Intrinsic Images by Entropy Minimization

How Peter Pan Really Lost His Shadow

Presentation By: Jordan Frank

Based on work by Finlayson, Hordley, Drew, and Lu [1,2,3]

Image from http://dvd.monstersandcritics.com/reviews/article_1273419.php/DVD_Review_Peter_Pan__Two-Disc_Platinum_Edition_

Colour Constancy

- Humans automatically remove the effect of lighting in visual perception. Our vision system is very robust to illumination.
- We can easily tell that the bricks in the sunlight are the same colour as the bricks in the shade.

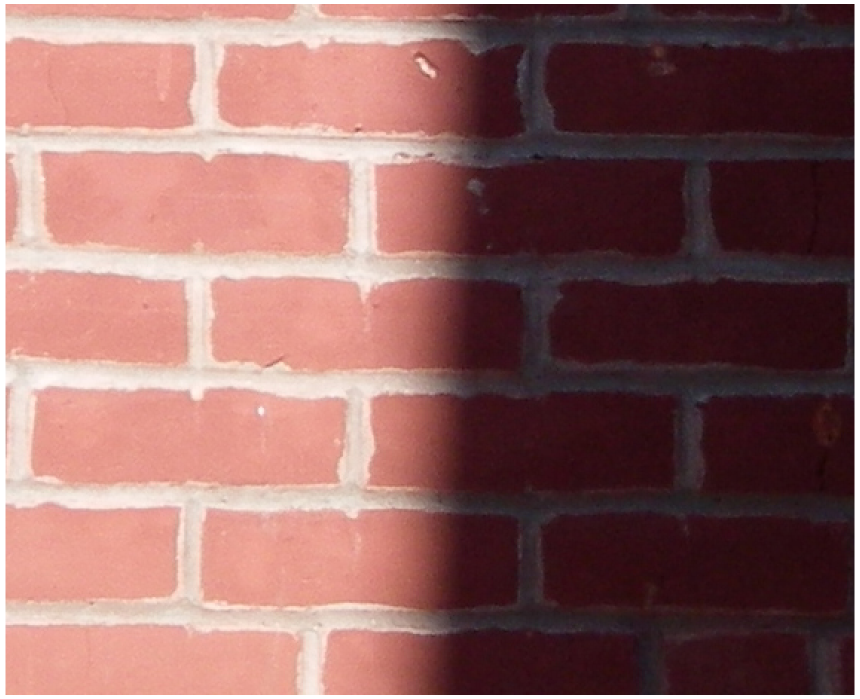

Colour Constancy

- However, this is an Ill-posed problem.
- We cannot tell whether differences in the colour that we detect are due to differences in the colour of the object, or differences in illumination.
- What if the bricks on the right really were darker? In fact, there are small differences in the colours, even within a single brick.

Invariant Image

- Goal is to produce an image that is invariant to effects of illumination.
- Motivation:
  - Add our own illumination.
  - Remove shadows! Almost every presentation prior to this one has talked about how shadows are problematic.

Cameras/Sensors

- RGB colour at a pixel results from an integral over the visible wavelength

$$R_k = \sigma \int E(\lambda) S(\lambda) Q_k(\lambda) d\lambda, k = R, G, B \tag{1}$$

σ – Lambertian shading
$E(\lambda)$ – illumination spectral power distribution
$S(\lambda)$ – surface spectral reflectance
$Q_k(\lambda)$ – camera sensitivity

Cameras/Sensors

- For convenience we assume that camera sensitivity is exactly a Dirac delta function

$$Q_k(\lambda) = q_k \delta(\lambda - \lambda_k)$$

- $q_k = Q_k(\lambda_k)$ is the strength of the sensor.
- So (1) reduces to

$$R_k = \sigma E(\lambda_k) S(\lambda_k) q_k$$

More Approximations

- Supposing that lighting can be approximated by Planck's law, with Wien's approximation [4] we get

$$R_k = \sigma I k_1 \lambda_k^{-5} e^{-\frac{k_2}{T\lambda}} S(\lambda_k) q_k \qquad (2)$$

- k_1, k_2 are constants, temperature T characterizes the lighting colour and I gives the overall light intensity.

Removing I and σ

- We can effectively remove the effect of Lambertian shading and illumination from (2) by dividing to get the band-ratio 2-vector chromaticities $\boldsymbol{c}$,

$$c_k = R_k / R_p,$$

where p is one of the channels and k=1,2 indexes over the remaining responses.

Log Chromaticities

- Log chromaticities are independent of illumination intensity, and translate under change of illumination colour.
- $(\ln R/G, \ln B/G)$ under one light becomes $(\ln R/G, \ln B/G)+(a,b)$ under a second light.
- More importantly, the translational term for different illuminants can always be written as $(\alpha a, \alpha b)$ where a,b are constants and α depends on illumination.

Invariance, finally

- Illumination change translates log chromaticities in the same direction.
- Therefore, the coordinate axis orthogonal to the direction of illumination variation, $y = -(a/b)x$, records only illuminant invariant information.

Example (simulated)

Measured log chromaticities

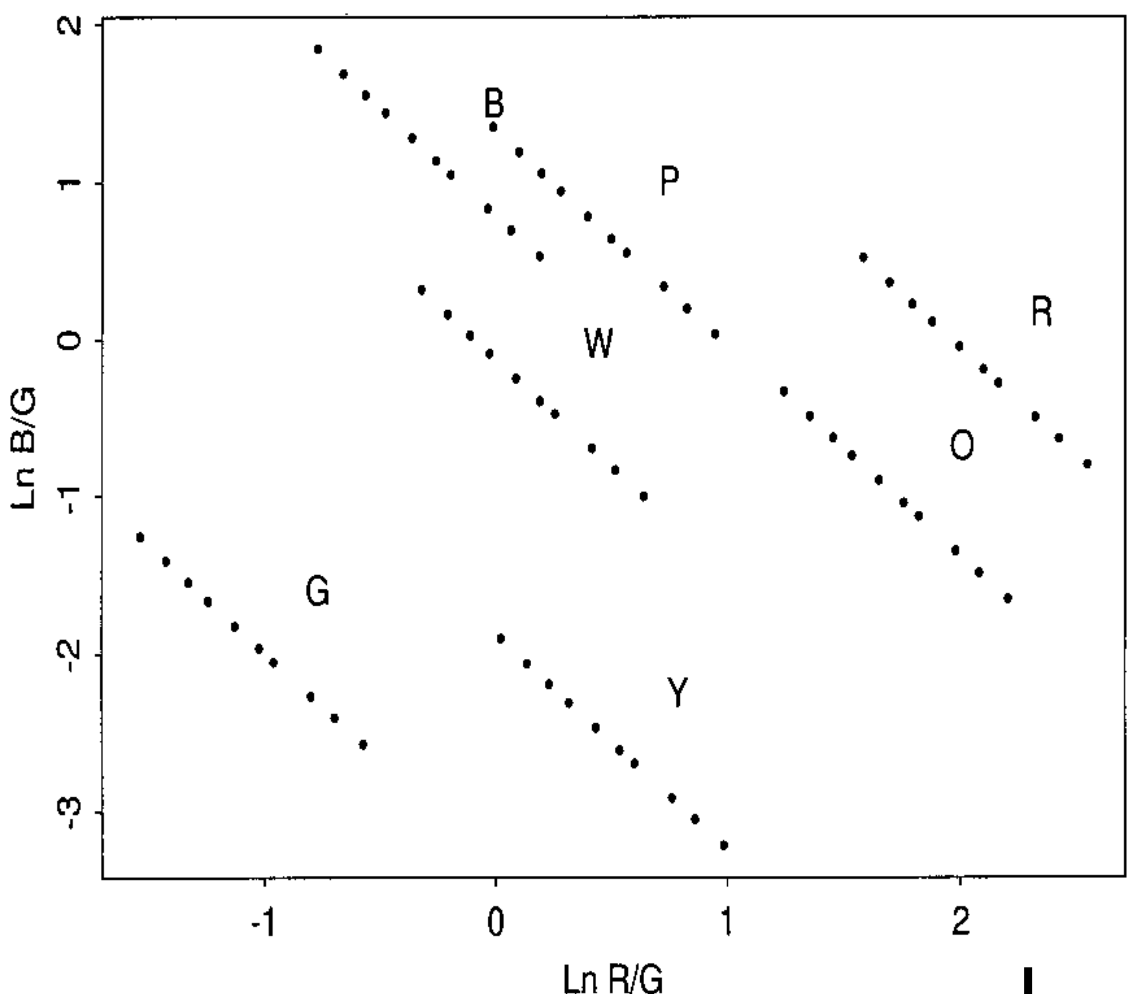

After rotation

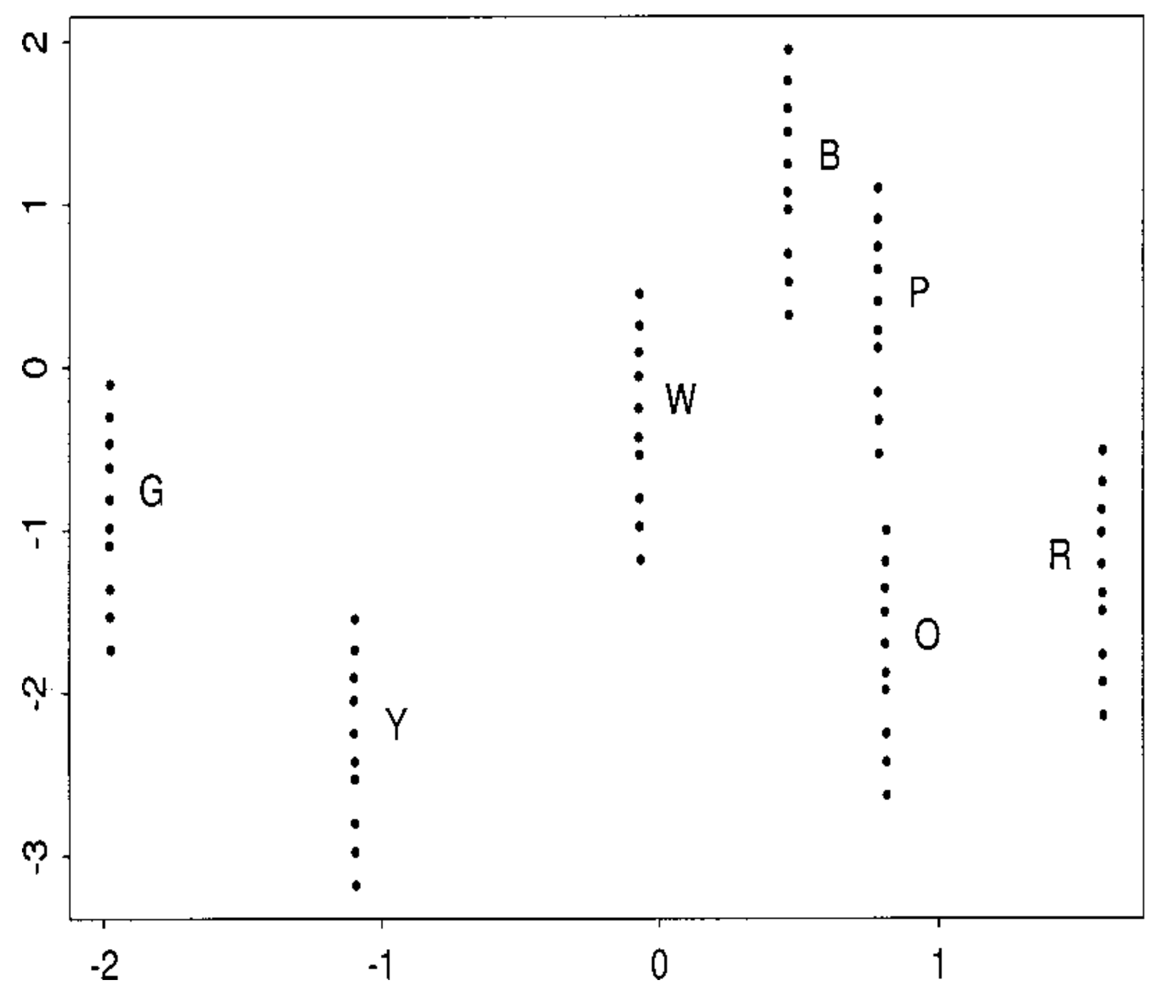

Figures from [3]

Great in theory, but...

- There are a number of problems with this approach.
  - Cameras do not have Dirac delta responses.
  - Noise!
  - How do we determine how much to rotate.

Narrowband, not Dirac

Actual measurements from a SONY DXC-930 camera

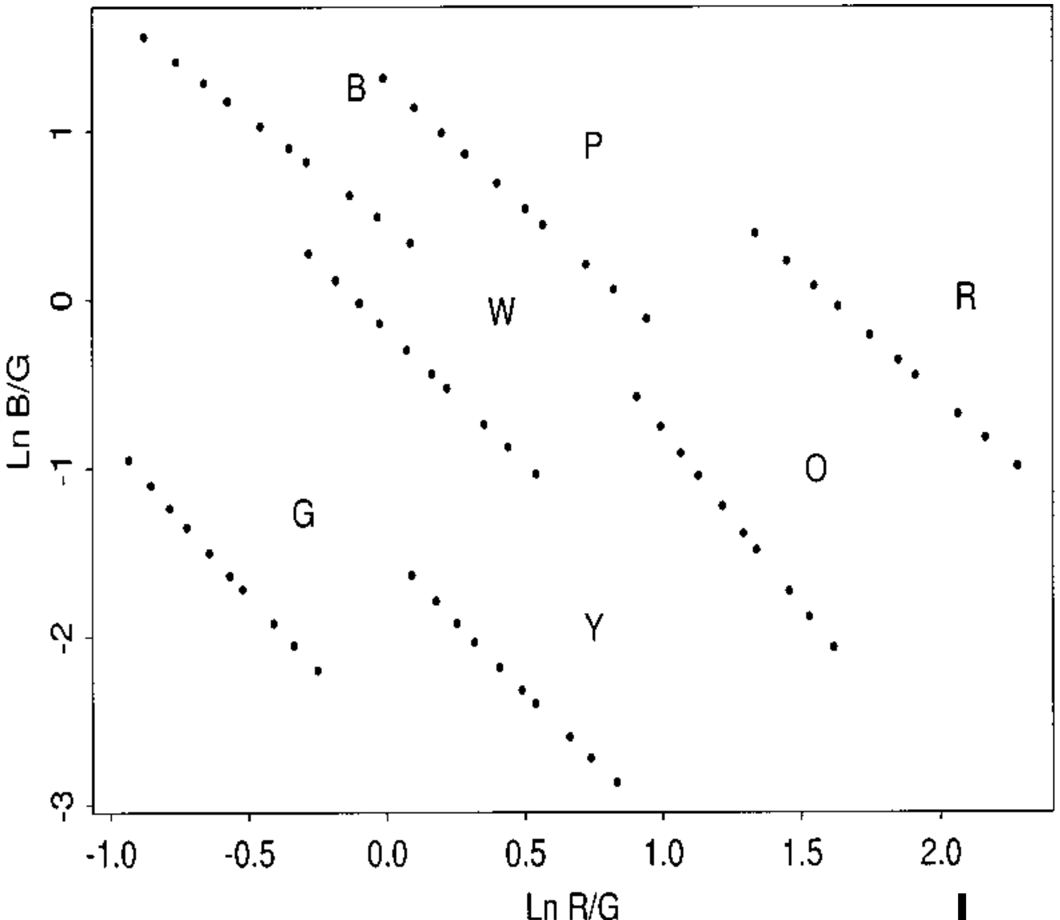

Rotated

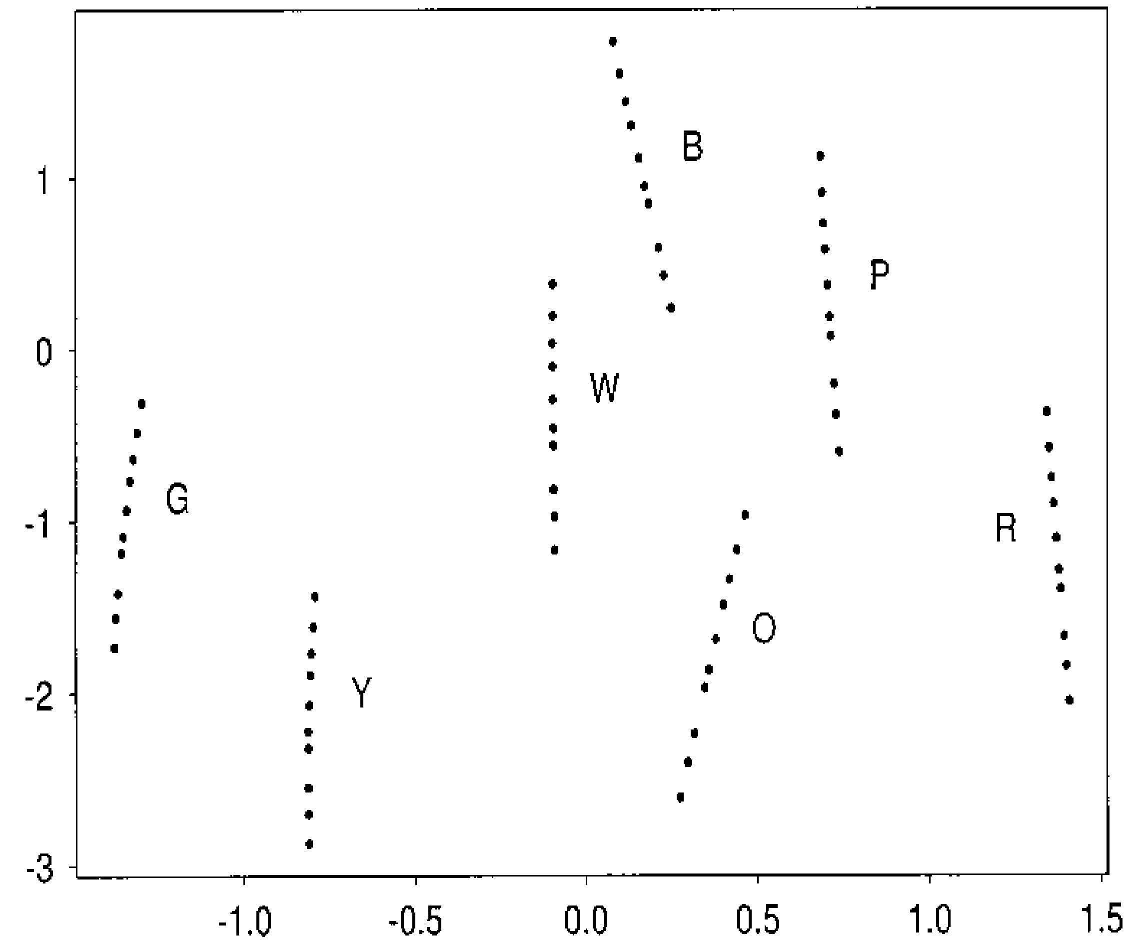

So we are still doing pretty well.

Figures from [3]

How to calibrate

- We can carefully calibrate our cameras to determine the best rotation angle through inspection of log chromaticities.

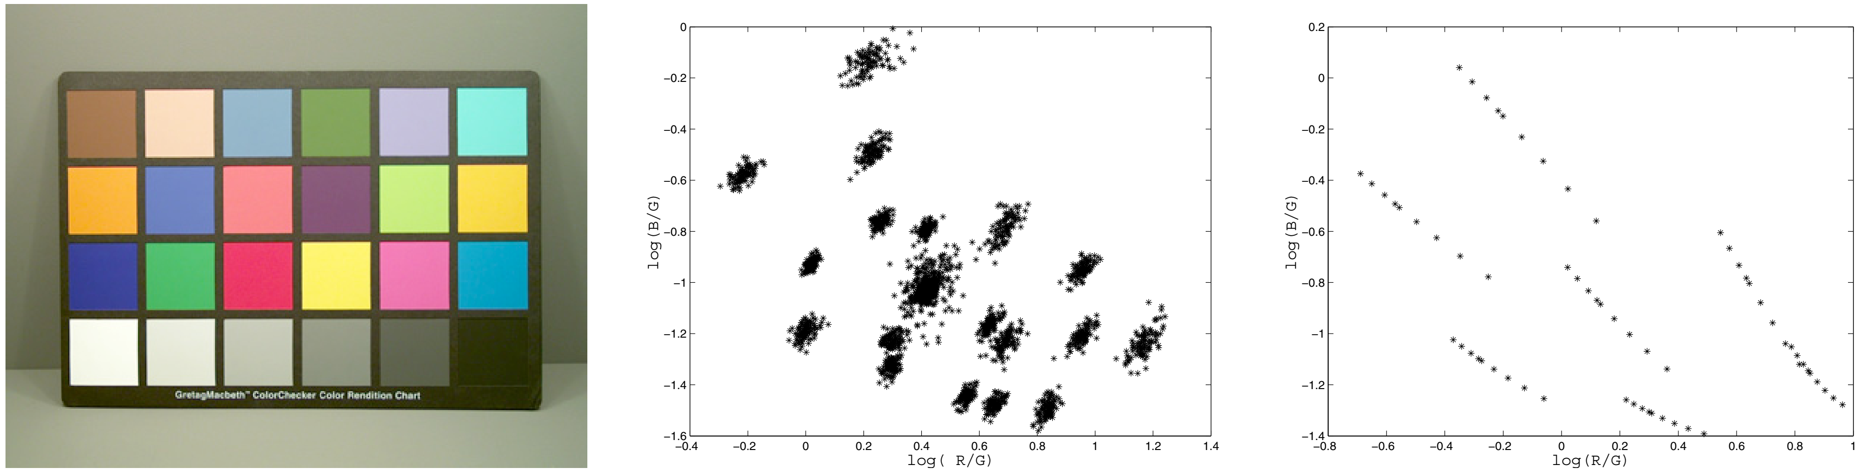

(a) (b) (c)

Fig. 2. (a): Macbeth ColorChecker Chart image under a Planckian light. (b): Log-chromaticities of the 24 patches. (c): Median chromaticities for 6 patches, imaged under 14 different Planckian illuminants.

Figure and caption from [1]

How to calibrate

- Manual calibration is time-consuming, and requires numerous images under various lighting conditions.
- We would like to find a way to find the desired parameters from just one image, without even knowing the camera that was used.

Intuition

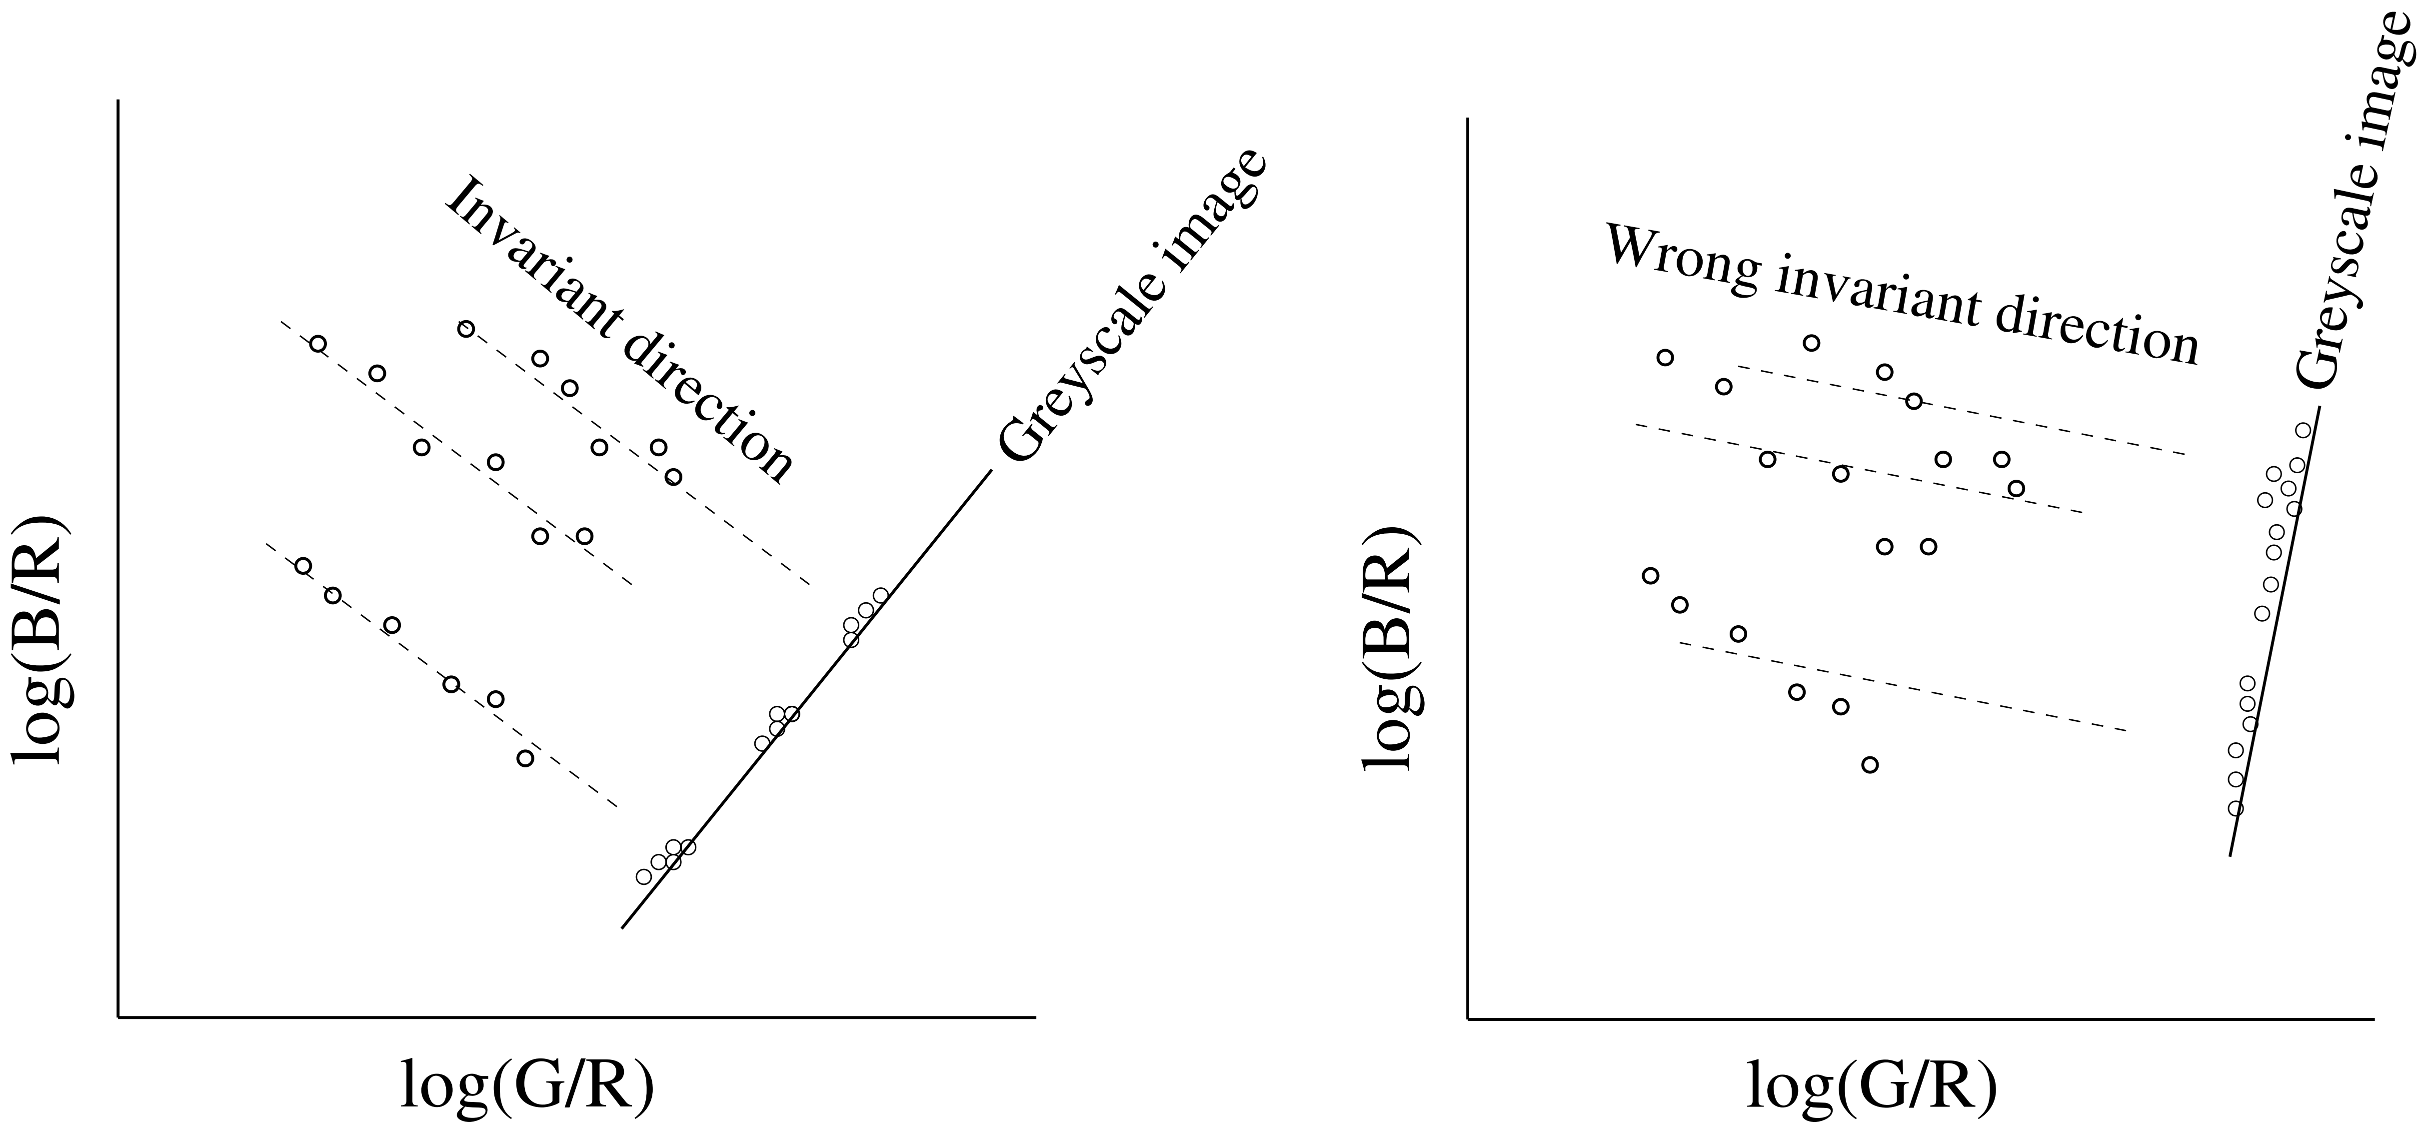

And so what might be a good way to determine the best invariant direction?
(Hint: we learned about it in this course)

Figures from [1]

Entropy to the Rescue

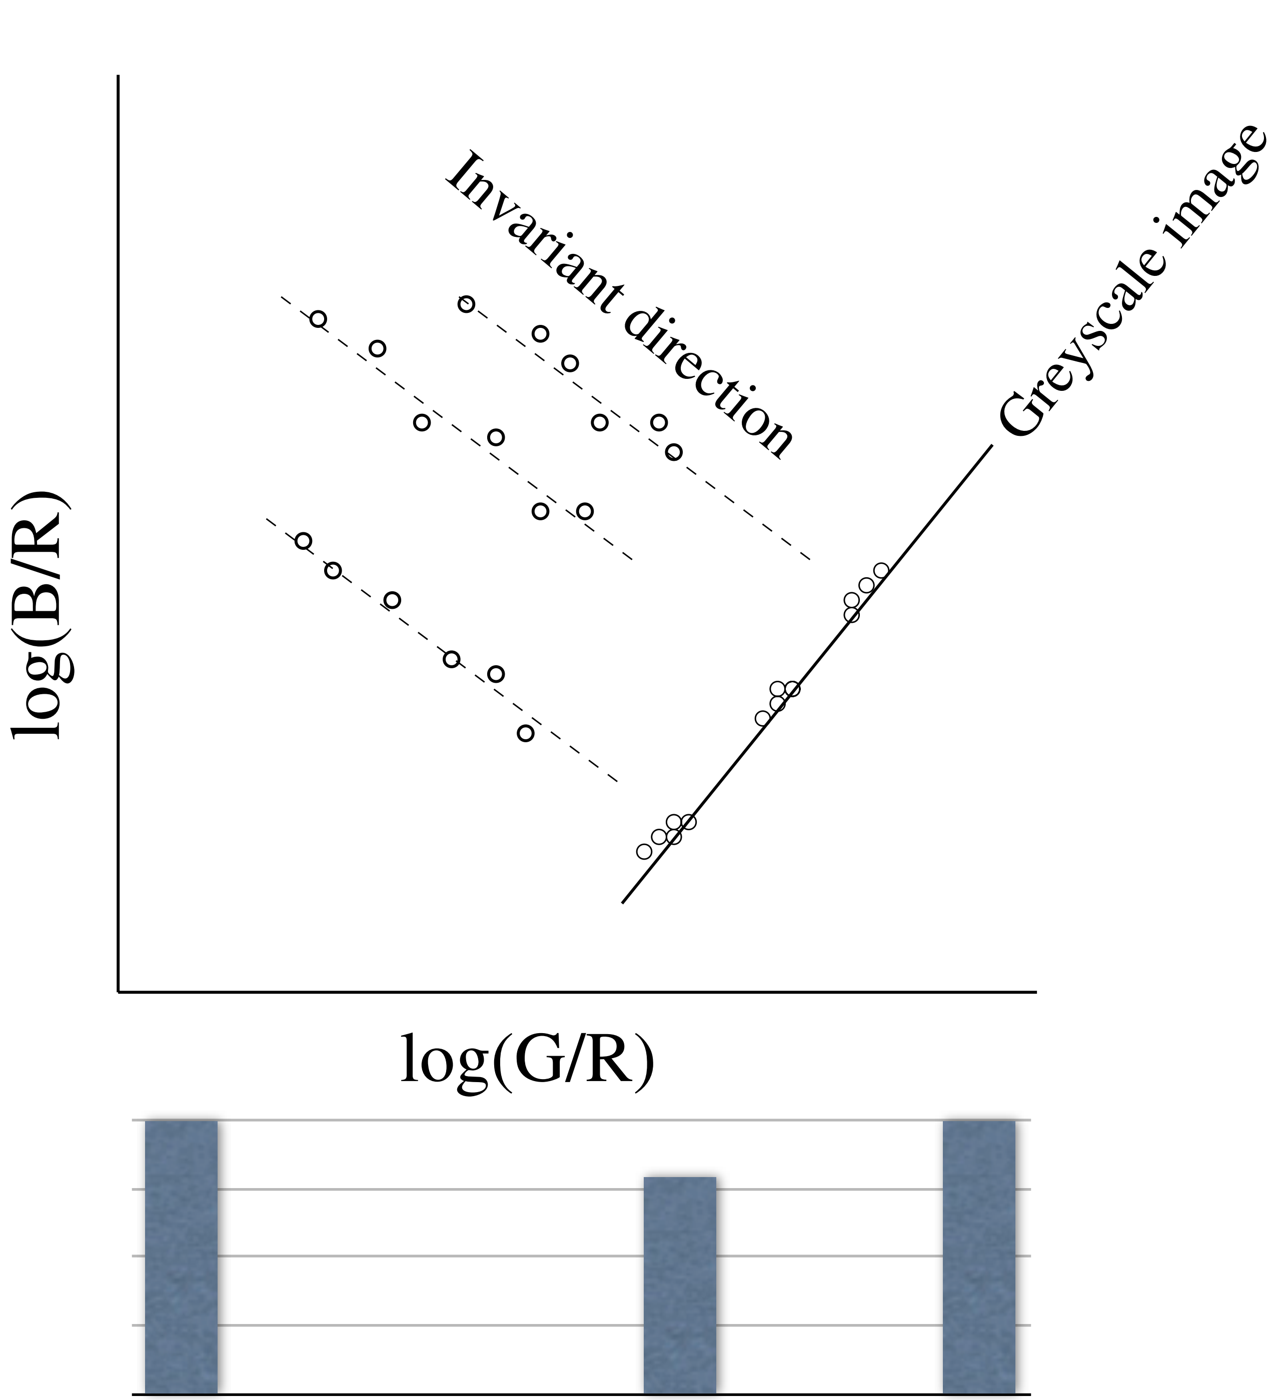

Low Entropy

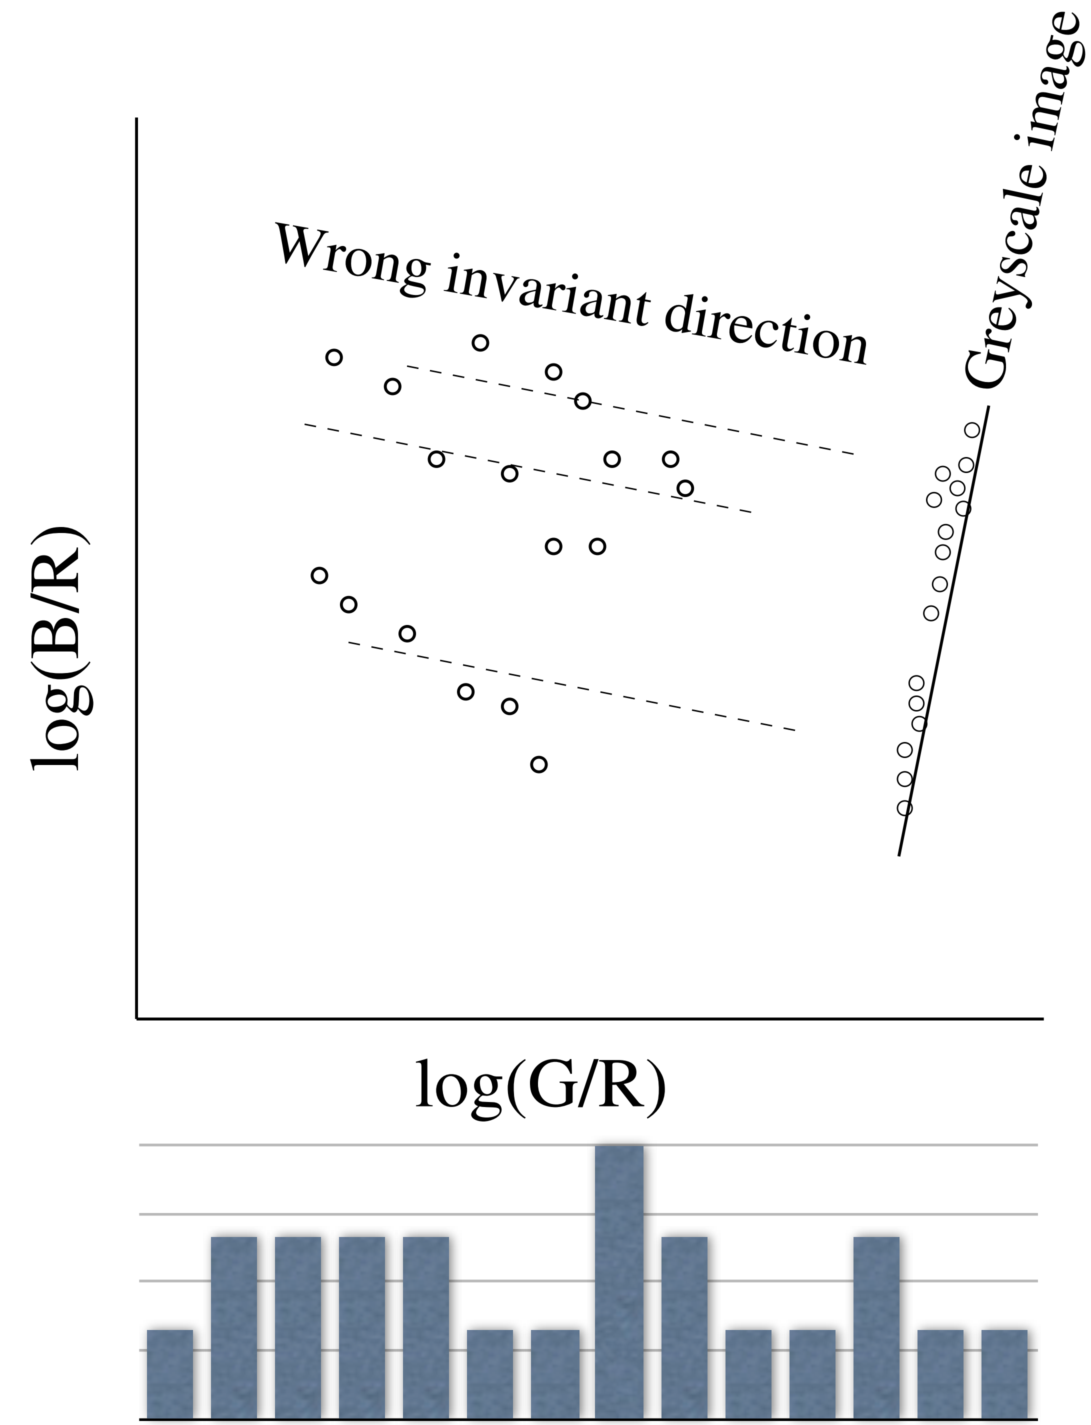

High Entropy

Figures from [1]

Entropy Minimization

Algorithm:

1. Form a 2D log-chromaticity representation of the image.
2. for $\theta = 1..180$
   a) Rotate by θ and take projection onto x-axis
   b) Calculate entropy
   c) Keep track of θ that minimizes entropy

Problems/Solutions

- How to calculate entropy?
- Noise in the data?
- How to go back from rotated data to a usable image?

Problems/Solutions

- How to calculate entropy?
  Solution: Create a histogram, compute bin widths using Scott's Rule:
  $\text{bin\_width} = 3.49\ std(\text{projected\_data})\ N^{1/3}$
- Noise in the data?
- How to go back from rotated data to a usable image?

Problems/Solutions

- How to calculate entropy?
  Solution: Create a histogram, compute bin widths using Scott's Rule:
  bin_width = 3.49 *std*(projected_data) $N^{1/3}$
- Noise in the data?
  Solution: Use only the middle 90% of the projected data.
- How to go back from rotated data to a usable image?

Problems/Solutions

- How to calculate entropy?
  Solution: Create a histogram, compute bin widths using Scott's Rule:
  $\text{bin\_width} = 3.49\ \mathit{std}(\text{projected\_data})\ N^{1/3}$
- Noise in the data?
  Solution: Use only the middle 90% of the projected data.
- How to go back from rotated data to a usable image?
  Solution: See the paper, not easy!

Example

Original Image

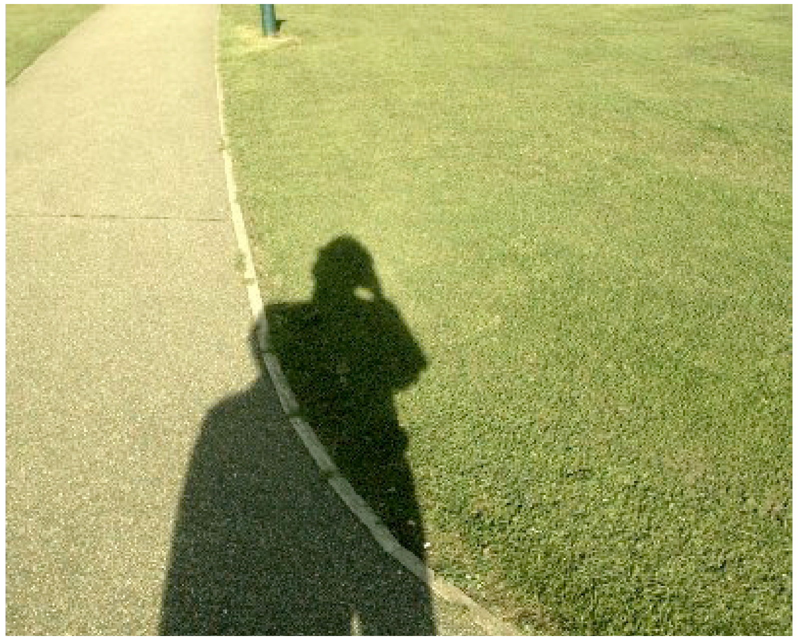

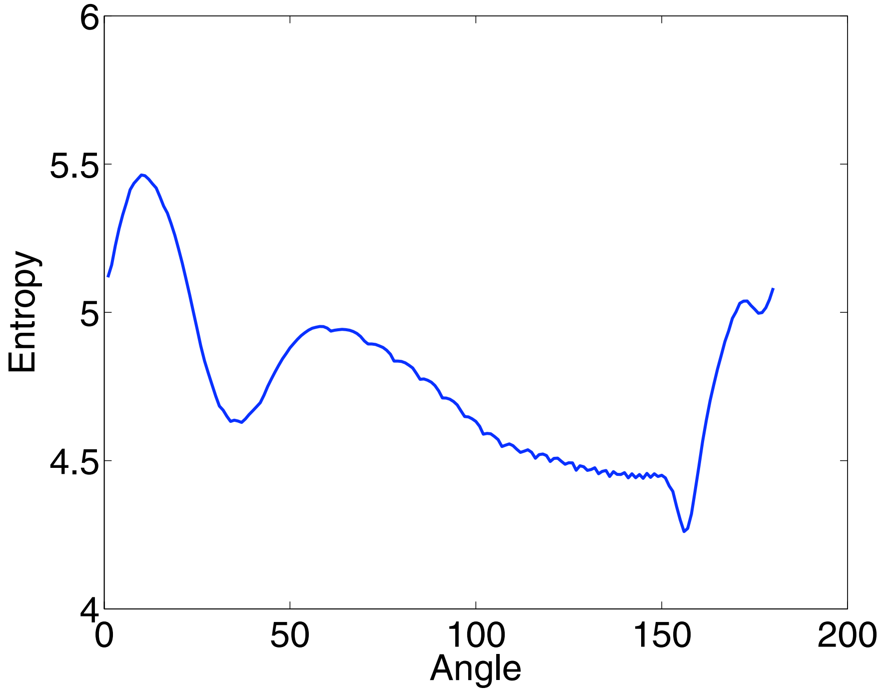

Min-Entropy Projection

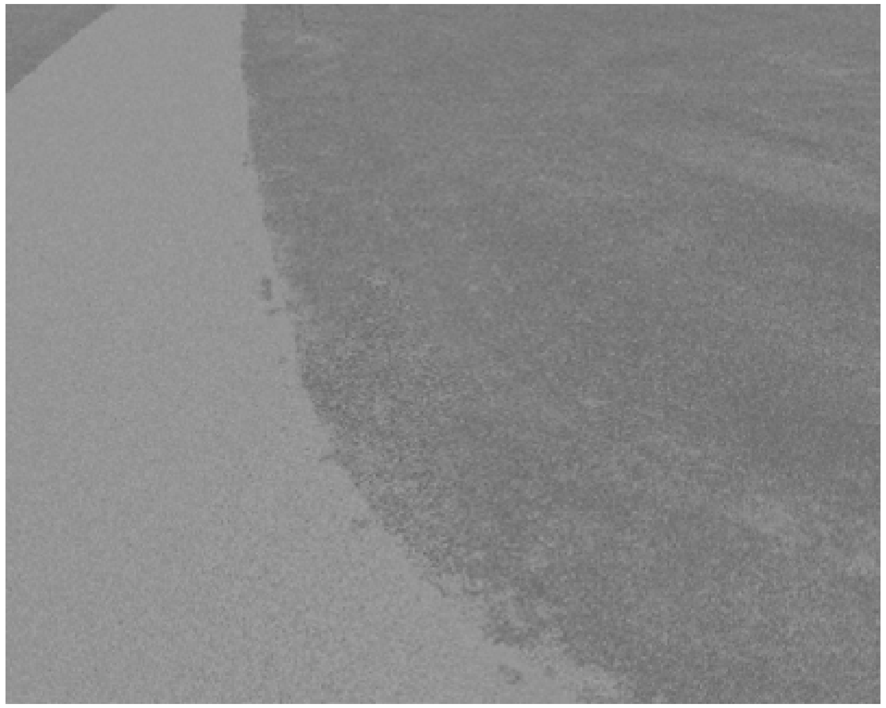

Figures from [1]

Removing Shadows

- Now that we have an image without shadows, how can we use this to remove shadows from the original image?

Image from http://collectibles.about.com/library/priceguides/blpgDSpeterpan903.htm

Removing Shadows

- Need to reintegrate the illumination invariant image into the original image.
- Create an edge map for the Mean-Shift processed original image as well as the invariant image.
- If the magnitude of the gradient in the invariant image is close to zero where the gradient in log response in the original image is high, then this is evidence of a shadow edge.

Removing Shadows

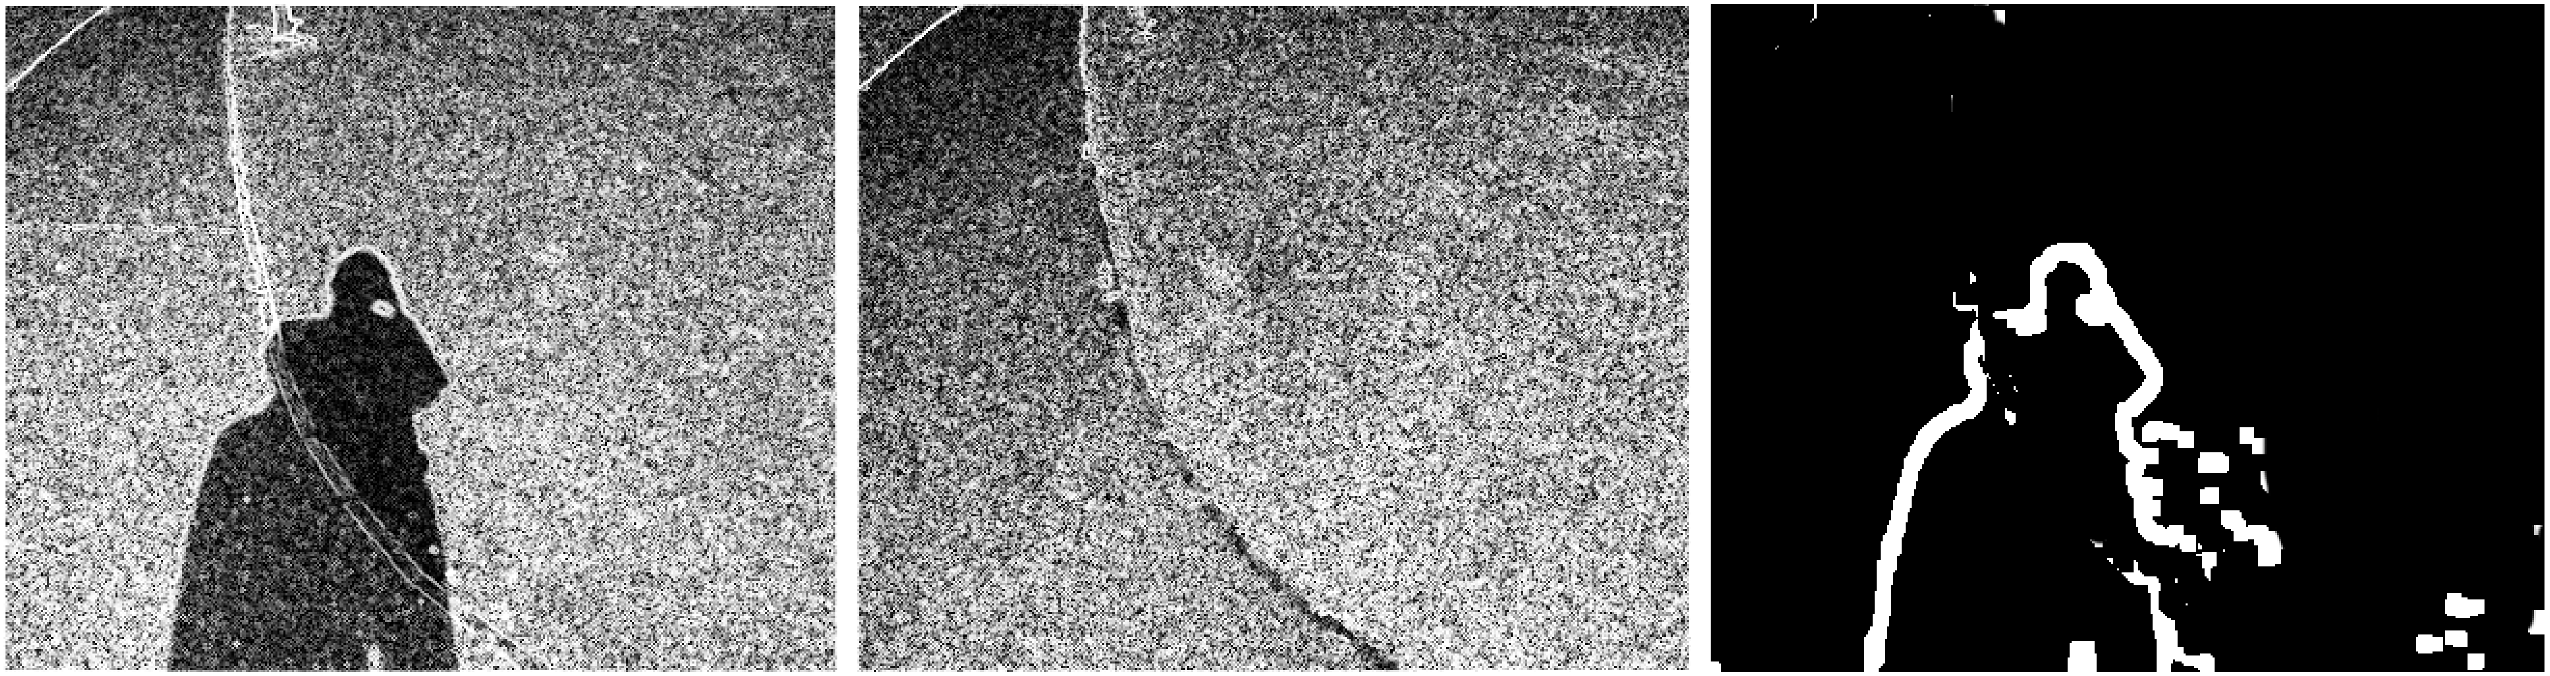

Left to right, Edges in the original image, edges in the invariant image, recovered shadow edge.

Figures from [2]

Removing Shadows

- All edges in the image that are not shadow edges are indicative of material changes. There are no sharp changes due to illumination and so shadows have been removed.
- Reintegrating the gradient gives a log response image which does not have shadows. For details, see the paper.
- To get back to a realistic image, we simply have to add artificial illumination.

Results (Paper)

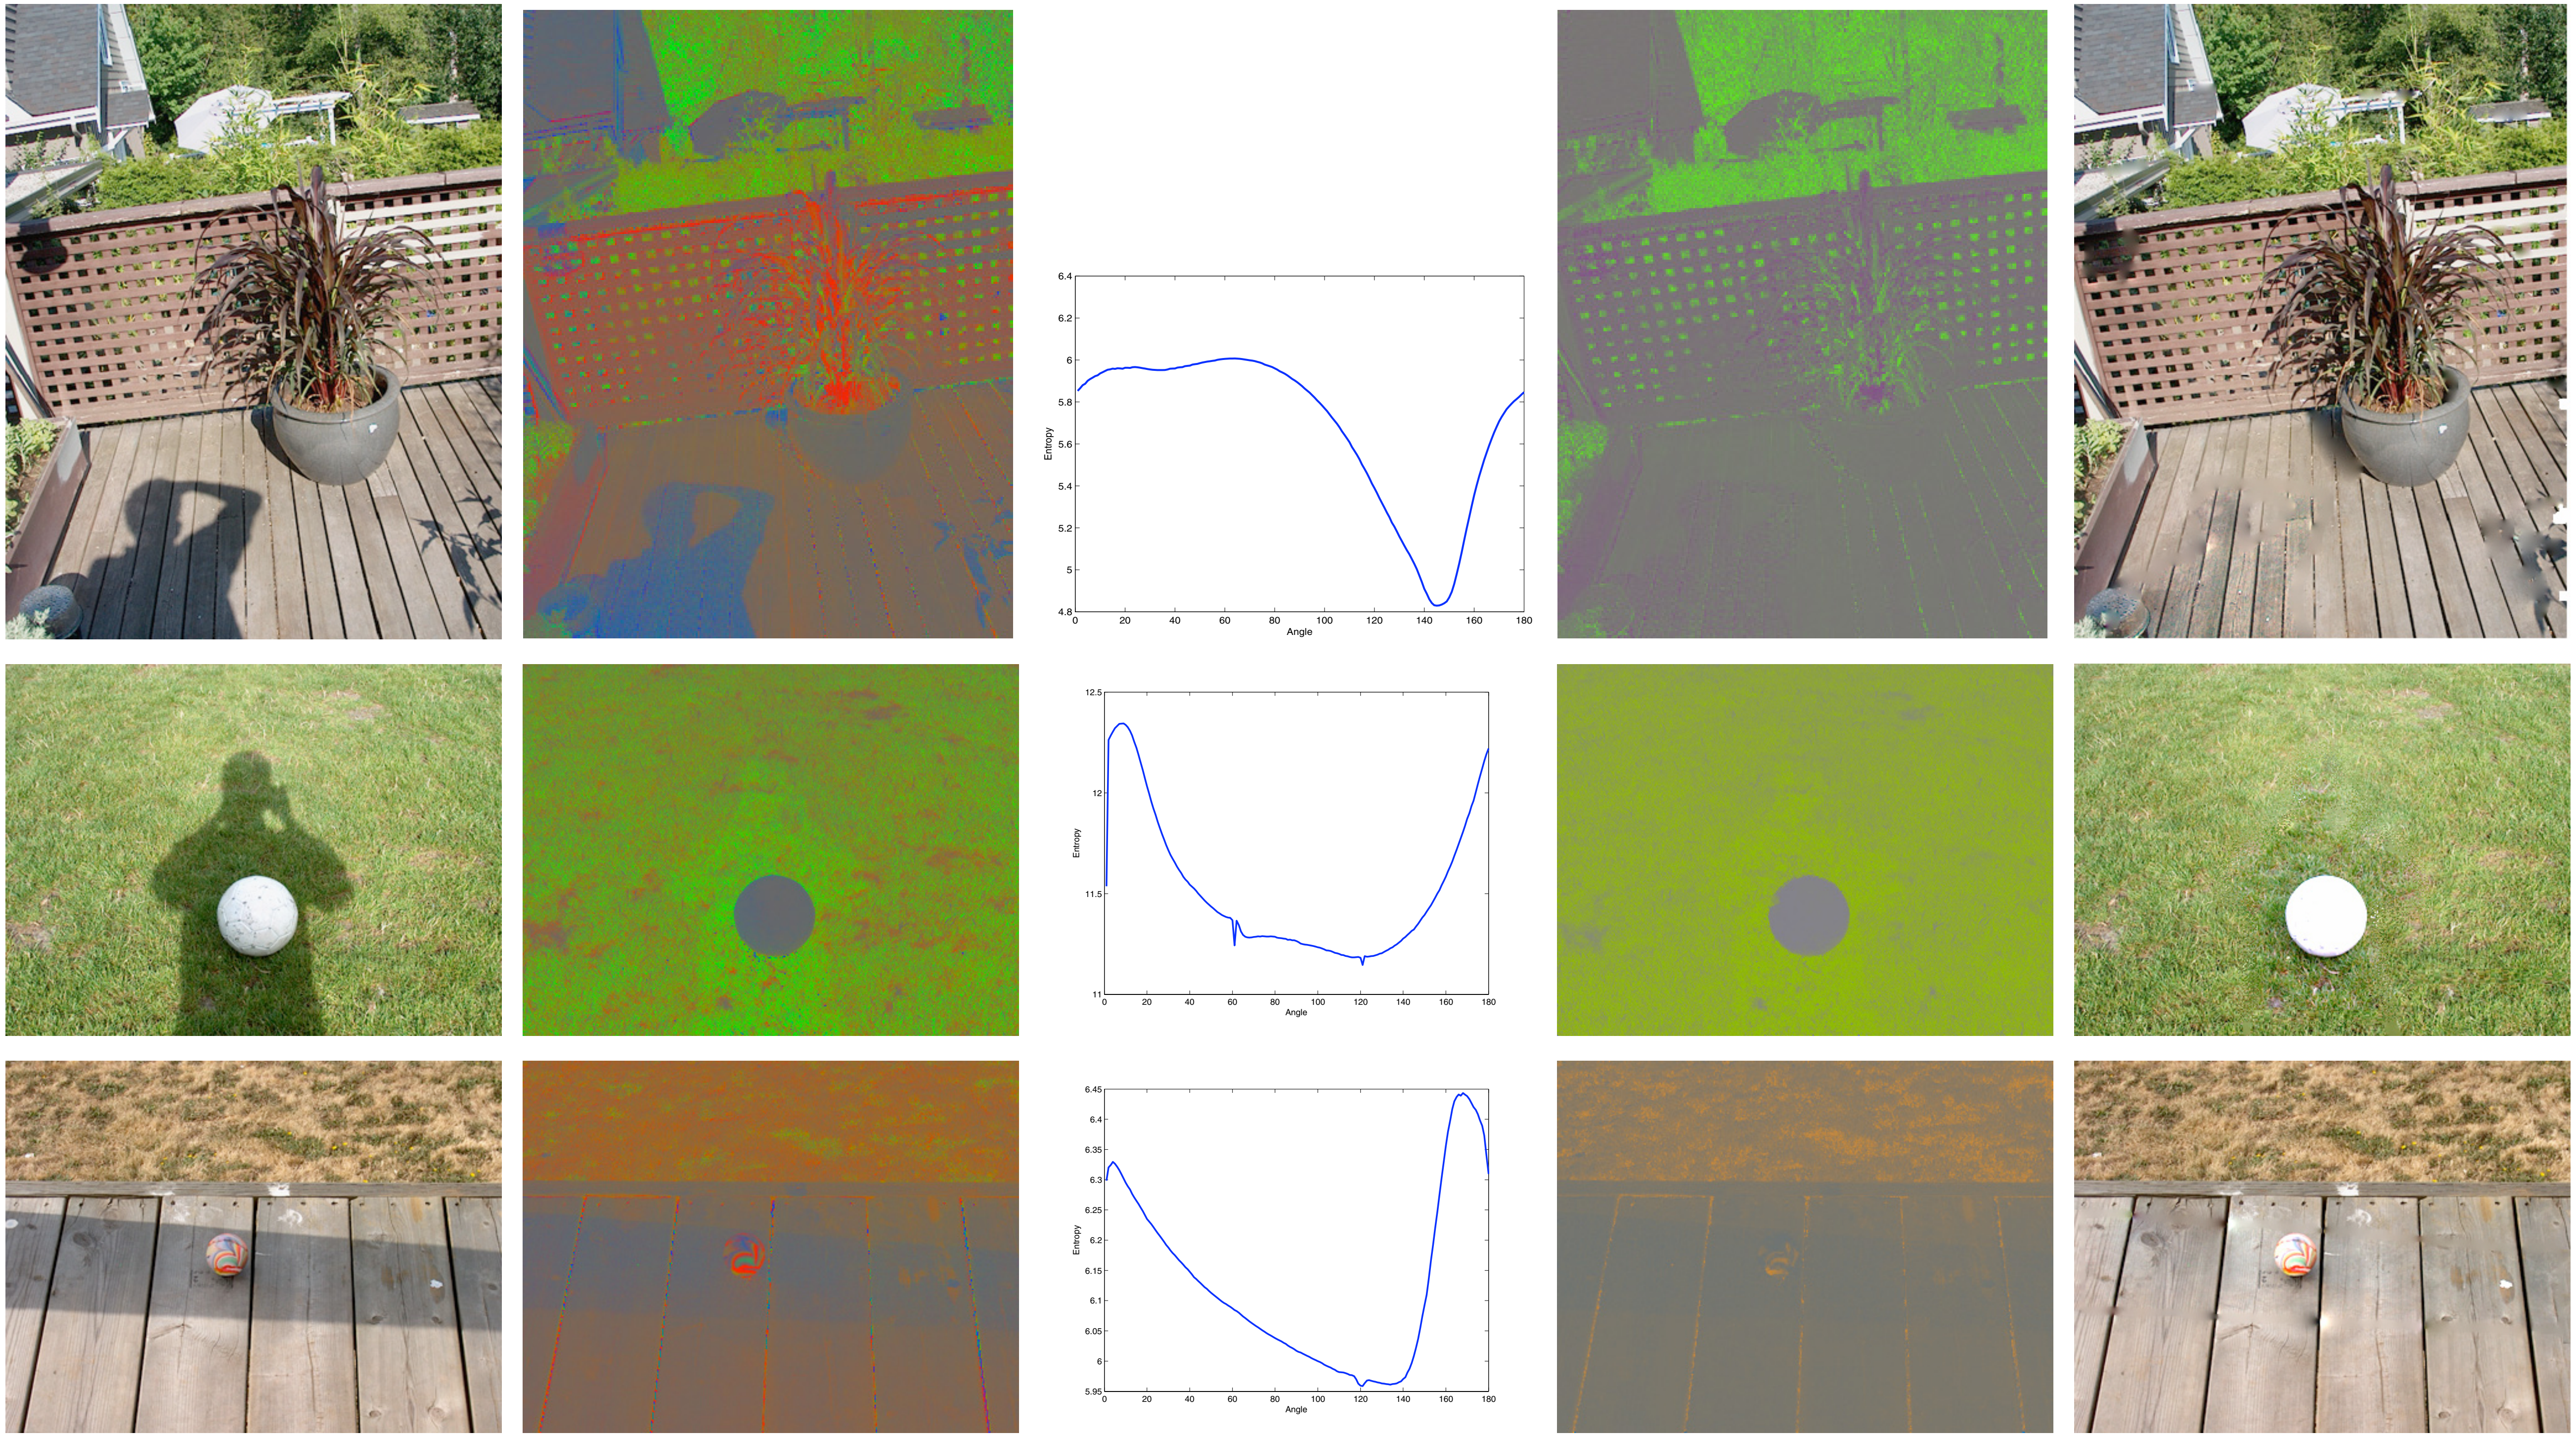

Figures from [3]

Results (Paper)

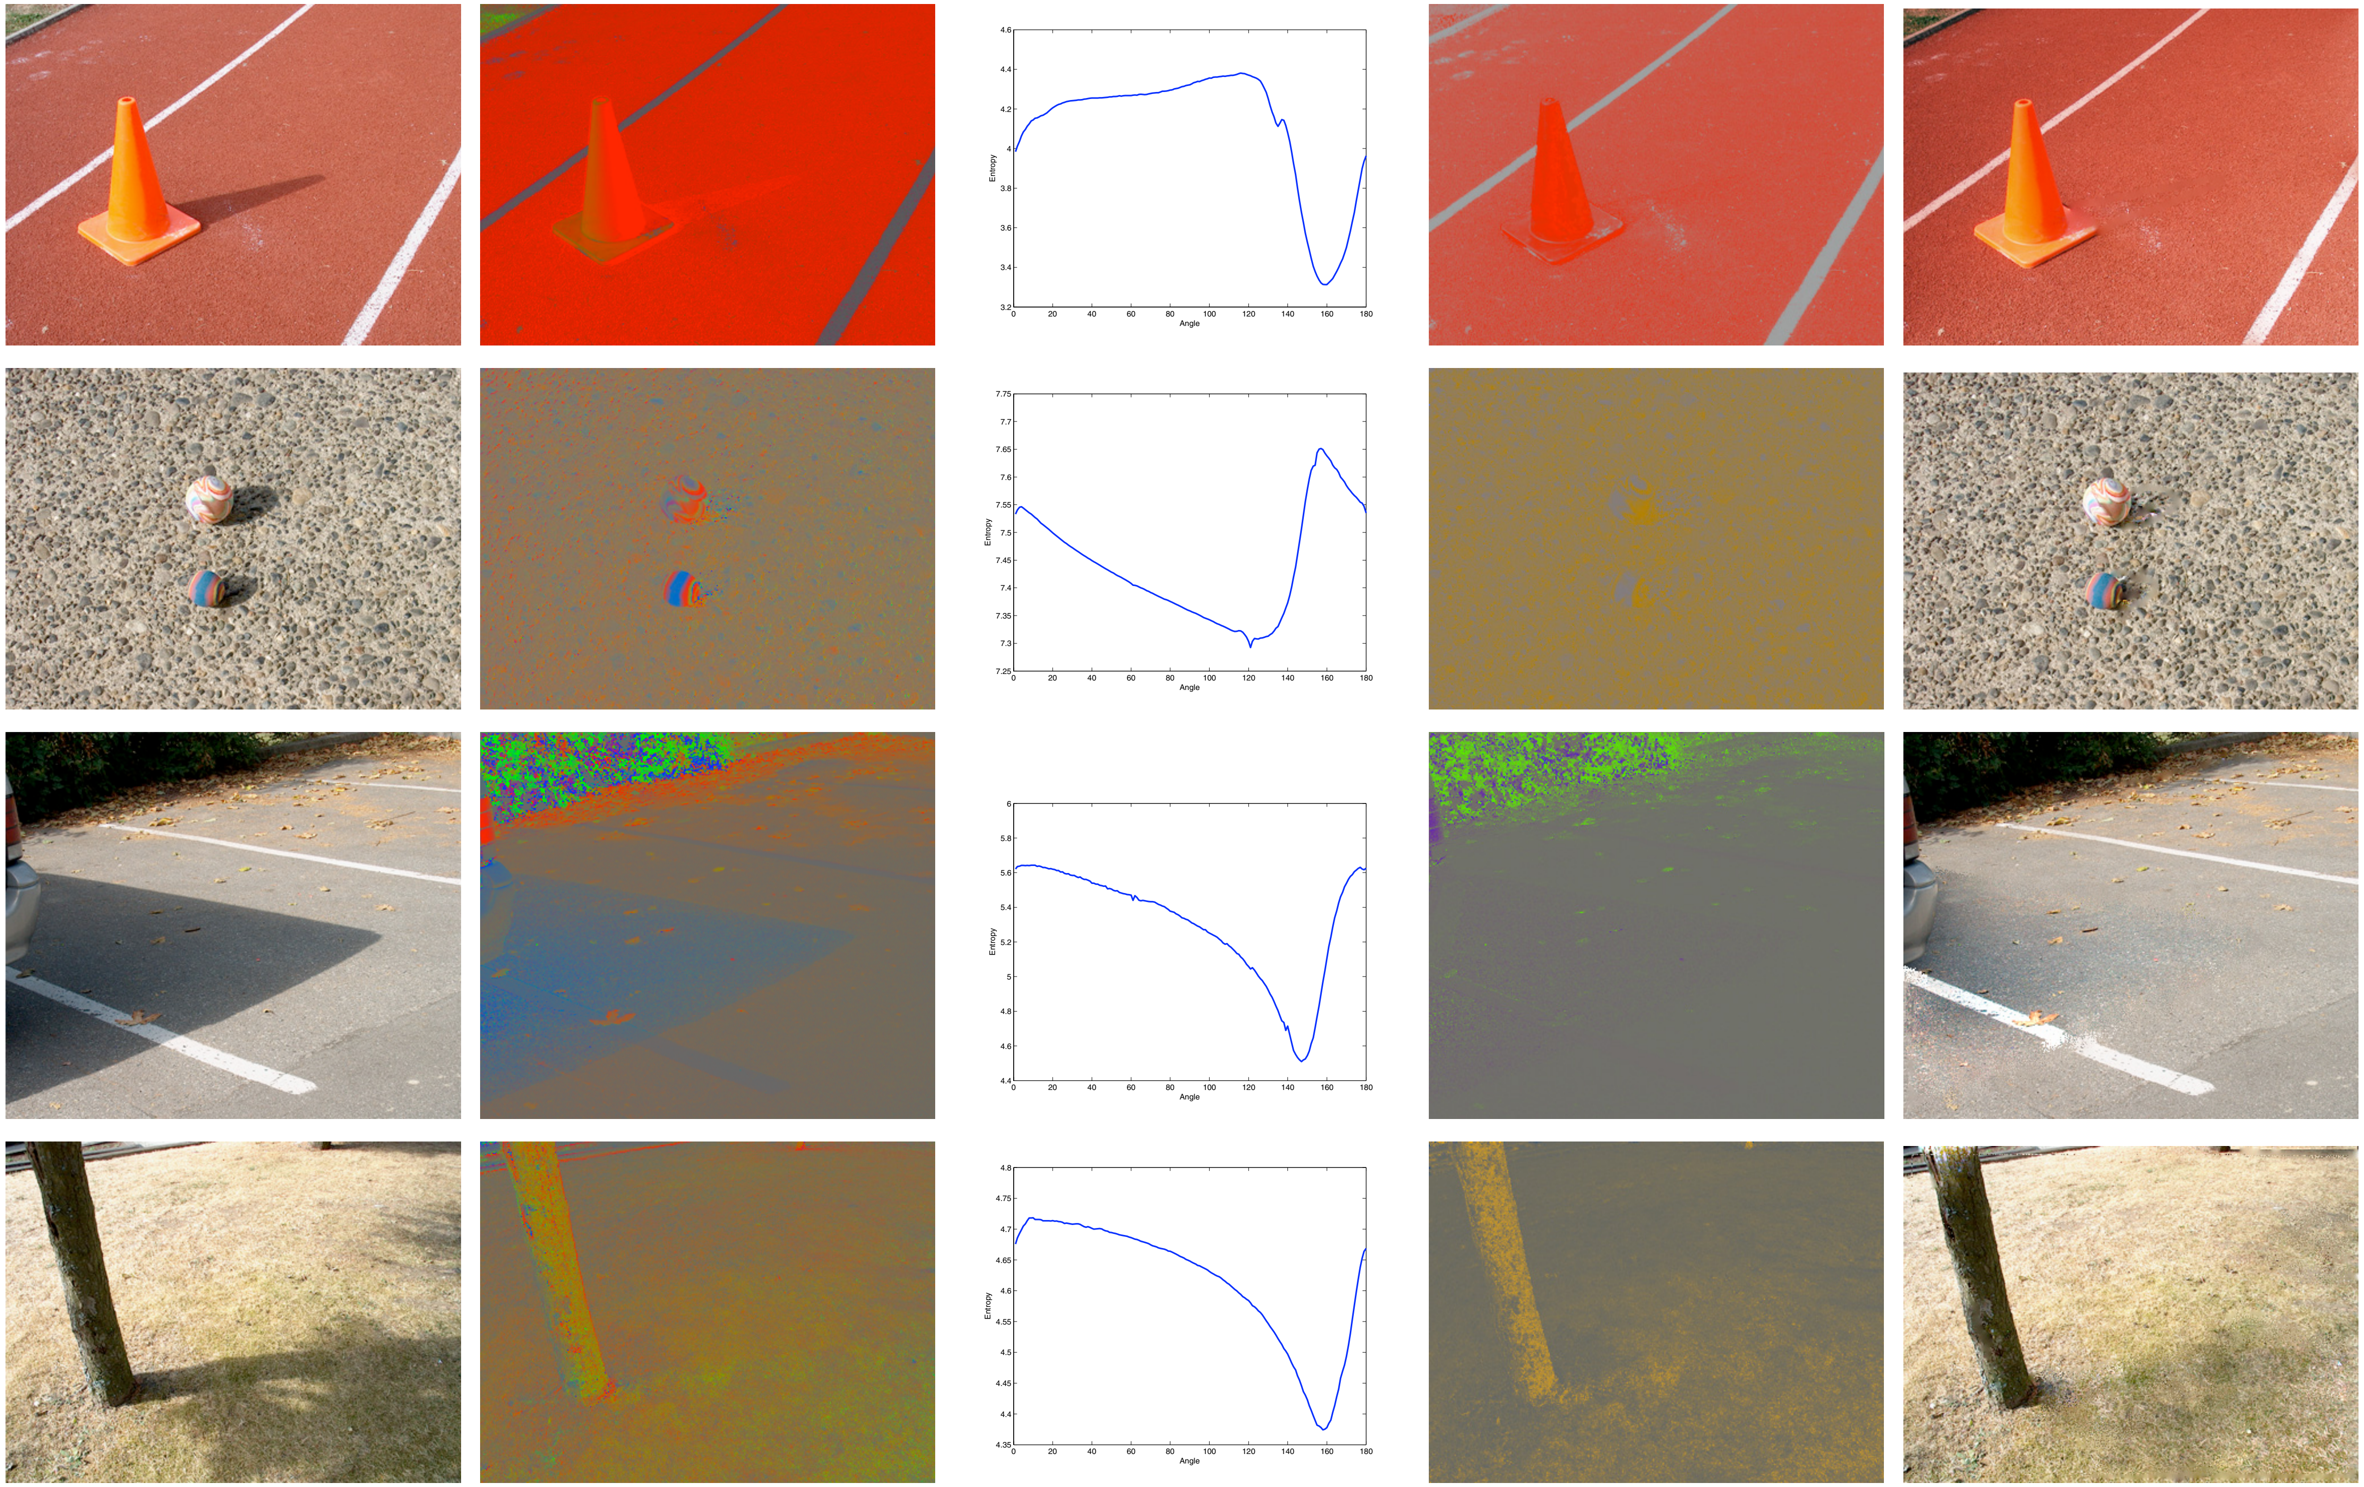

Figures from [3]

Results (Mine)

- The cameras that were used by the authors of the paper are professional quality. They claim in the paper that the methods work on all of the cameras that they tested, but the question is whether they actually tested consumer-grade cameras.
- So I went out on a sunny day and snapped a few shots with my Pentax Optio WPi 6.0 megapixel camera.

Results (Mine)

- My results weren't quite as spectacular.

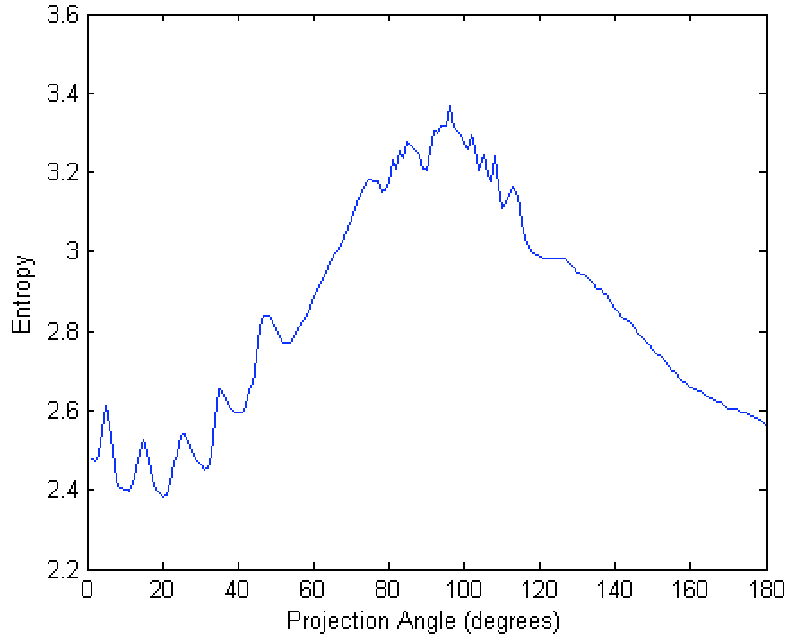

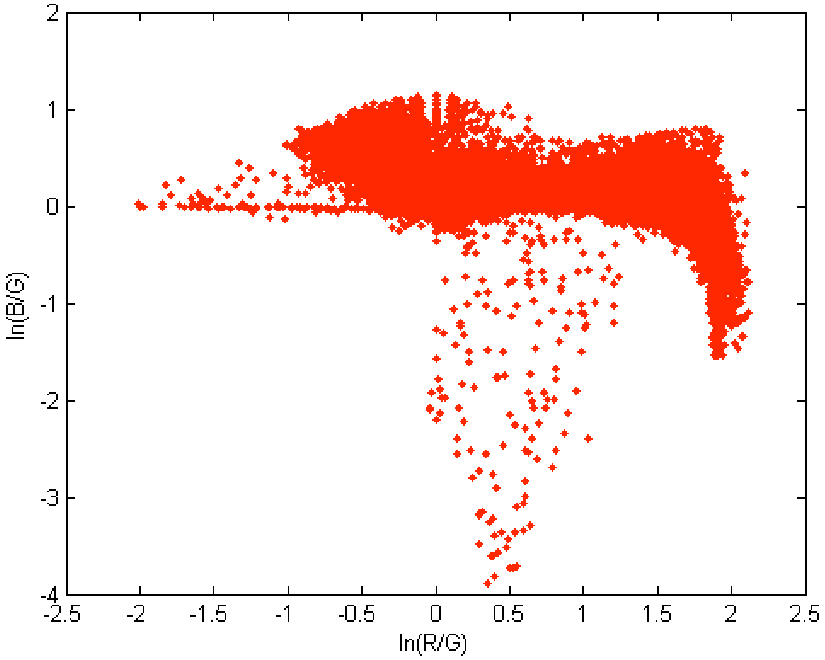

Results (Mine)

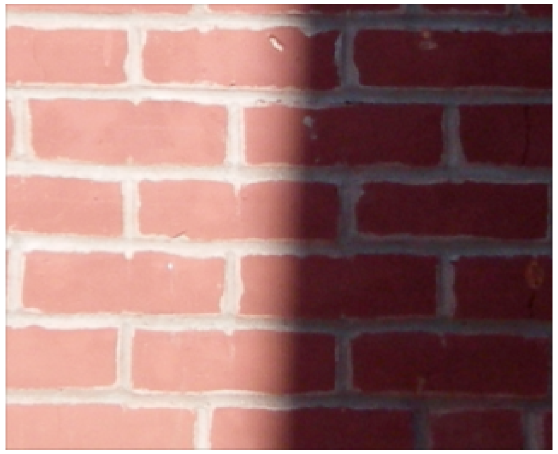

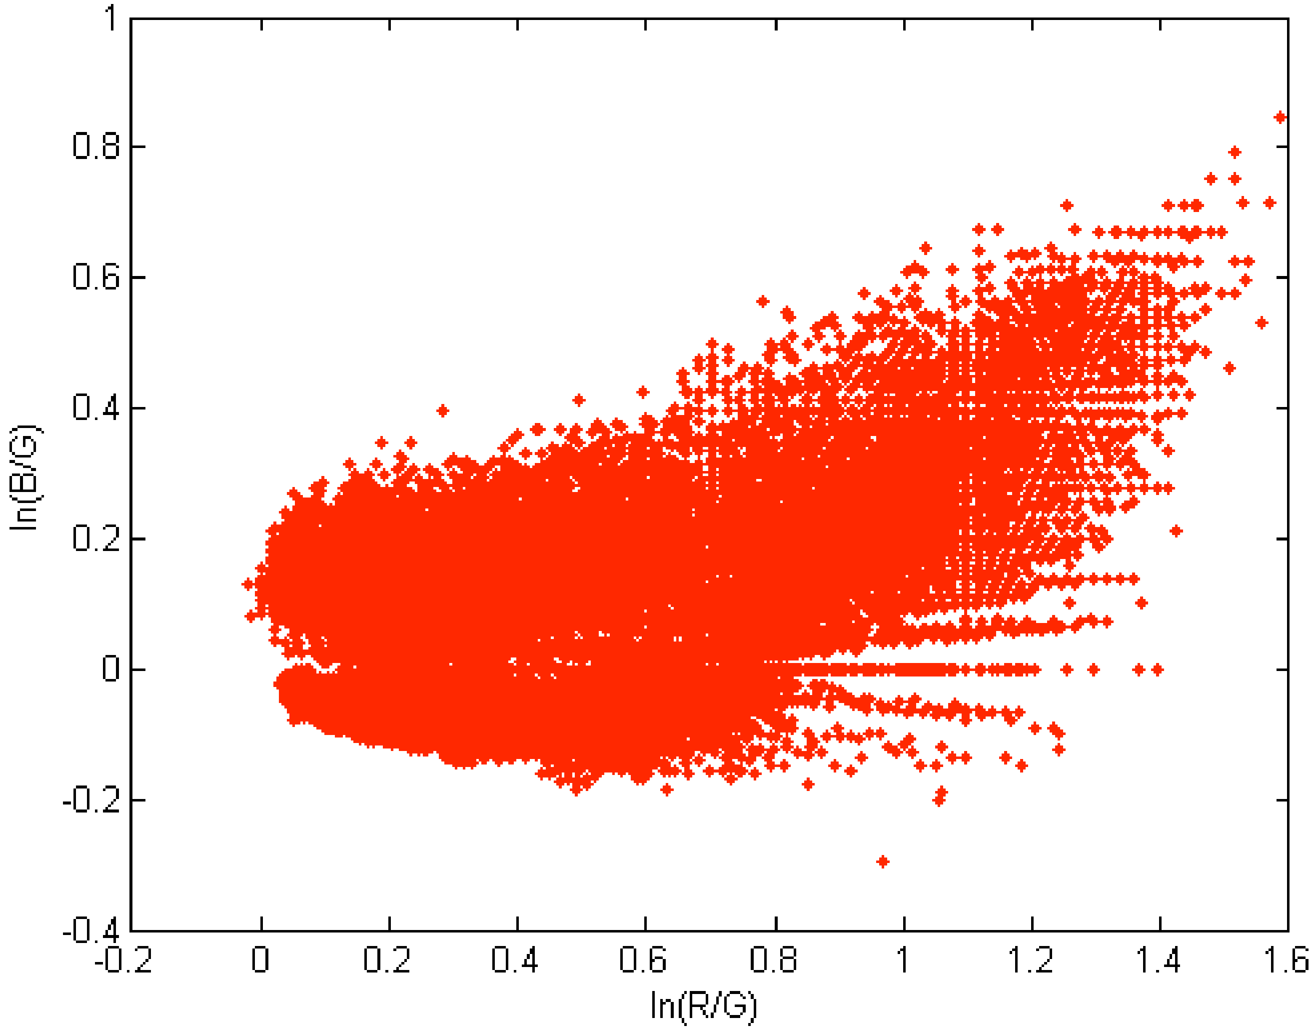

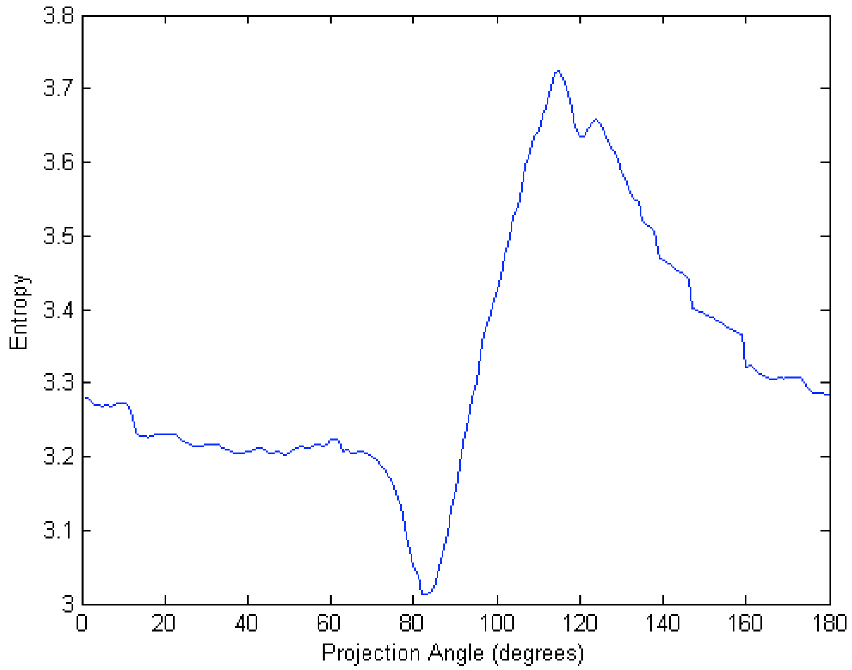

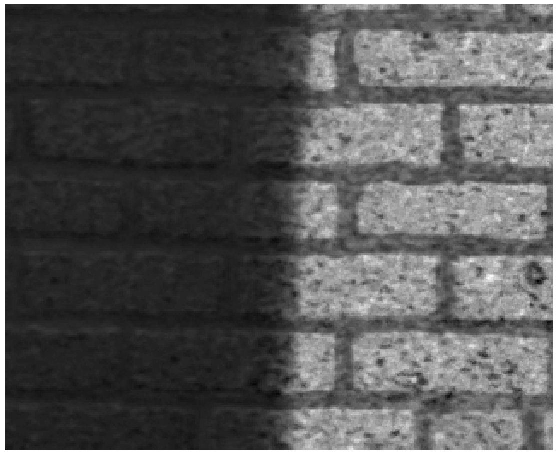

Questions?

“If he thought at all, but I don't believe he ever thought, it was that he and his shadow, when brought near each other, would join like drops of water, and when they did not he was appalled. He tried to stick it on with soap from the bathroom, but that also failed. A shudder passed through Peter, and he sat on the floor and cried.”

“Peter Pan : The Story of Peter and Wendy”
J. M. Barrie

References

- [1] Graham D. Finlayson, Mark S. Drew, and Cheng Lu, "Intrinsic Images by Entropy Minimization", European Conference on Computer Vision, Prague, May 2004. Springer Lecture Notes in Computer Science, Vol. 3023, pp. 582-595, 2004. http://www.cs.sfu.ca/~mark/ftp/Eccv04/.
- [2] Graham D. Finlayson, Steven D. Hordley, and Mark S. Drew, "Removing Shadows from Images", European Conference on Computer Vision, ECCV'02 Vol.4, Lecture Notes in Computer Science Vol. 2353, pp. 823-836, 2002. http://www.cs.sfu.ca/~mark/ftp/Eccv02/shadowless.pdf
- [3] Graham D. Finlayson and Steven D. Hordley, “Color constancy at a pixel”, Journal of the Optical Society of America, Optics, Image Science, and Vision, Volume 18, Issue 2, February 2001, pp.253-264.
- [4] G. Wyszecki and W.S. Stiles, “Color Science: Concepts and Methods, Quantitative Data and Formulas”, Wiley, New York, 2nd edition, 1982.