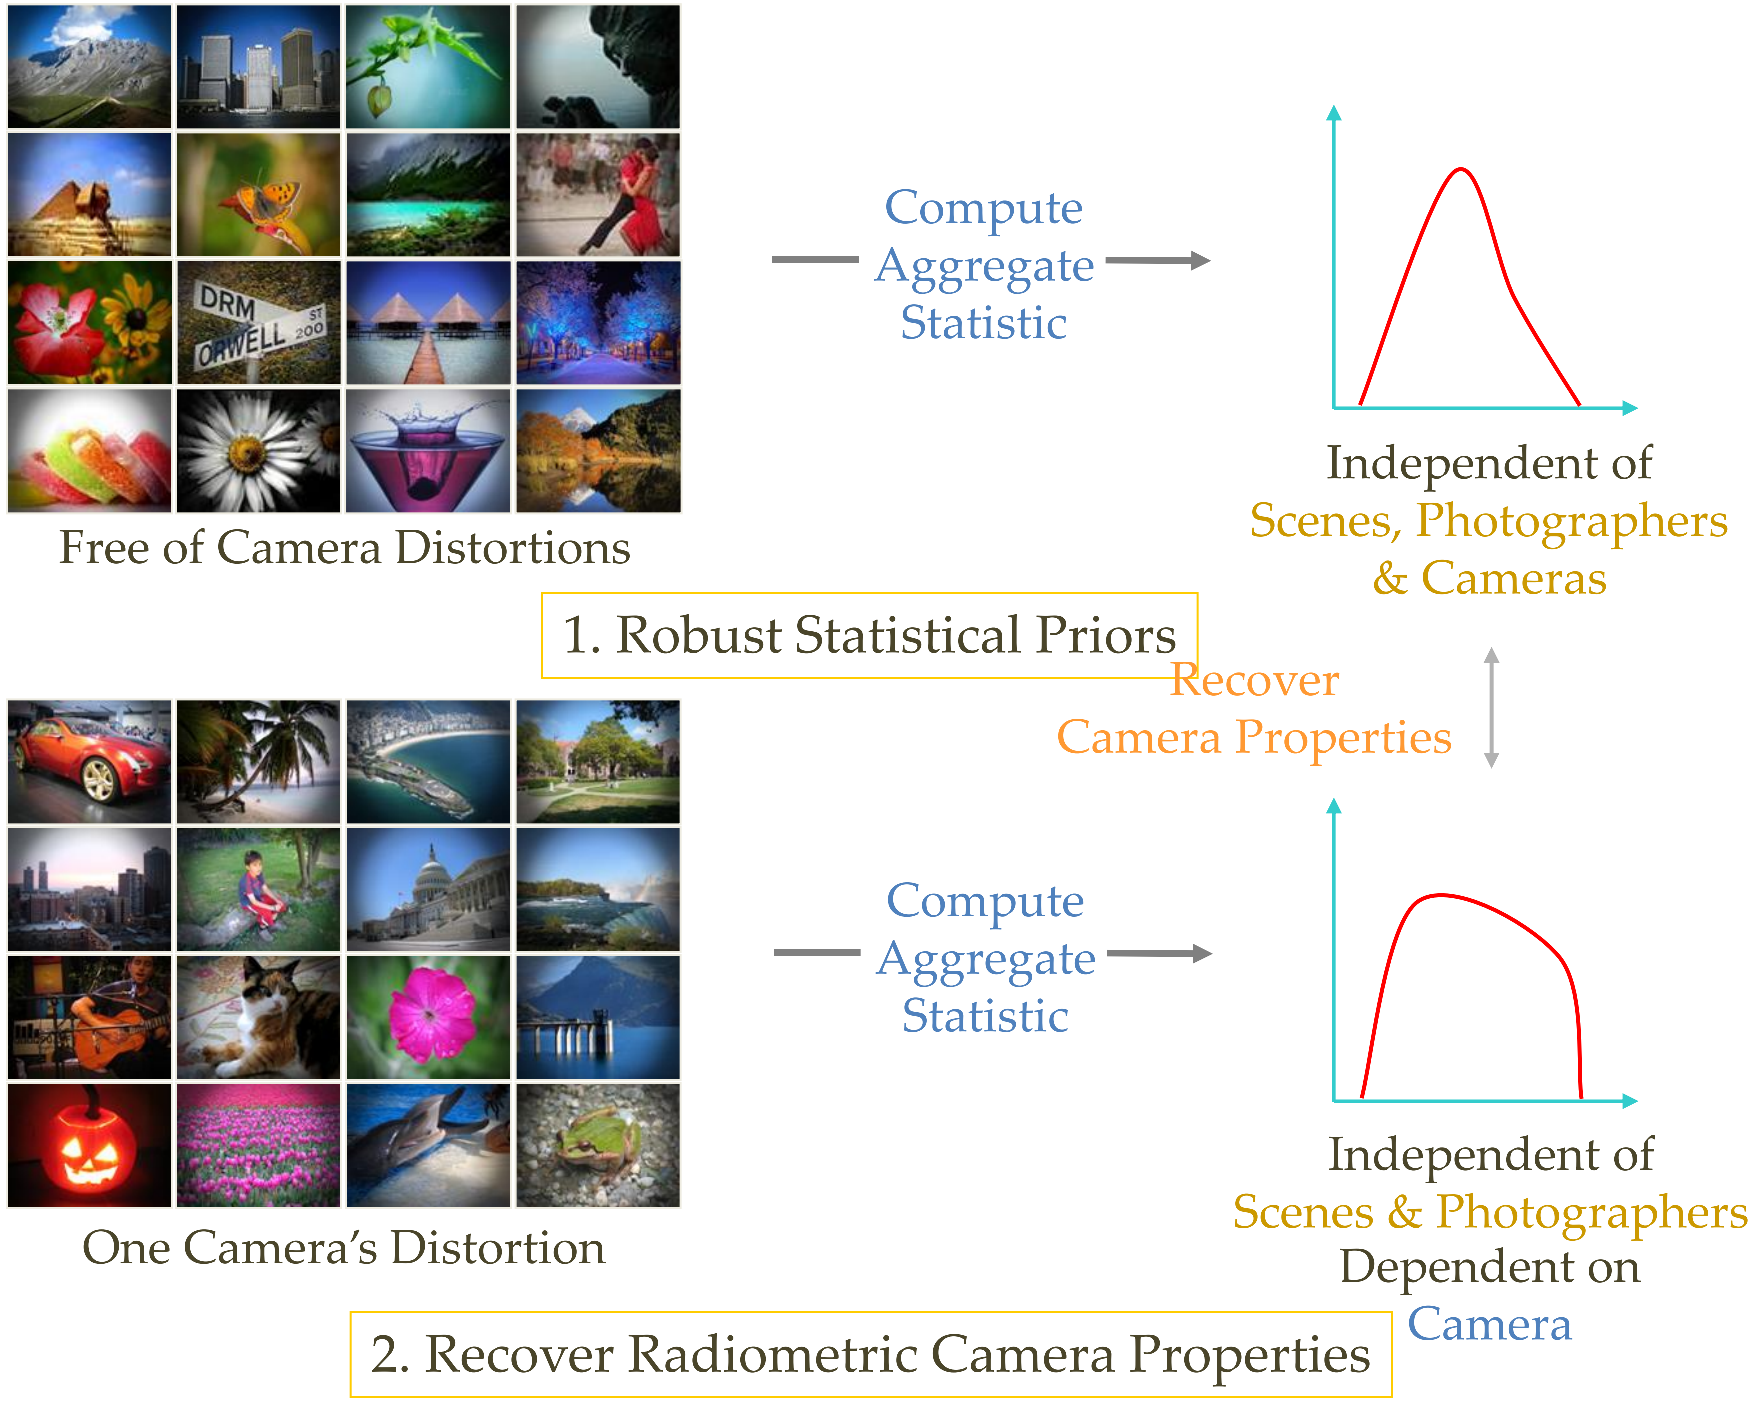
Free of Camera Distortions
Compute
Aggregate
Statistic
Independent of
Scenes, Photographers
& Cameras
1. Robust Statistical Priors
Recover
Camera Properties
One Camera's Distortion
Compute
Aggregate
Statistic
Independent of
Scenes & Photographers
Dependent on
Camera
2. Recover Radiometric Camera Properties

Related Work: Large Photo Collections

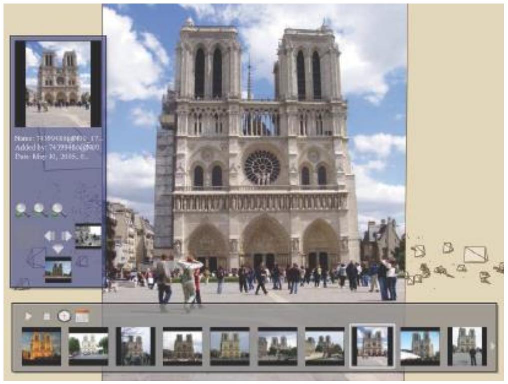

Photo Tourism
Snavely et al. '06

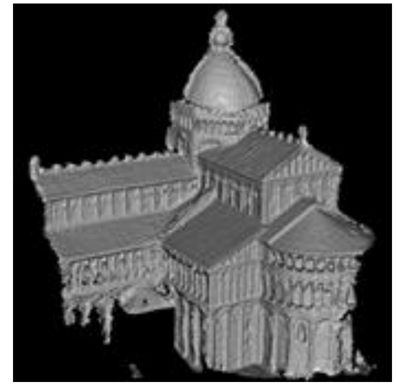

Internet Stereo
Goesele et al. '07

Object Insertion
Lalonde et al. '07

Hole Filling
Hays et al. '07

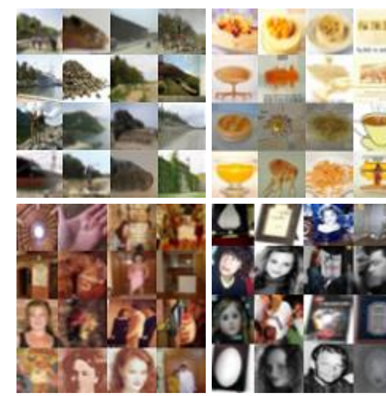

Recognition
Torralba et al. '07

Recover Camera Properties

Related Work: Image Statistics

- Natural Image Statistics
  - 1/f amplitude spectrum fall-off *Burton & Moorhead '87, Field '87*
  - Sparsity of image derivatives *Olshausen & Field '96, Simoncelli '97*
  - Bias in gradient orientations *Switkes et al. '78, Baddeley '97*

- Exploit priors for
  - Scene recognition *Baddeley. '97, Torralba & Oliva '03*
  - Super-resolution *Tappen et al. '03*
  - Deriving intrinsic images *Weiss '01*
  - Image denoising *Roth et al. '05*
  - Removing camera shake *Fergus et al. '06*

- Priors attempt to describe statistics of individual photographs

---

- Our Priors describe aggregate statistics of many photographs

Camera Model Centric Photo Collections

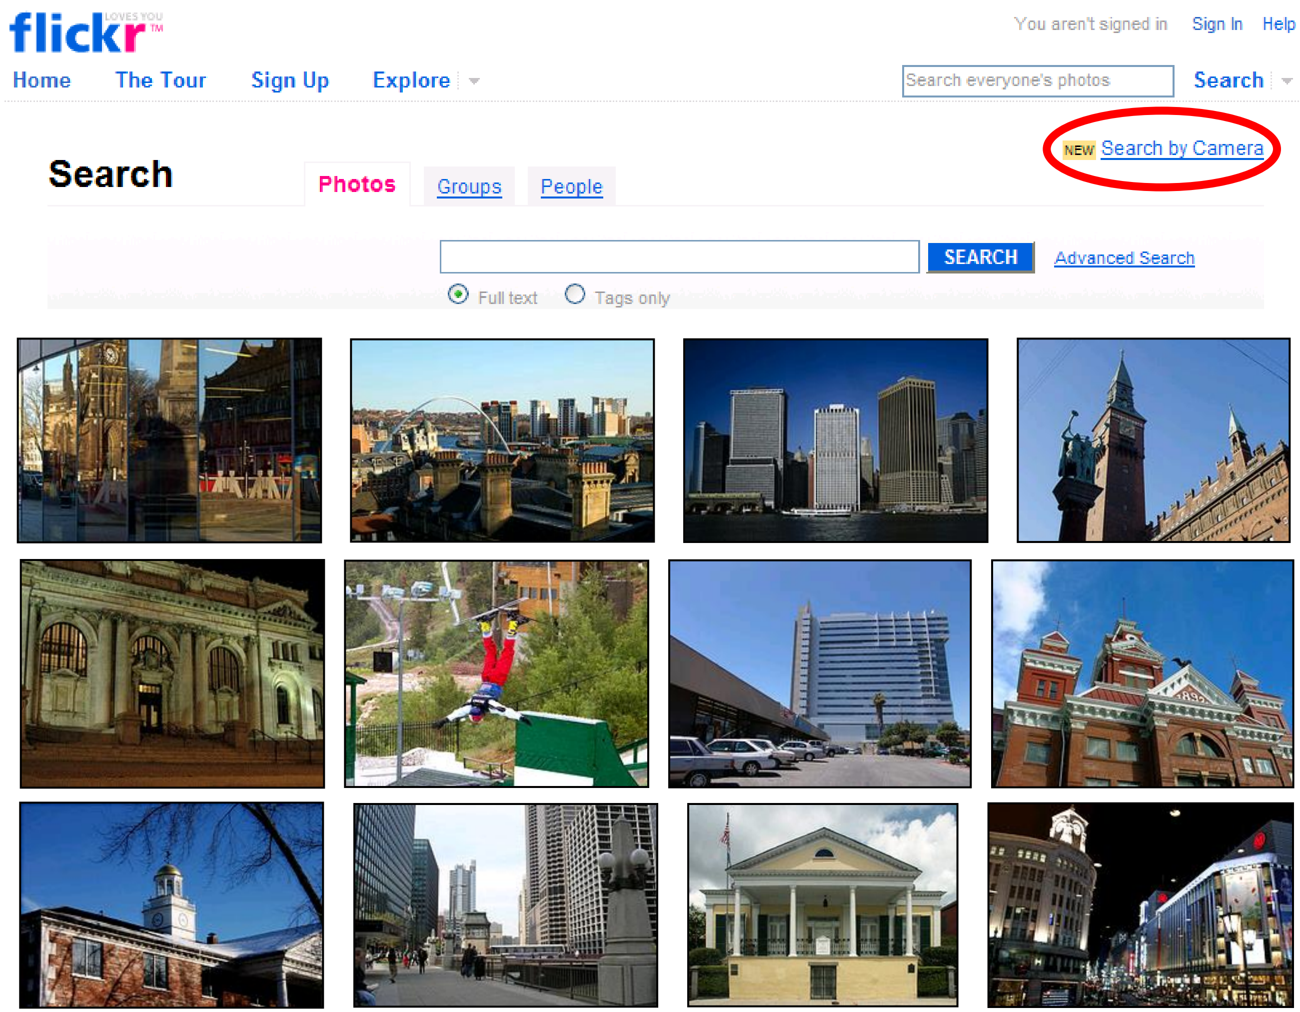

Point-and-Shoot Camera Models

Canon S1IS

Exif Tags






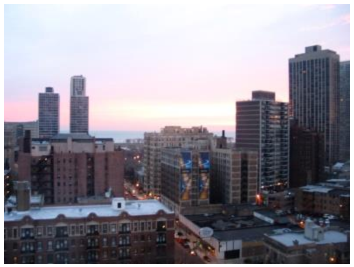

Cropped







Photoshopped



Portrait Mode


Flash


Canon S1IS

Focal Length: 5.8 mm F-Number: 2.8

Focal Length: 5.8 mm F-Number: 4

Exif Tags

Focal Length: 58 mm F-Number: 3

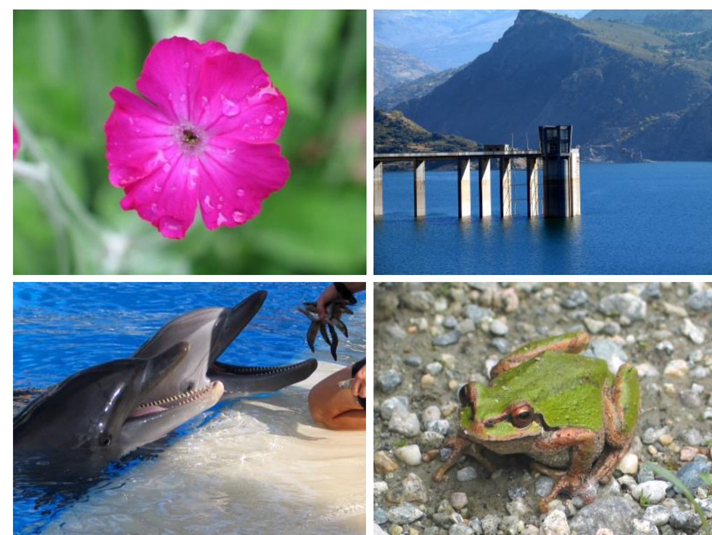

Focal Length: 58 mm F-Number: 4

Camera Distortion Free

Training Set

Compute
Aggregate
Statistic

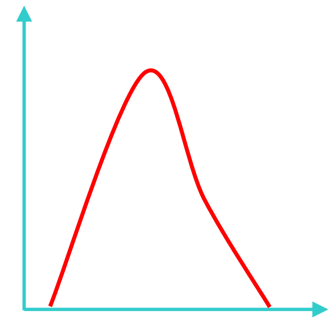

Independent of
Scenes, Photographers
& Cameras

1. Robust Statistical Priors

Creating the Training Set

Canon S1 IS

Camera Response

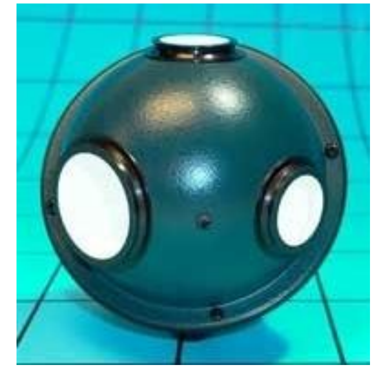

Vignetting

Remove camera-specific properties

Training Set

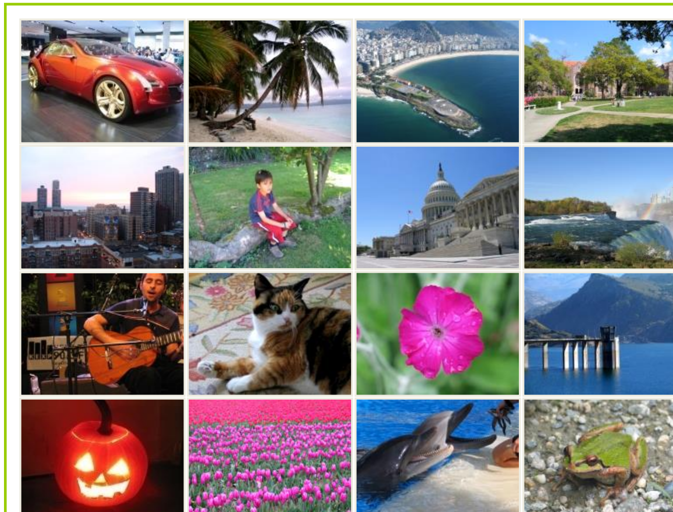

Camera Distortion Free

Radiometric Camera Properties

- Properties of specific *camera models*
  - Camera response function
  - Vignetting for different lens settings

- Properties of specific *camera instances*
  - Bad pixels on the detector

Related Work: Response Estimation

- Multiple images
  - Varying camera exposures
    Mann & Picard '95, Debevec & Malik '97, Mitsunaga & Nayar '99, Grossberg & Nayar '03

  - Combinations of illuminations
    Manders et al. '04

- Single image
  - High order Fourier correlations
    Farid '01
  - Intensity statistics at edges
    Lin et al. '04, '05

- Fully automatic, robust estimation
- Do not need access to the camera

The Gradient Prior

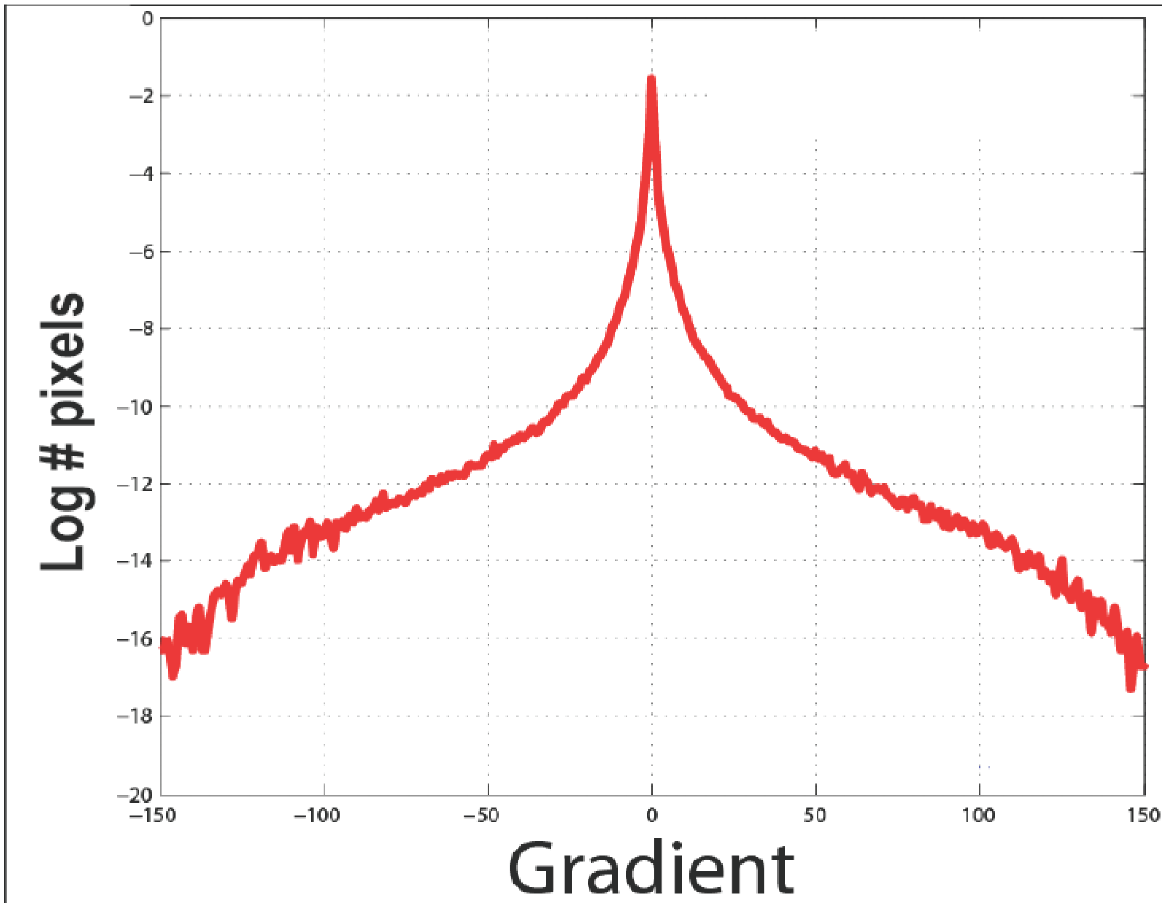

Fergus et al. '06

Joint Histogram of Irradiances at Neighboring Pixels (Linearized Images)

Joint Histogram of Irradiances at Neighboring Pixels (Linearized Images)

Red Channel

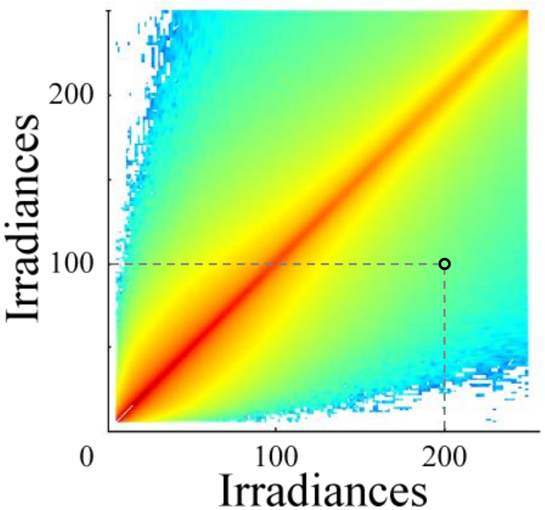

Green Channel

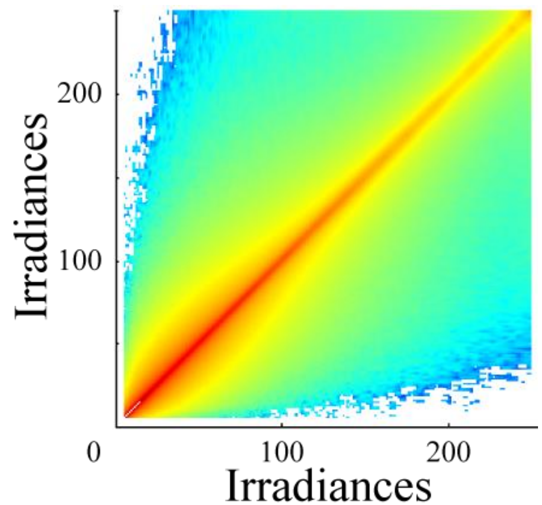

Blue Channel

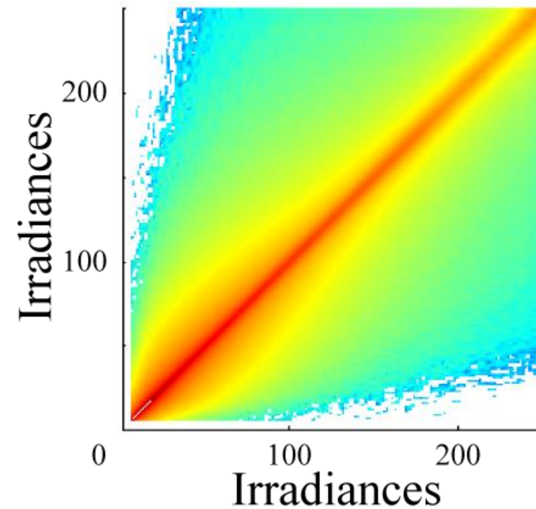

Canon S1IS   Focal Length: 5.8 mm   F-Number: 4.5   15,550 Images

Joint Histogram of Irradiances at Neighboring Pixels (Linearized Images)

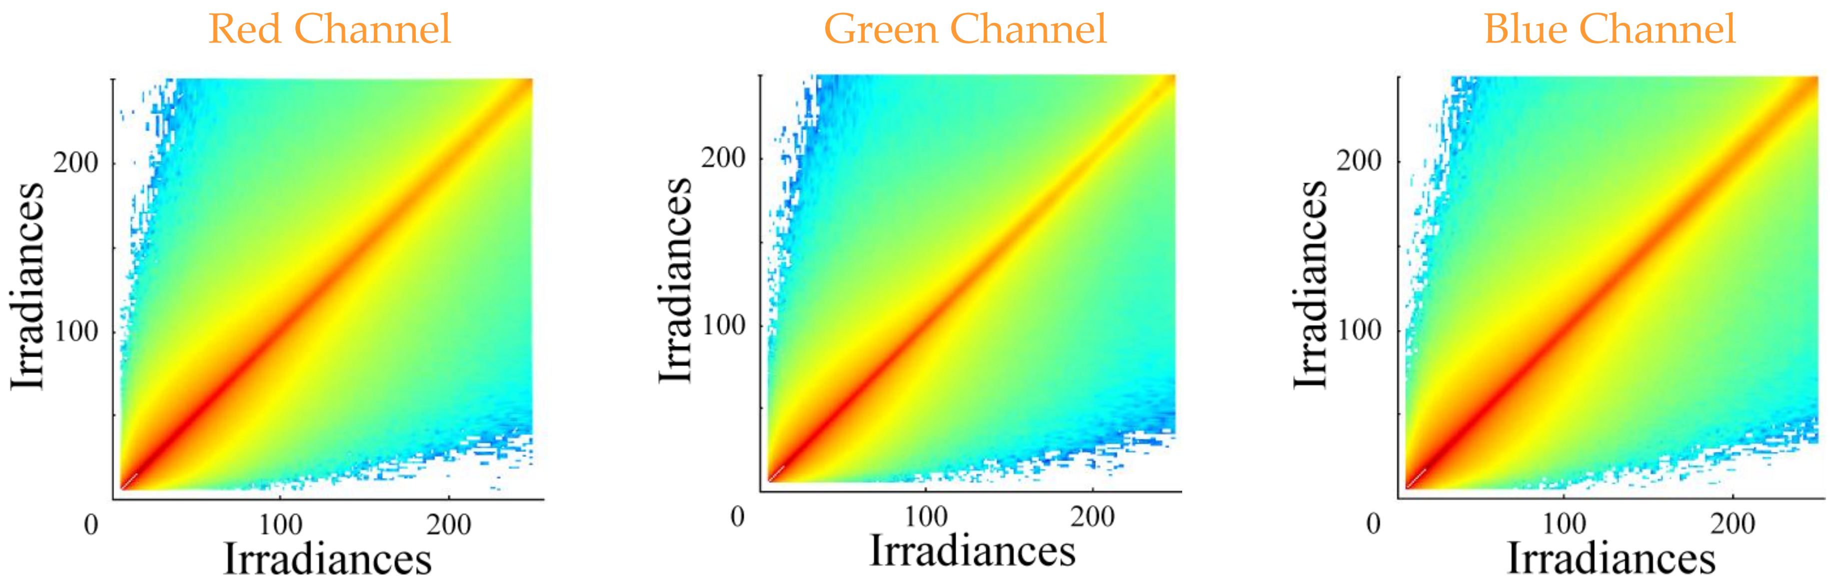

Canon S1IS   Focal Length: 5.8 mm   F-Number: 4.5   15,550 Images

- Joint Histograms are very similar across camera models
  - Especially for smallest focal length and largest f-number
  - KL Divergence between corresponding histograms of Canon S1IS and Sony W1 cameras –

Prior: Joint Histograms of any one camera model

Red: 0.059   Green: 0.081   Blue: 0.068

Estimating Camera Response Function

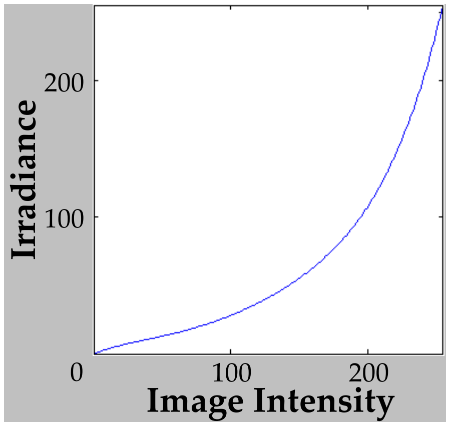

Inverse Response Guess

Compute

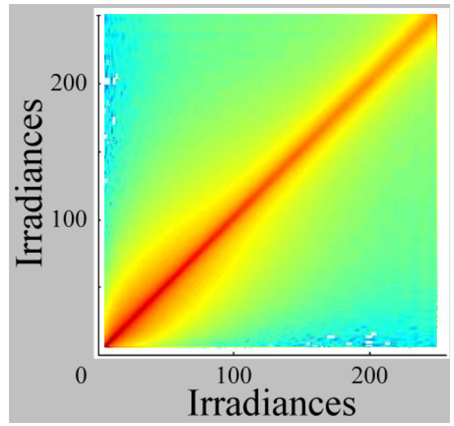

Estimated Joint Histogram

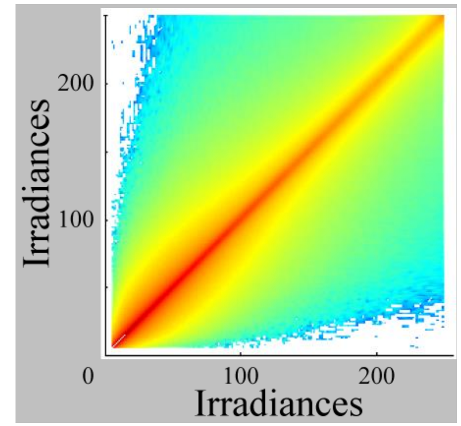

Prior Joint Histogram

Similarity (KL Divergence)

Optimization

Quality of Inverse Response Guess

$$R(i) = 255 * \sum_{k=1}^{D} \alpha_k \left(\frac{i}{255}\right)^k$$

$R(i)$	Irradiance corresponding to intensity i
α	Polynomial coefficients
D	Polynomial degree (5)

Estimating Camera Response Function

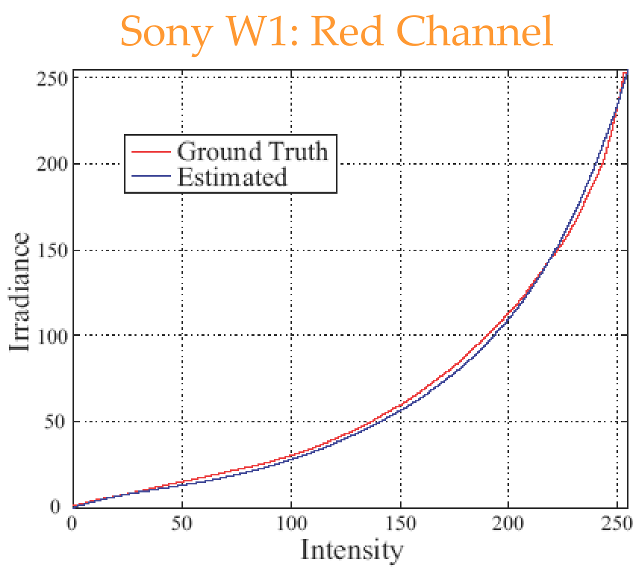

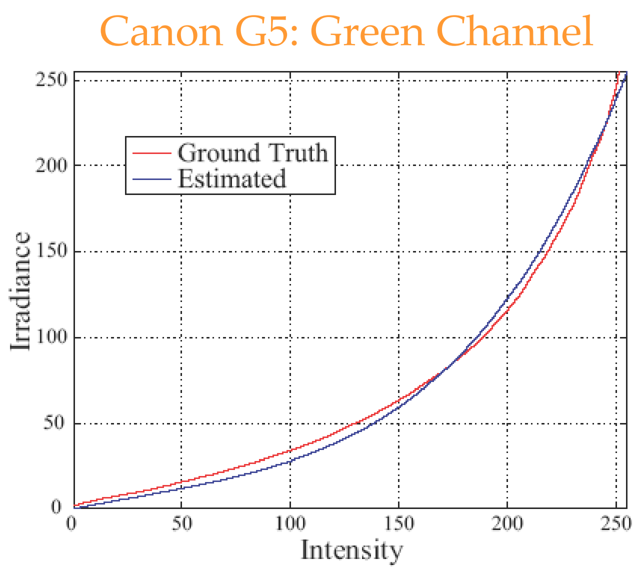

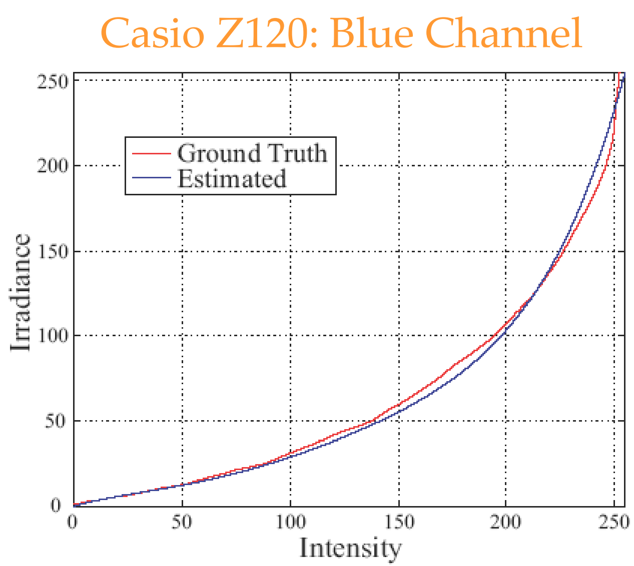

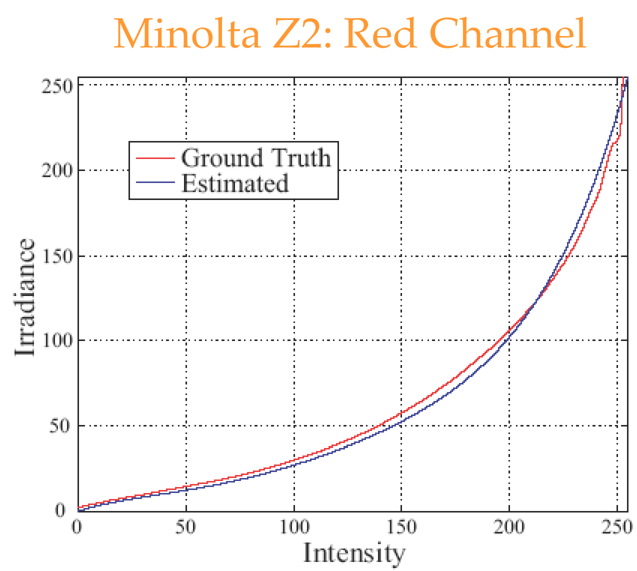

Estimating Camera Response Function

	Sony W1		Canon G5		Casio Z120		Minolta Z2	
	RMS %	Mean %	RMS %	Mean %	RMS %	Mean %	RMS %	Mean %
Red	1.344	1.131	1.759	1.554	2.269	1.176	2.226	1.701
Green	1.993	1.498	0.865	0.748	2.521	2.071	2.743	2.011
Blue	1.164	0.771	2.523	2.244	2.051	1.403	2.653	2.200

We need ~ 50 photographs to get estimates with RMS Error ~ 2%

Radiometric Camera Properties

- Properties of specific *camera models*
  - Camera response function
  - Vignetting for different lens settings

- Properties of specific *camera instances*
  - Bad pixels on the detector

Related Work: Vignetting

- Integrating sphere

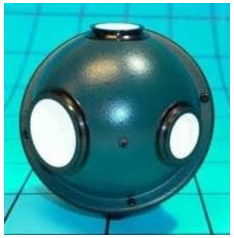

- Multiple images
  - Known illuminant at different image locations
    Stumpfel et al. '04
  - Overlapping images of an arbitrary scene
    Goldman & Chen '05, Litvinov & Schechner '05, Jia & Tang '05
- Single image
  - Iterative segmentation and vignetting estimation
    Zheng et al. '06
  - Distribution of radial gradients
    Zheng et al. '08

- Fully automatic, robust linear estimation
- Do not need access to the camera

Torralba & Oliva. '02

Newlyweds

The Graduate

Salavon

What does the average of a group of photographs
with the same lens setting look like?

Canon S1IS Focal Length: 5.8 mm F-Number: 4.5 195/15,550 Images

Images are linearized and have no vignetting

Spatial Distribution of Average Luminances

Average Log(Luminance) of 15,500 Images
Canon S1 IS
Focal Length: 5.8 mm
F-Number: 4.5

Average Log(Luminance) of 13,874 Images
Canon S1 IS
Focal Length: 5.8 mm
F-Number: 2.8

ed out particula

Spatial Distribution of Average Luminances

Prior

- Have a vertical gradient
- No horizontal gradient

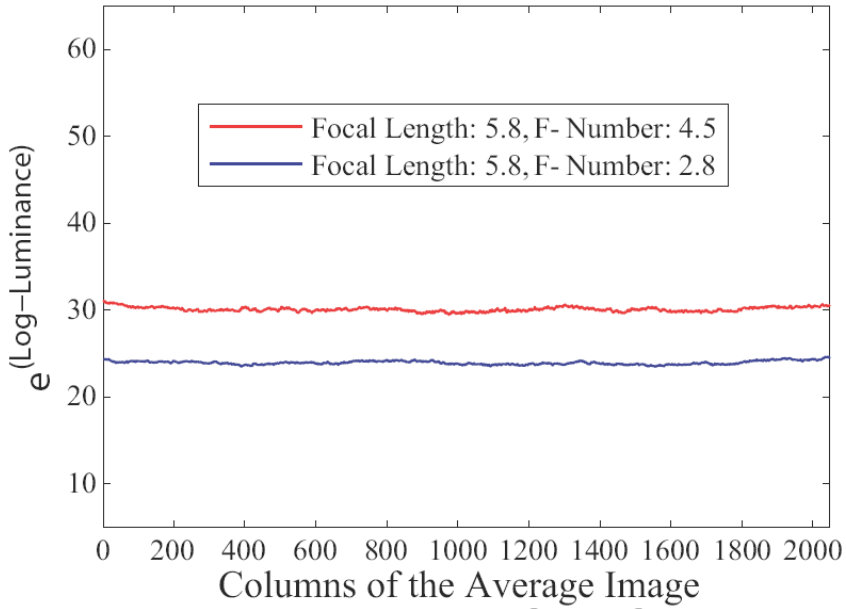

Focal Length: 5.8 mm
F-Number: 4.5

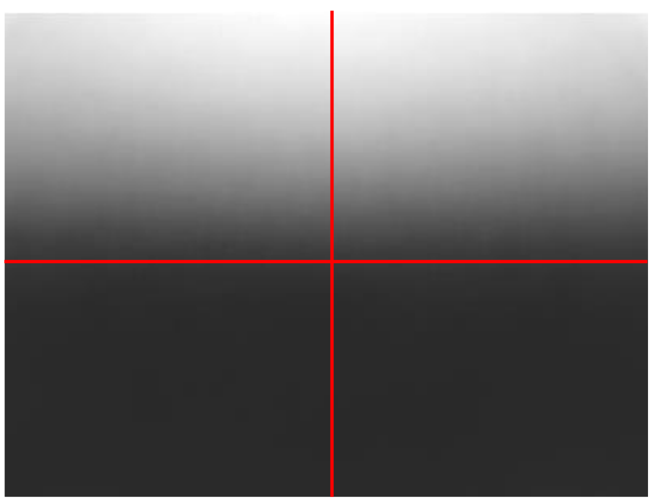

Focal Length: 5.8 mm
F-Number: 2.8

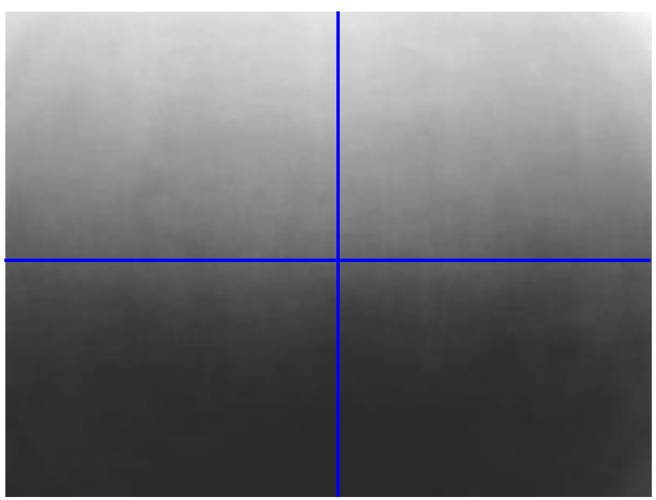

Estimating Vignetting for a Lens Setting

- Use estimated response function to linearize images
- Compute average log-luminance image

Prior: In the *absence of vignetting*, average log-luminance image

  - Has a vertical gradient
  - No horizontal gradient

What if photographs *have vignetting*?

Estimating Vignetting for a Lens Setting

Average Log(Luminance) of 15,500 Images
Canon S1 IS
Focal Length: 5.8 mm
F-Number: 4.5

Average Log(Luminance) of 13,874 Images
Canon S1 IS
Focal Length: 5.8 mm
F-Number: 2.8

Estimating Vignetting for a Lens Setting

$$m_i(x,y) =$$

Measured image luminance

Vignetting

Image luminance when no vignetting

$$\frac{1}{N}\sum_i \log(m_i(x,y)) = \log(v(x,y)) + \log(L_i(x,y))$$

$$\frac{1}{N}\sum_i \log(l_i(x,y))$$

N = Number of photographs

$$M(x,y)$$

Known

Unknown

Unknown

$$M(x,y) = V(x,y) + L(y)$$

Prior: $L(x,y)$

- Has a vertical gradient
- No horizontal gradient

$$V(x,y) = \sum_{k}^{D} \beta_k r^k$$

r = radial distance to (x,y)

β= Polynomial coefficients

D = Polynomial degree (9)

$$M(x,y) = \sum_{k}^{D} \boxed{\beta_k} r^k + \boxed{L(y)}$$

Estimate Vignetting Linearly

Estimating Vignetting for a Lens Setting

Average Log(Luminance) of 15,500 Images
Canon S1 IS
Focal Length: 5.8 mm
F-Number: 4.5

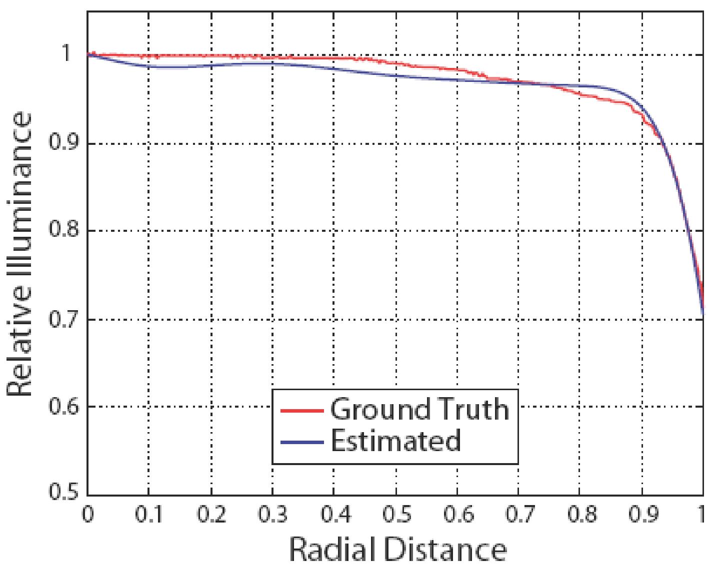

Average Log(Luminance) of 13,874 Images
Canon S1 IS
Focal Length: 5.8 mm
F-Number: 2.8

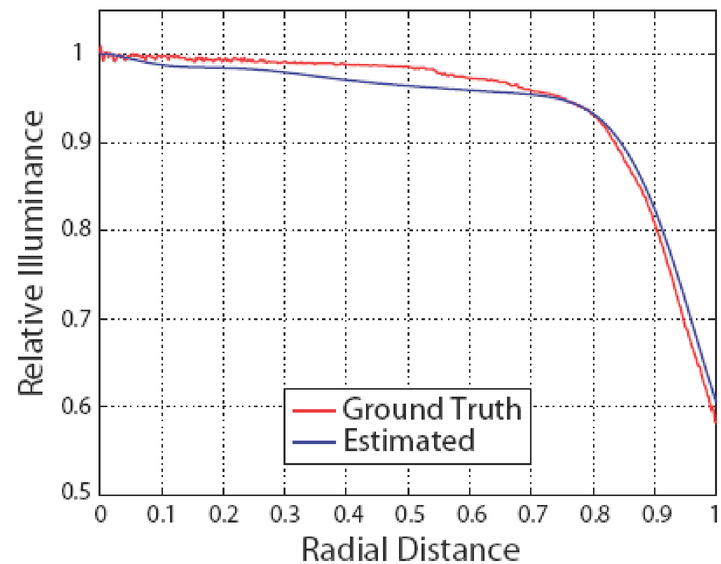

Vignetting Estimation Results

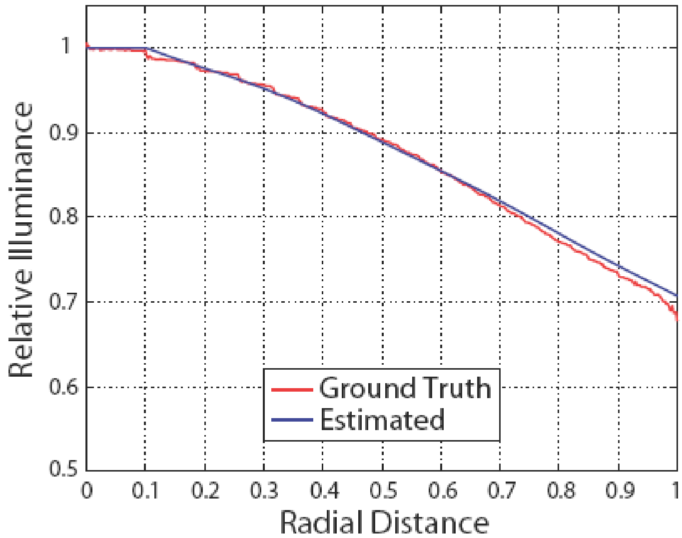

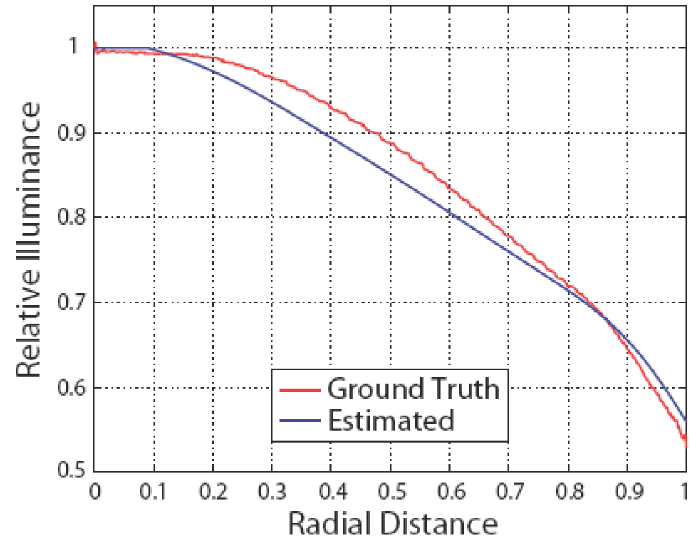

Sony W1 Focal Length: 7.9 mm, F-Number: 5.6 Focal Length: 7.9 mm, F-Number: 2.8

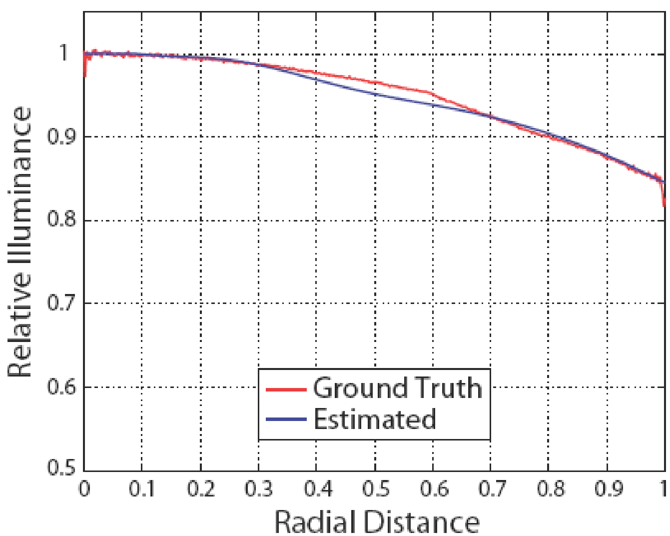

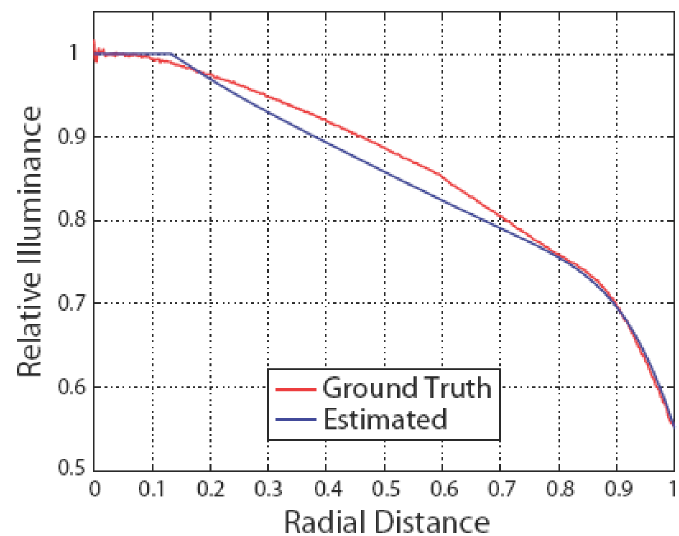

Canon G5 Focal Length: 7.2 mm, F-Number: 4 Focal Length: 7.2 mm, F-Number: 2

Vignetting Estimation Error

	Canon S1IS Focal Length: 5.8 mm		Sony W1 Focal Length: 7.9 mm		Canon G5 Focal Length: 7.18 mm	
	F/4.5	F/2.8	F/5.6	F/2.8	F/4.0	F/2.0
RMS %	0.989	1.399	0.594	2.324	0.664	1.723
Mean %	0.895	1.221	0.460	1.980	0.484	1.398
Median %	0.990	1.183	0.317	1.966	0.296	1.084

We need ~ 3000 photographs to get estimates with RMS Error ~ 2%

Estimated Vignetting

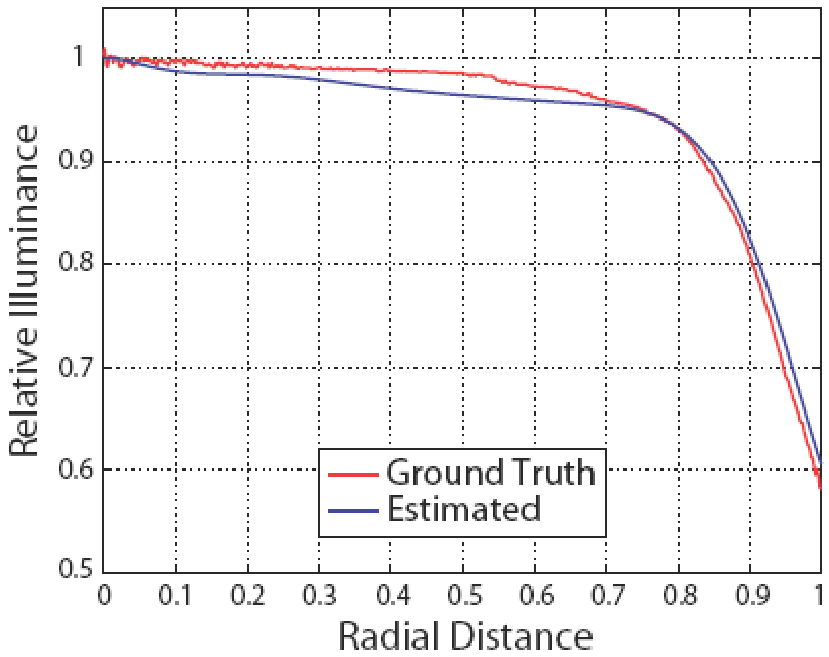

Canon S1IS

Focal Length: 5.8 mm, F-Number: 2.8

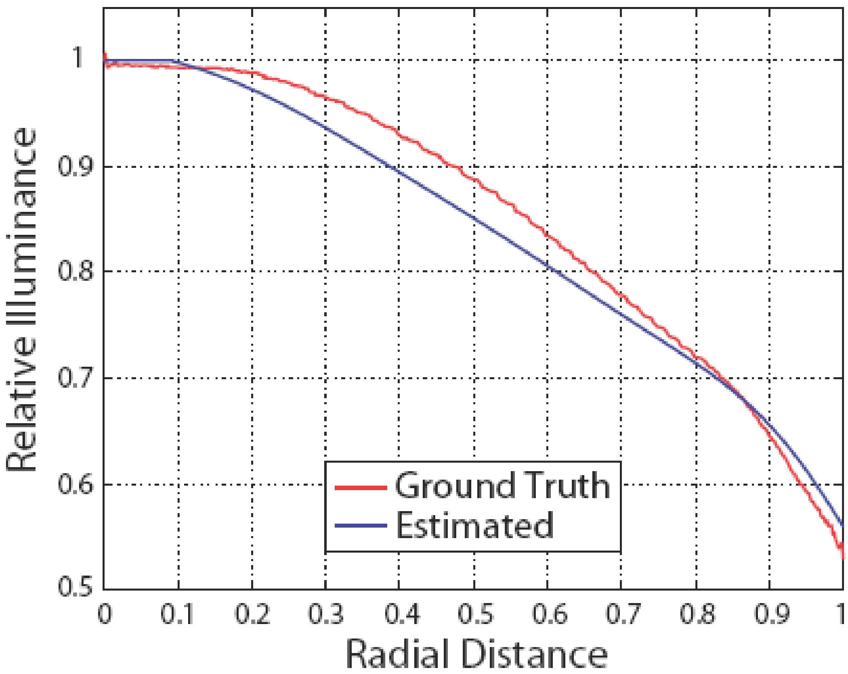

Sony W1

Focal Length: 7.9 mm, F-Number: 2.8

Canon S1IS Focal Length: 5.8 mm, F-Number: 2.8

Canon S1IS Focal Length: 5.8 mm, F-Number: 2.8

w/ Vignetting Correction

Radiometric Camera Properties

- Properties of specific *camera models*
  - Camera response function
  - Vignetting for different lens settings

- Properties of specific *camera instances*
  - Bad pixels on the detector

Identifying Bad Pixels on a Camera Detector

Group images by camera instance (Flickr username)

Prior: Average image should be smooth

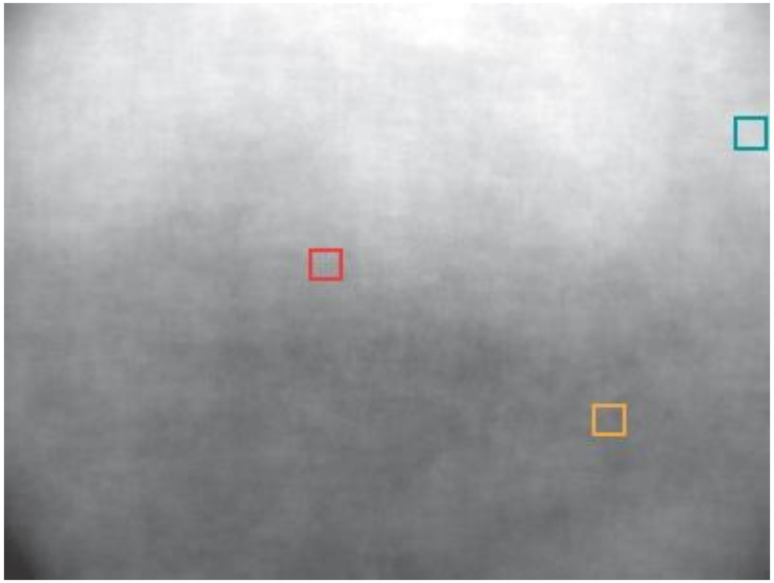

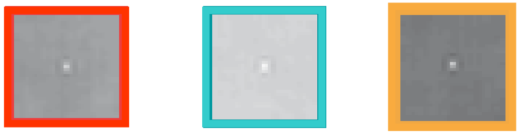

Camera Model	# of Cameras	Mean Defects	Median Defects
Canon G5	15	2.2	1
Canon SD 300	13	1.08	1
Sony W1	13	0.384	0

Priors and Camera Properties

Robust Statistical Priors

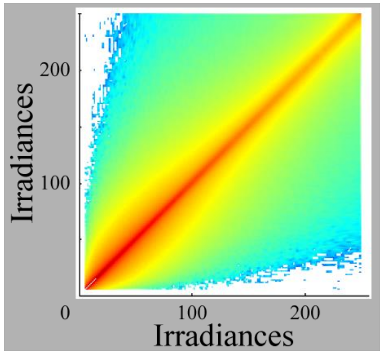

Joint Histogram of Irradiances

Spatial Distribution of Average Luminances

Recover Radiometric Camera Properties

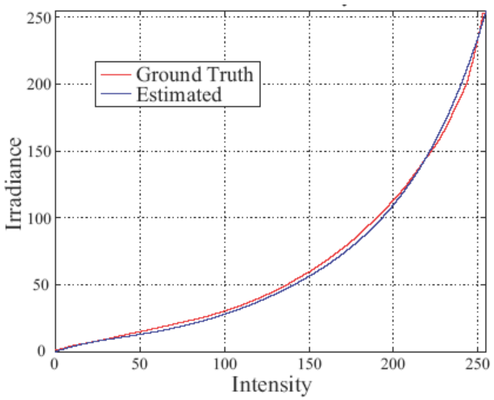

Response Function

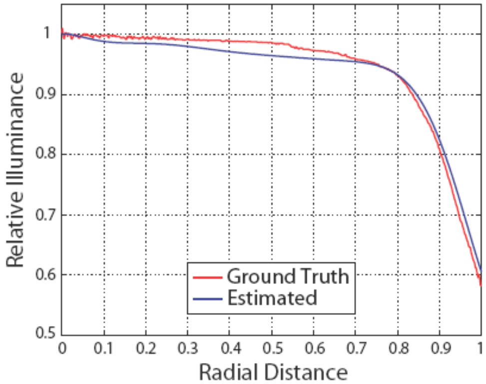

Vignetting

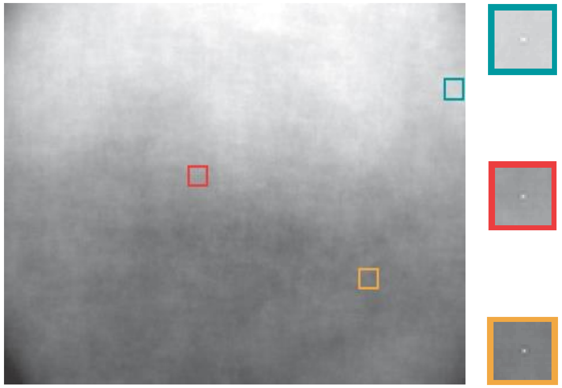

Bad Pixels

Discussion

- Need a large number of photographs

- Fully automatic , do not need access to camera

- Other priors
  - Distribution of gradients
  - Statistics of Fourier coefficients
  - Higher order joint statistics

Other camera properties
  - Radial distortion
  - Chromatic aberration
  - Varying lens softness

- Database of camera properties
  - Fully automatic
  - Zero cost

PTLens, DxO

- Information about scenes and photographers

Priors and Camera Properties

Robust Statistical Priors

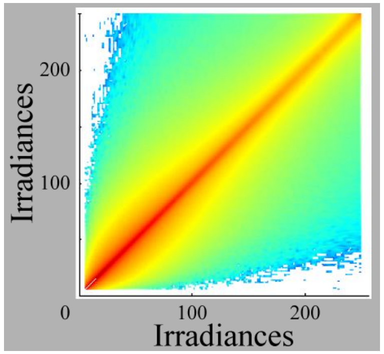

Joint Histogram of Irradiances

Spatial Distribution of Average Luminances

Recover Radiometric Camera Properties

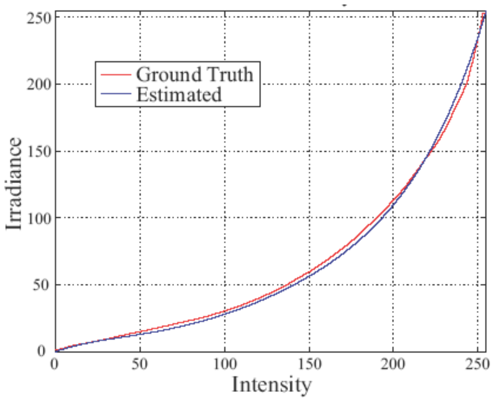

Response Function

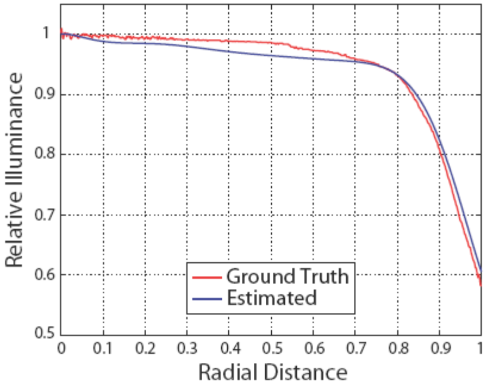

Vignetting

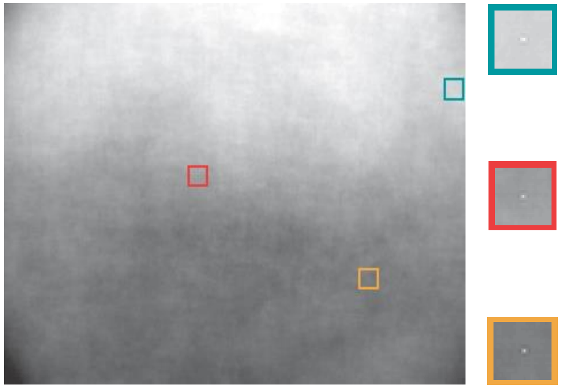

Bad Pixels

Estimating Camera Response Function

	Sony W1		Canon G5		Casio Z120		Minolta Z2	
	Our	Lin et al.	Our	Lin et al.	Our	Lin et al.	Our	Lin et al.
Red	1.344	2.587	1.759	2.553	2.269	1.518	2.226	4.914
Green	1.993	1.243	0.865	3.396	2.521	1.155	2.743	3.237
Blue	1.164	1.783	2.523	2.154	2.051	3.053	2.653	3.292

Vignetting in Canon S1IS Cameras

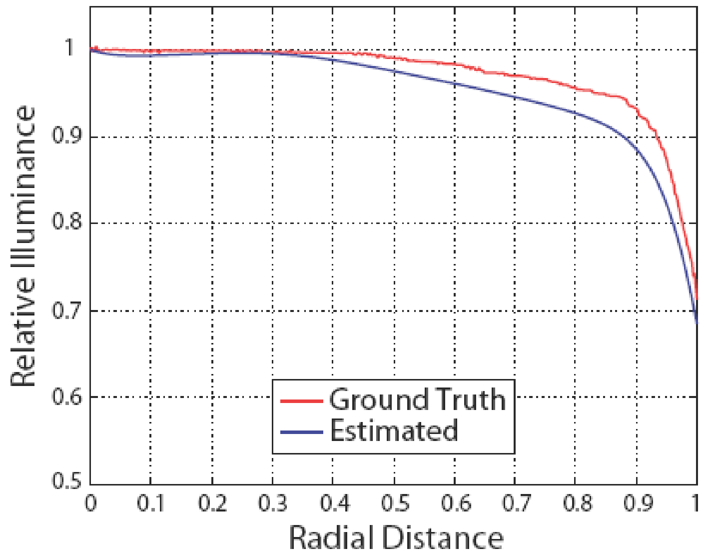

Focal Length: 5.8 mm, F-Number: 4.5

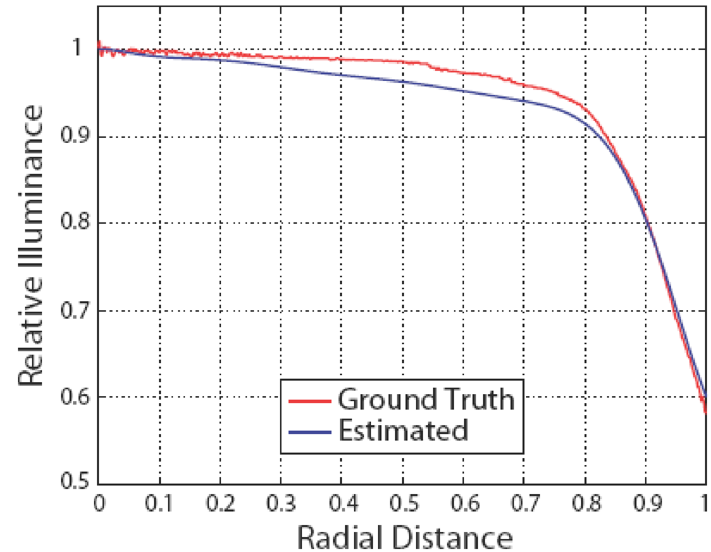

Focal Length: 5.8 mm, F-Number: 2.8

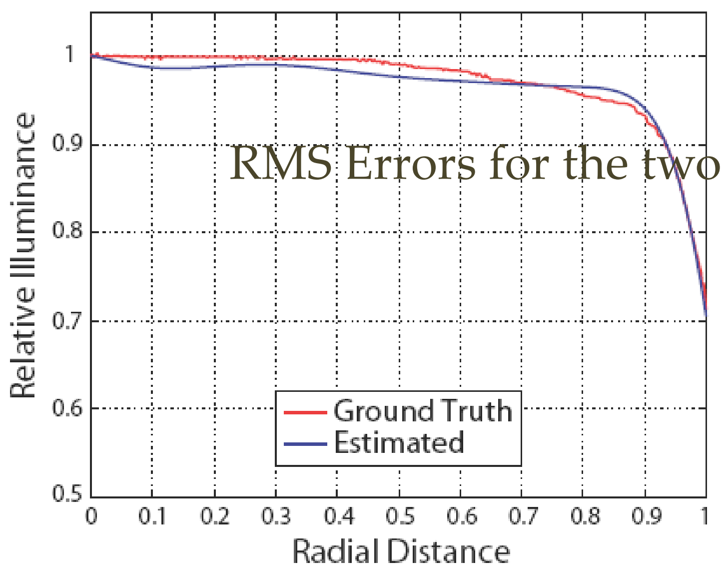

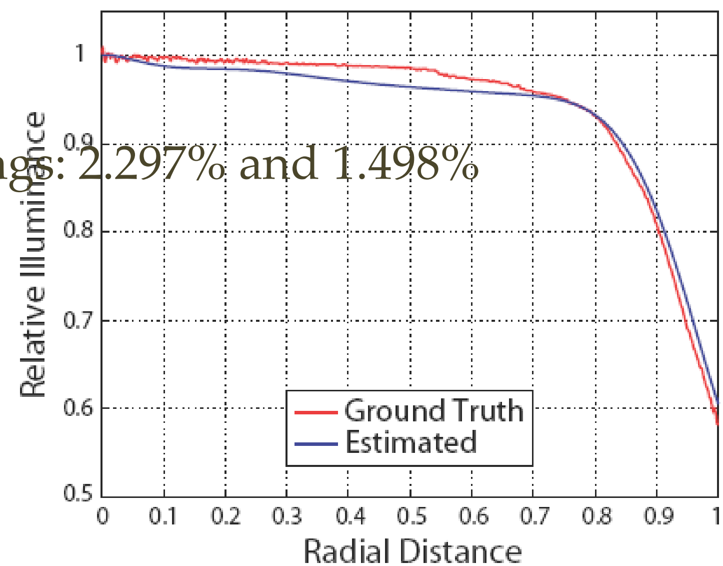

RMS Errors for the two settings: 2.297% and 1.498%

Estimating Camera Response Function

How many images do we need?

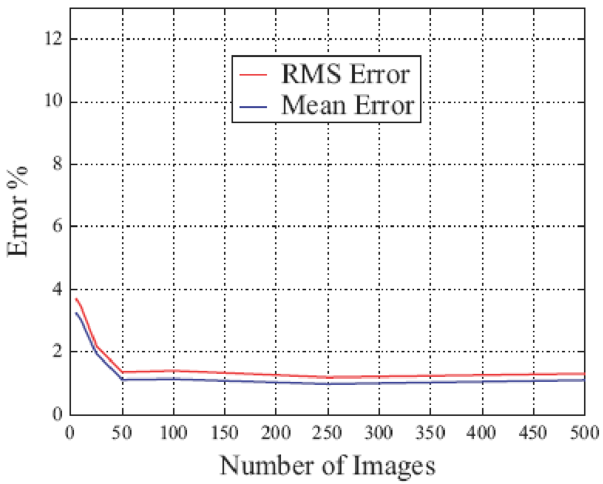

Sony W1: Red Channel

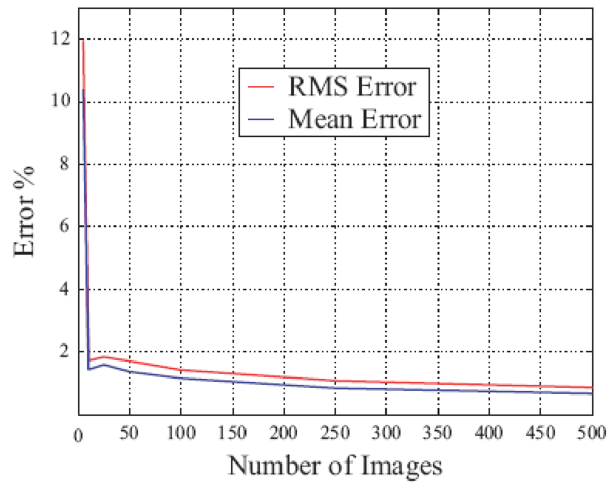

Canon G5: Green Channel

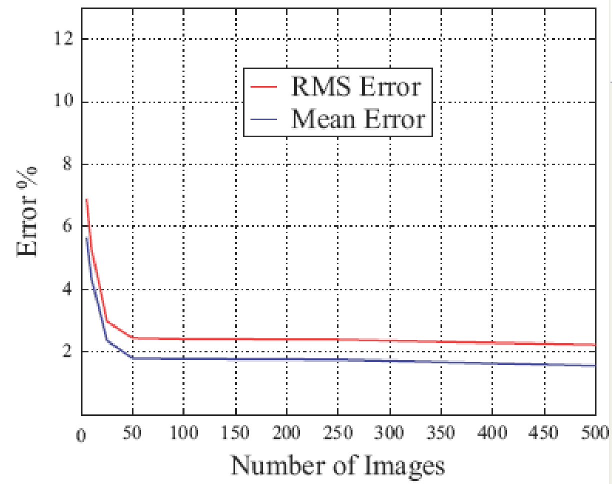

Casio Z120: Blue Channel

We need ~50 photographs to get estimates with RMS Error < 2%

Vignetting Estimation Results

How many images do we need?

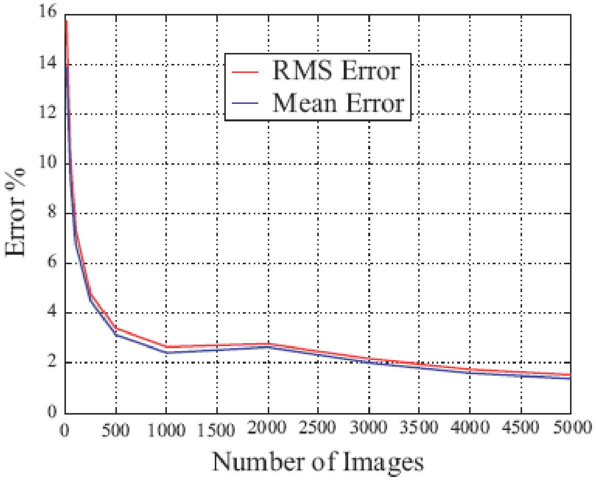

Canon S1IS
Focal Length: 5.8 mm
F Number: 4.5

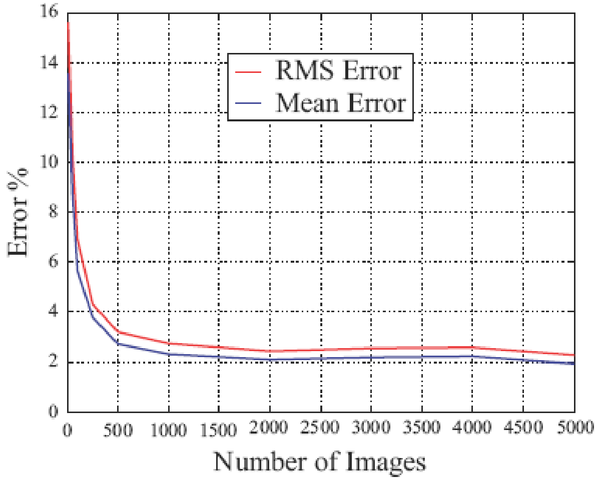

Sony W1
Focal Length: 7.9 mm
F Number: 2.8

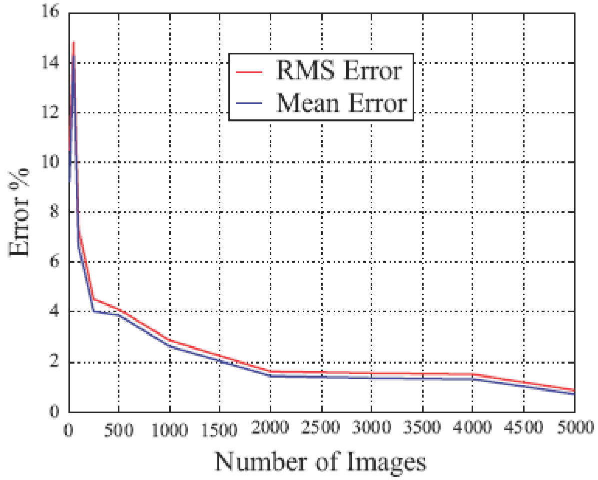

Canon G5
Focal Length: 7.2 mm
F Number: 4.0

We need ~3000 photographs to get estimates with RMS Error < 2%

Reference

- S. Kuthirummal et al., “Priors for large photo collections and what they reveal about cameras,” ECCV 2008.
- http://www1.cs.columbia.edu/~sujit/

- T. Mitsunaga et al., “Radiometric self calibration,” CVPR 1999.
- S. Lin et al., “Radiometric calibration using a single image,” CVPR 2004.

Oral: 3.8% (62), Overall: 31.6% (506)

Li Shen (MSRA), Ping Tan (Univ. of Singapore), Stephen Lin (MSRA)

International Conference on Computer Vision and Pattern Recognition (CVPR) 2008, Poster

INTRINSIC IMAGE DECOMPOSITION WITH NON-LOCAL TEXTURE CUES

Intrinsic images

Input = reflectance × illumination

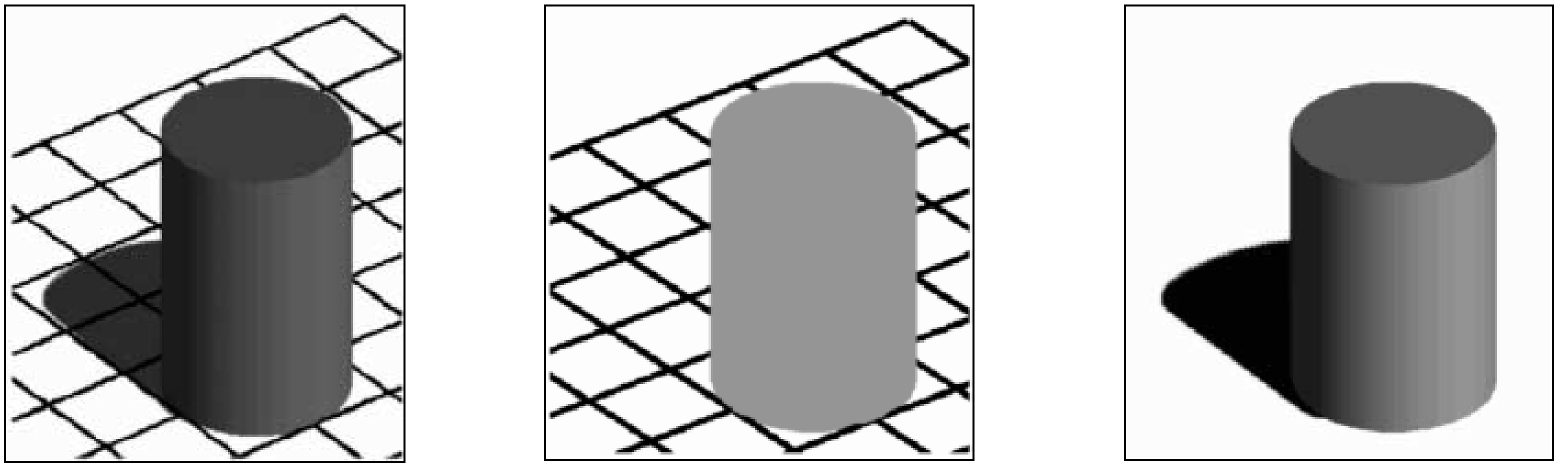

a b c

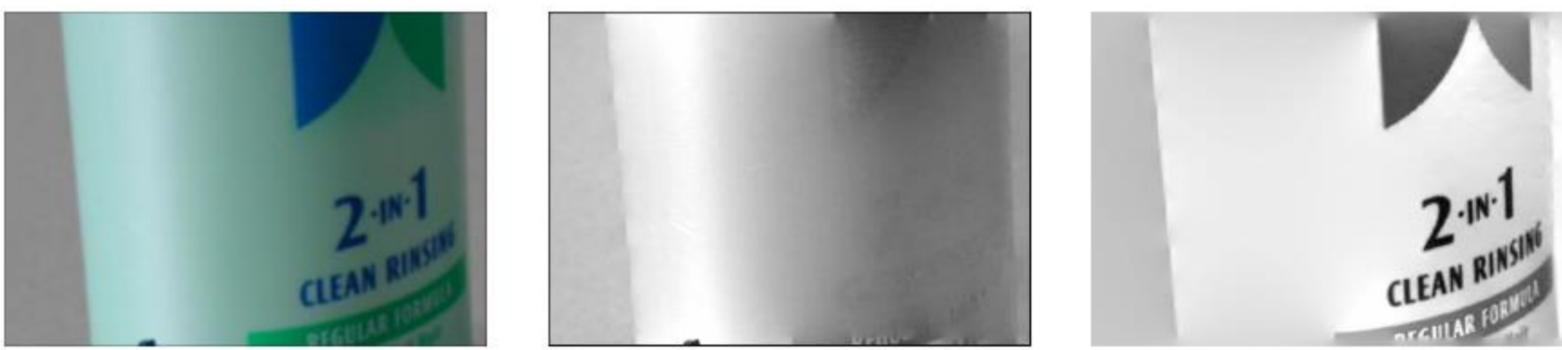

Original Image Shape Image Reflectance Image

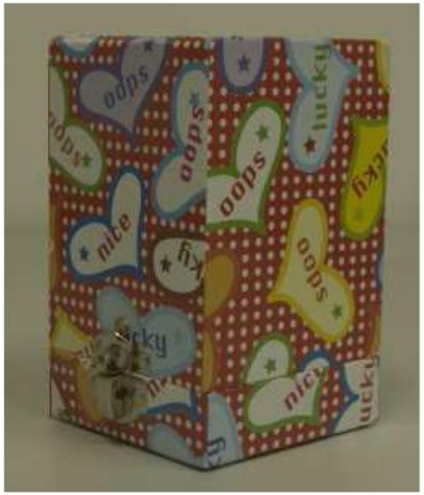

Result - Shading

Result - Reflectance

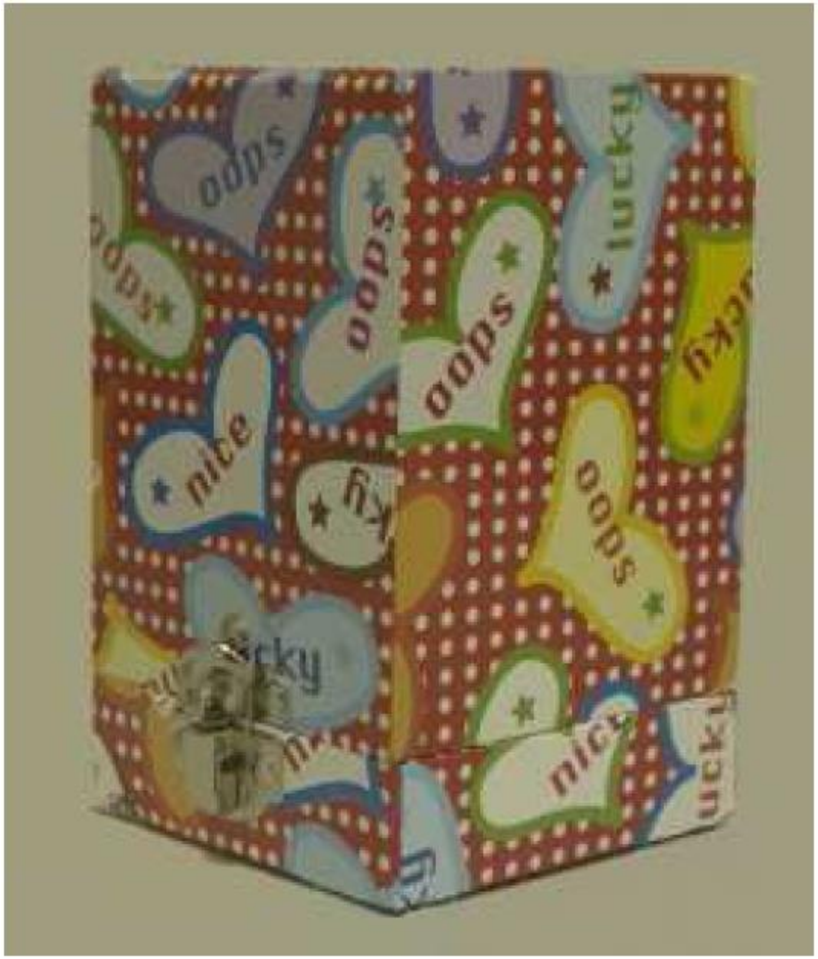

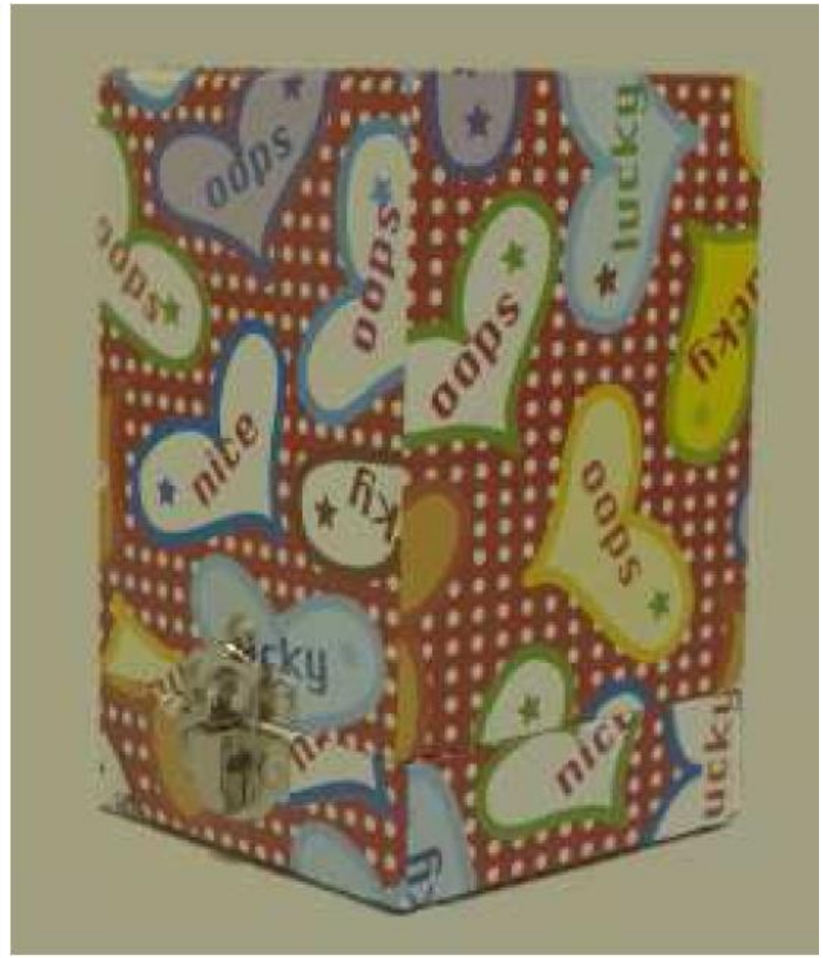

Result - Shading

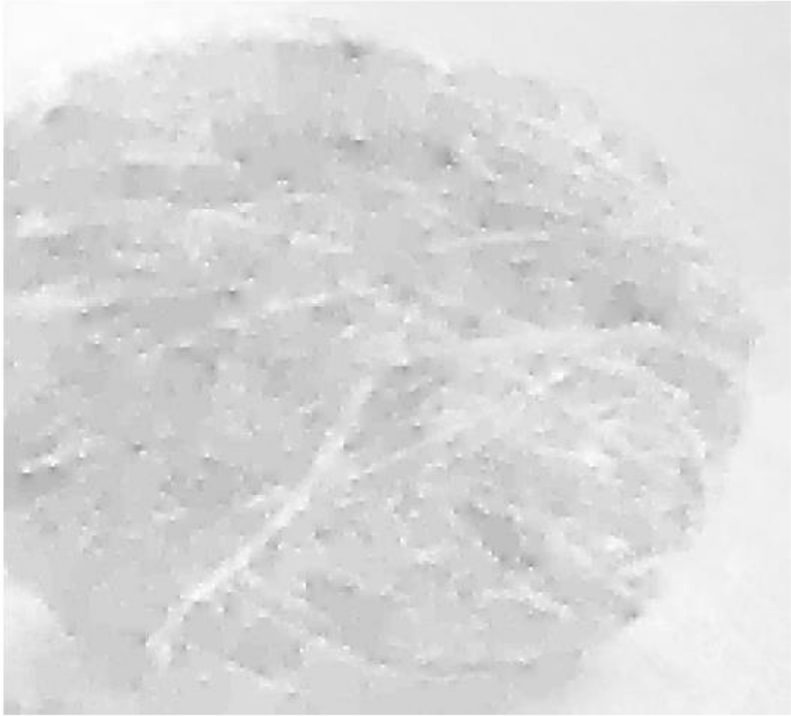

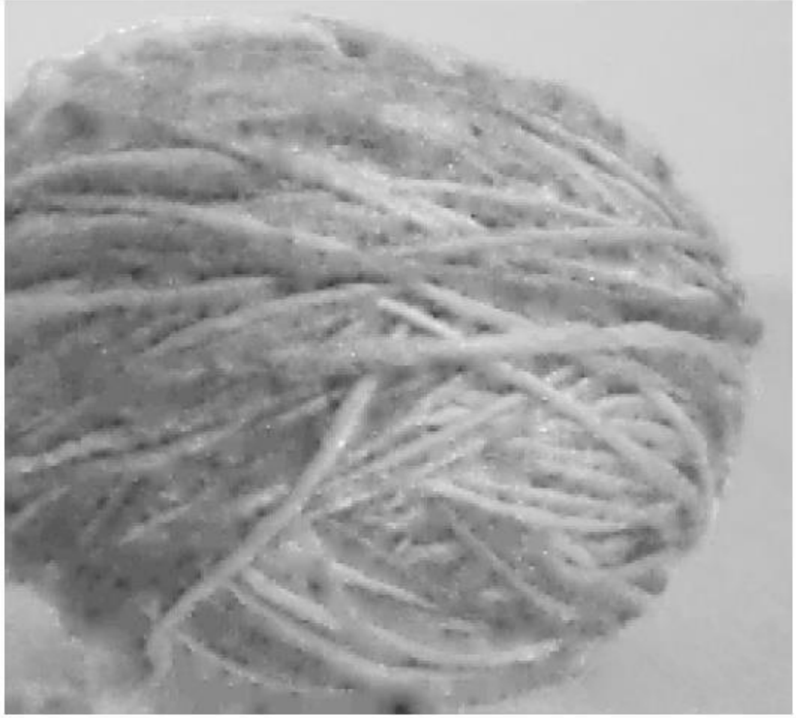

Result - Reflectance

Result - Shading

Result - reflectance

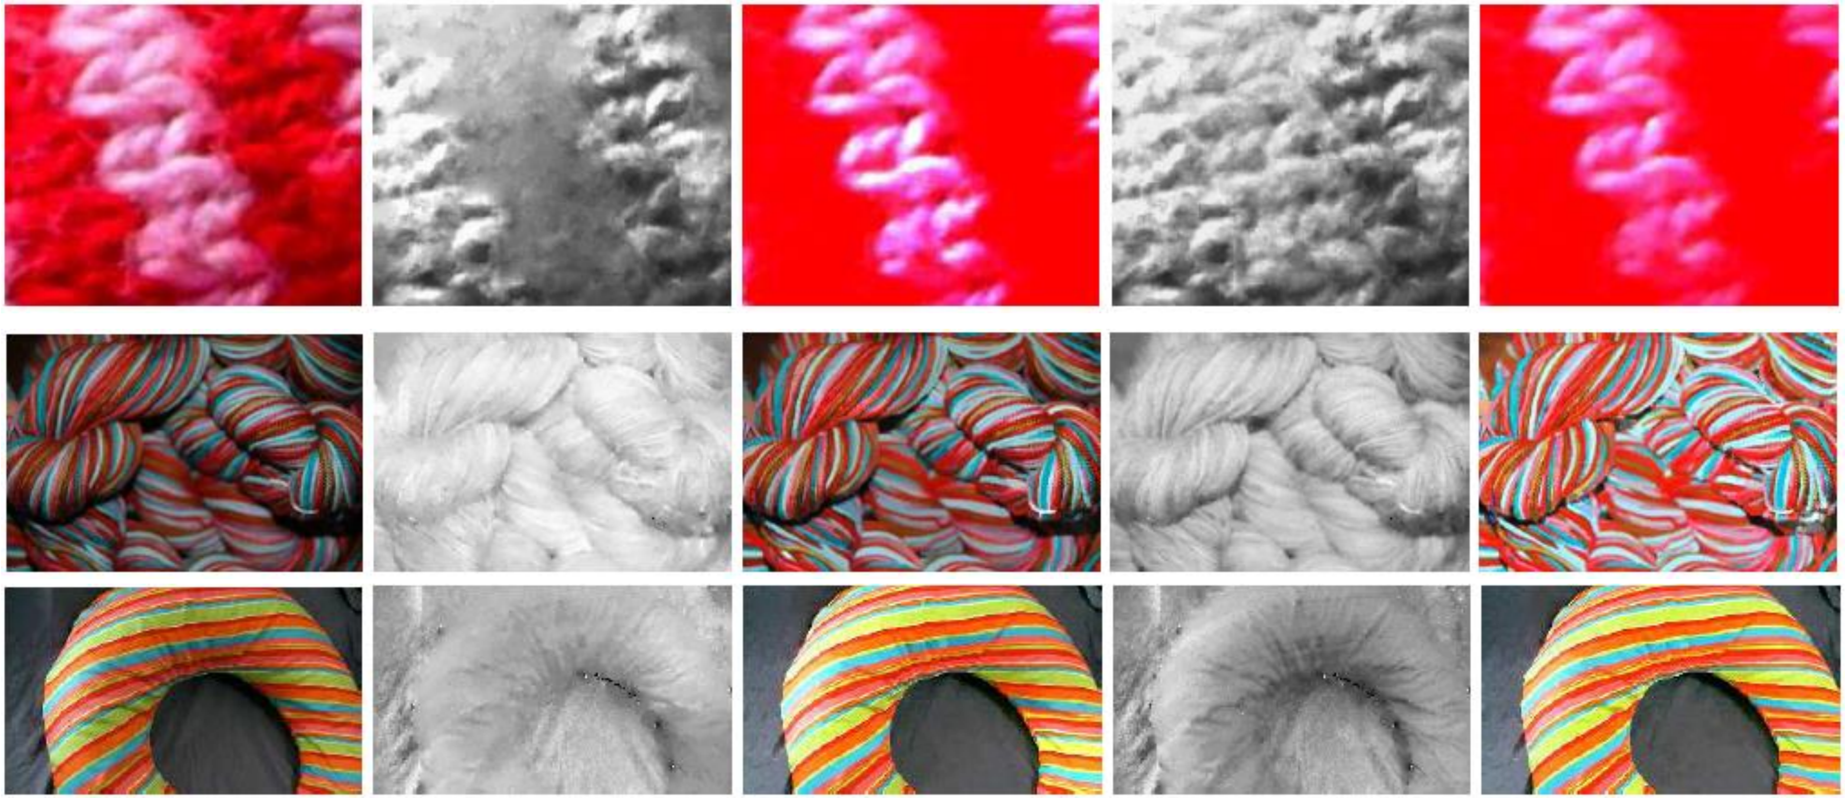

References

- L. Shen et al., “Intrinsic image decomposition with non-local texture cues,” CVPR 2008.

- M. F. Tappen et al., “Recovering intrinsic images from a single image,” TPAMI, 2005.
- Y. Weiss, “Deriving intrinsic images from image sequences,” ICCV 2001.