From Pixels to Physics: Probabilistic Color De-rendering

Ying Xiong[1], Kate Saenko[2], Trevor Darrell[2] and Todd Zickler[1]

[1]Harvard University and [2]UC Berkeley

Introduction

- Many computer vision algorithms assume that pixel values are linearly related to scene radiance.
- We often want to apply these algorithms to non-linear consumer images that are shared online.

Photometric stereo

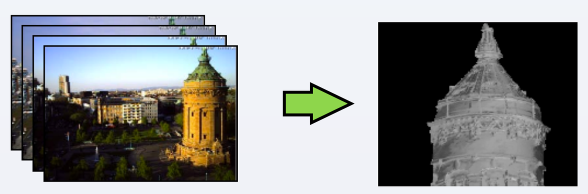

[Ackerman et al., 2010]

Shadow removal

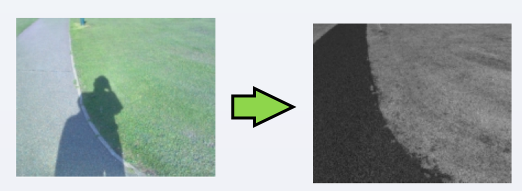

[Finlayson et al., 2002]

Multi-exposure HDR imaging

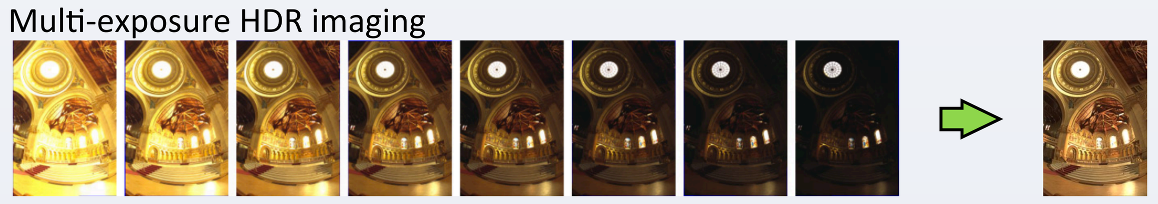

[Debevec and Malik, 1997]

De-hazing

[He et al., 2009]

De-glossing

[Mallick et al., 2006]

De-blurring

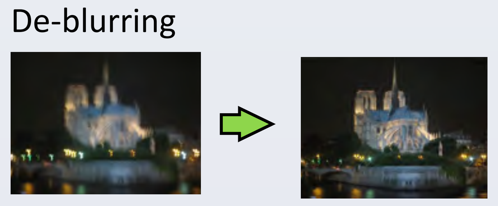

[Whyte et al., 2011]

Intrinsic images

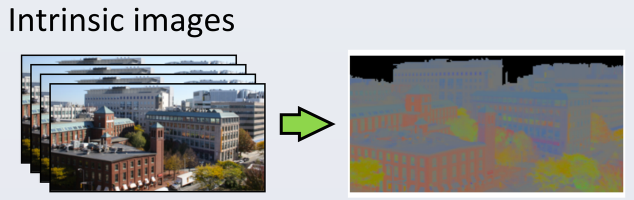

[Sunkavalli et al., 2008]

Lighting/Weather estimation and reflectometry

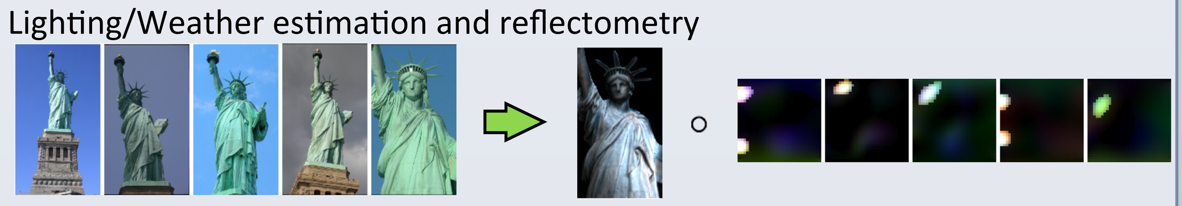

[Haber et al., 2009]

Problems and Objectives

Problems: Consumer cameras render colors to make visually pleasing images for narrow-gamut displays.

- Physical accuracy is lost.
- Many-to-one mapping cannot be deterministically undone.

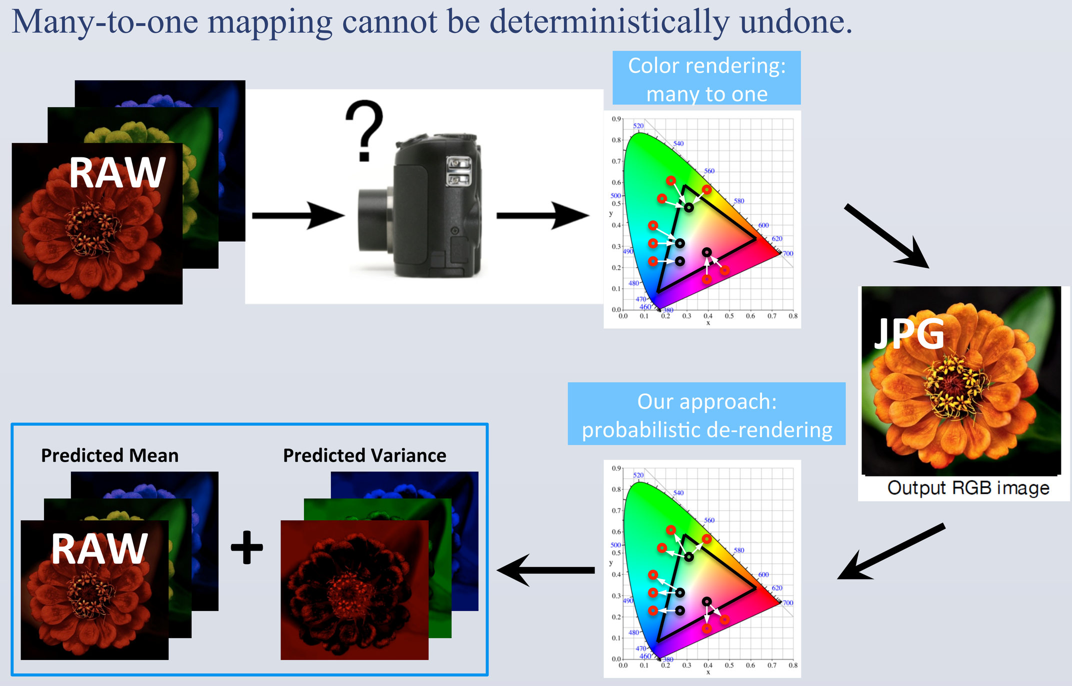

- De-rendering: Inferring linear (RAW) values from non-linear (JPEG) ones.
- Use *probabilistic* models to account for information lost during color rendering.
- Probabilistic output can be directly applied in upstream applications.

Previous Works

- Per-channel radiometric response functions
  - permits online "self calibration"
  - cannot recover out-of-gamut chromaticities

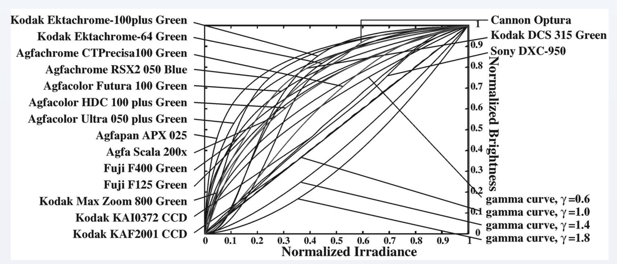

[Grossberg and Nayar, 2004]

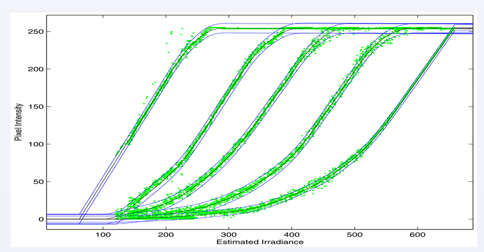

[Pal et al., 2004]

- Cross-channel radiometric calibration
  - requires offline calibration
  - can recover out-of-gamut chromaticities, but is still one-to-one

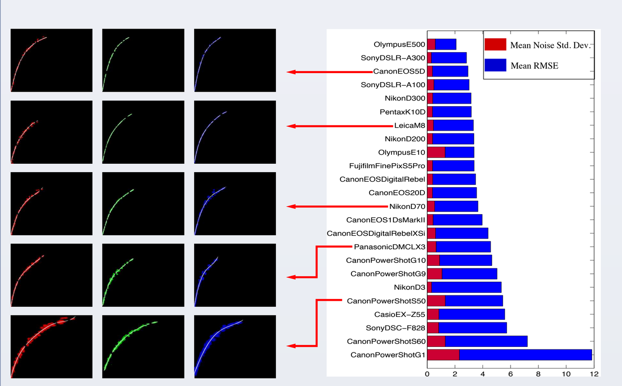

[Chakrabarti et al., 2009]

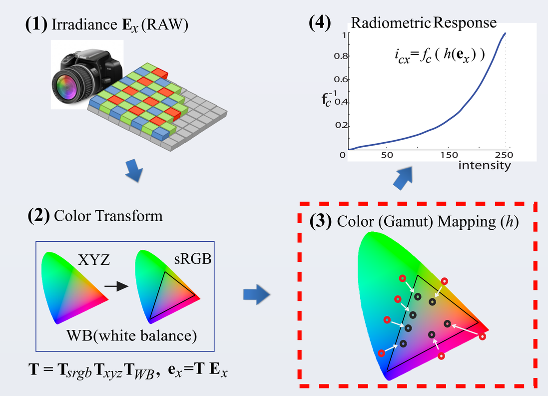

[Lin et al., 2011]

Forward (rendering) Model

The rendering pipeline of a common customer digital camera:

Spectral Irradiance $E(\lambda)$ ⇨ Spectral Filters $\{\pi^c(\lambda)\}_{c=R,G,B}$ ⇨ Scene Color $\mathbf{x}=\{x^c\}$ ⇨ Sensor Saturation $\min(1,x^c)$ ⇩ RAW Value $\tilde{\mathbf{x}}=\{\tilde{x}^c\}$ ⇦ White Balance W ⇦ Color Transform C ⇦ Color Rendering $f(\cdot)$ ⇦ Rendered Color (sRGB) $\mathbf{y}=\{y^c\}$

Many-to-one effects due to tone-mapping and saturation

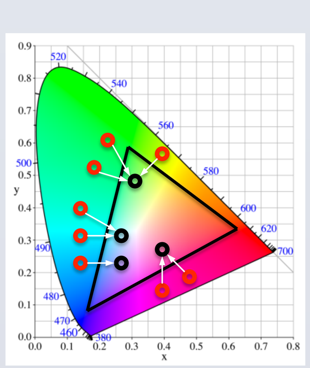

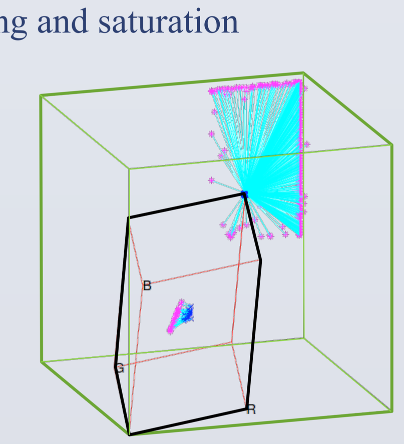

DMC-LX3, Panasonic Inc.

Non-linear and even non-monotonic effect for out-of-gamut chromaticities.

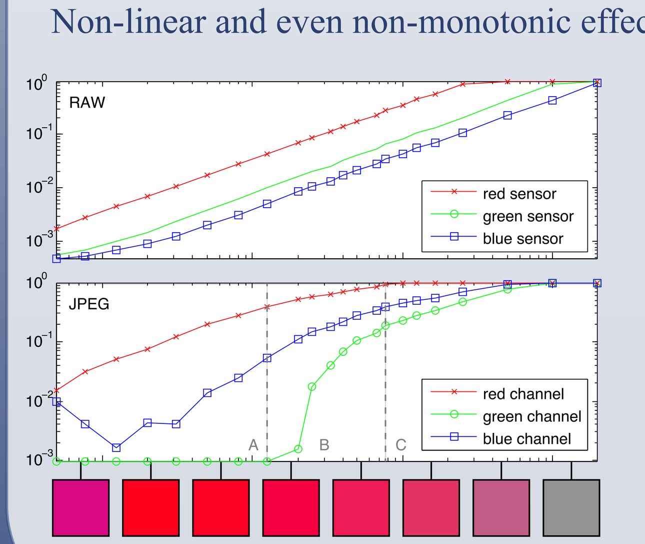

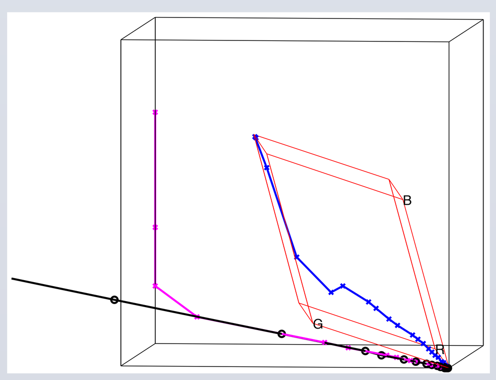

Inverse (de-rendering) Model

- The inverse process is one-to-many: use Gaussian Process (GP) regression to produce a *distribution* instead of a single point estimate

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

- The input noise are non-stationary: use local GP regression

$$p_x(\mathbf{x}|\mathbf{y}) = \prod_c p_{GP}(x^c|\mathcal{D}_{N(\mathbf{y})}, \mathbf{y})$$

We quickly collect thousands of RAW/JPEG color pairs by imaging a color checker under many different illuminations. (Done once for each imaging mode of each camera.)

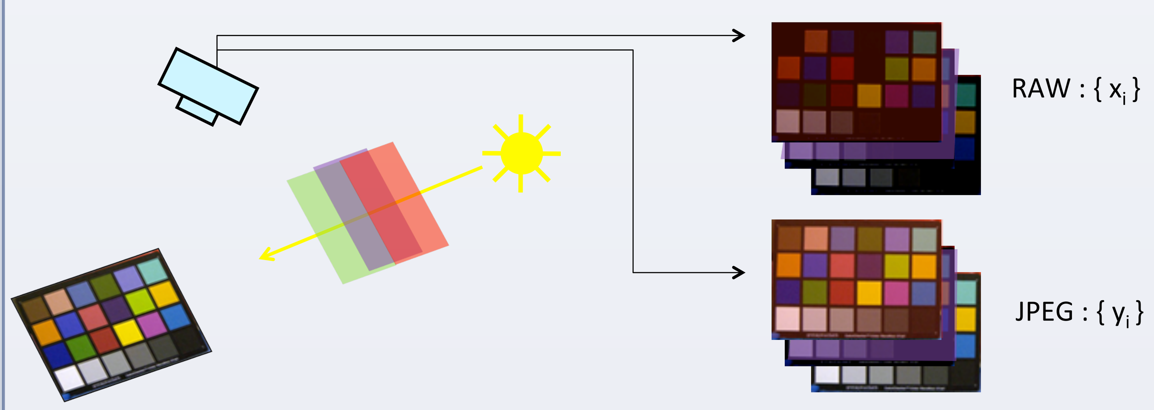

Variance visualization in a chromaticity slice of output sRGB color space. (Panasonic DMC-LX3)

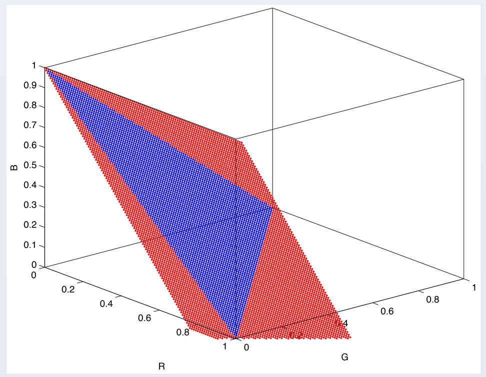

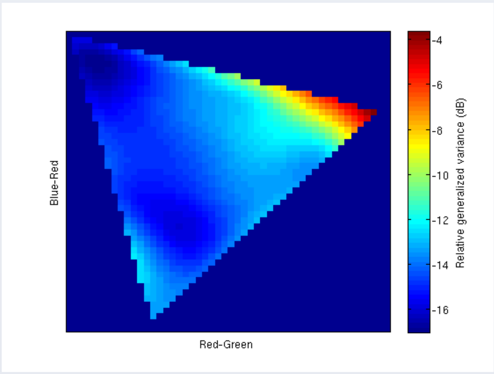

Application I: Probabilistic multi-exposure imaging

Combine measurements at different exposure to produce HDR and wide-gamut images.

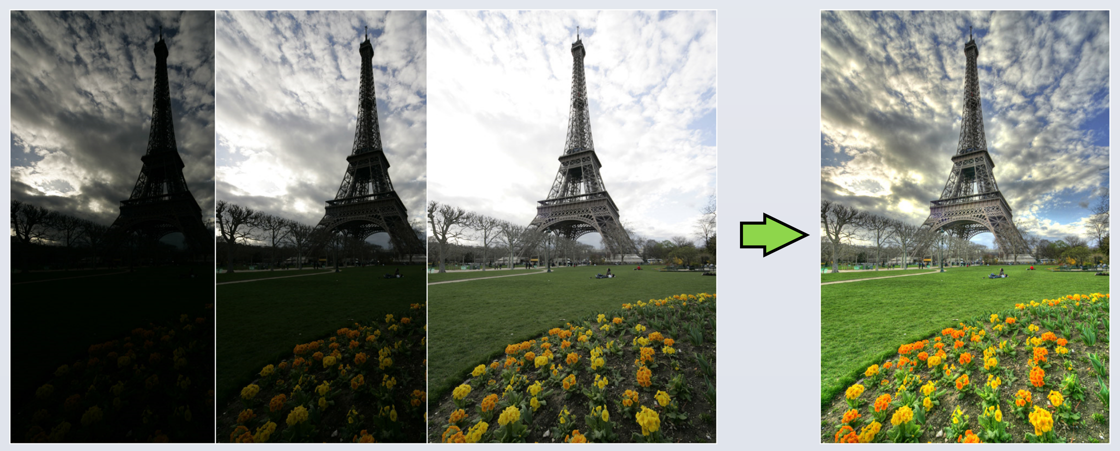

Traditional HDR approach: weigh all measurements equally

$$\mathbf{x_0} = \frac{1}{n}\sum_i \frac{\alpha_0}{\alpha_i}\mathbf{x_i} = \frac{1}{n}\sum_i \frac{\alpha_0}{\alpha_i} f^{-1}(\mathbf{y_i})$$

Probabilistic approach: weigh measurements according to confidence

$$p_{x_0}(\mathbf{x_0}|\mathbf{y_1},\ldots,\mathbf{y}_N) = \prod_i p_{x_0}(\mathbf{x_0}|\mathbf{y_i}) = \prod_i \frac{\alpha_i}{\alpha_0} p_{x_i}\left(\frac{\alpha_i}{\alpha_0}\mathbf{x_0}|\mathbf{y_i}\right)$$

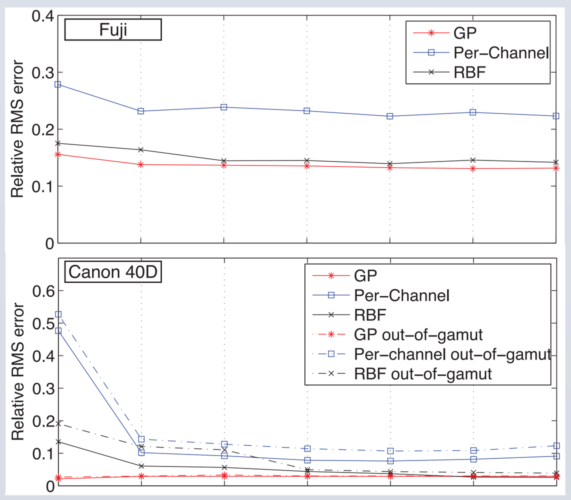

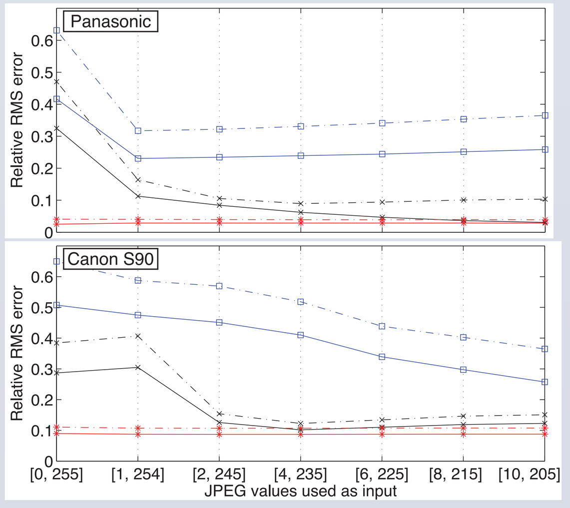

Application II: Probabilistic photometric stereo

Recover 3D surface by observing the object under different lighting condition.

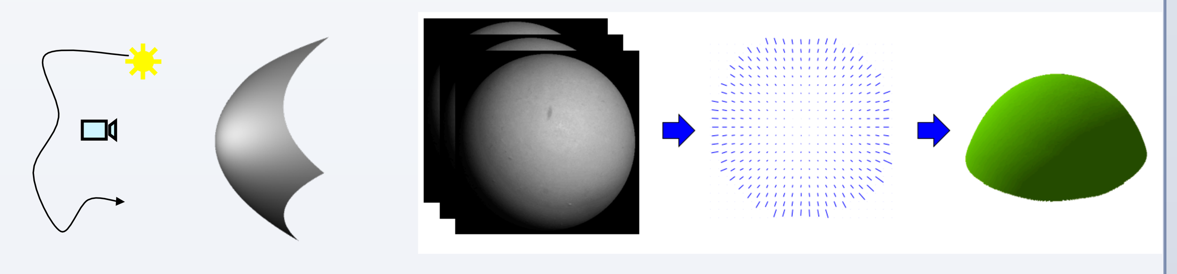

Traditional approach: least square solution

$$\mathbf{l}_i^T\mathbf{b} = \mu_i$$
$$\mathbf{b} = (\boldsymbol{L}^T\boldsymbol{L})^{-1}\boldsymbol{L}^T\boldsymbol{\mu}$$

Probabilistic approach: weighted least square solution

$$\mathbf{l}_i^T\mathbf{b} = \mu_i + \sigma_i\epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0,1)$$
$$\mathbf{b} = (\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L})^{-1}\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{\mu},$$
$$\boldsymbol{W} = \mathrm{diag}\{\sigma_i^{-2}\}_{i=1}^N$$

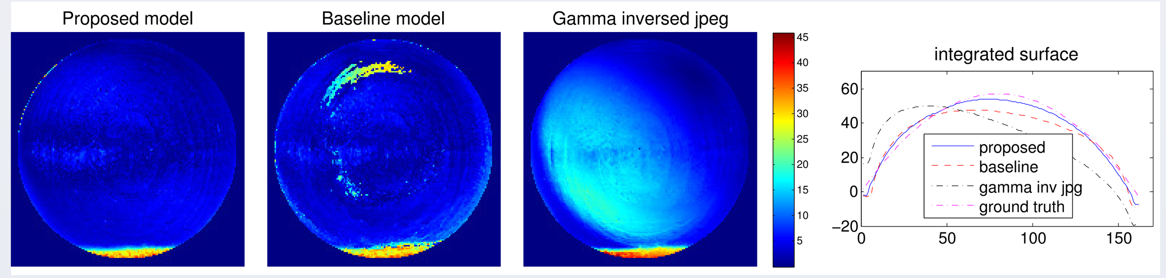

Conclusions

- Most computer vision algorithms requires RAW images, while most available Internet images exists in narrow-gamut sRGB.
- We propose a de-rendering algorithm to recover physical color values from reported sRGB (JPEG) pixel values.
- The proposed model is probabilistic, embracing the multivalued nature of the de-rendering map.
- The output distribution can be used in probabilistic version of many tradition computer vision applications, such as de-blurring, de-hazing, color constancy, image-based modeling, object recognition, etc.

References

[1] A. Chakrabarti, D. Scharstein, and T. Zickler. An empirical camera model for internet color vision. In BMVC, 2009.

[2] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. PAMI, 10(4):439–451, 1988.

[3] M. Grossberg and S. Nayar. Modeling the space of camera response functions. PAMI, 26 (10):1272–1282, 2004.

[4] H. Lin, S. J. Kim, S. Susstrunk, and M. Brown. Revisiting radiometric calibration for color computer vision. In ICCV, 2011.

[5] T. Mitsunaga and S. Nayar. Radiometric self calibration. In CVPR, 1999.

[6] C. Pal, R. Szeliski, M. Uyttendaele, and N. Jojic. Probability models for high dynamic range imaging. In CVPR, 2004.

[7] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec. High dynamic range imaging. Elsevier, 2006.

Acknowledgement

The authors thank Ayan Chakrabarti for many helpful discussions and for providing code from [1]. Authors Ying Xiong and Todd Zickler were supported by the NSF awards IIS-0905243, CRI-0708895, IIS-0546408. Authors Kate Saenko and Trevor Darrell were supported by DARPA contract W911NF-10-2-0059, by NSF awards IIS-0905647 and IIS-0819984, and from Toyota and Google.