

From Pixels to Physics

Probabilistic color de-rendering



Ying Xiong



Kate Saenko



Trevor Darrell



Todd Zickler



Harvard University



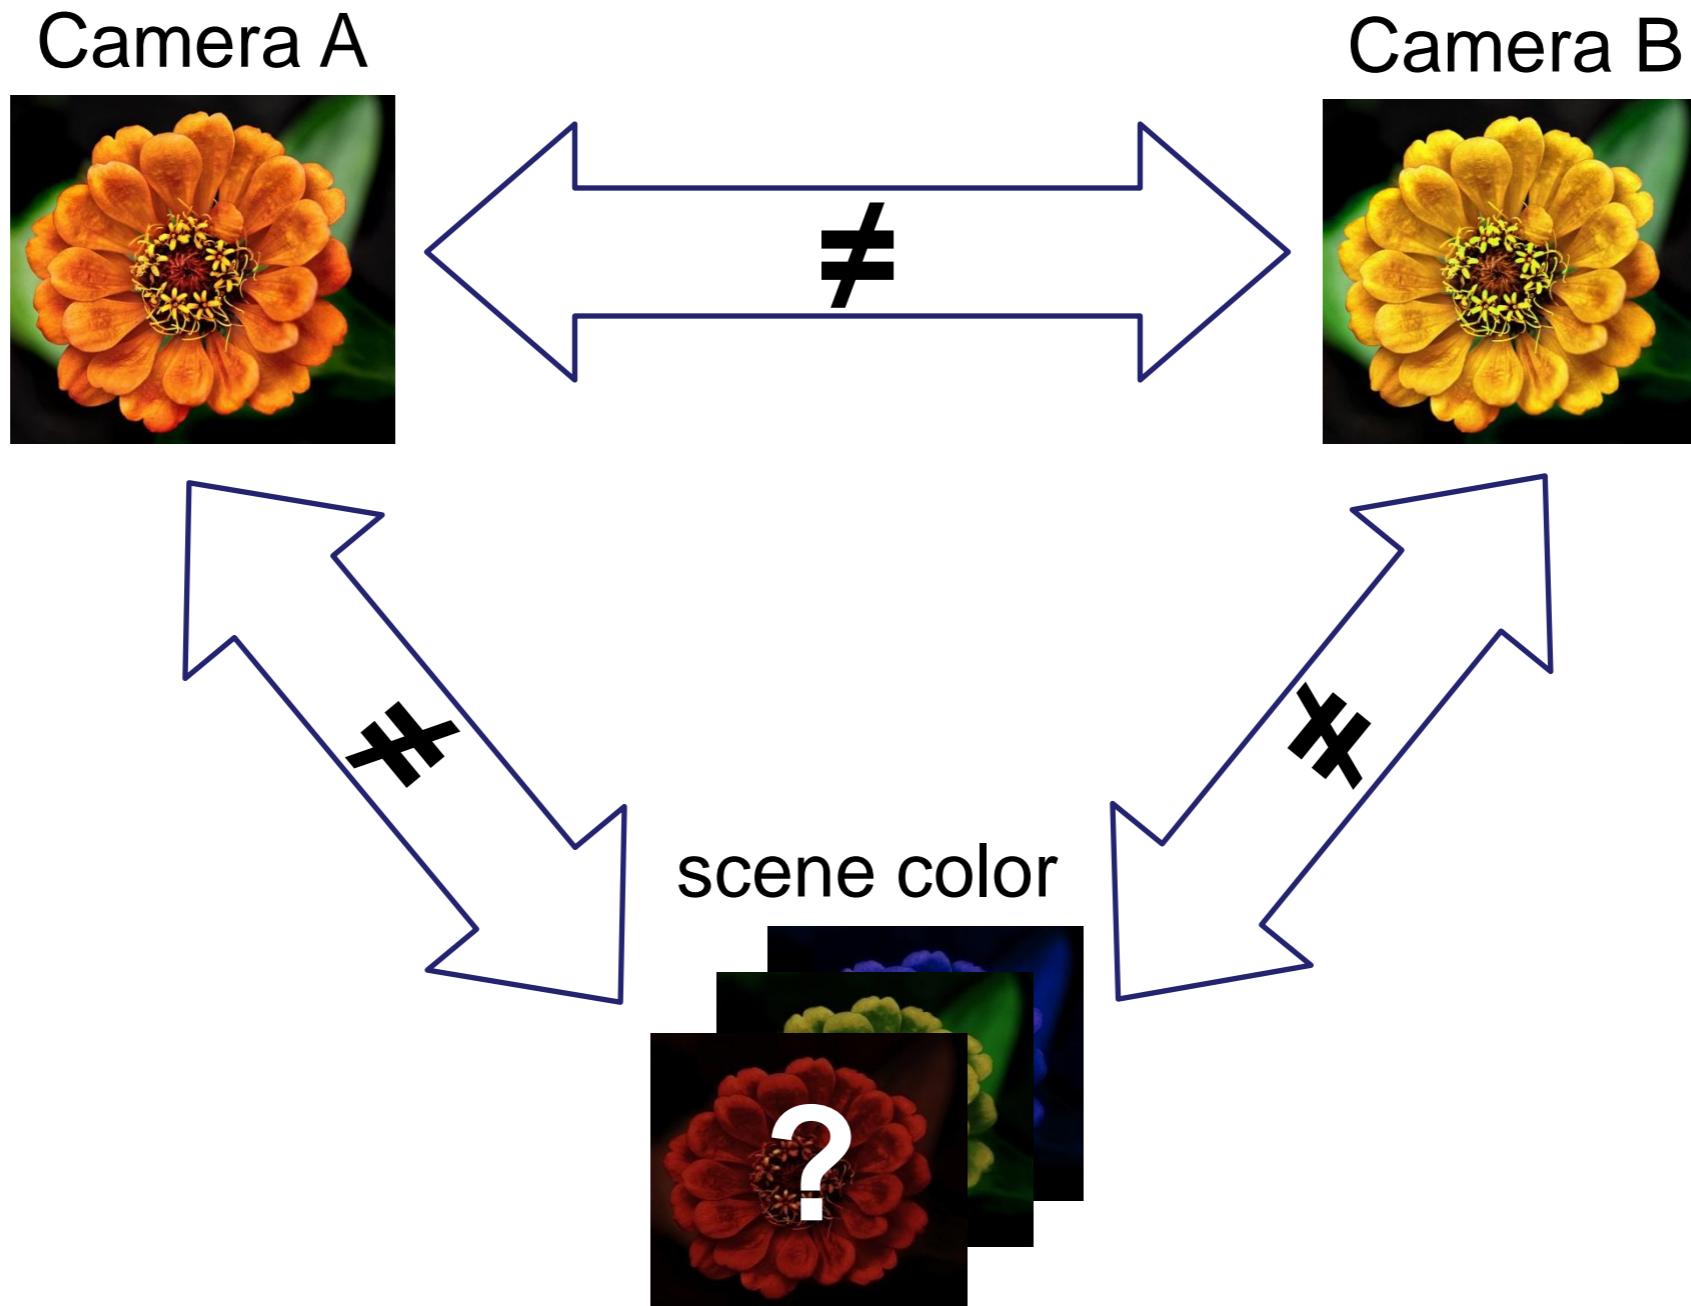
University of California, Berkeley

Colors in Internet vision



Photos courtesy of ImageNet

Colors in Internet vision



An approach: probabilistic derendering



rendered color values (ignoring
compression artifacts)

scene color

$$p\left(\begin{array}{c} \text{image} \\ \text{of flower} \end{array} \middle| \text{input image}\right)$$

values linearly pro-
portional to physical
scene colors

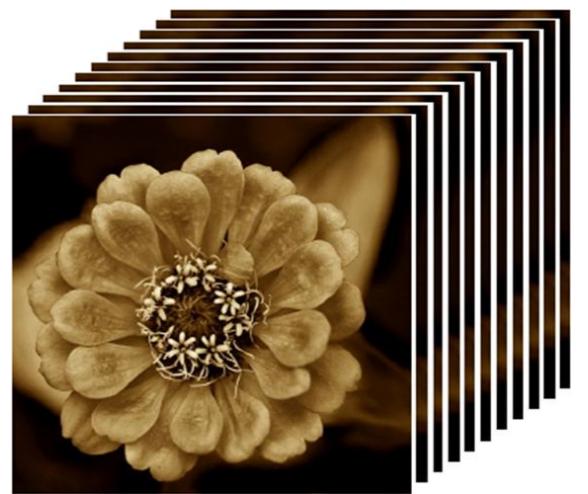
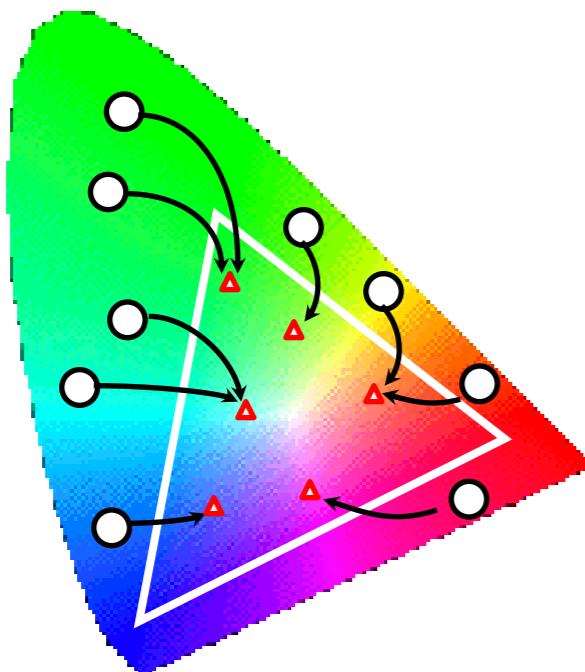
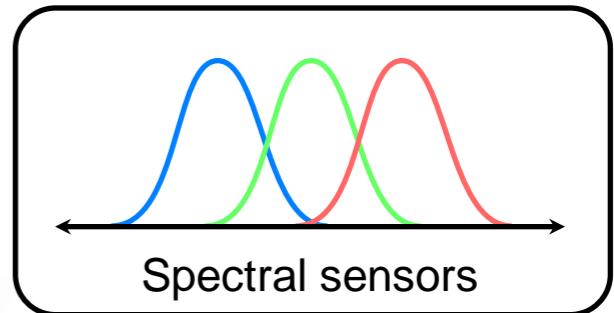
Problem: tone mapping is many-to-one

Rendered
Color



Tone
mapping

RAW



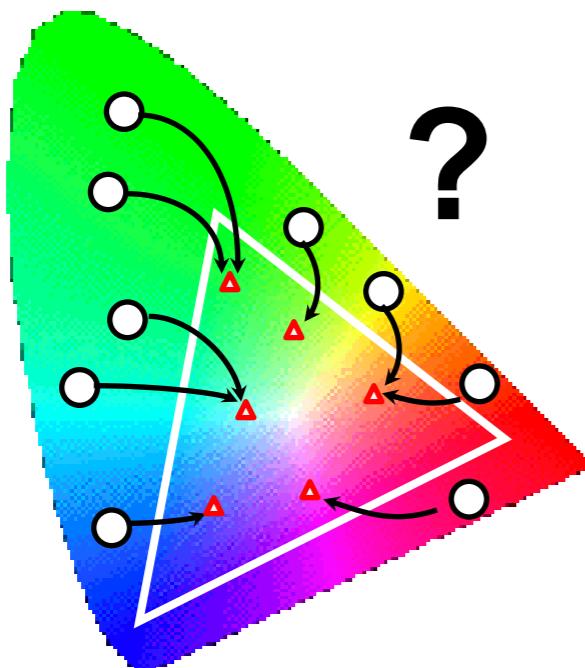
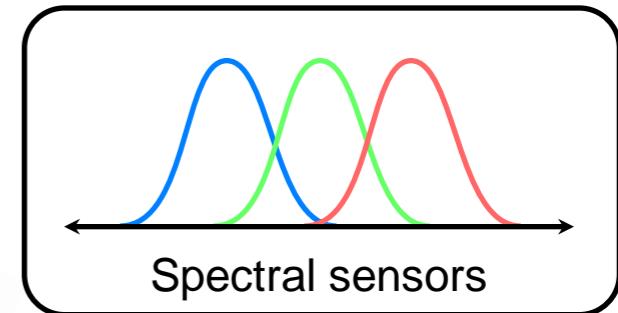
Spectral
Scene
Radiance

Problem: tone mapping is many-to-one

Rendered
Color



←
Tone
mapping



A Standard Default Color Space for the Internet - sRGB

Michael Stokes (Hewlett-Packard), Matthew Anderson (Microsoft),
Srinivasan Chandrasekar (Microsoft), Ricardo Motta (Hewlett-Packard)

Version 1.10, November 5, 1996

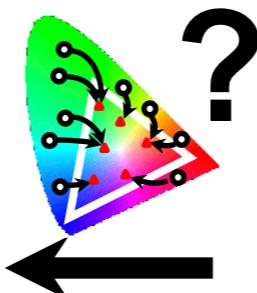
Introduction

Hewlett-Packard and Microsoft propose the addition of support for a standard color space, sRGB, within the Microsoft operating systems, HP products, the Internet, and all other interested vendors. The aim of this color space is to complement the current color management strategies by enabling a third method of handling color in the operating systems, device drivers

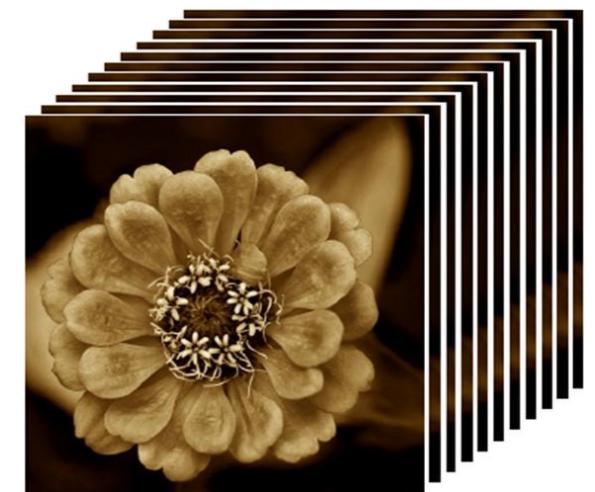
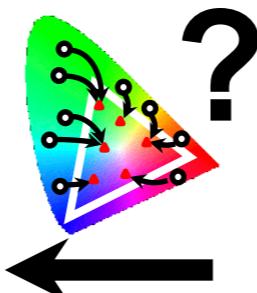


Problem: tone mapping is many-to-one

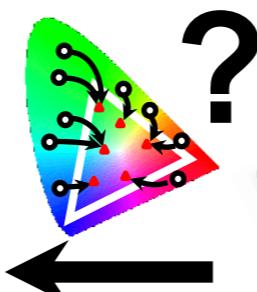
Rendered
Color



Rendered
Color



Rendered
Color



Spectral
Scene
Radiance

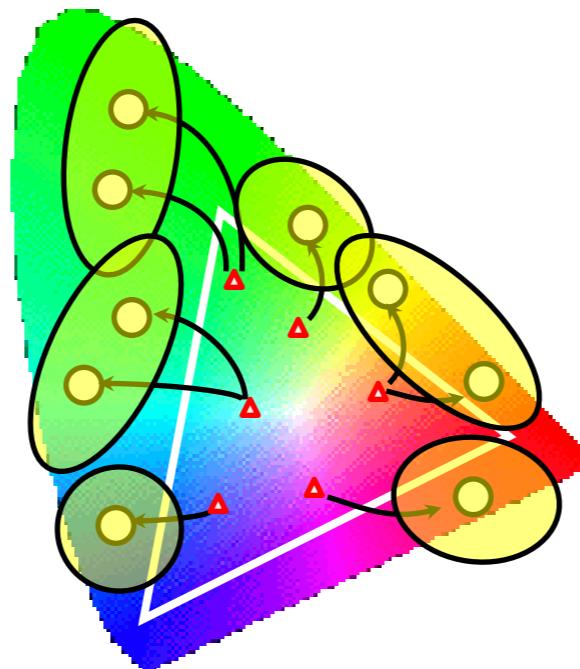
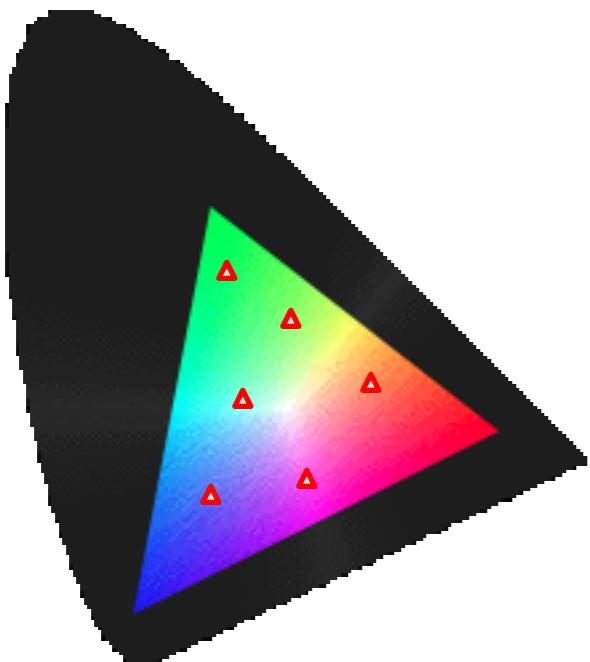
Our approach

Rendered
Color

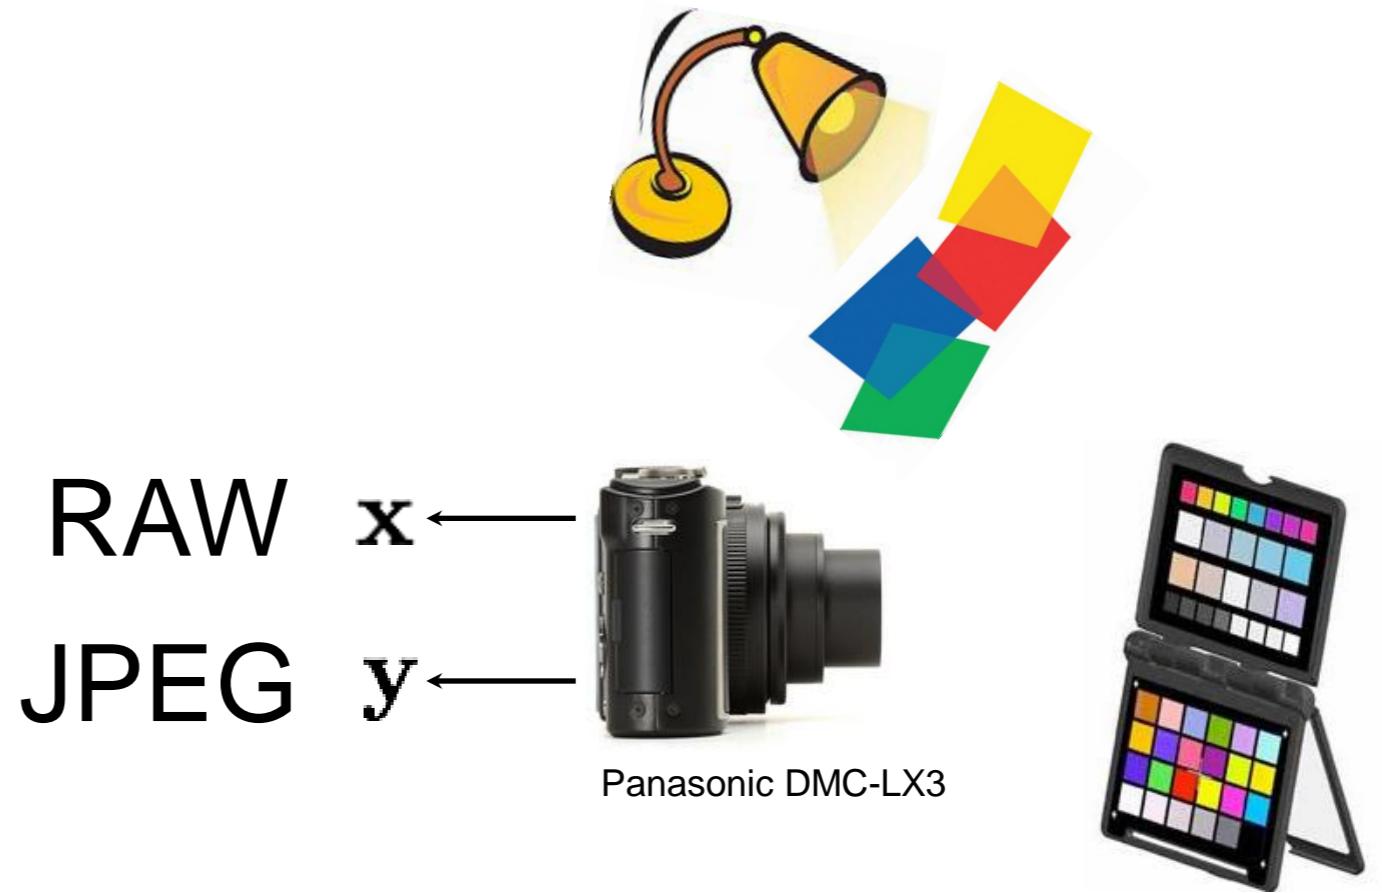


→
De-
rendering

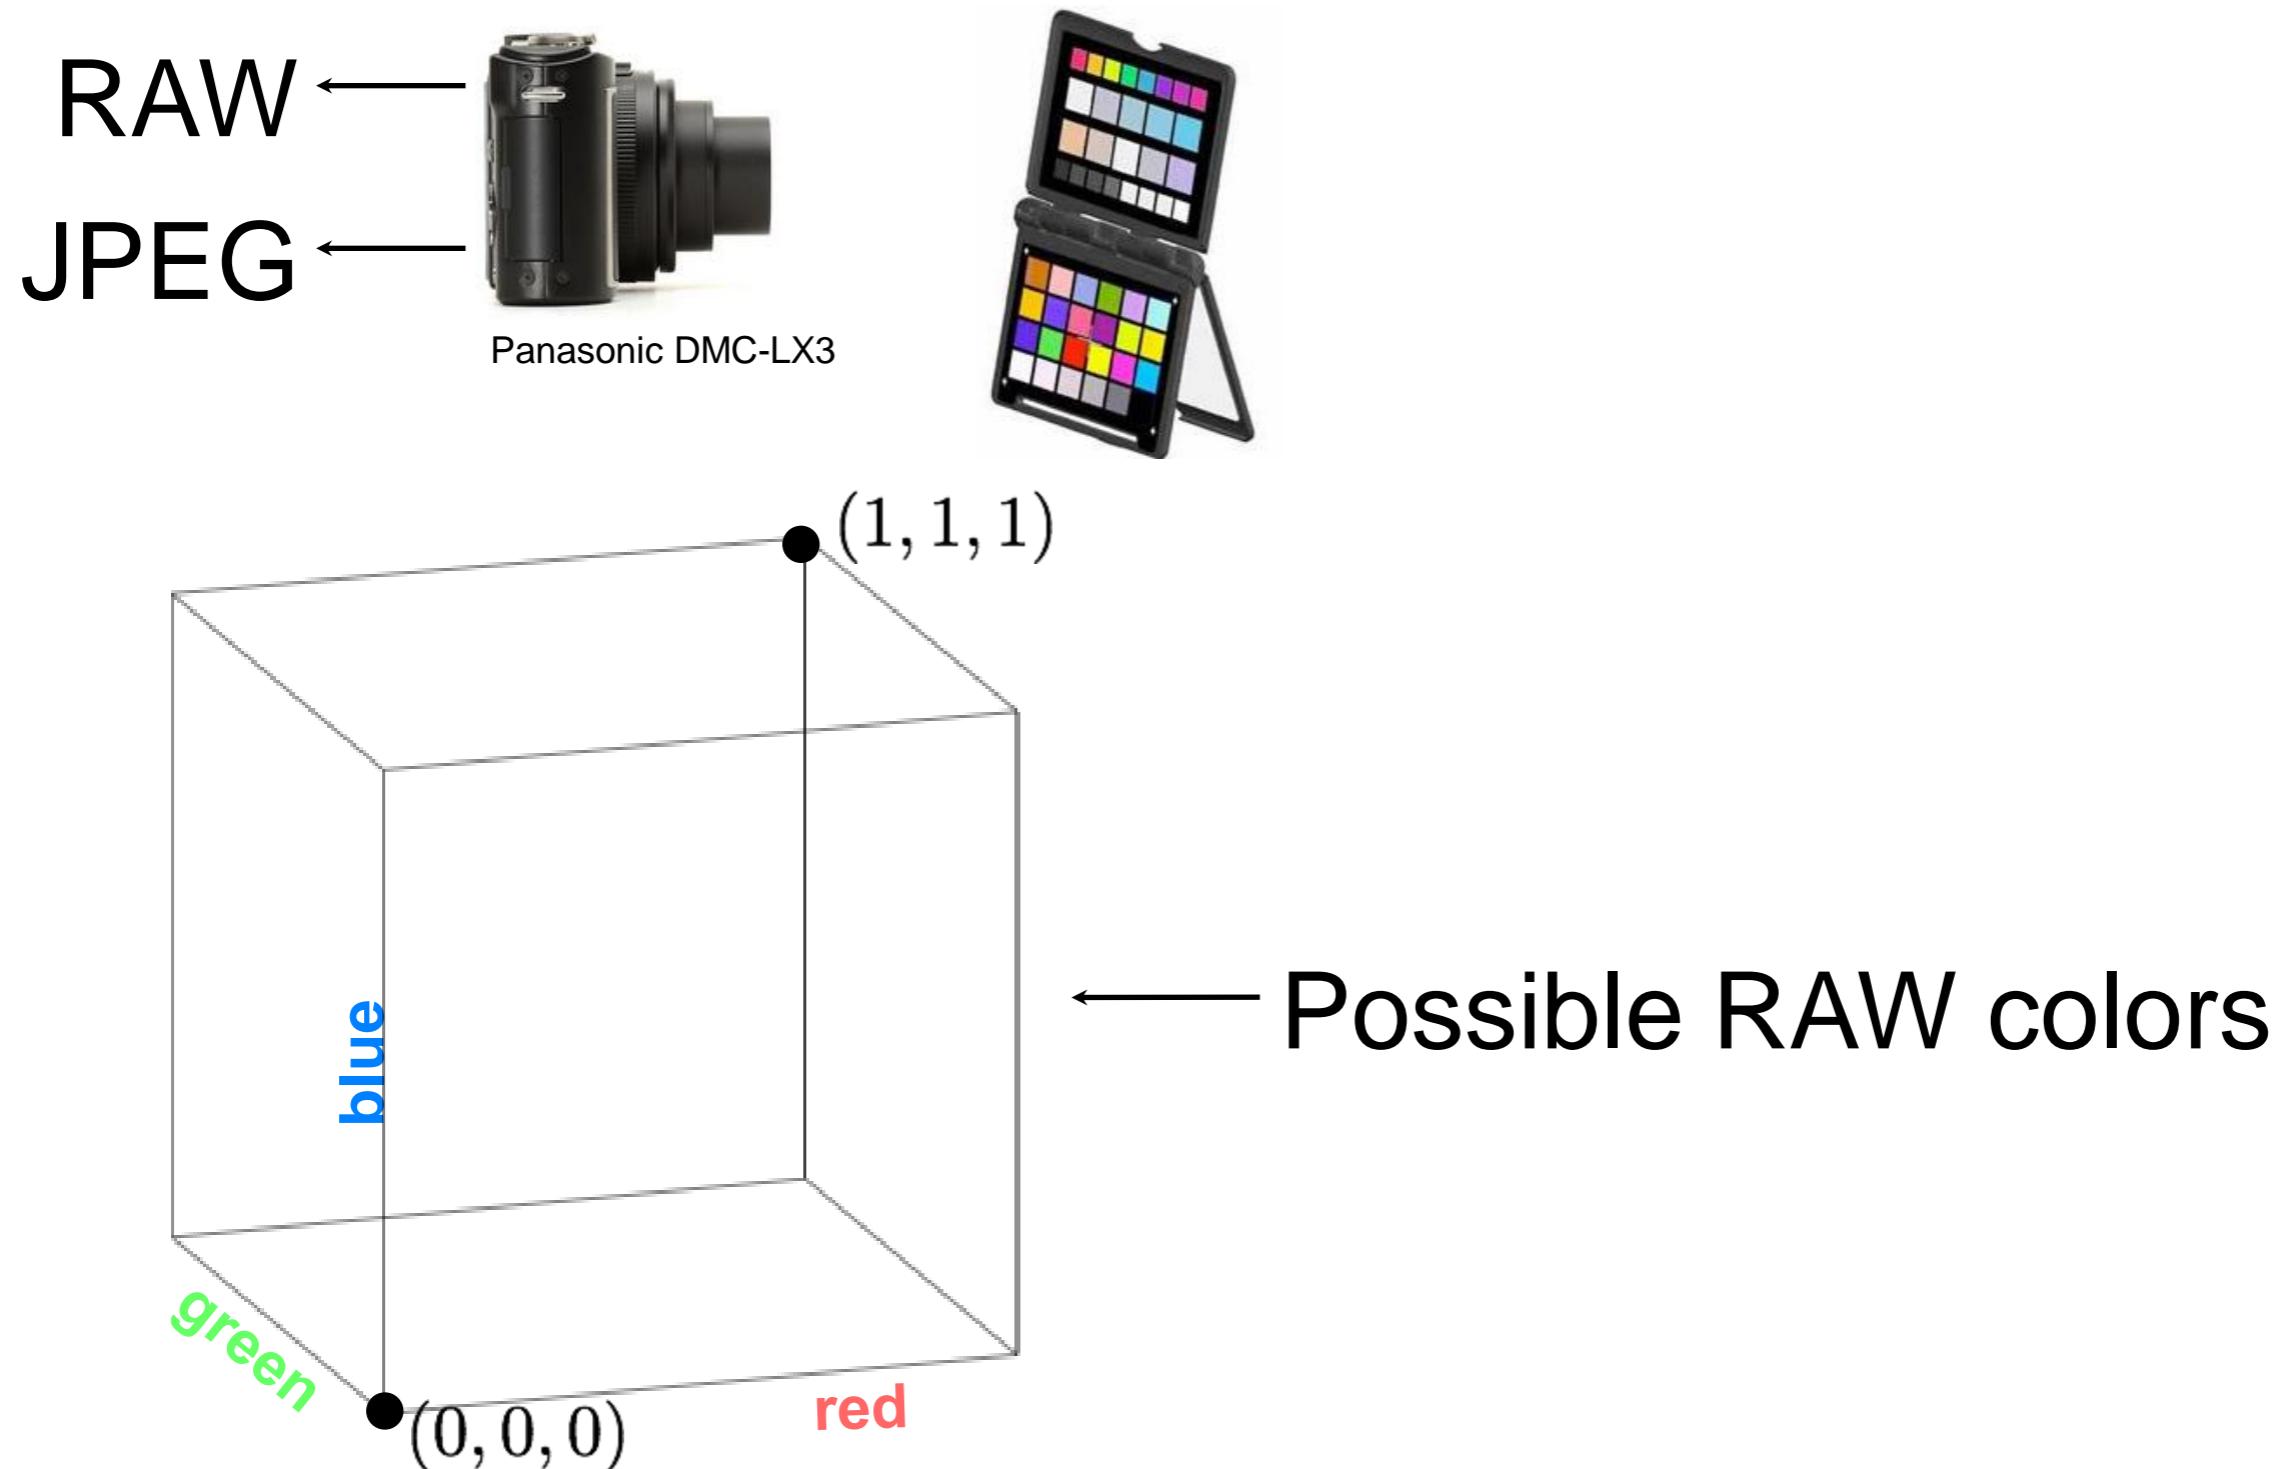
$P(\text{RAW} | \dots)$



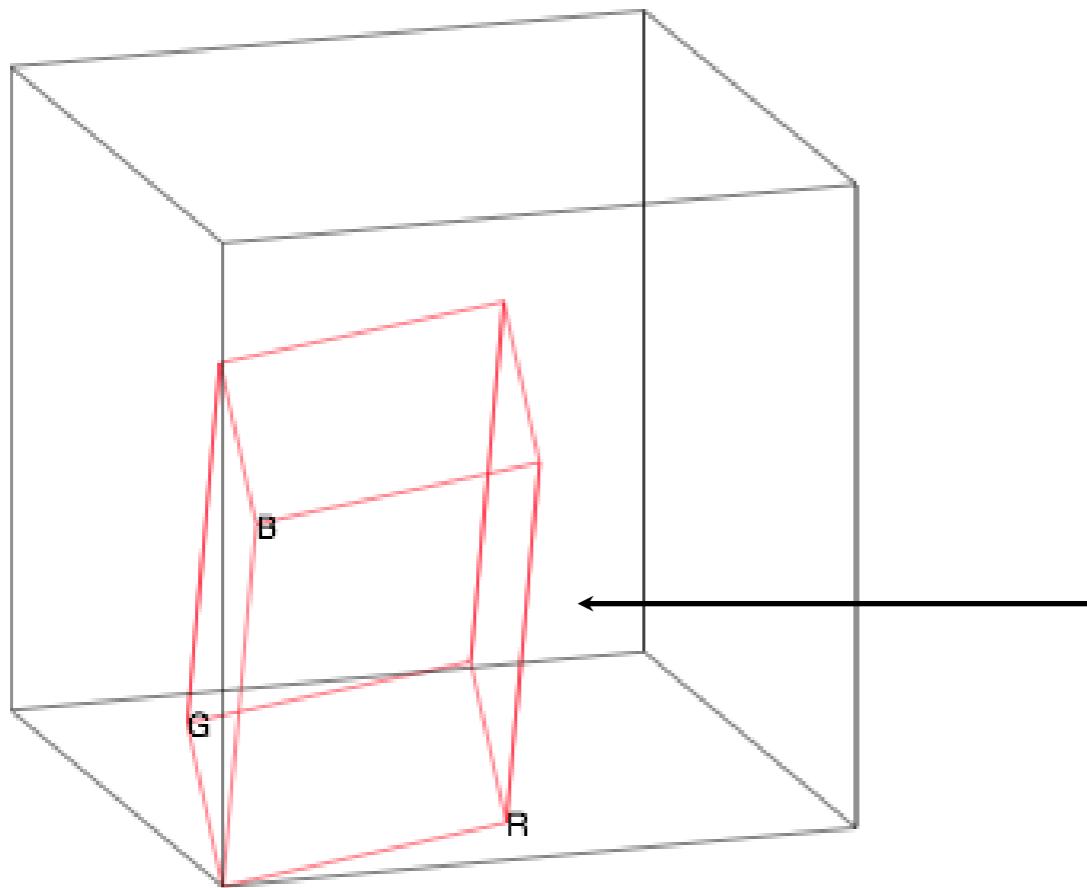
Some data



Some data

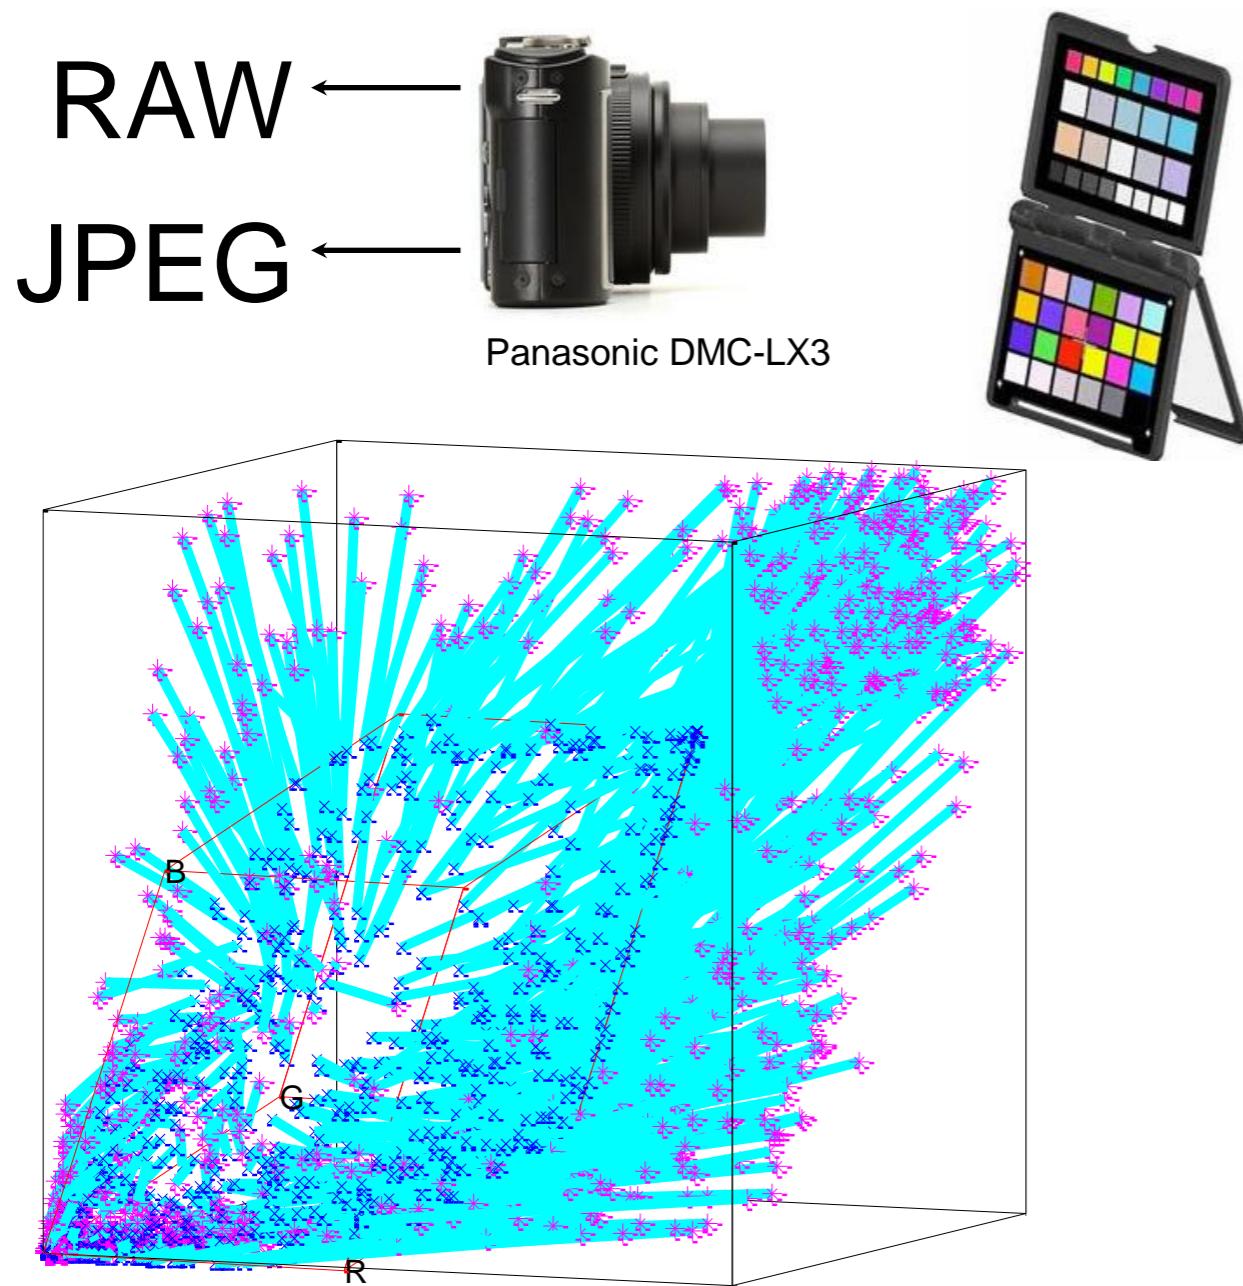


Some data

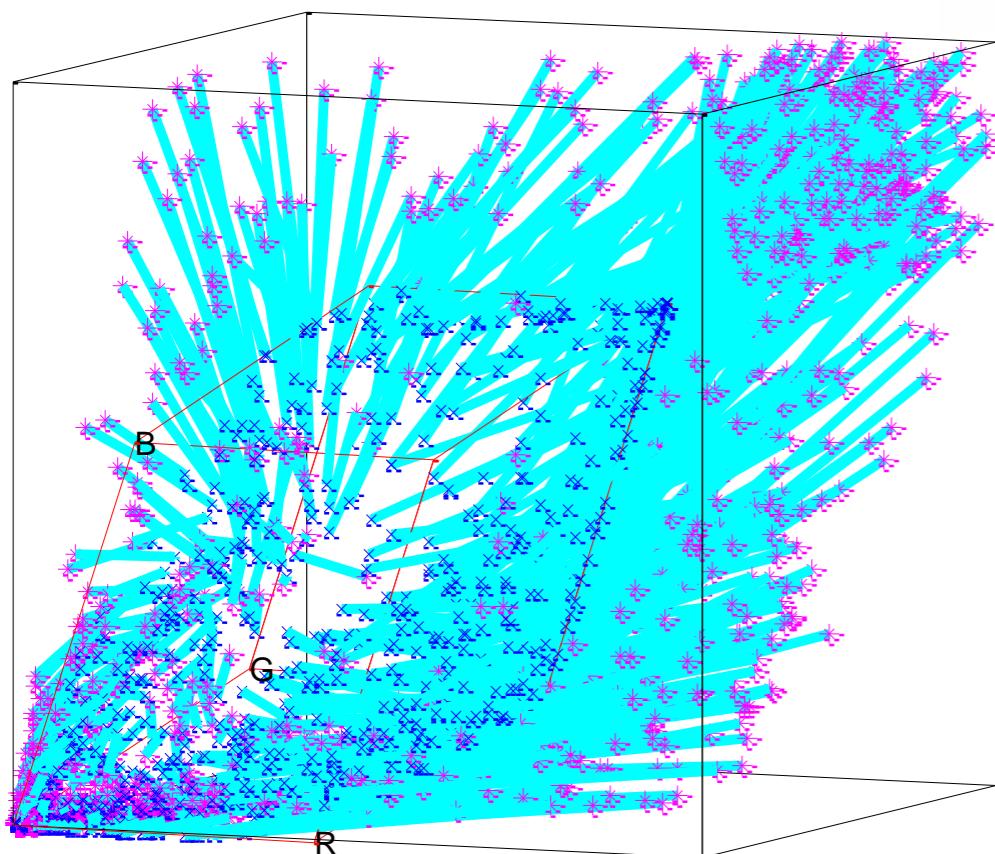


sRGB output gamut

Some data

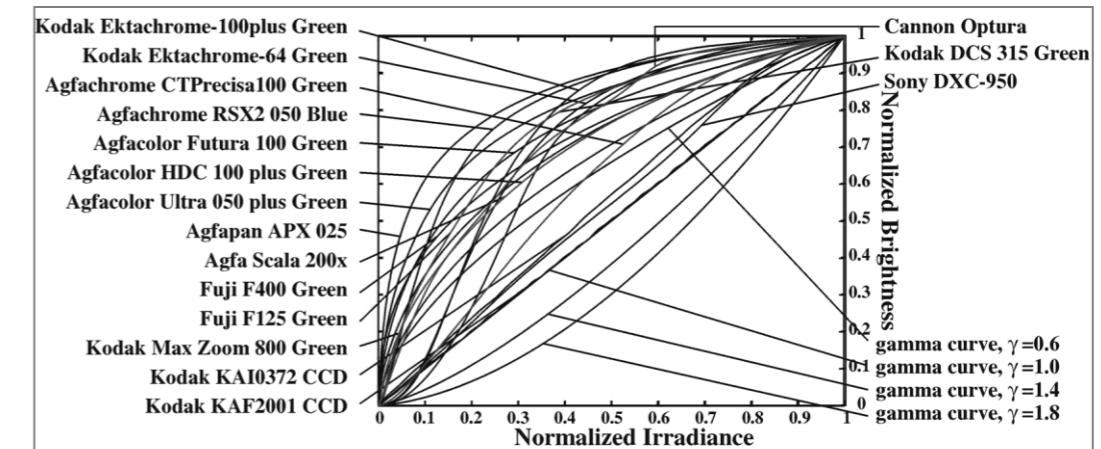


Some data



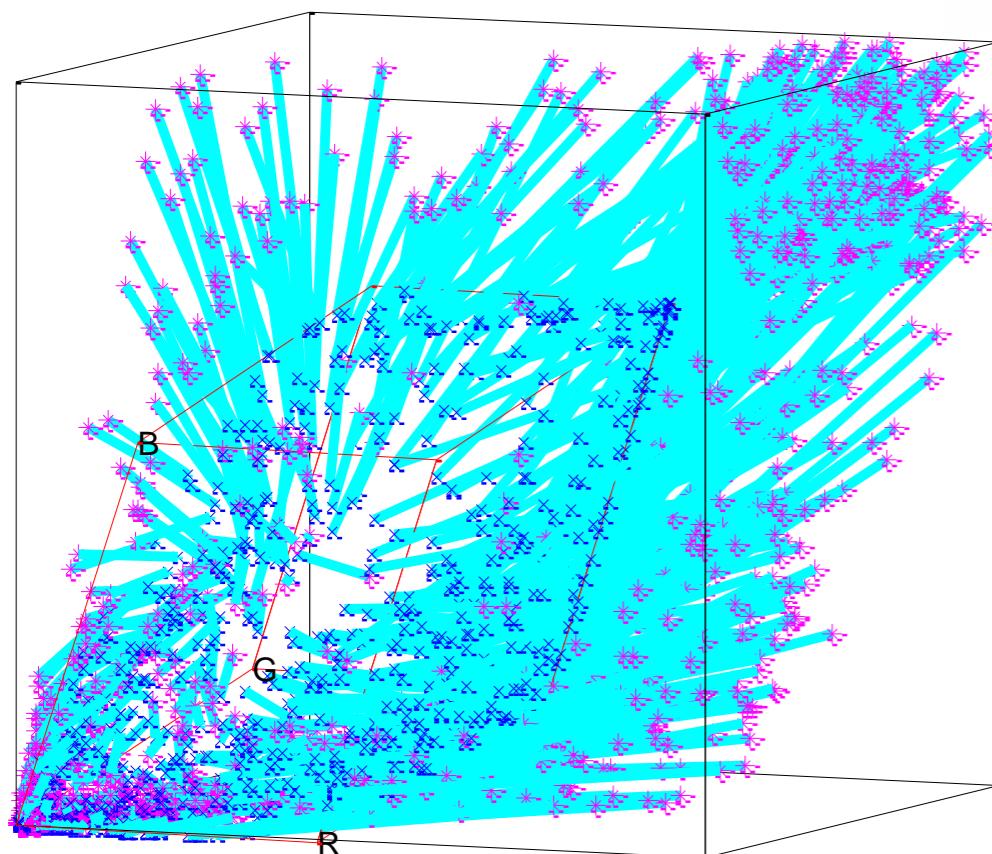
Related Work:

Per-channel
“radiometric response functions”

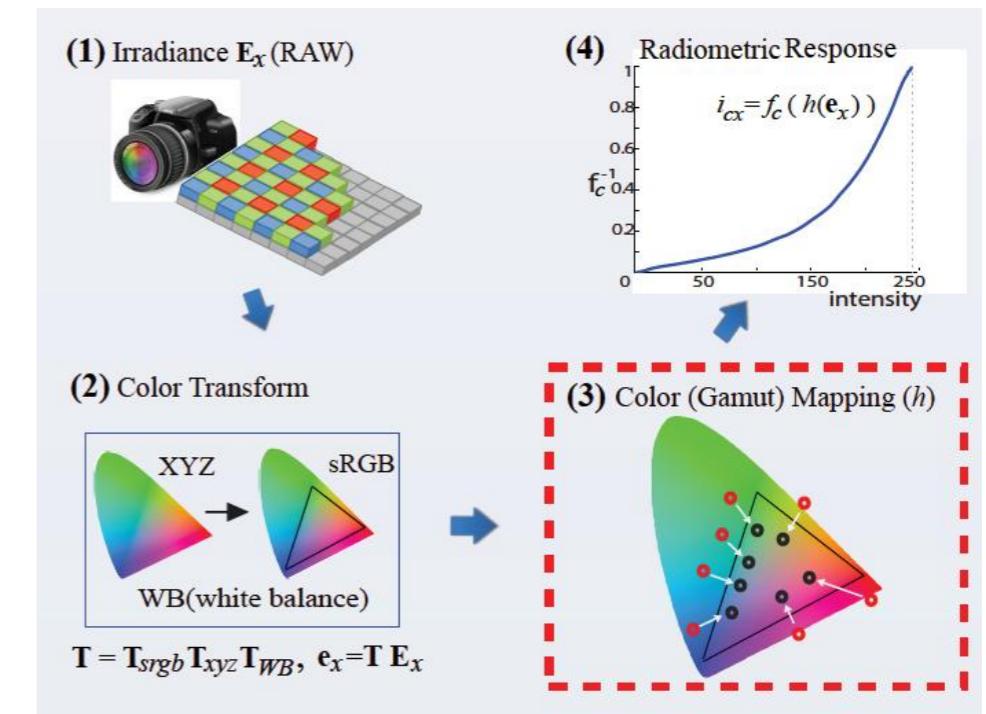


[Mann & Picard, 1995; Debevec & Malik, 1997; Mitsunaga & Nayar, 1999; Grossberg & Nayar, 2004; Pal et al., 2004; ...]

Some data

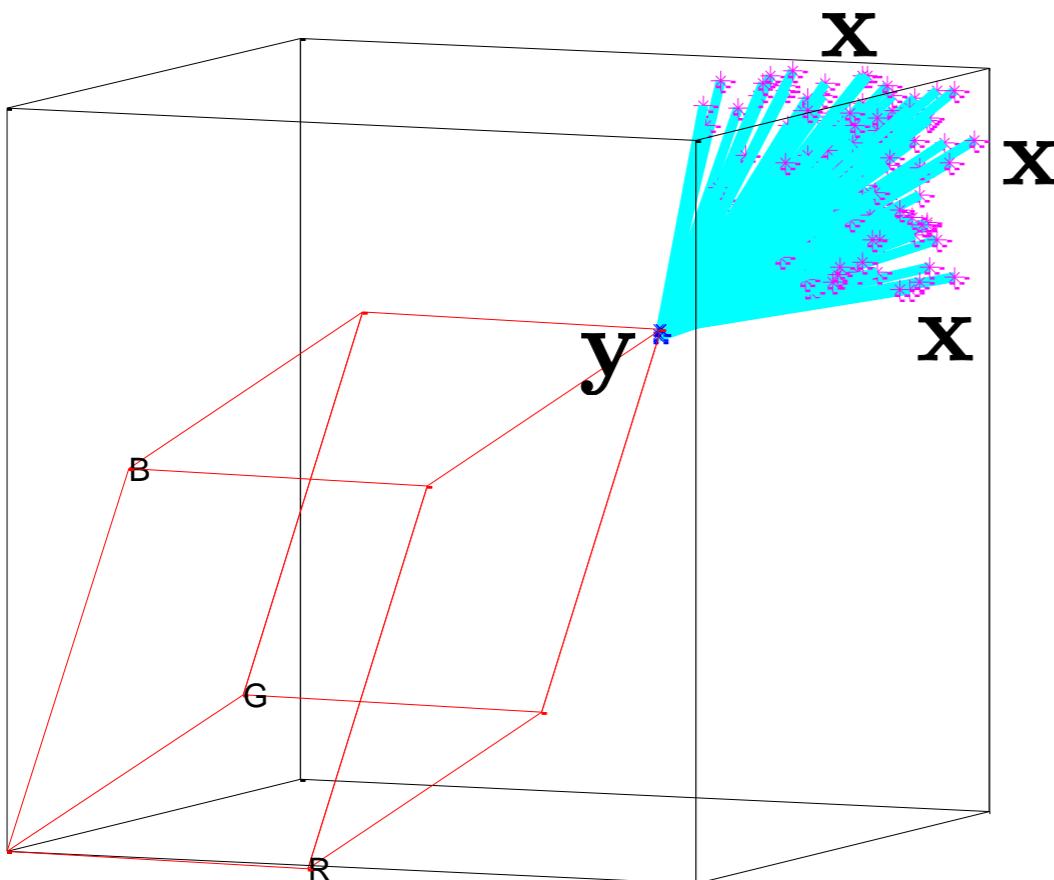


Related Work: Cross-channel deterministic maps

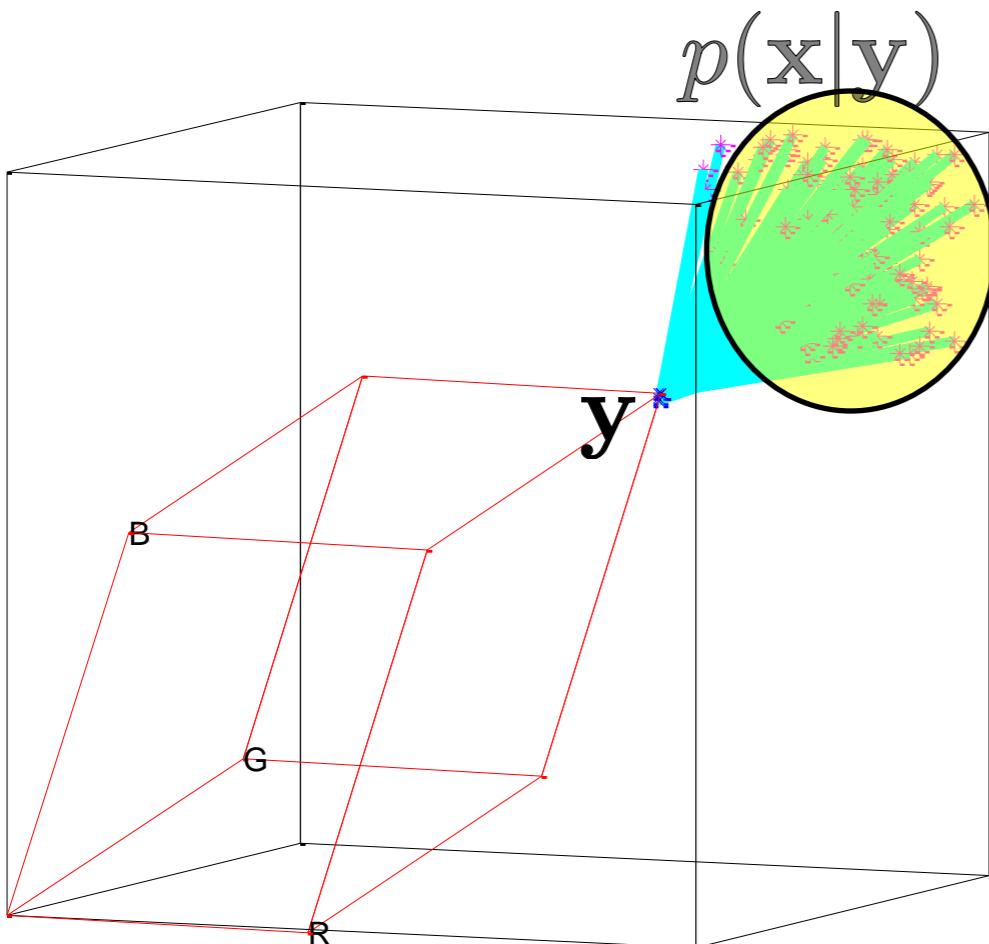


[Chakrabarti et al, 2009; Lin et al., 2011]

Our approach: local GP regression



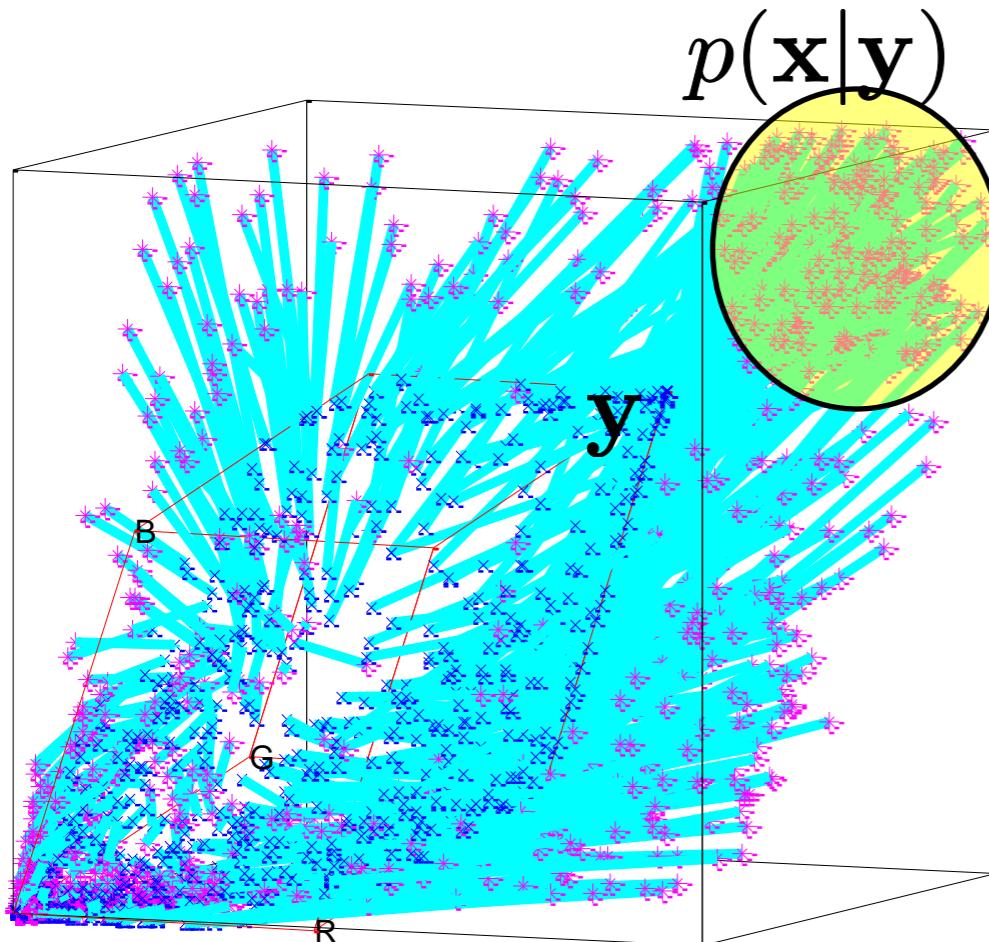
Our approach: local GP regression



Our approach: local GP regression

Training set:

$$\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i), i = 1 \dots N\}$$



$$\mathbb{R}^1 \longleftarrow \mathbb{R}^3$$

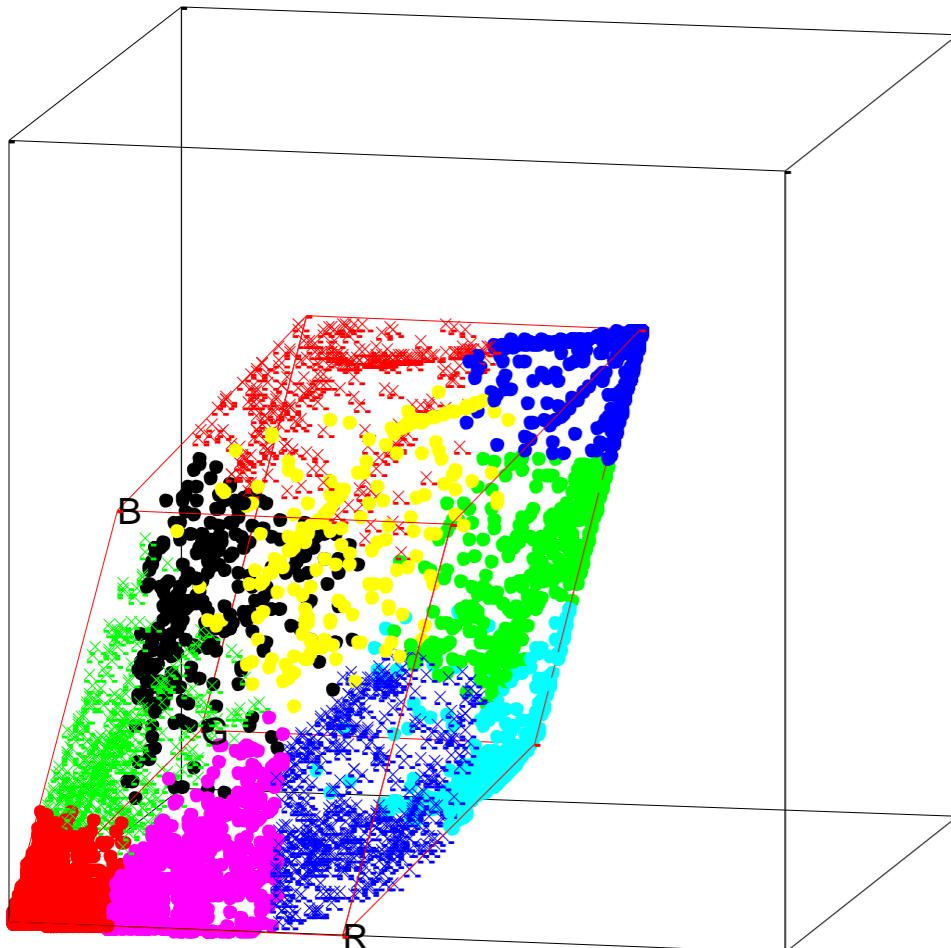
$$x_i^{\text{red}} = z^{\text{red}}(\mathbf{y}_i) + \epsilon_i, \quad \epsilon_i \propto \mathcal{N}(0, \sigma^2)$$

$$p(\mathbf{x}|\mathbf{y}) = \prod_{c=r,g,b} p_{\text{GP}}(x^c|\mathcal{D}, \mathbf{y})$$

Our approach: local GP regression

Training set:

$$\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i), i = 1 \dots N\}$$



$$\mathbb{R}^1 \longleftarrow \mathbb{R}^3$$

$$x_i^{\text{red}} = z^{\text{red}}(\mathbf{y}_i) + \epsilon_i, \quad \epsilon_i \propto \mathcal{N}(0, \sigma^2)$$

$$p(\mathbf{x}|\mathbf{y}) = \prod_{c=r,g,b} p_{\text{GP}}(x^c | \mathcal{D}_{N(\mathbf{y})}, \mathbf{y})$$

[Urtasan & Darrell, 2008]
[Rasmussen and Z. Ghahramani, 2002]
[McHutchon and Rasmussen, 2011]

Potential Applications



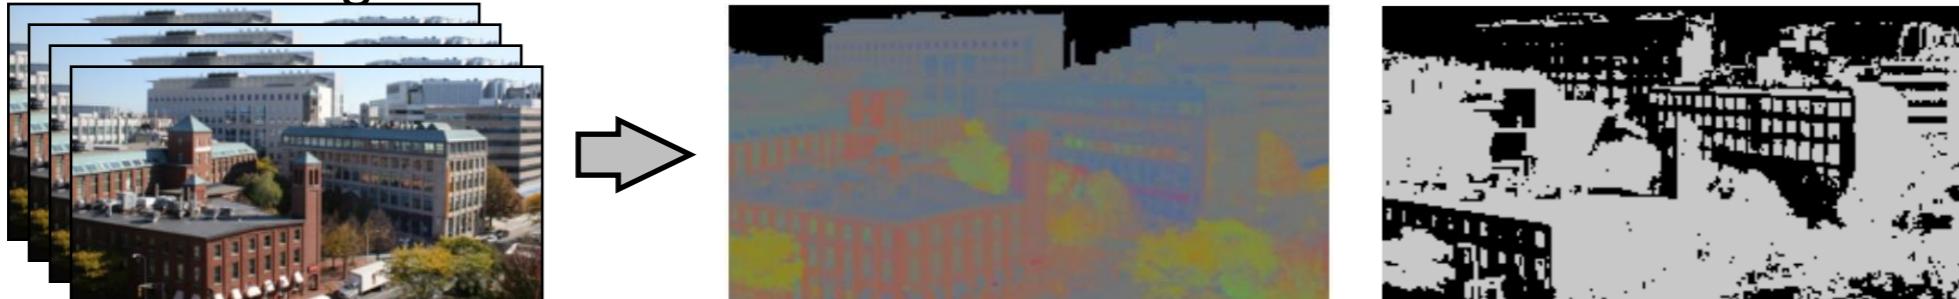
Potential Applications

Recognition



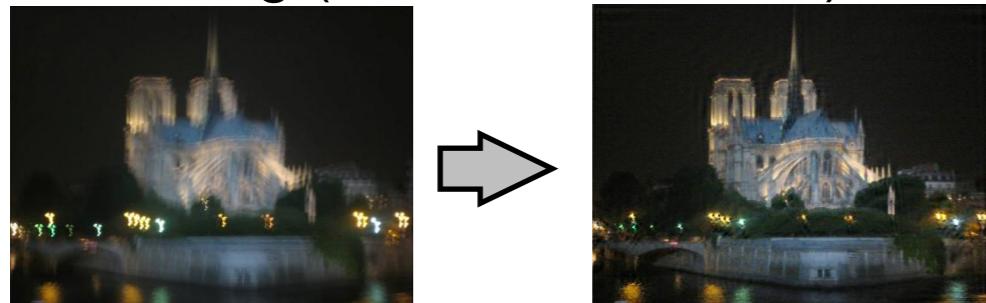
[Nilsback and Zisserman, 2006]

Intrinsic images



[Sunkavalli et al., 2008]

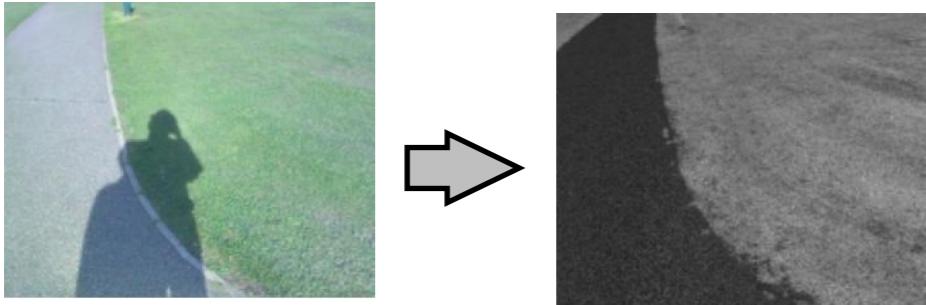
Deblurring (blind or non-blind)



[Whyte et al., 2011]

Potential Applications

Shadow removal



[Finlayson et al., 2002]

Glare removal



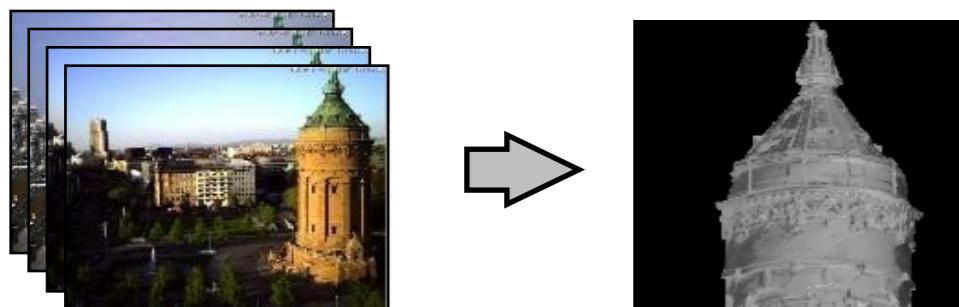
[Mallick et al., 2006]

Multi-exposure HDR imaging



[Debevec and Malik, 1997]

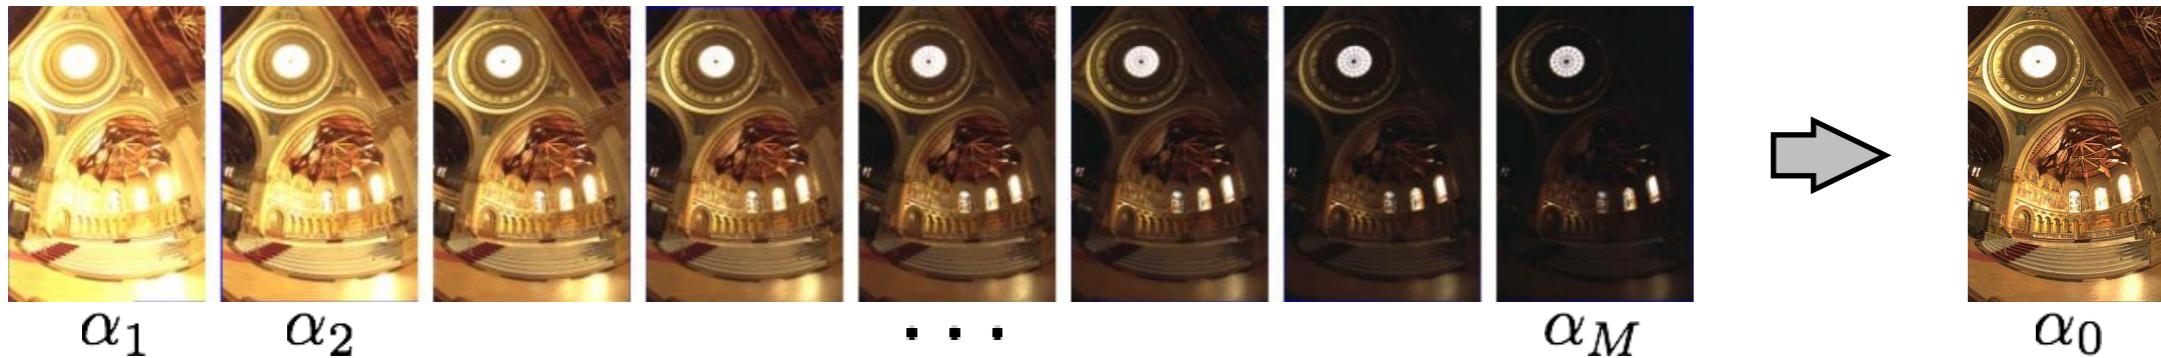
Photometric stereo



[Ackerman et al., 2010]

Probabilistic multi-exposure imaging

Multi-exposure imaging



Conventional approach:
weight all measurements equally

$$\begin{aligned}\mathbf{x}_0 &= \frac{1}{n} \sum_i \frac{\alpha_0}{\alpha_i} \mathbf{x}_i \\ &= \frac{1}{n} \sum_i \frac{\alpha_0}{\alpha_i} f^{-1}(\mathbf{y}_i)\end{aligned}$$

[Debevec and Malik, 1997]

[Mitsunaga and Nayar, 1999]

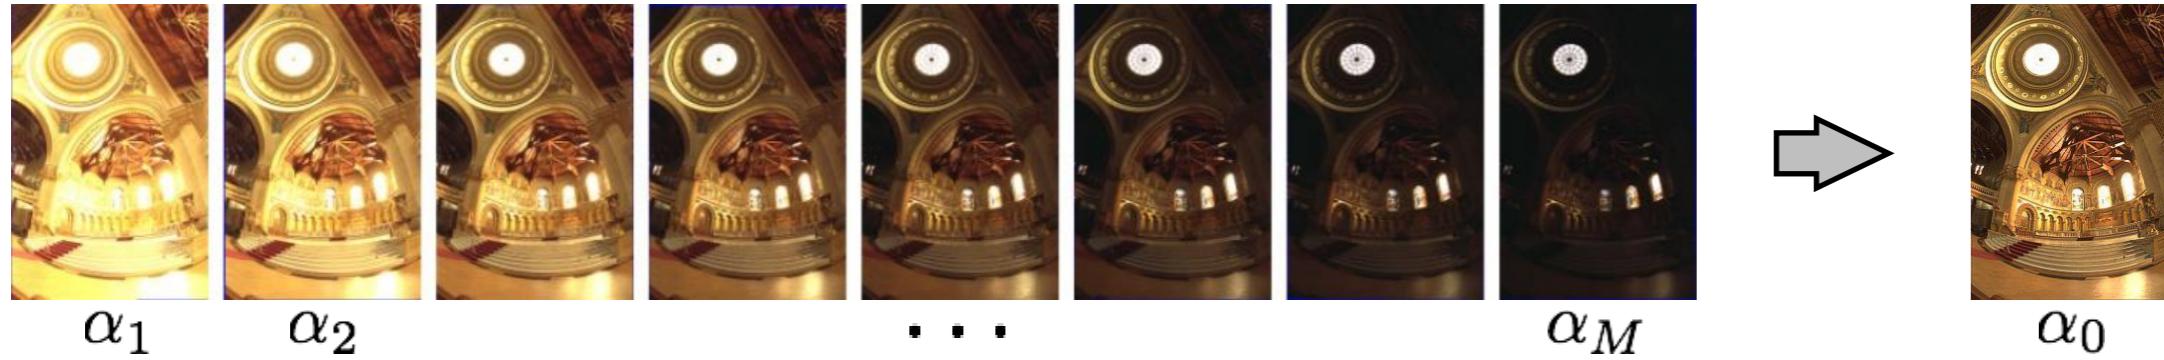
[Robertson et al. 2003]

Probabilistic approach:
combine estimates according to variance

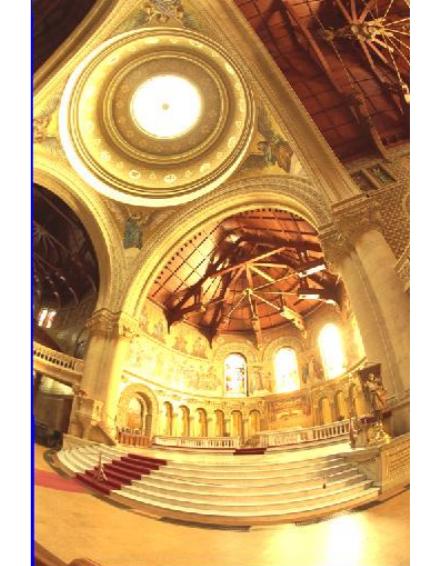
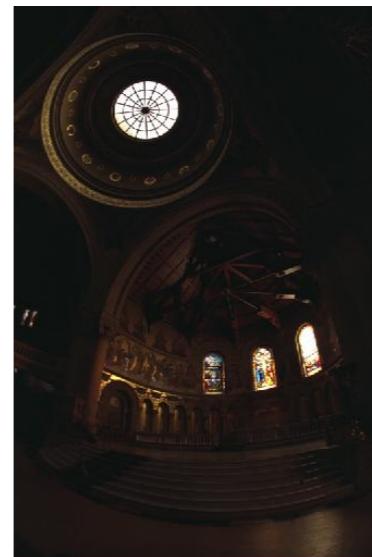
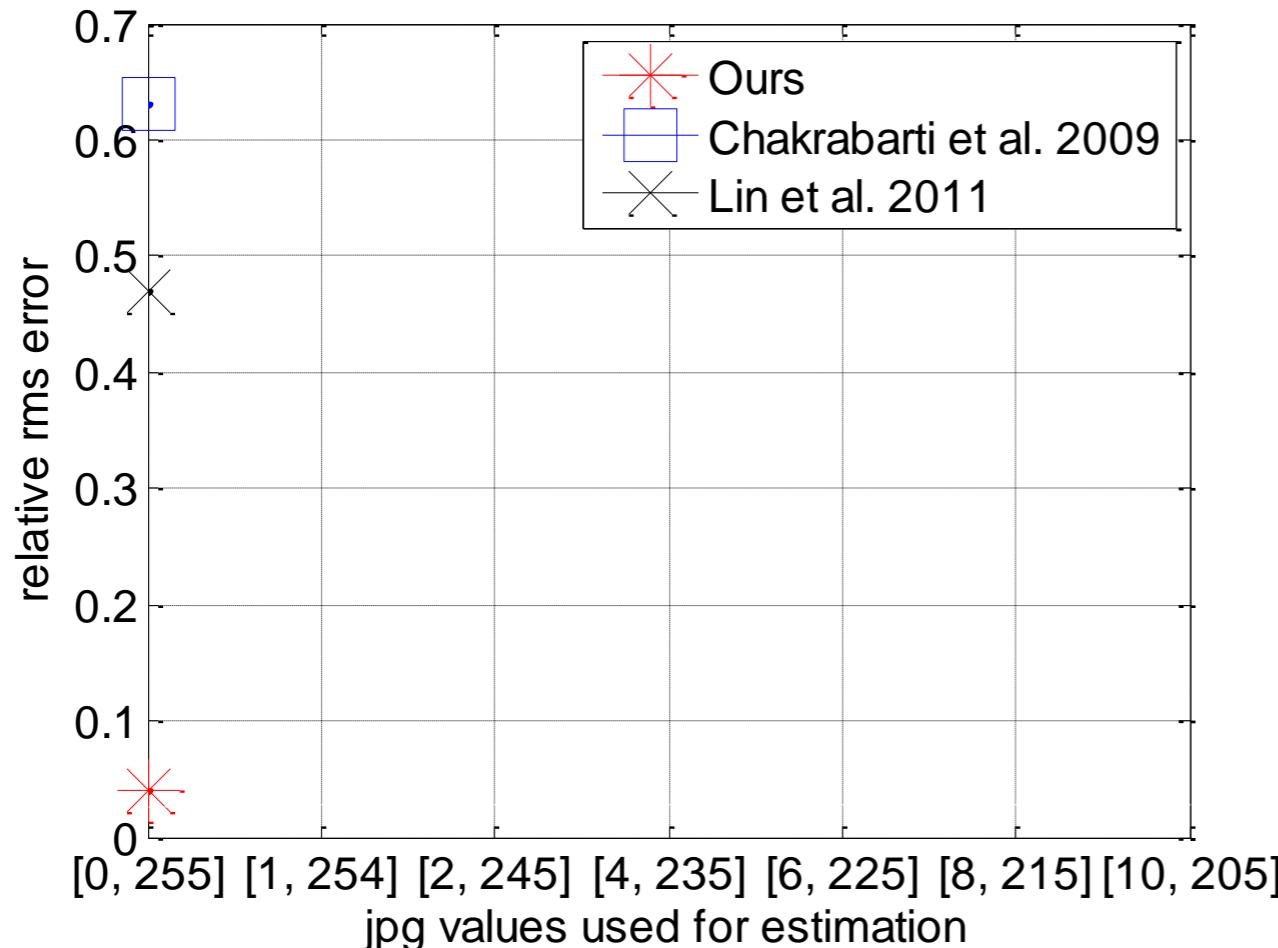
$$\begin{aligned}p(\mathbf{x}_0 | \mathbf{y}_1 \dots \mathbf{y}_N) &= \prod_i p_{x_0|y_i}(\mathbf{x}_0 | \mathbf{y}_i) \\ &= \prod_i \frac{\alpha_i}{\alpha_0} p_{x_i|y_i} \left(\frac{\alpha_i}{\alpha_0} \mathbf{x}_0 | \mathbf{y}_i \right)\end{aligned}$$

Probabilistic multi-exposure imaging

Multi-exposure imaging



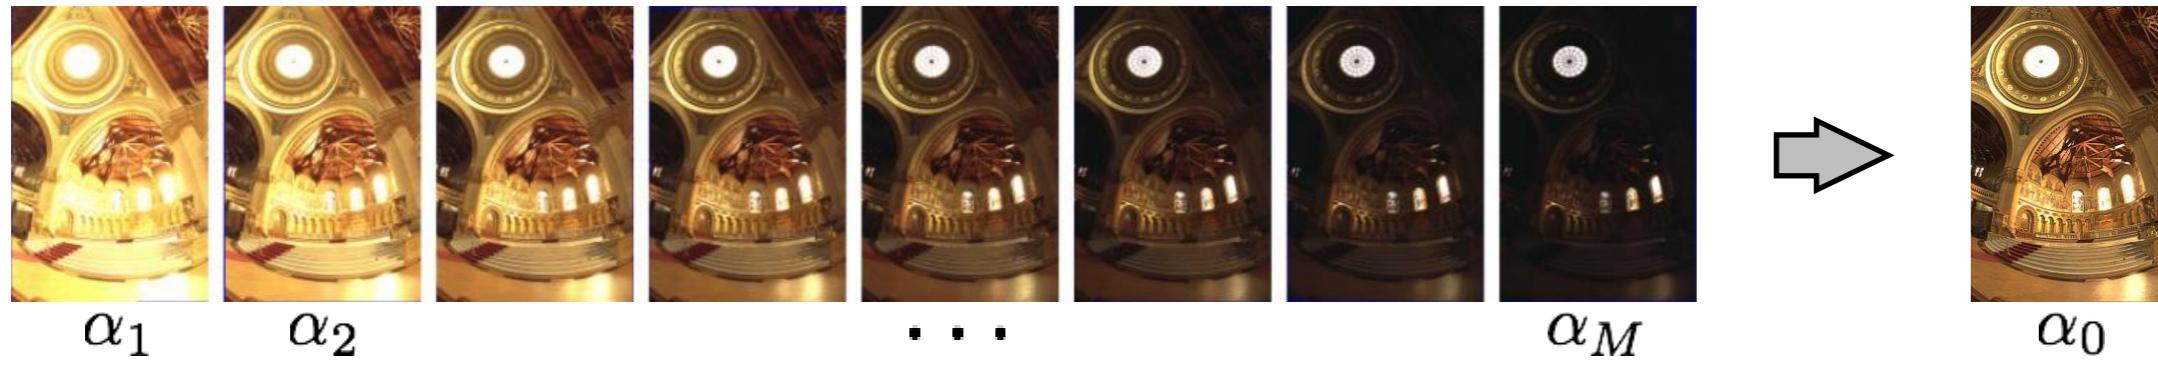
Panasonic DMC-LX3



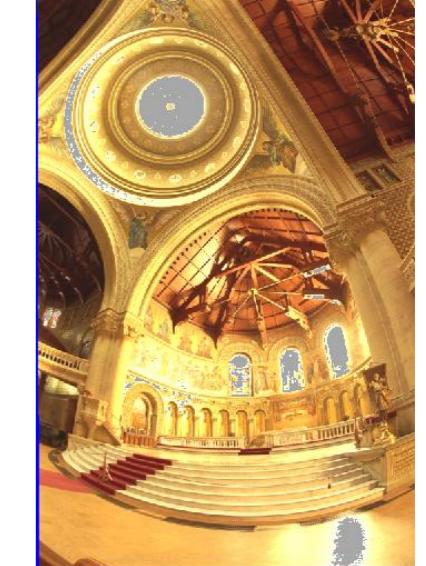
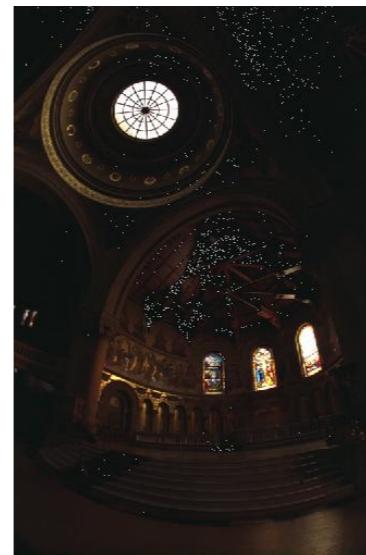
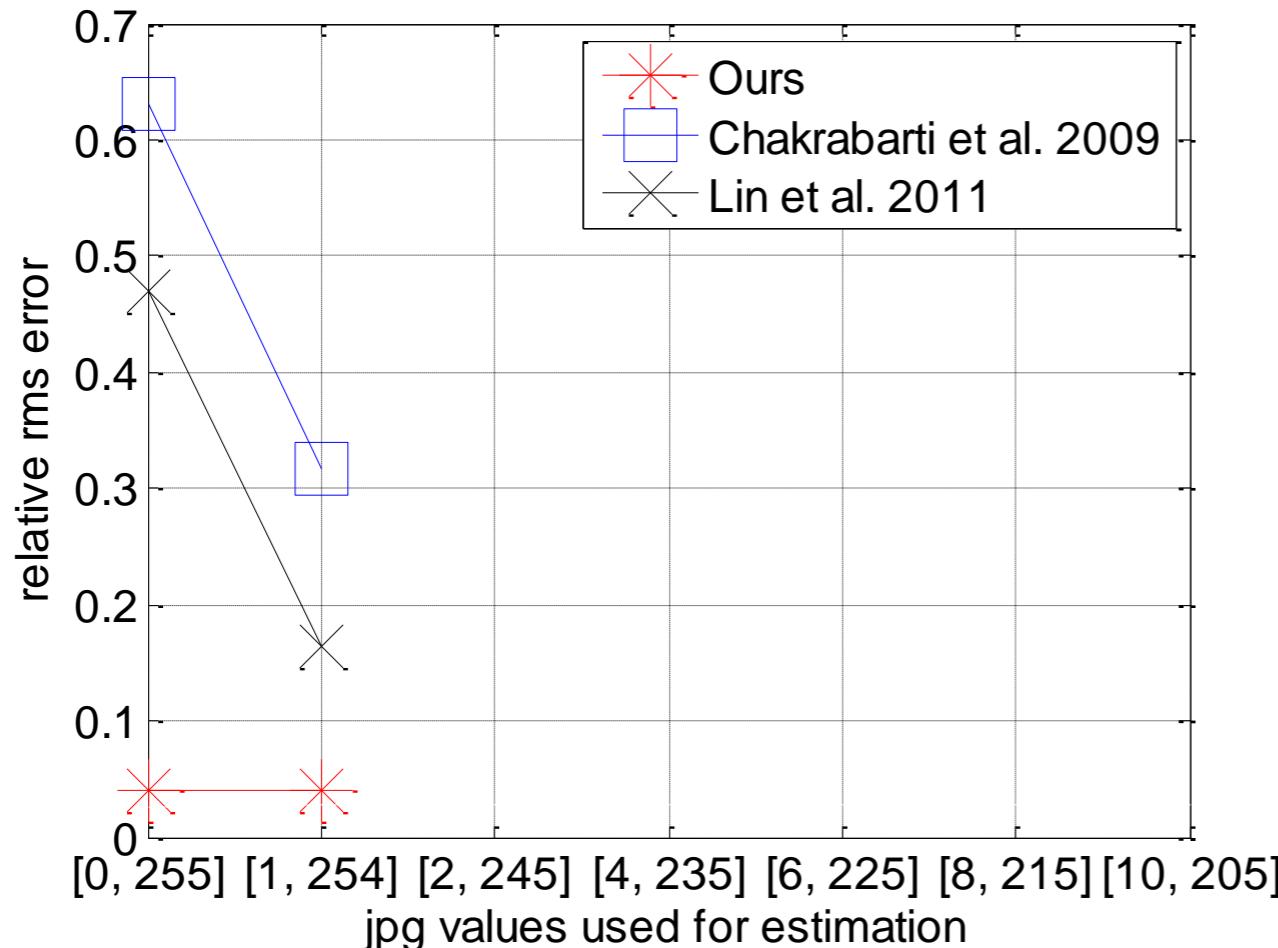
[0 , 255]

Probabilistic multi-exposure imaging

Multi-exposure imaging



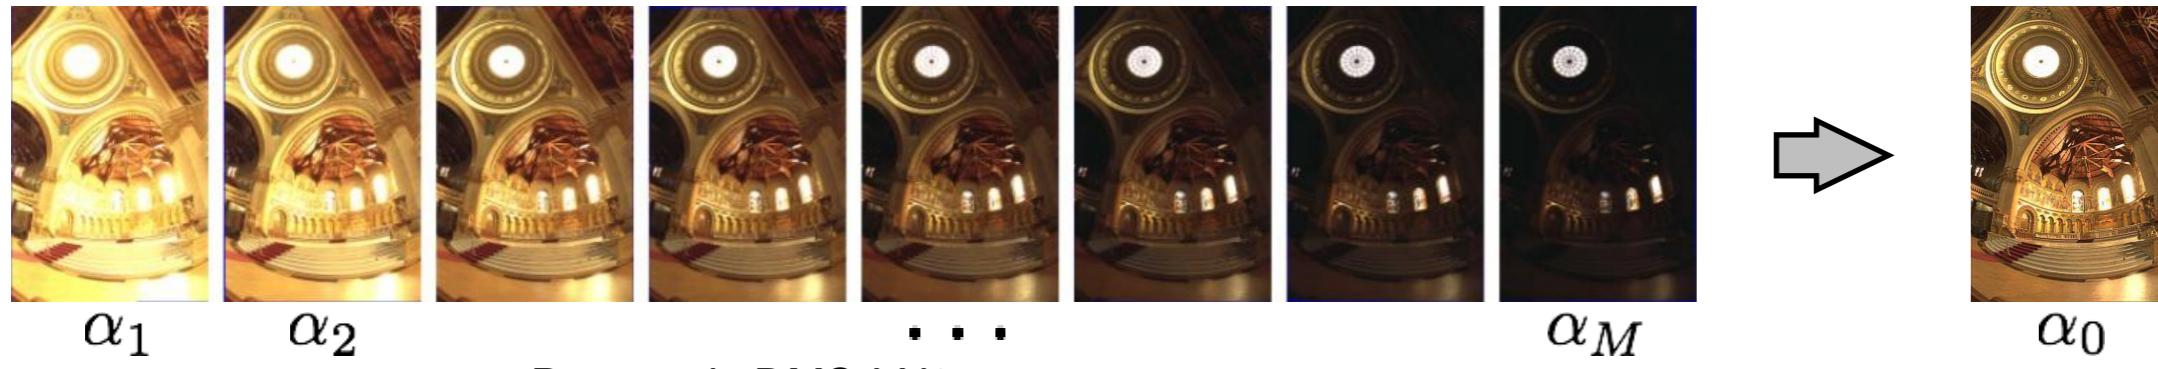
Panasonic DMC-LX3



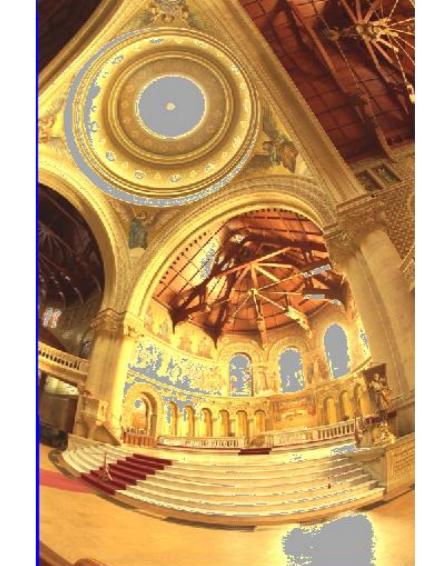
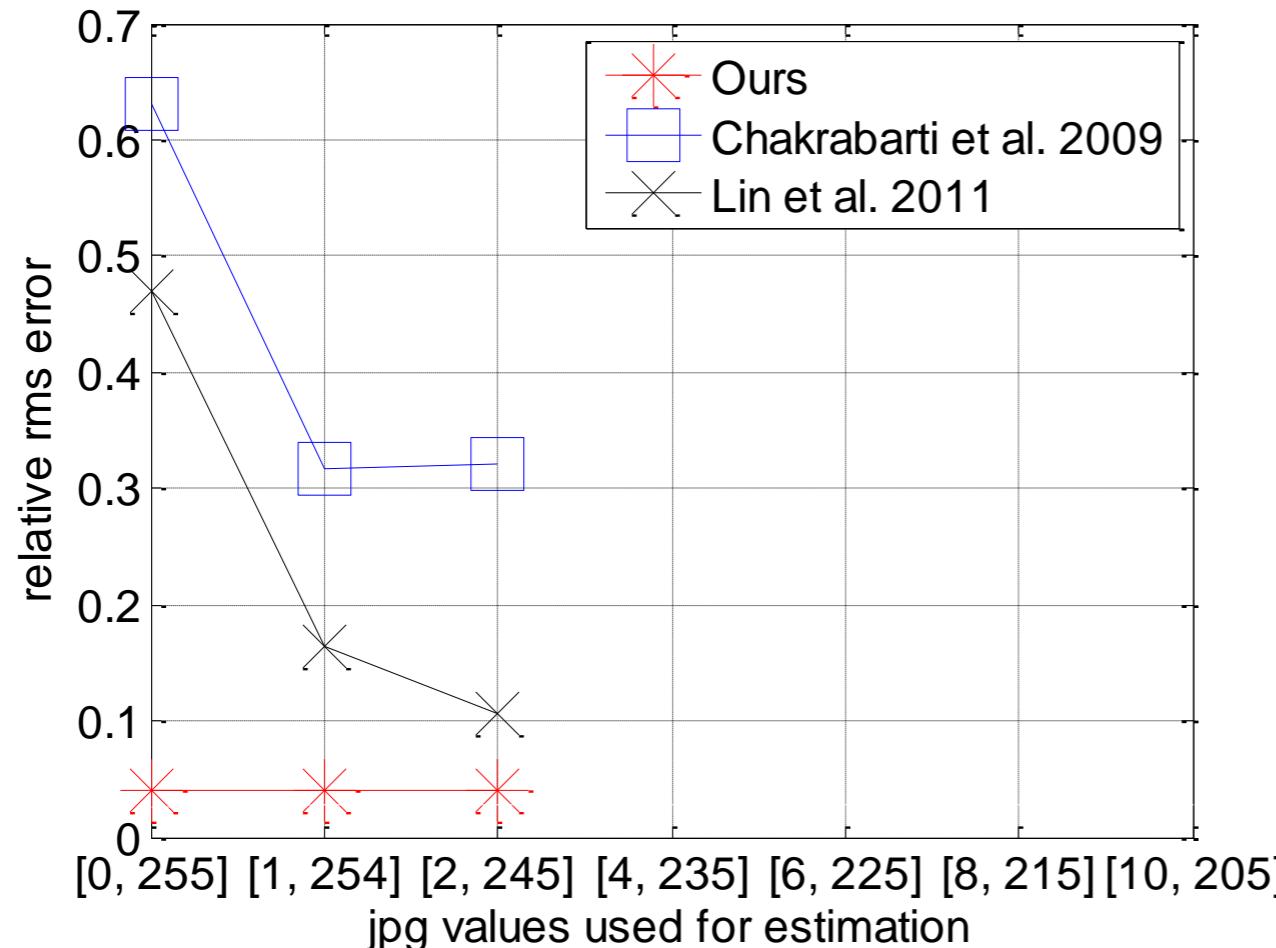
[1 , 254]

Probabilistic multi-exposure imaging

Multi-exposure imaging



Panasonic DMC-LX3

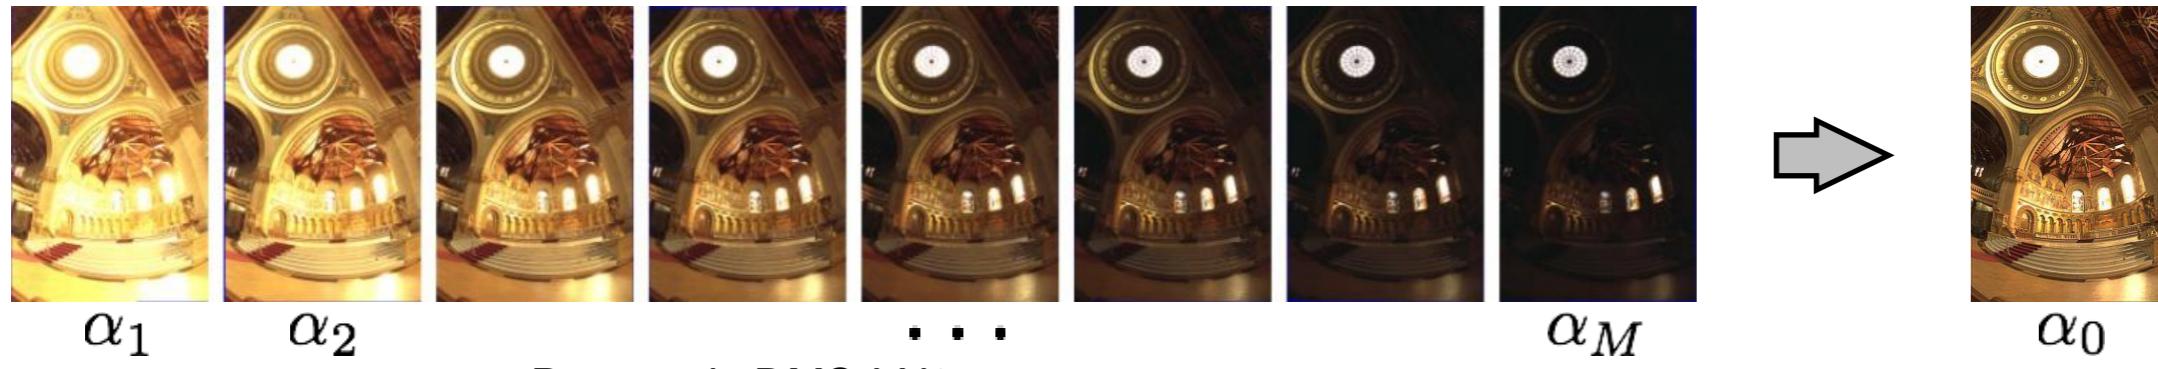


[2 ,

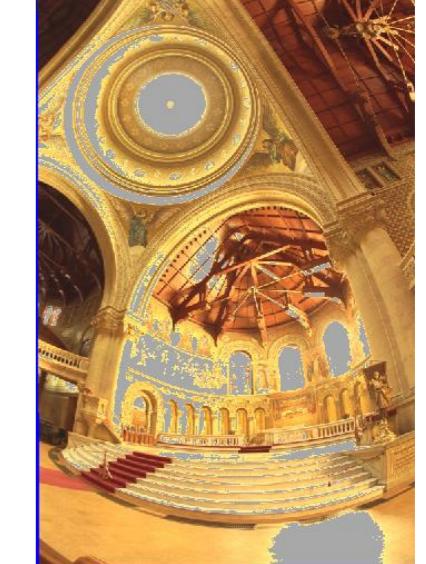
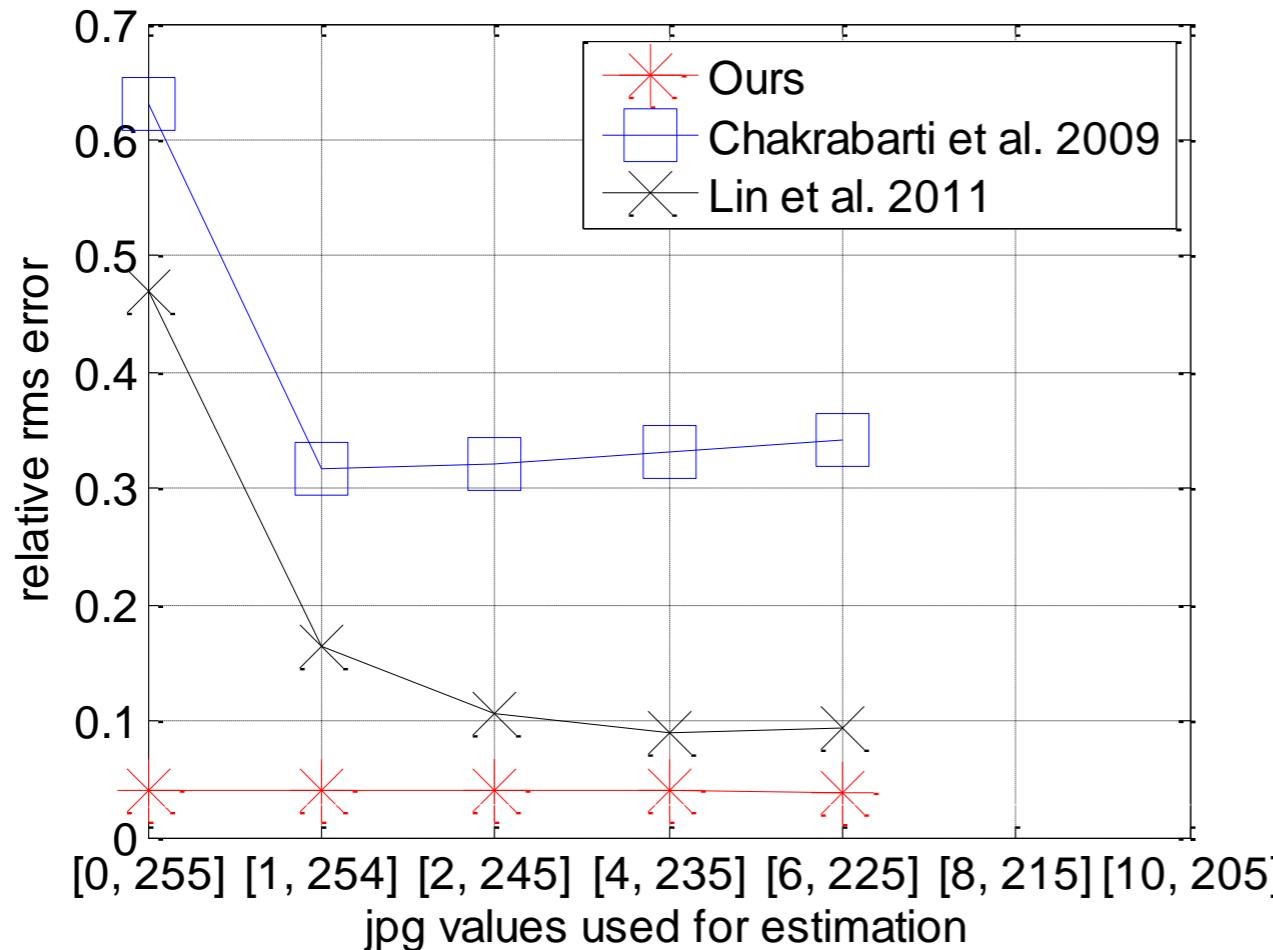
245]

Probabilistic multi-exposure imaging

Multi-exposure imaging



Panasonic DMC-LX3



[6 ,

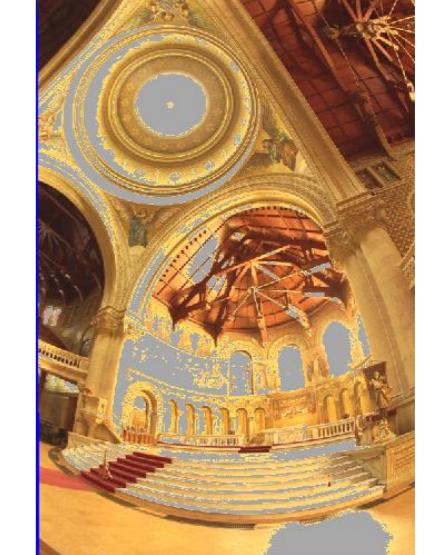
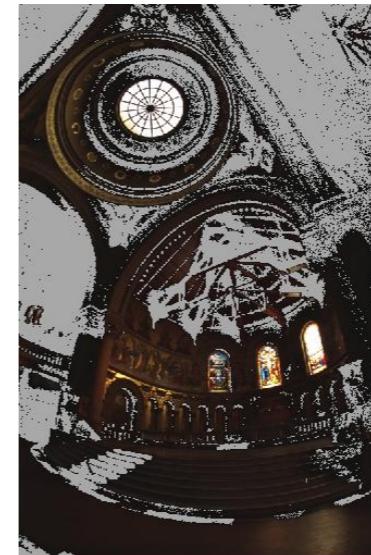
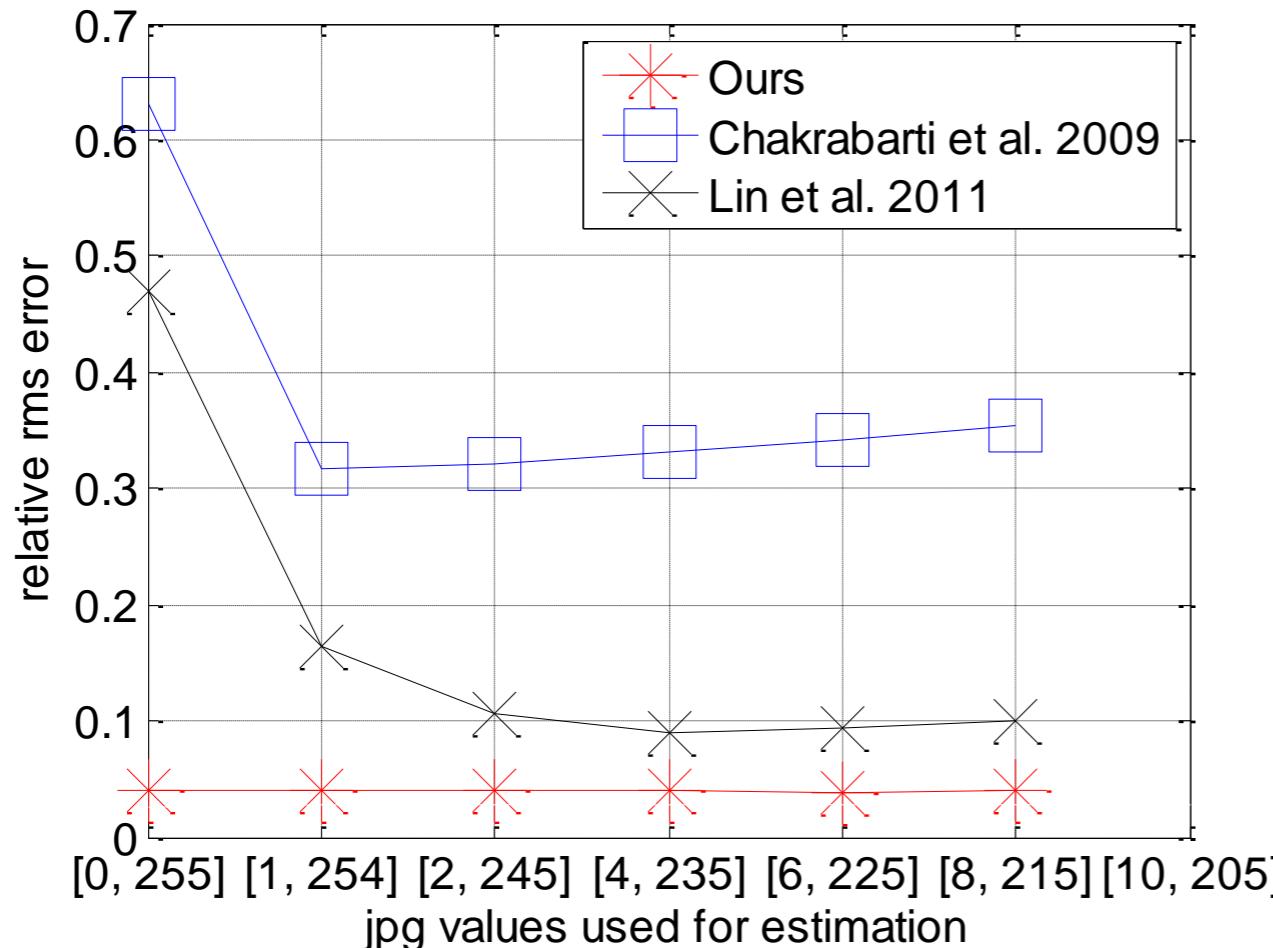
225]

Probabilistic multi-exposure imaging

Multi-exposure imaging



Panasonic DMC-LX3

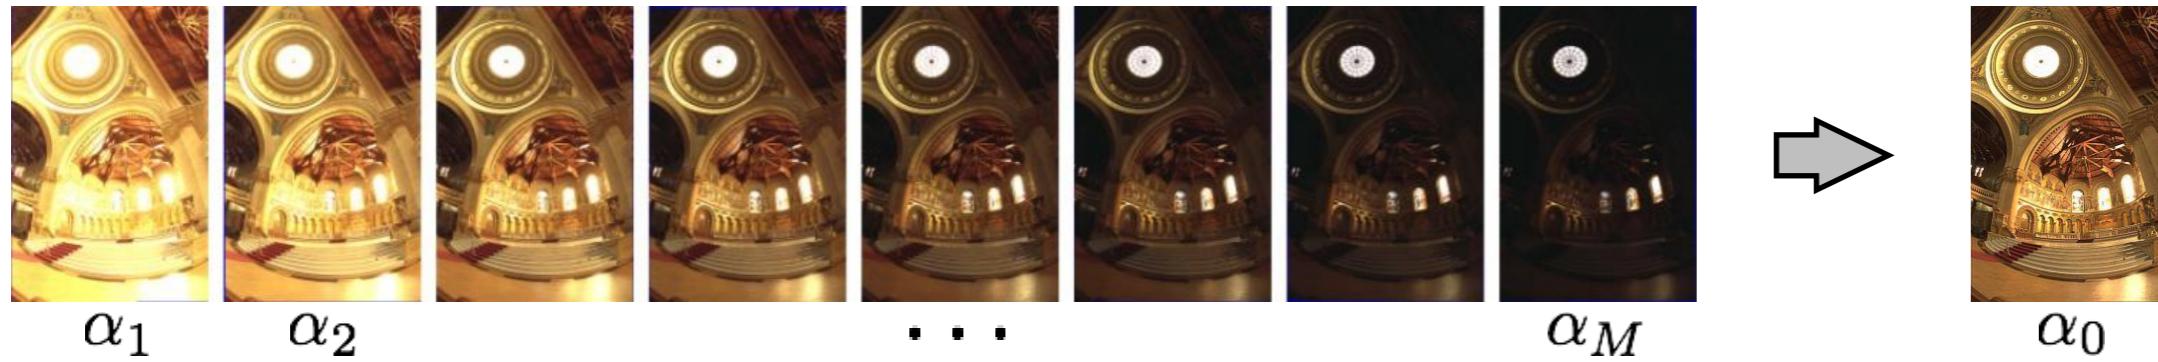


[8 ,

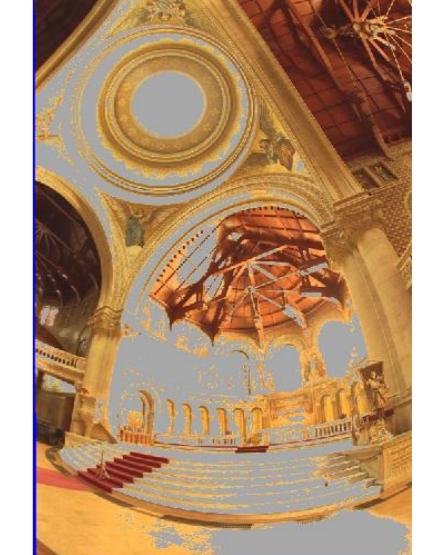
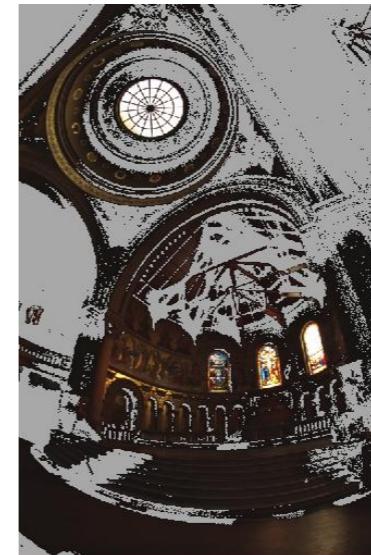
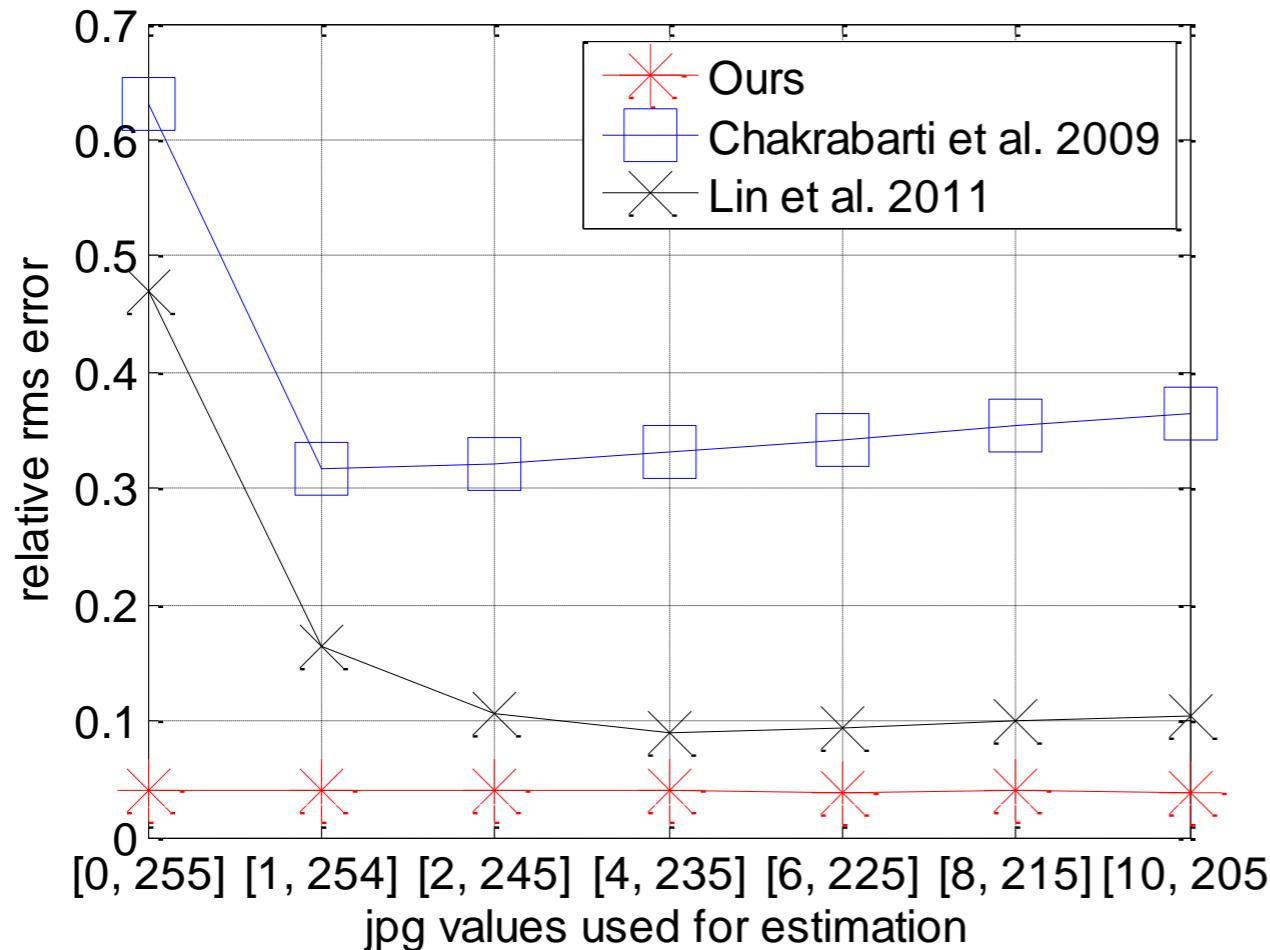
215]

Probabilistic multi-exposure imaging

Multi-exposure imaging



Panasonic DMC-LX3

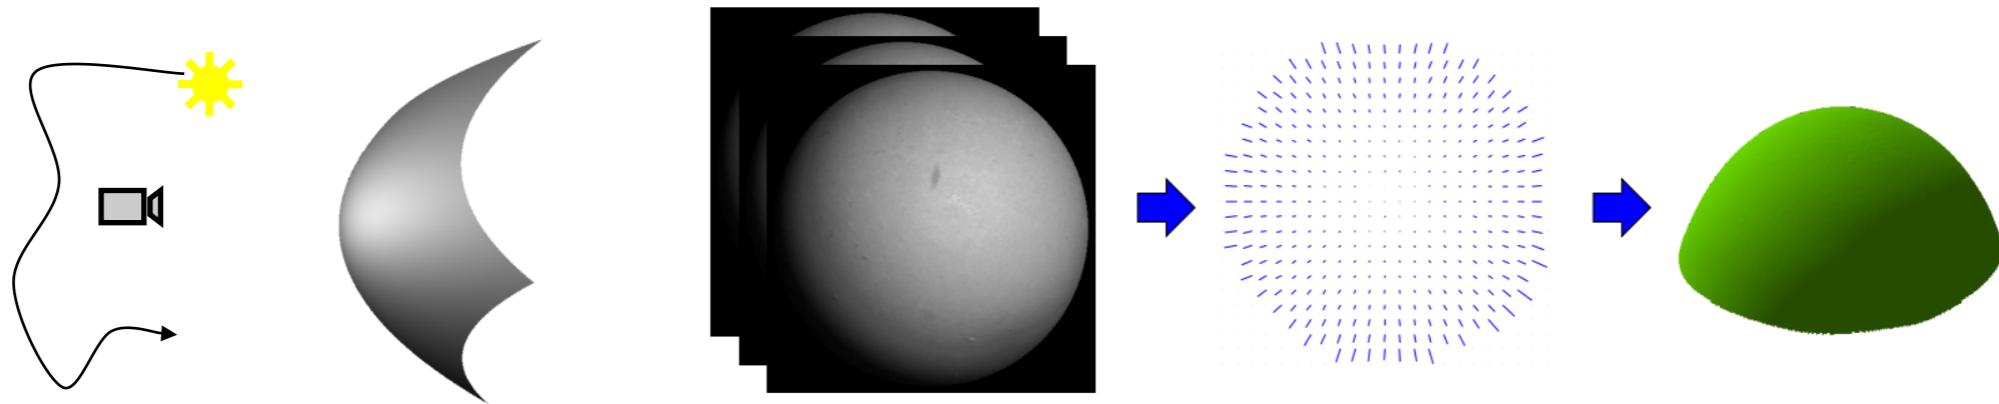


[10 ,

205]

Probabilistic photometric stereo

(Calibrated, Lambertian) Photometric stereo



Conventional approach:
weight all measurements equally

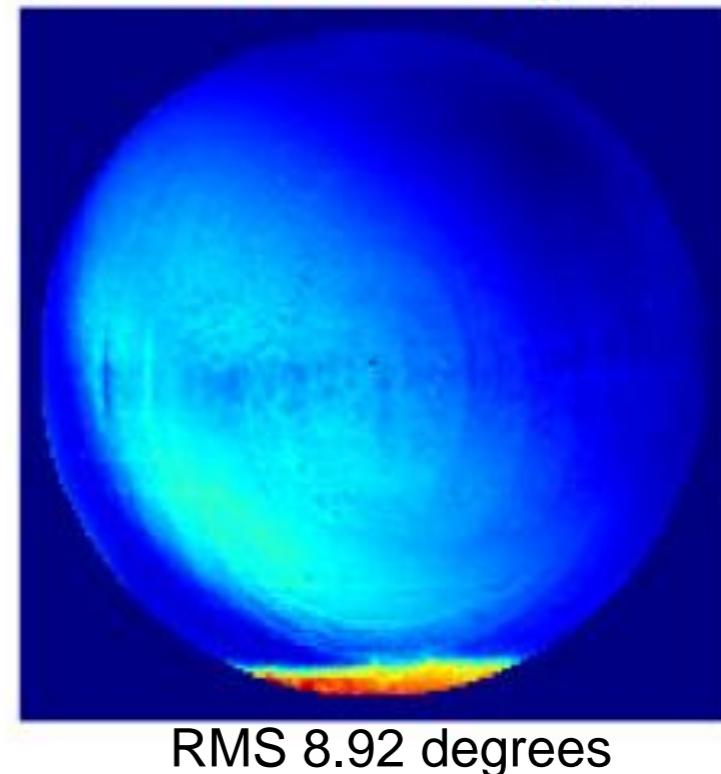
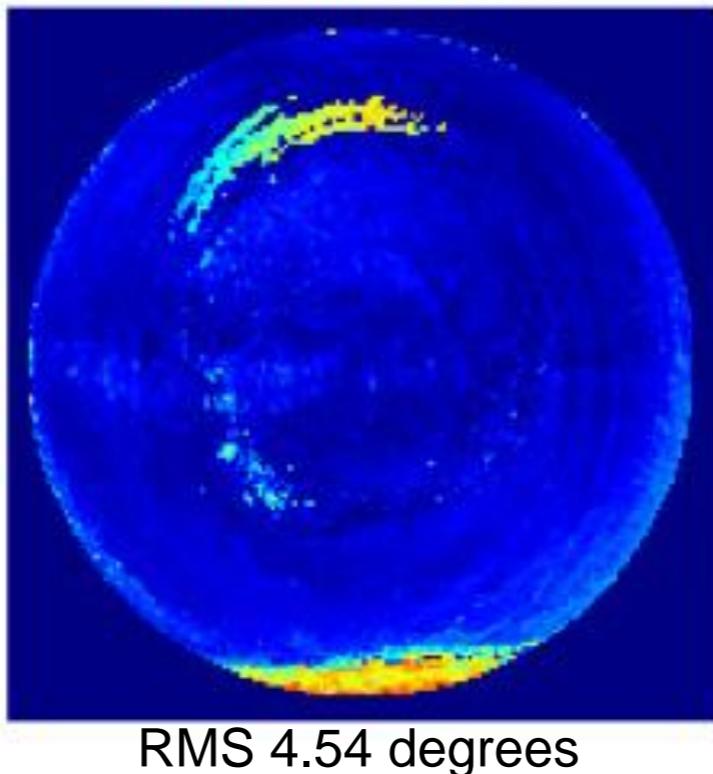
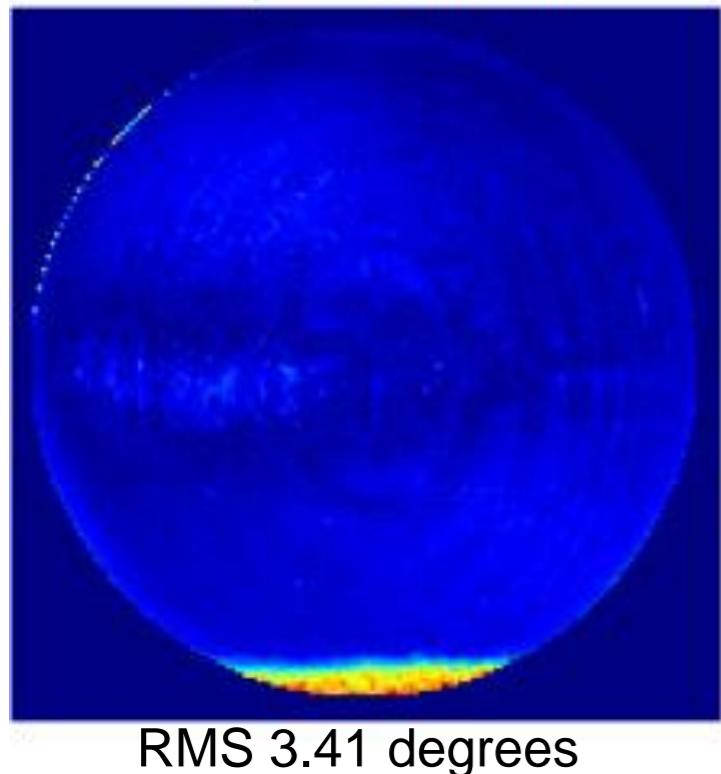
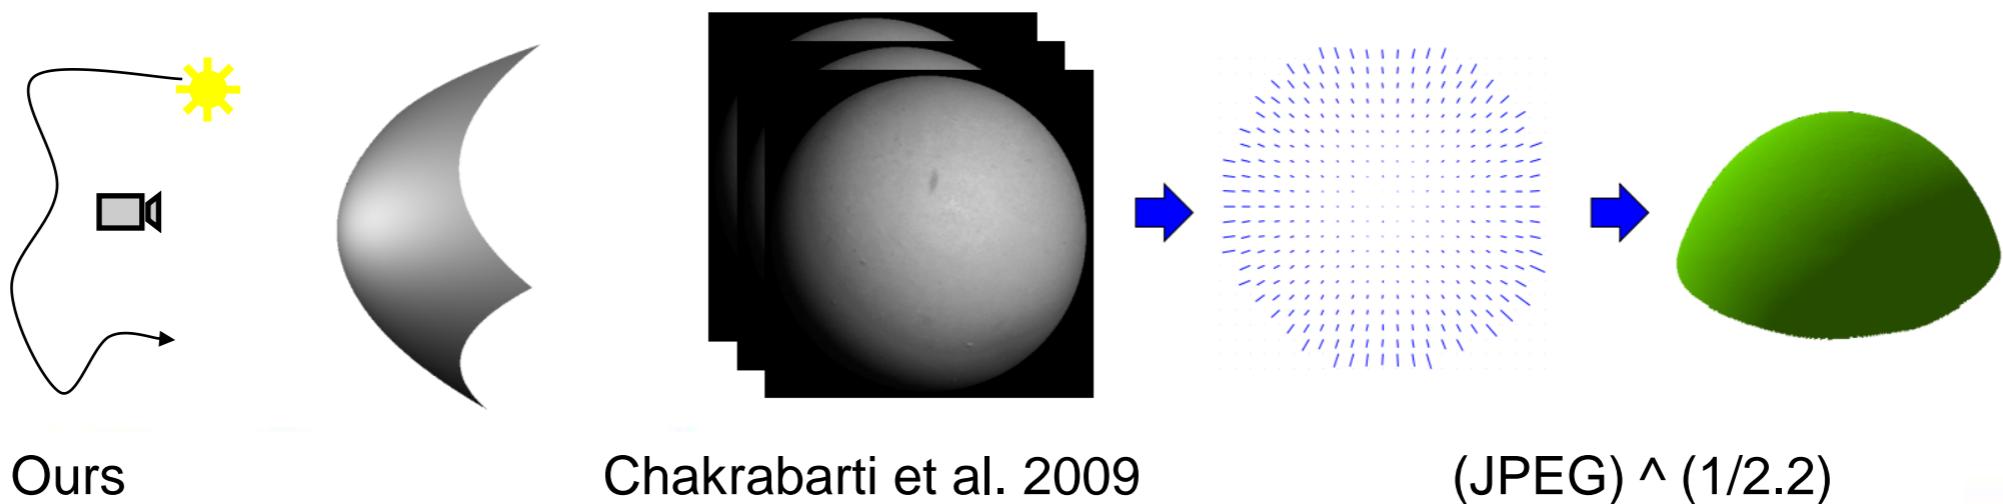
$$\begin{aligned}\mathbf{l}_j^\top \mathbf{b} &= E_j \\ \mathbf{b} &= (\mathbf{L}^\top \mathbf{L})^{-1} \mathbf{L}^\top \mathbf{E}\end{aligned}$$

Probabilistic approach:
combine estimates according to variance

$$\begin{aligned}\mathbf{l}_j^\top \mathbf{b} &= E_j, \quad E_j \sim \mathcal{N}(\mu_j, \sigma_j^2) \\ \mathbf{b} &= (\mathbf{L}^\top \mathbf{W} \mathbf{L})^{-1} \mathbf{L}^\top \mathbf{W} \mathbf{E}\end{aligned}$$

Probabilistic photometric stereo

(Calibrated, Lambertian) Photometric stereo



Angular Error (degree)



Summary

- Color-derendering with a measure of uncertainty
- Improves performance for probabilistic adaptations of HDR imaging and photometric stereo
- Should be useful for probabilistic adaptations of other tasks as well (deblurring, dehazing, intrinsic images, stitching and matching, color constancy, etc.)