Natural Image Statistics: Foundations and Applications

Tania Pouli
Douglas Cunningham
Erik Reinhard

Eurographics 2013 Tutorial

Syllabus

Introduction - Tania (5 mins)

Foundations - Tania (25 mins)

Fourier statistics - Douglas (30 mins)

Wavelets - Erik (10 mins)

Color statistics - Erik (15 mins)

Discussion - All (5 mins)

Introduction

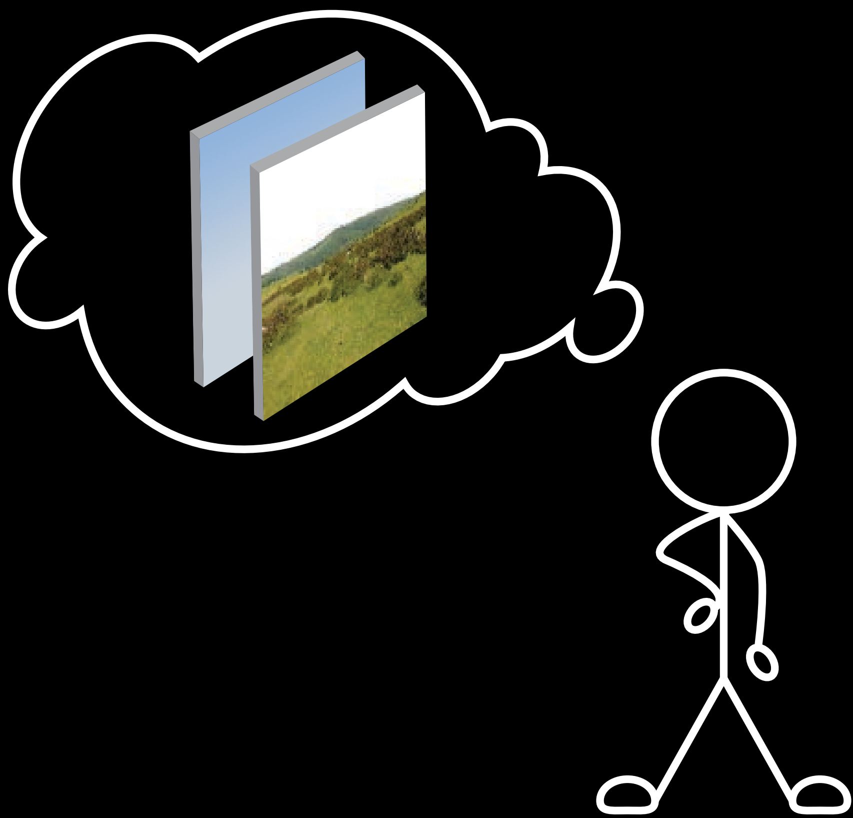

Introduction

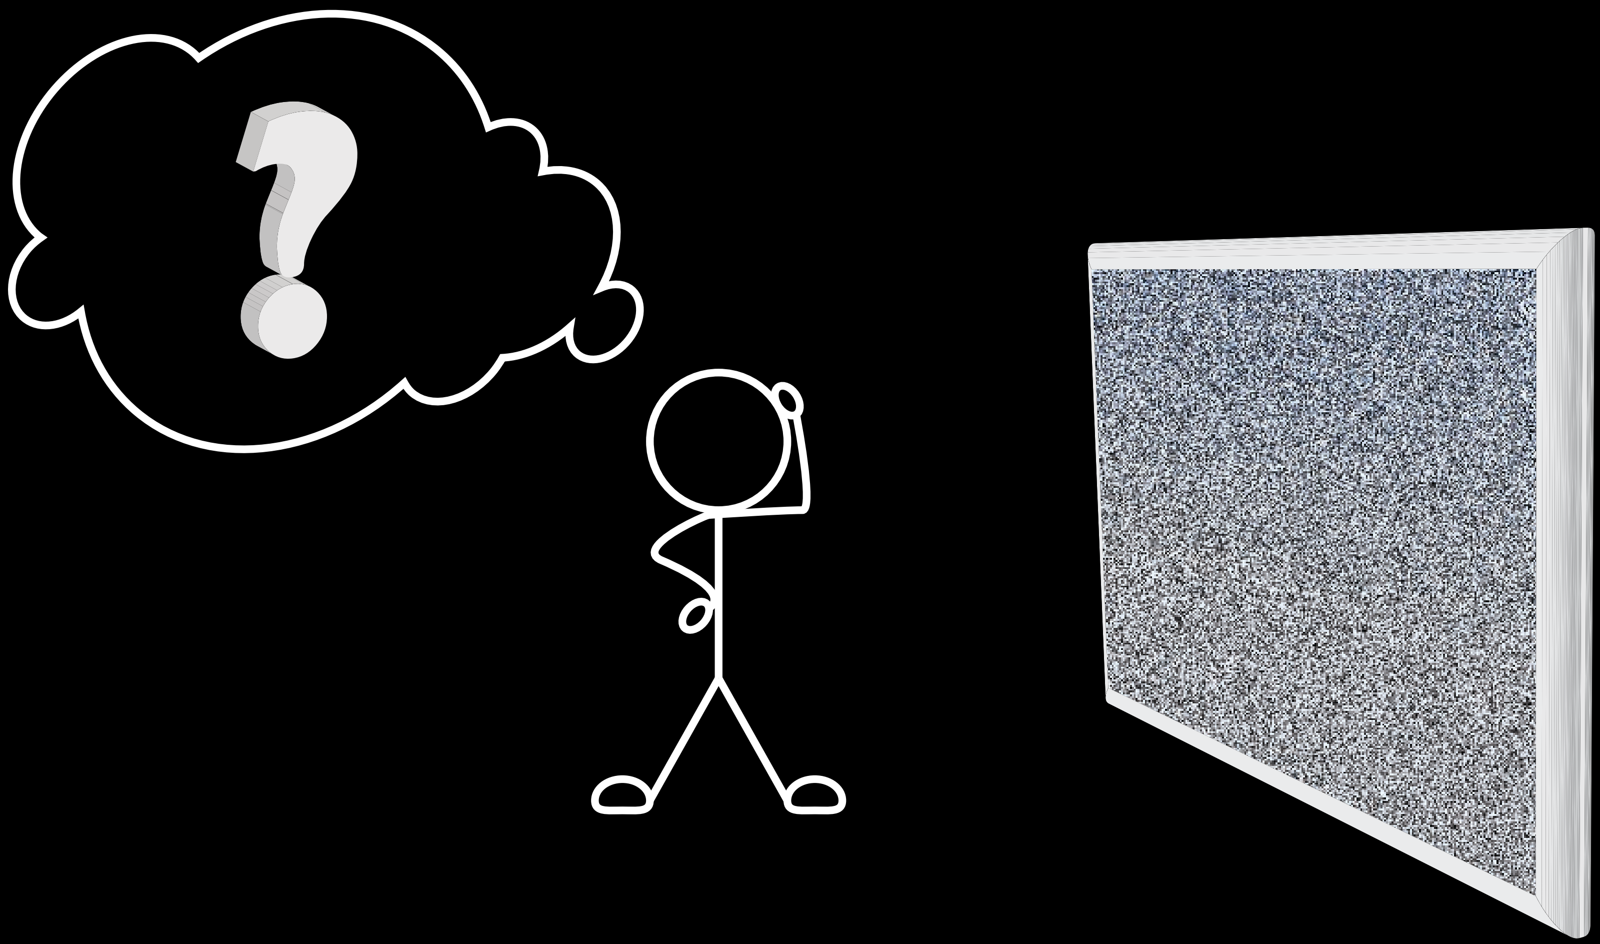

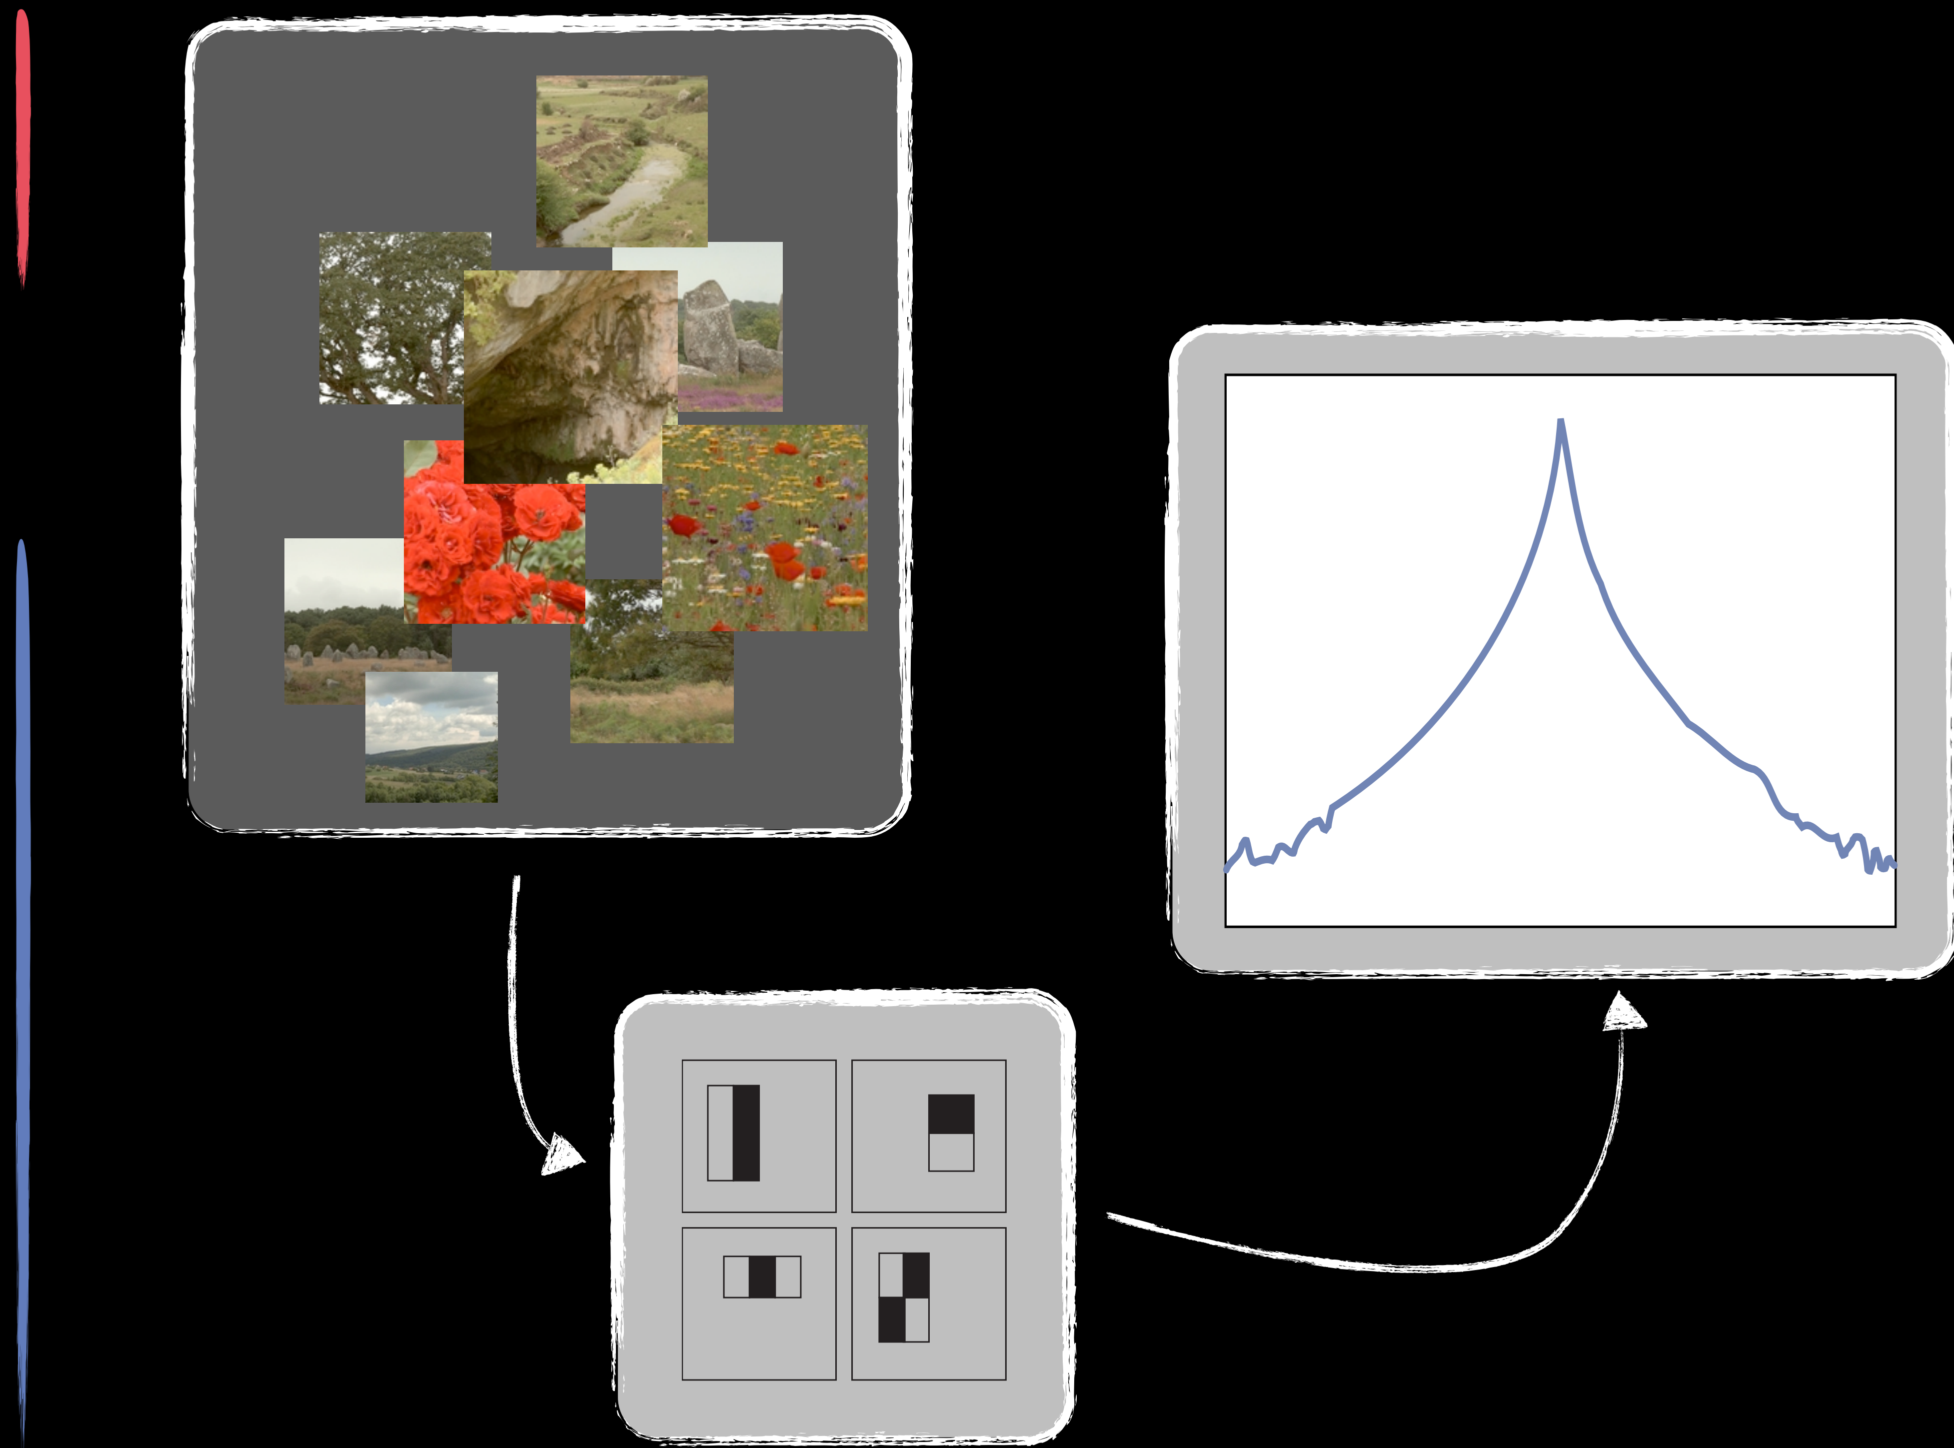

Statistics in Graphics

Types of Statistics

- **First order**
  - Each pixel viewed independently
- **Second Order**
  - Relations between pairs of pixels
- **Higher Order**
  - How does a pixel relate to more than one other pixel in the image?

Foundations

Tania Pouli

First Order

- Each pixel considered independently
  - Location invariant
  - Easy to compute & interpret
- First order statistics:
  - histogram moments
  - contrast*

Intensity Histograms

- How often does each intensity value occur?
- Histograms can be very important in data analysis, image analysis, and visualization

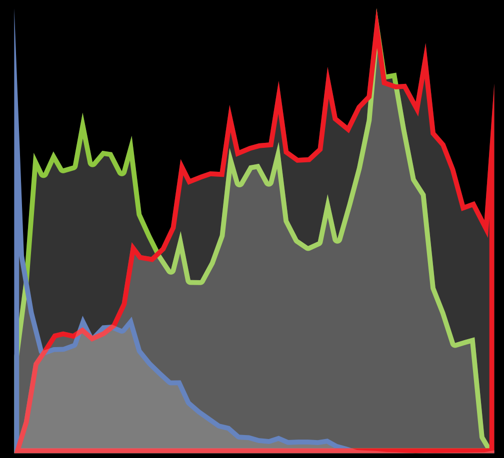

Individual frequency of occurrence

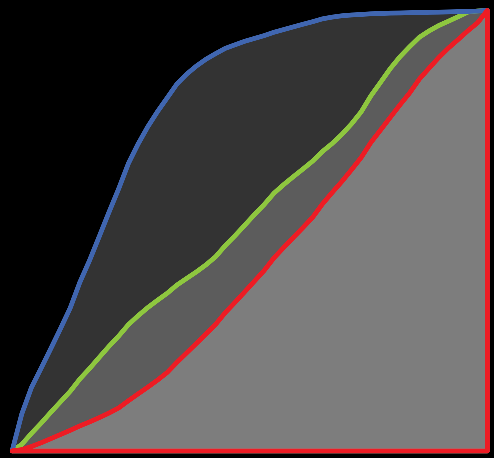

Cumulative frequency of occurrence

Histogram Moments

- They give information about the shape of the distribution
  - A measure of Gaussianity
  - $$m_k = \sum_{p=1}^{N} \frac{(I(p) - c)^k}{N}$$

Histogram Moments

- **1st moment:** mean
- **2nd moment:** variance
- **3rd moment:** relates to skewness $S = \frac{m_3}{\sigma^3}$
- **4th moment:** relates to kurtosis $\kappa = \frac{m_4}{\sigma^4}$

Intensity Histograms

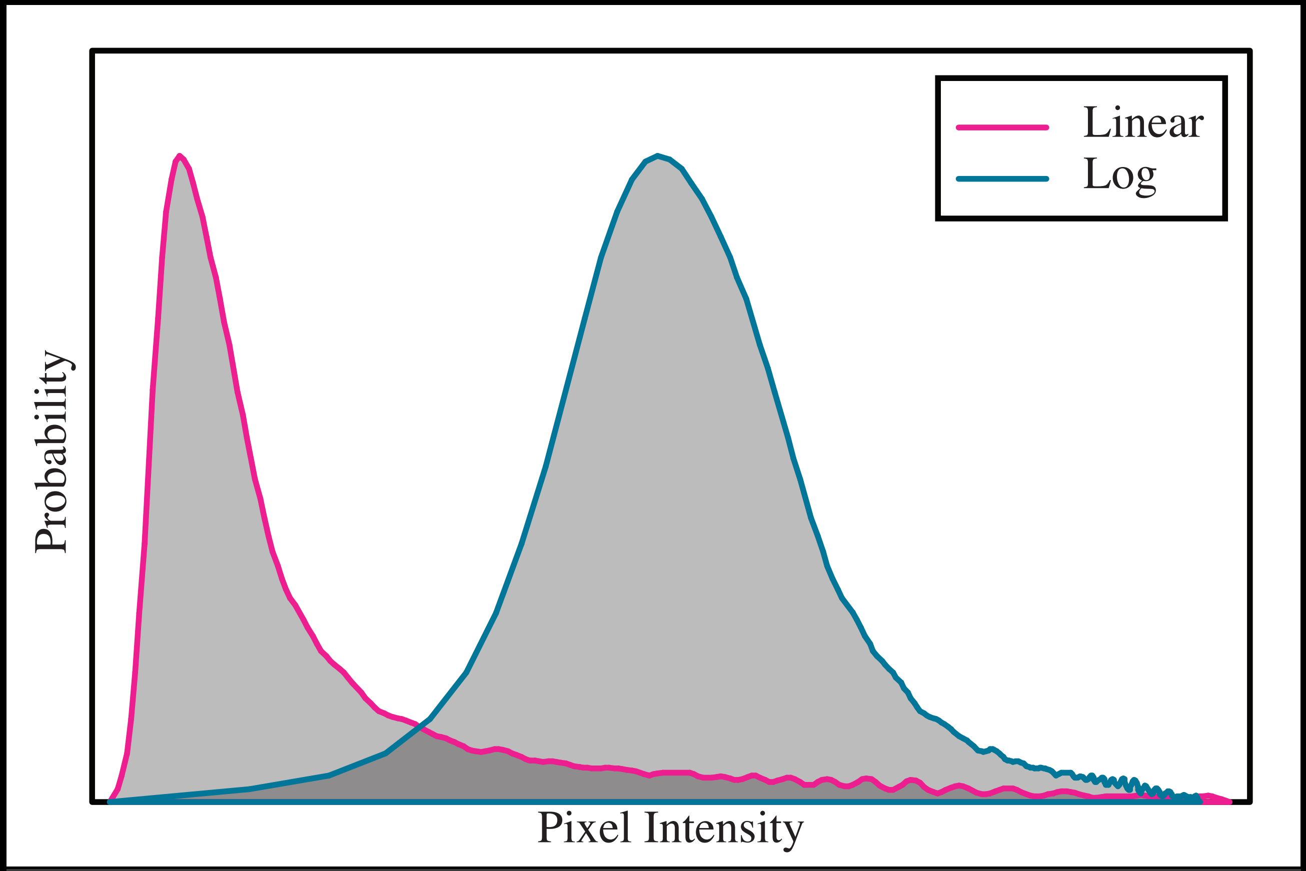

Local contrast in natural images: normalization and coding efficiency - Brady & Field (2000)

Intensity Histograms

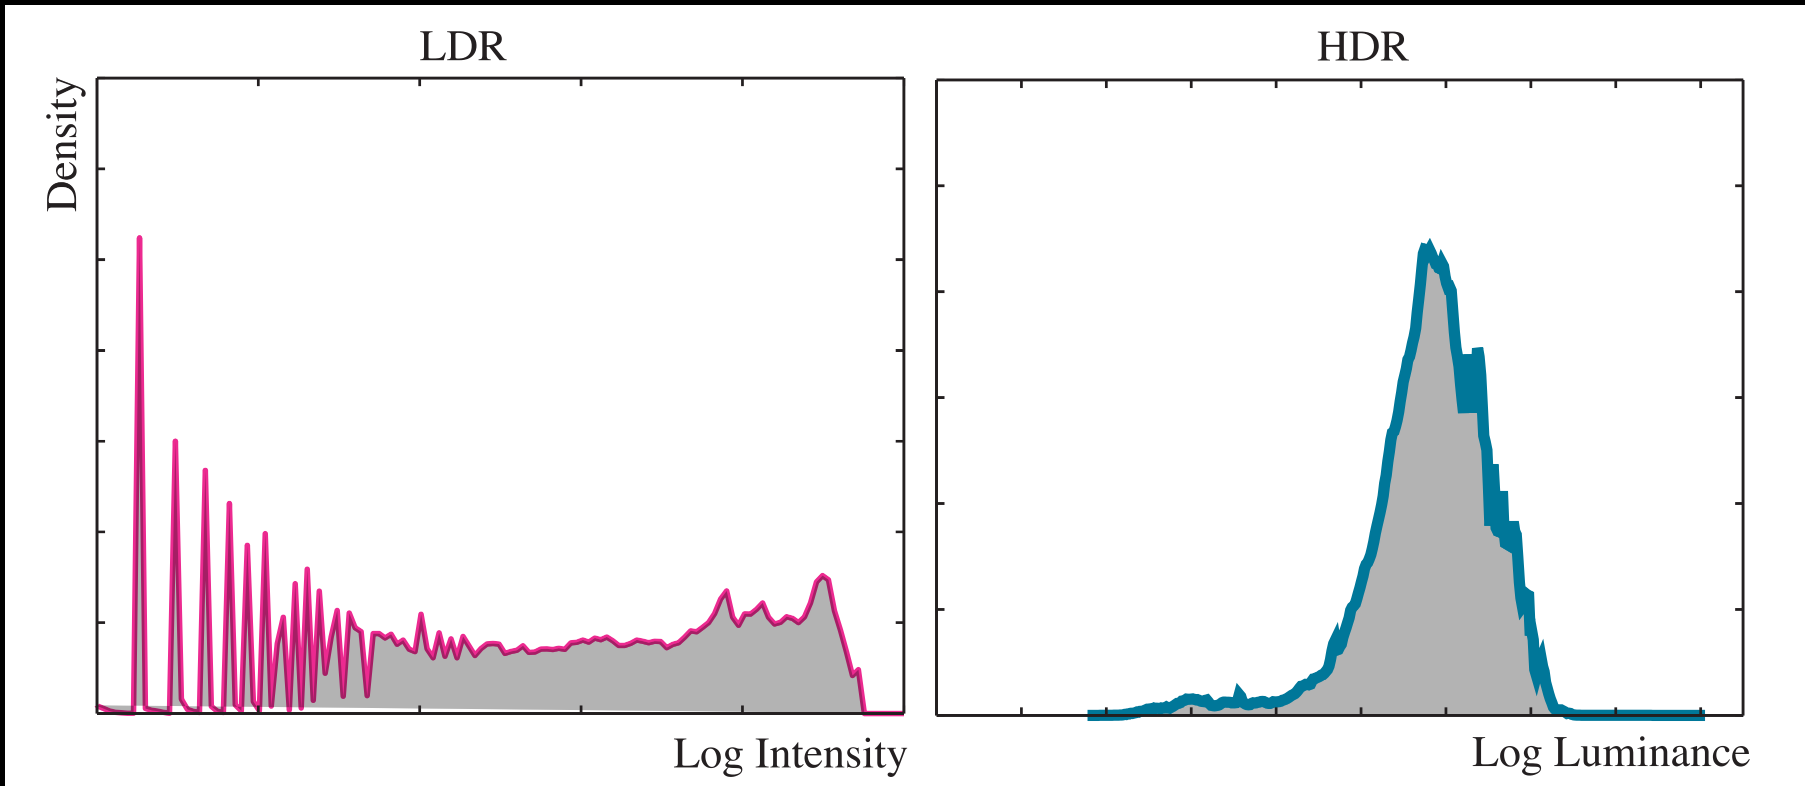

LDR vs HDR histograms for the same scenes - Pouli et al. (2010)

Histograms relate to surface properties:

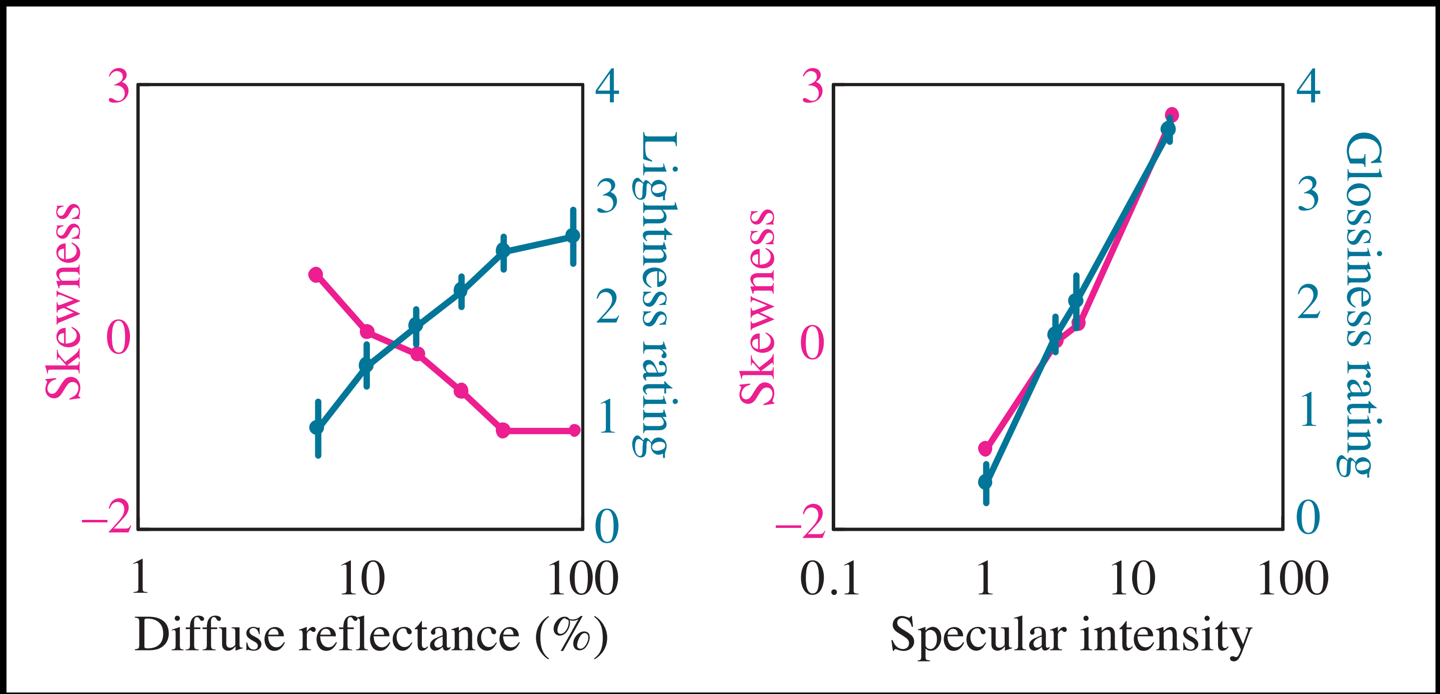

Image statistics and the perception of surface qualities - Motoyoshi et al. (2007)

Histogram Equalization

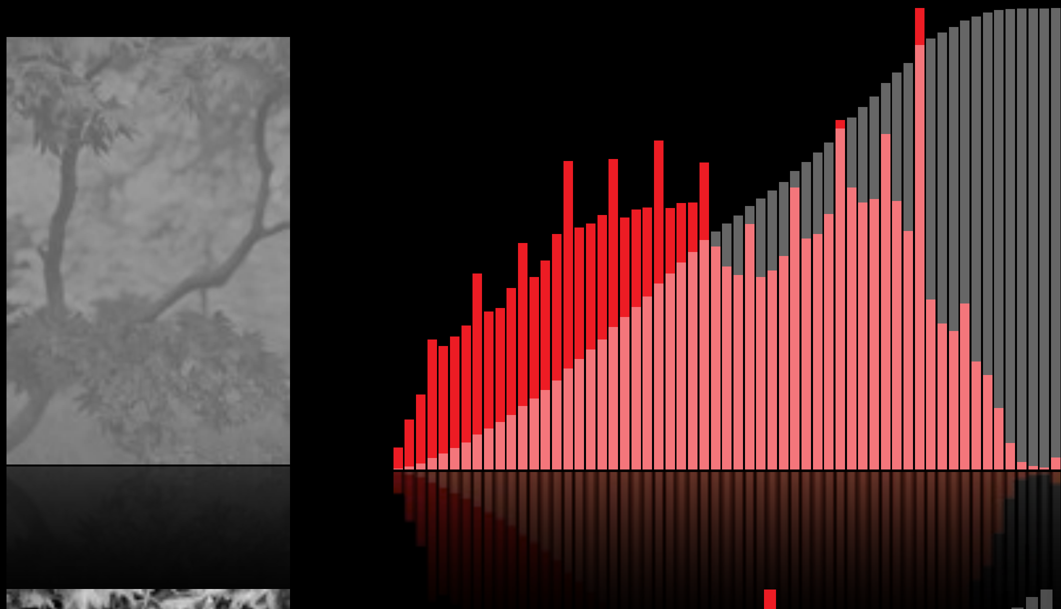

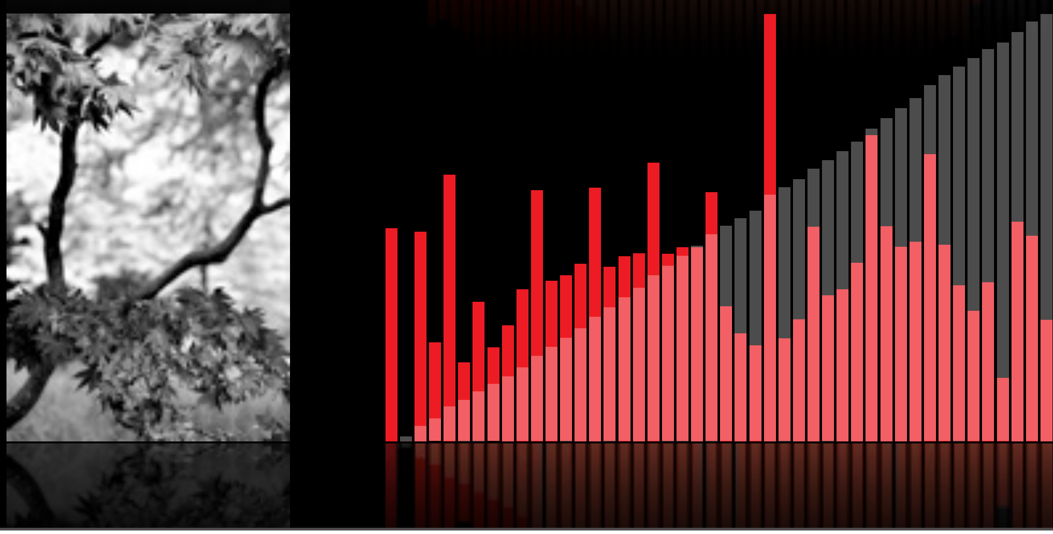

- Images do not use intensity values evenly
- Equalizing the values in the **cumulative** histogram can improve contrast

Histogram Matching

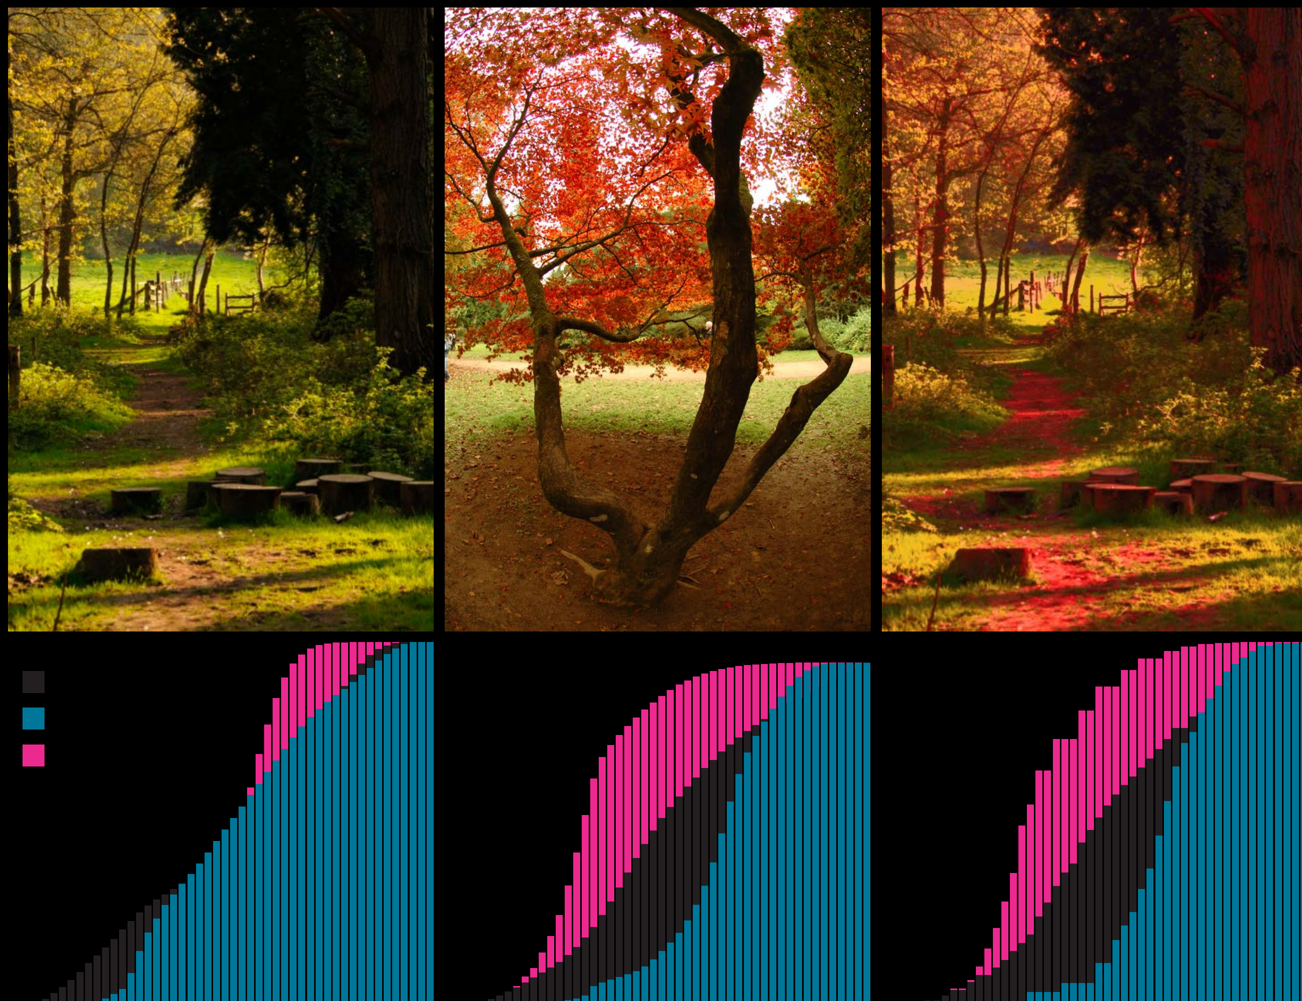

Source Target Result

Is 1st Order Enough?

- Simple to compute and interpret BUT...
- No spatial information
- No information on relations between pixels

- We need 2nd/higher order statistics for that!

Fourier Statistics

Douglas Cunningham

Spectral Slope

Fourier Transform Basics

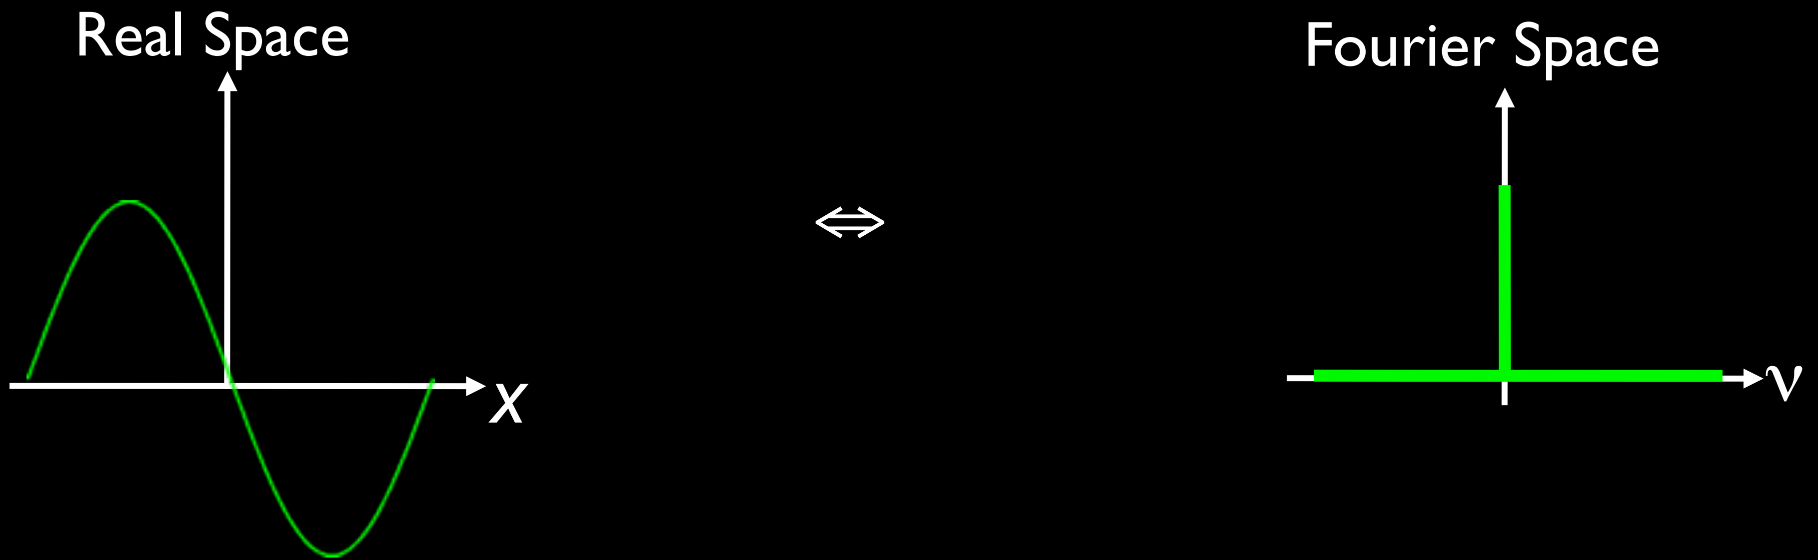

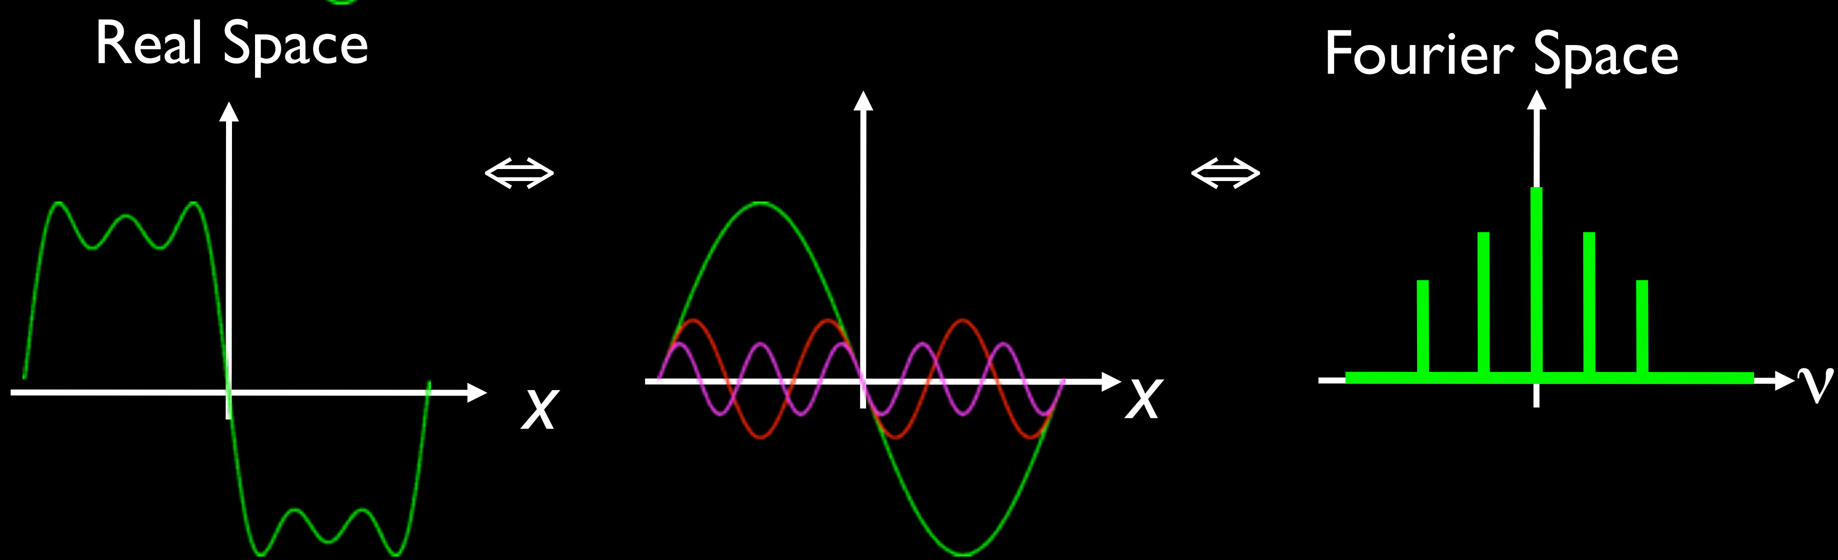

Spectral Slope

Fourier Transform Basics

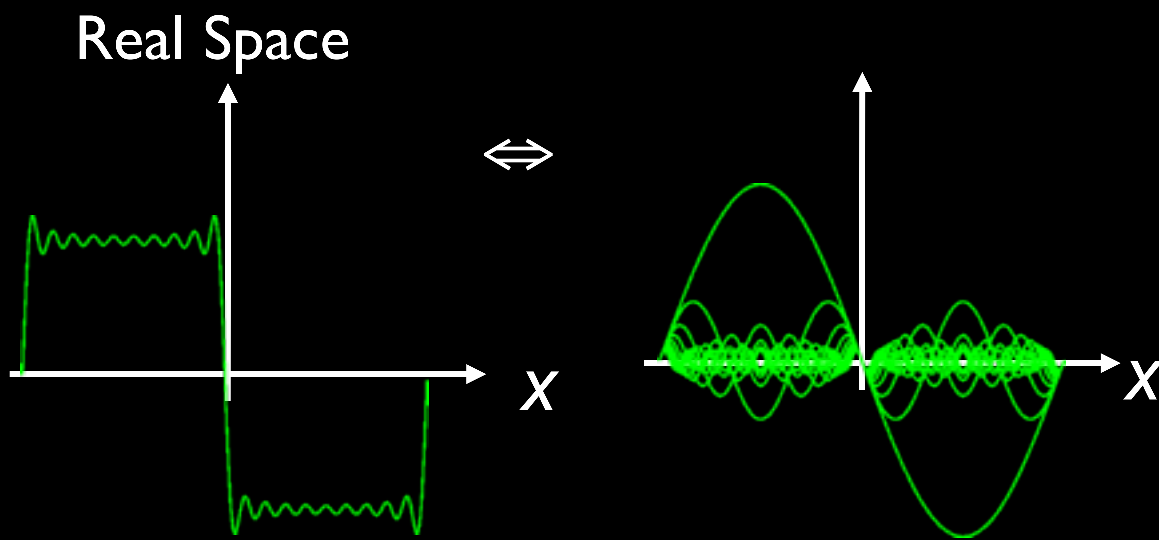

Edge Effects

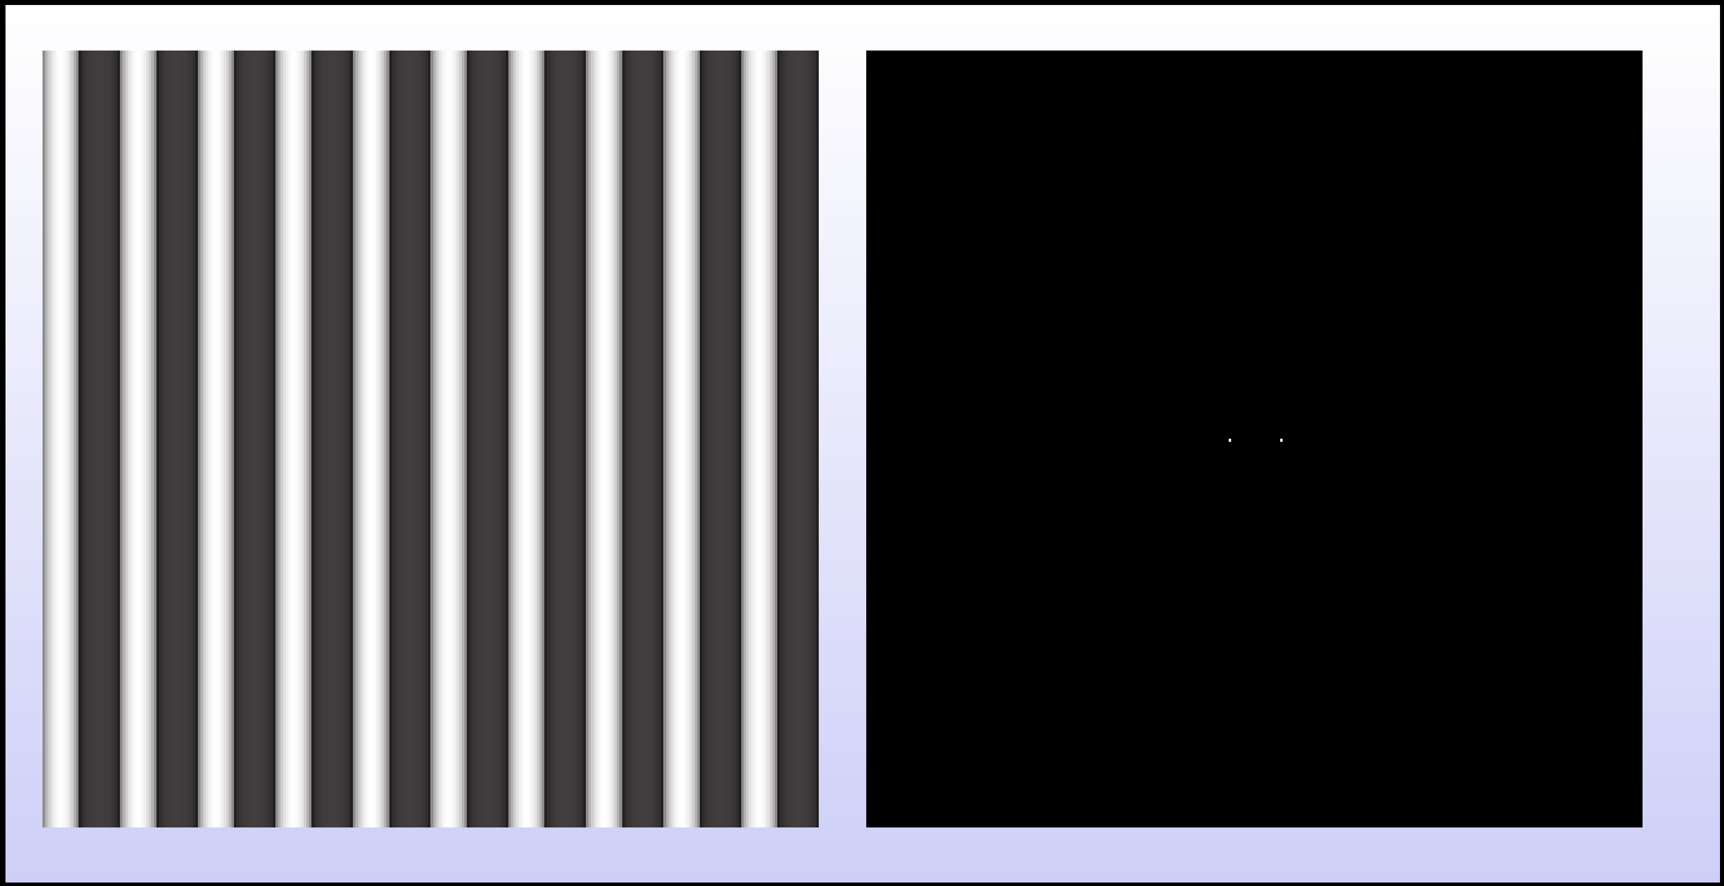

Cosine

Two spikes

Sum of Sines...

$I(x,y) = \sin(2\pi(n_x/2 - x)$ 1 sine

$I(x,y) = \sin(2\pi(n_x/2 - x) + \sin(\pi(n_x/2 - x)$ 2 sines

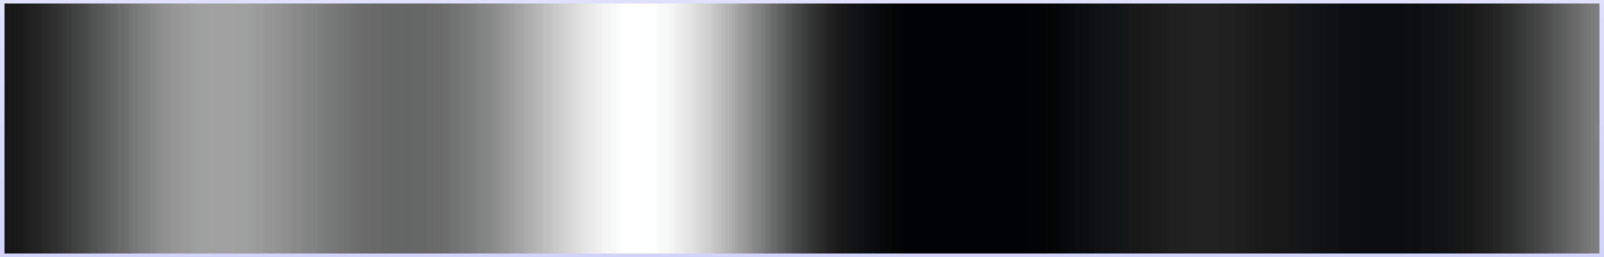

$I(x,y) = \sin(2\pi(n_x/2 - x) + \sin(\pi(n_x/2 - x) + \sin(\pi/2(n_x/2 - x)$ 3 sines

...Approximates an Edge

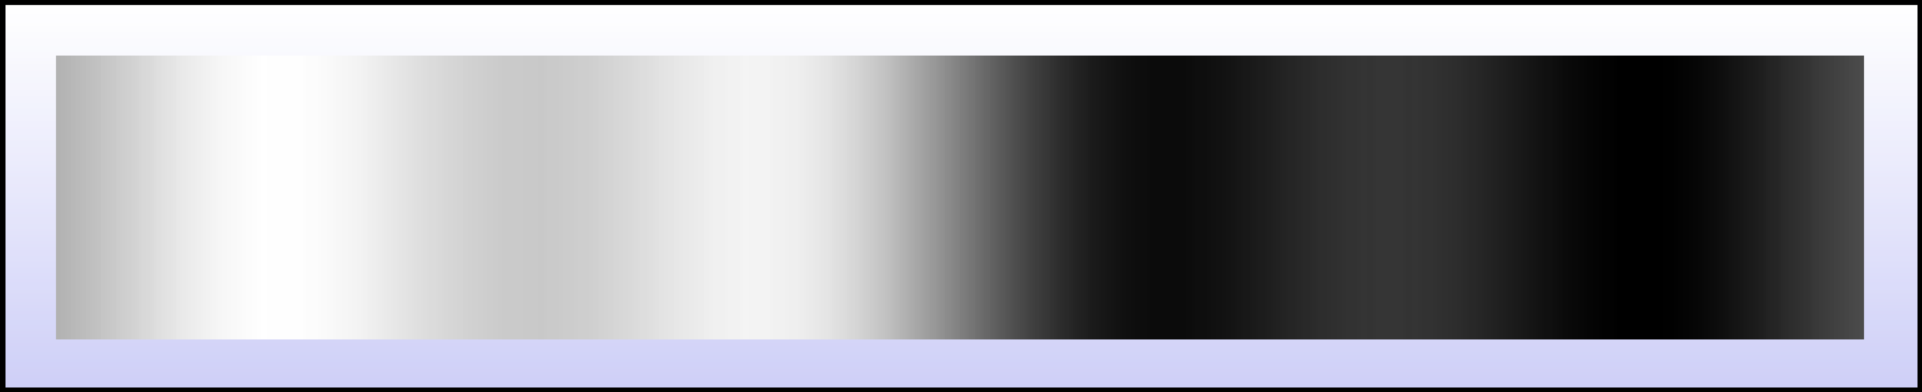

Power Spectrum

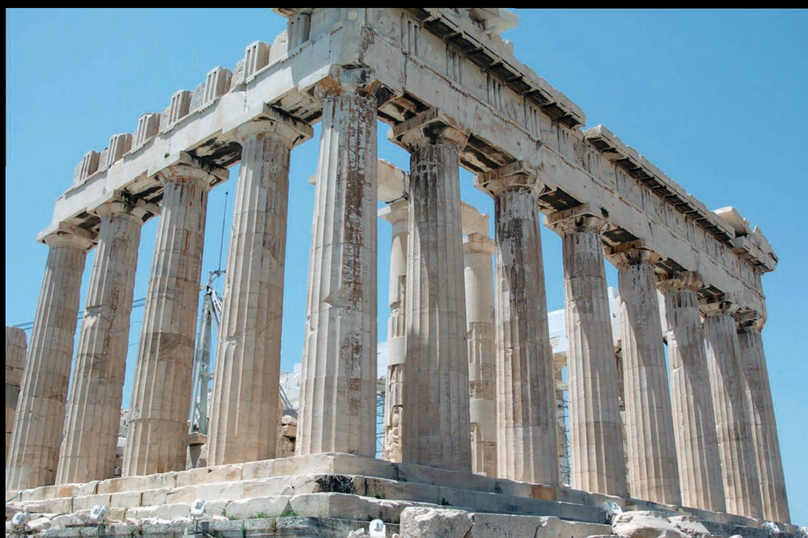

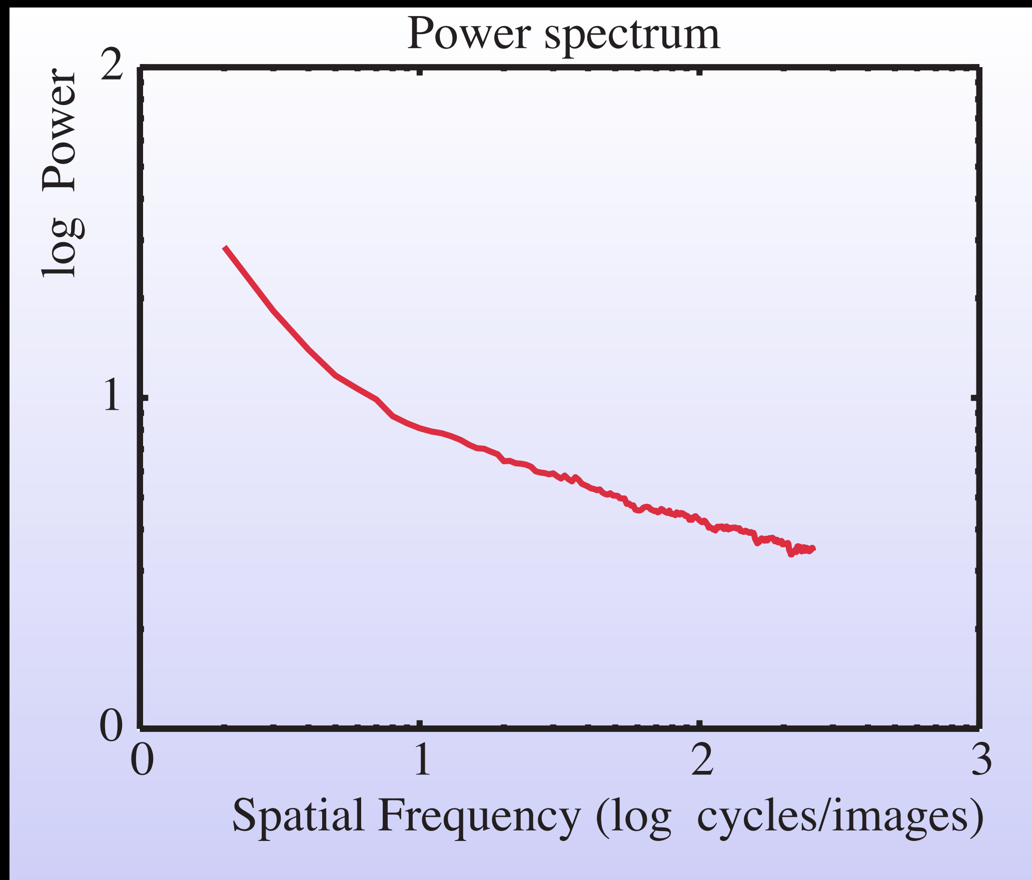

Spectral Slope

- Spectra of individual images varies
- AVERAGE spectra follow power law:

$$A \approx \frac{1}{f^{\beta}}$$

- Humans most sensitive to slopes between 2.8 and 3.2

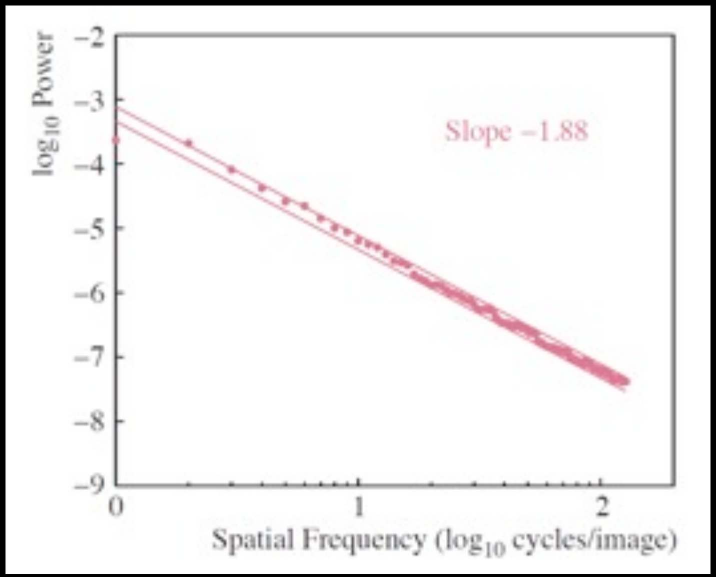

Spectral Slope

Study	# of Images	β±1sd
Burton G. J., Moorhead I. R. 1987	19	2.1±0.24
Dong D., Atick J. 1995	320	2.30
Dror R. O., Adelson E. H., Willsky A. S. 2001	95	2.29
Field D. J. 1987	6	2.0
Field D. J. 1993	85	2.20
Field D. J., Brady N. 1997	20	2.20±0.28
van Hateren J. 1992	117	2.13±0.36
Huang J., Mumford D. 1999	216	1.96
Pàrraga C. A., Brelstaff G., Troscianko T. 1998	29	2.22±0.26
Reinhard E., Shirley P., Troscianko T. 2001	133	1.88±0.42
Ruderman D., Bialec W. 1994	45	1.81
van der Schaaf A., van Hateren J. 1996	276	1.88±0.42
Thomson M., Foster D. 1997	82	2.38
Tolhurst D. J., Tadmory Y., Chiao T. 1992	135	2.4±0.26
Torralba A., Oliva A. 2003	12,000	2.08
Webster M., Miharaya E. 1997	48	2.26

Spectral Slope

- Spectral slope is related to autocorrelation
- Increasing slope increases coarseness
- Self similar (fractal)

β=0.8 β=1.6

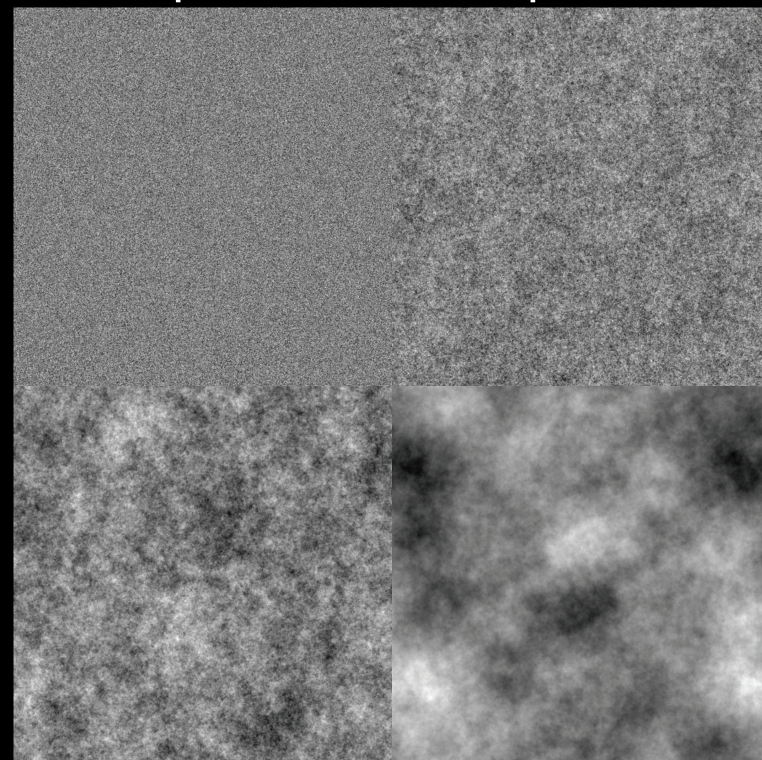

β=2.4 β=3.2

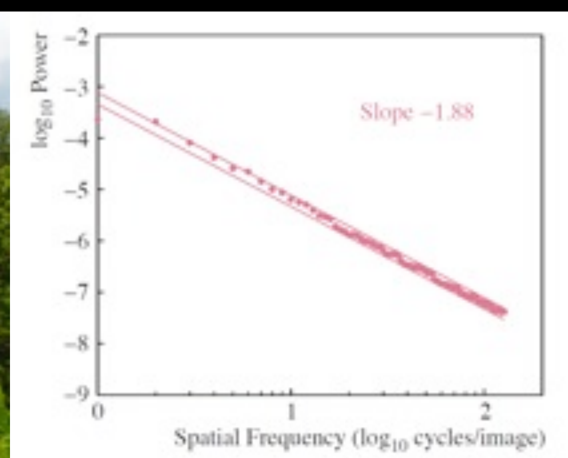

Edge Effects

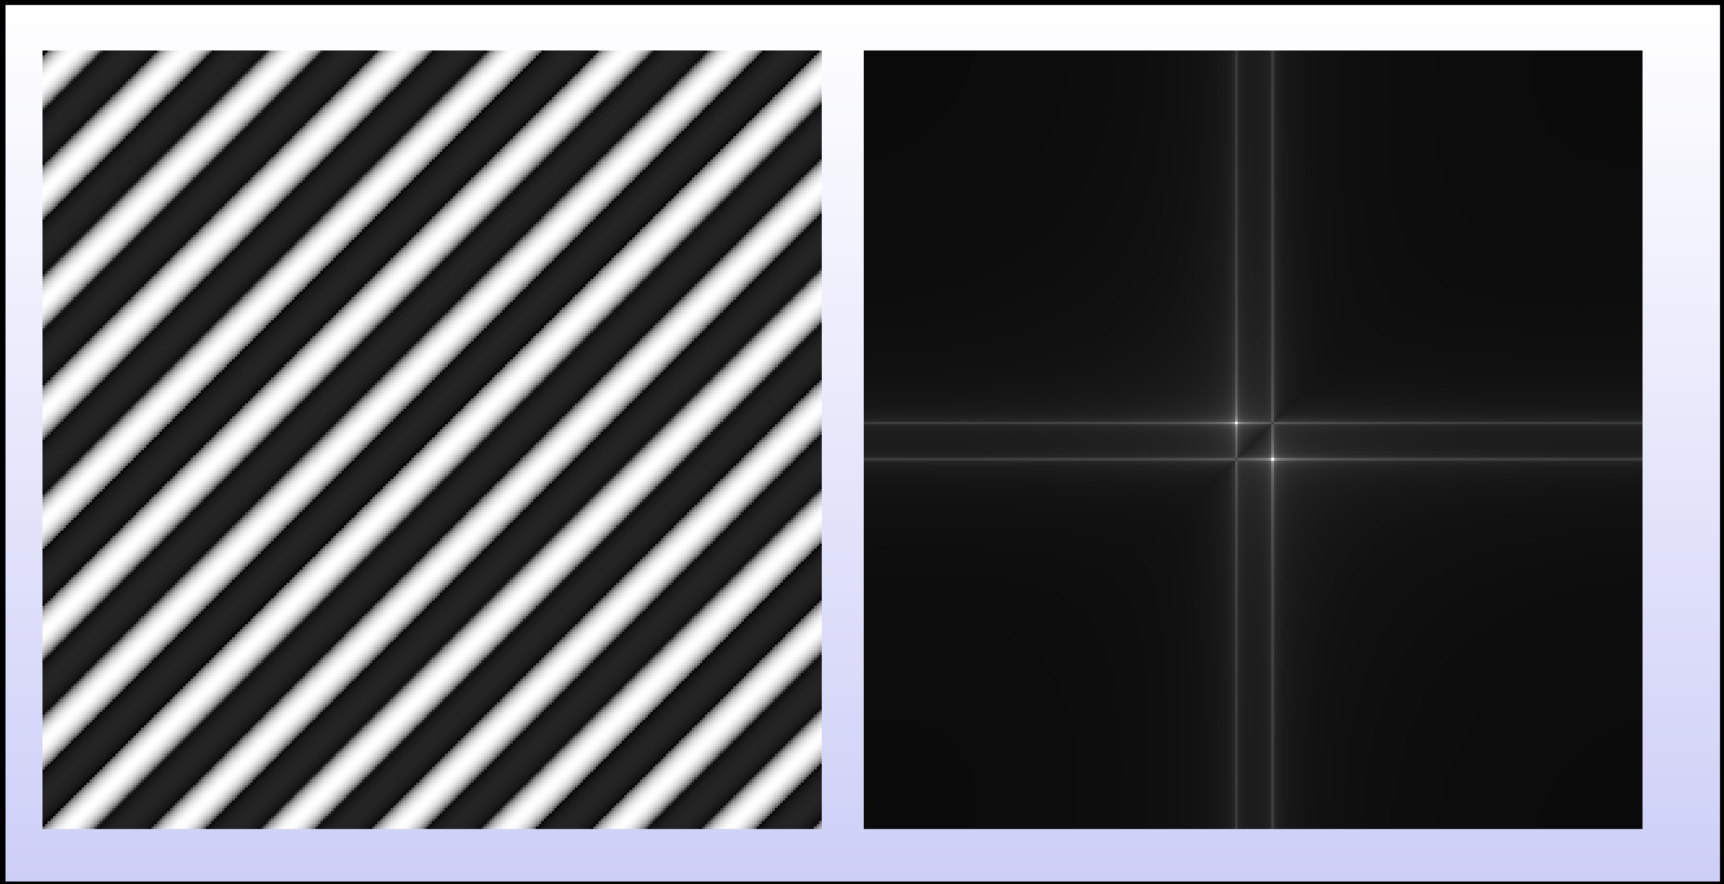

Cosine   Edge Effects

Windowing

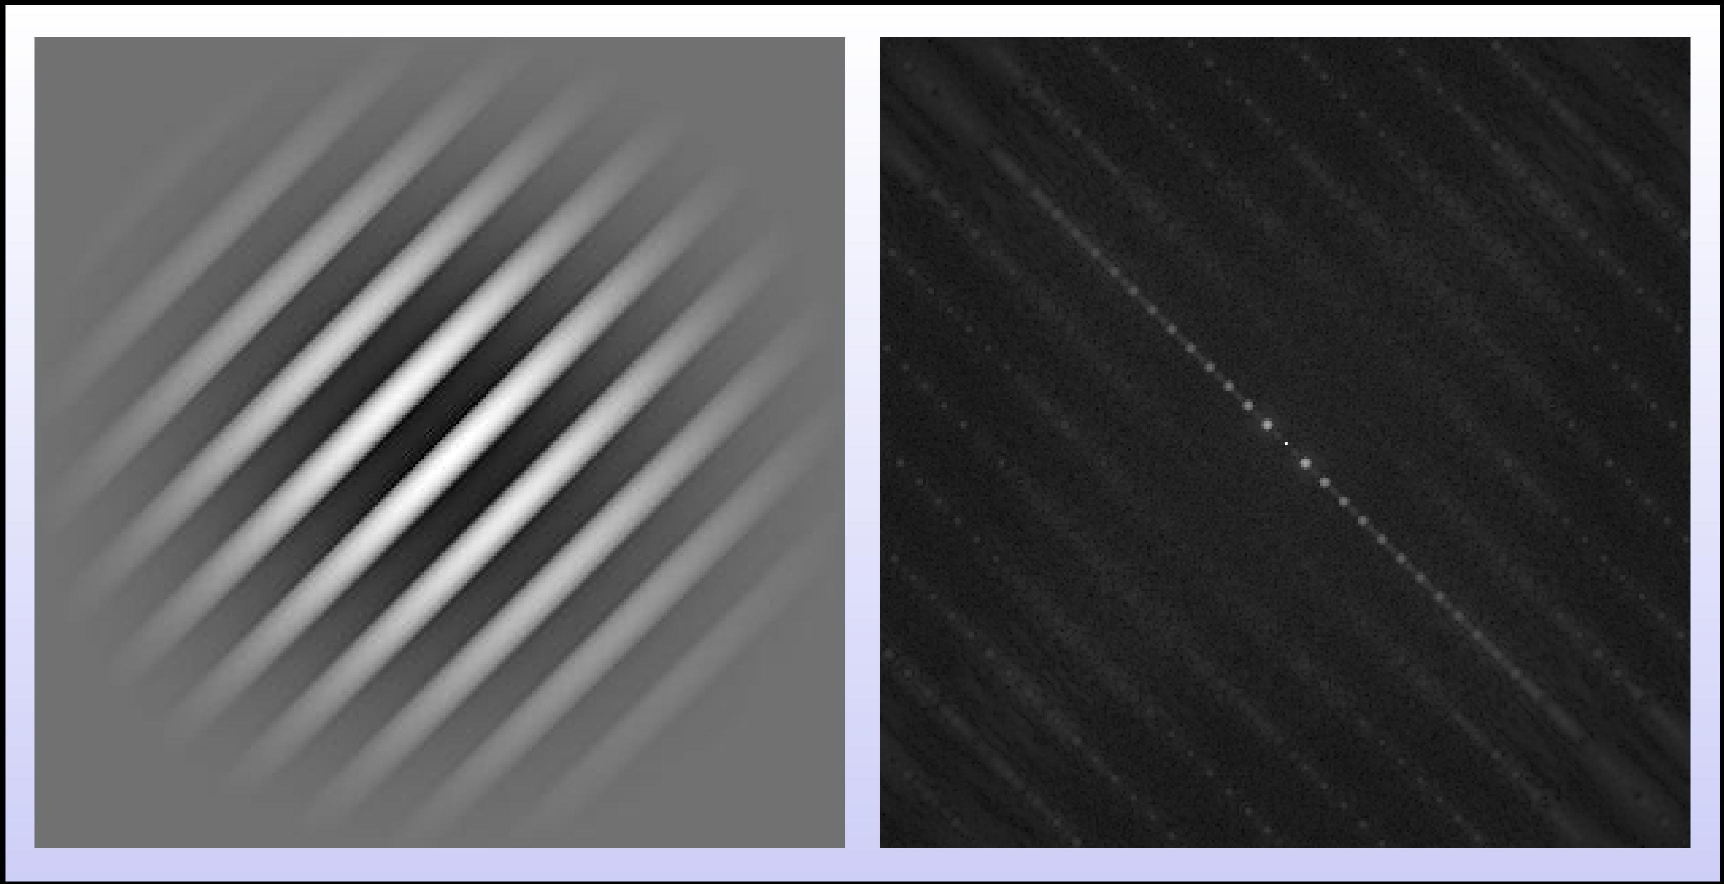

Angular Power Spectrum

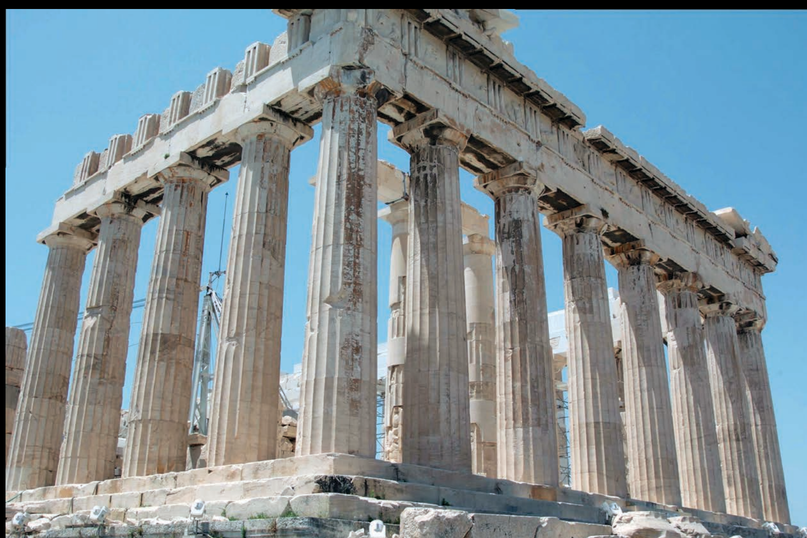

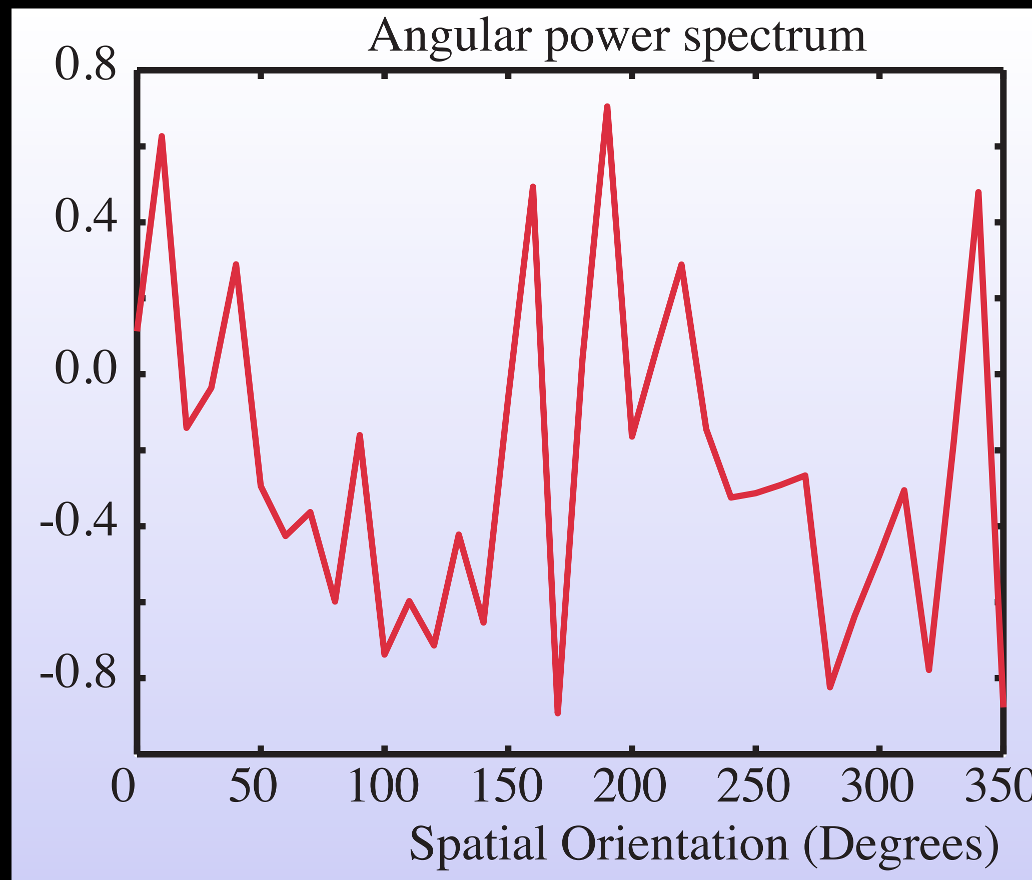

Applications: **Scene Classification**

- Spectral Slope differs by scene

Forest (2.15)

Close-up (2.23)

distant meadow (2.4)

Webster & Miyahara (1997)

Applications:
Scene Classification

- and orientation

Horizontal	Oblique	Vertical
1.98	2.02	2.22
	Natural	

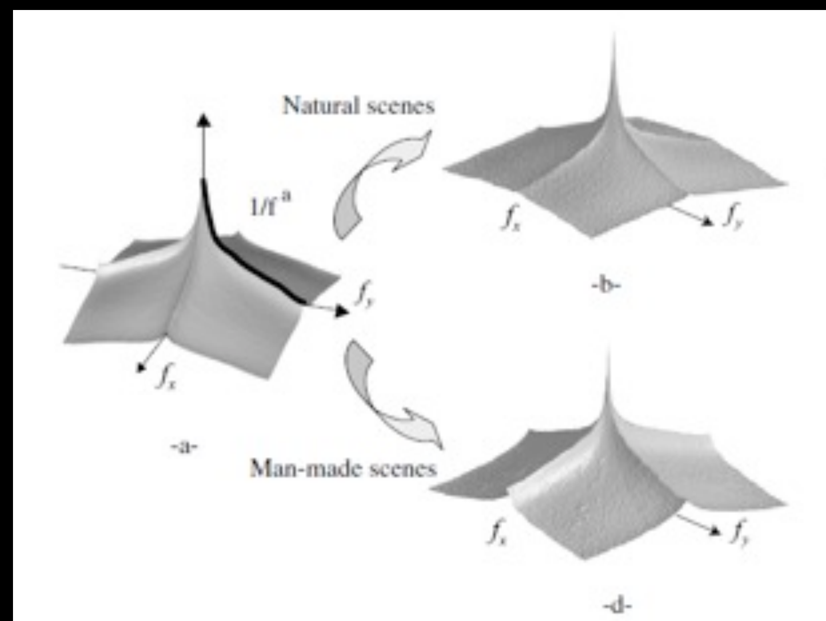

Torralba & Oliva (2003)

Applications:
Scene Classification

- and orientation
- and orientation by scene

	Horizontal	Oblique	Vertical
Natural	1.98	2.02	2.22
Man Made	1.83	2.37	2.07

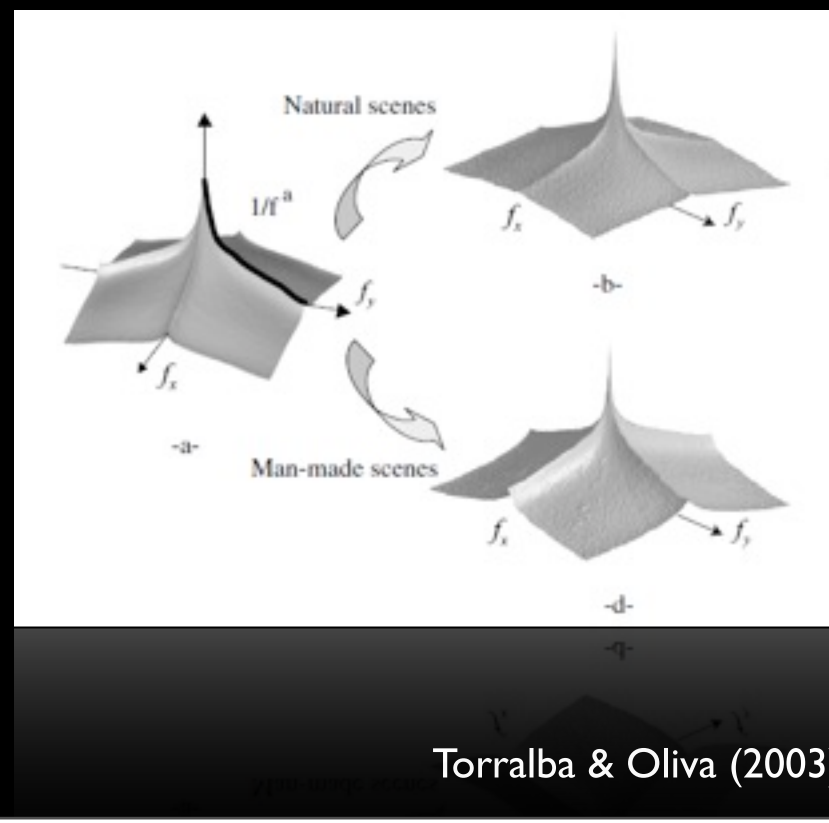

Torralba & Oliva (2003)

Applications: **Scene Classification**

- and orientation by scene

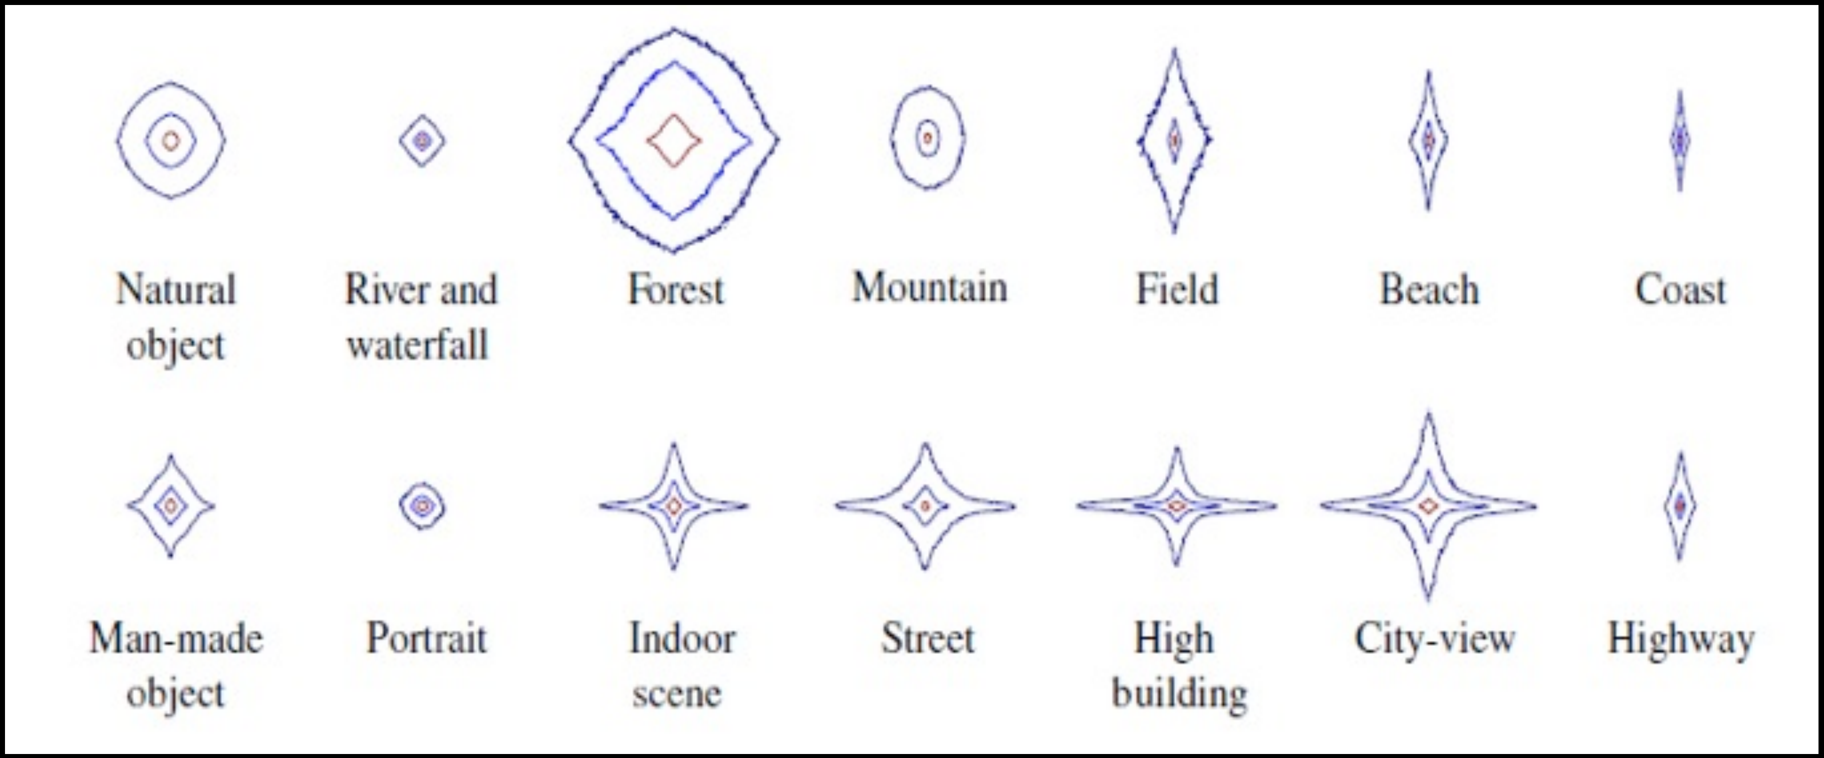

Torralba & Oliva (2003)

Applications:
Scene Synthesis

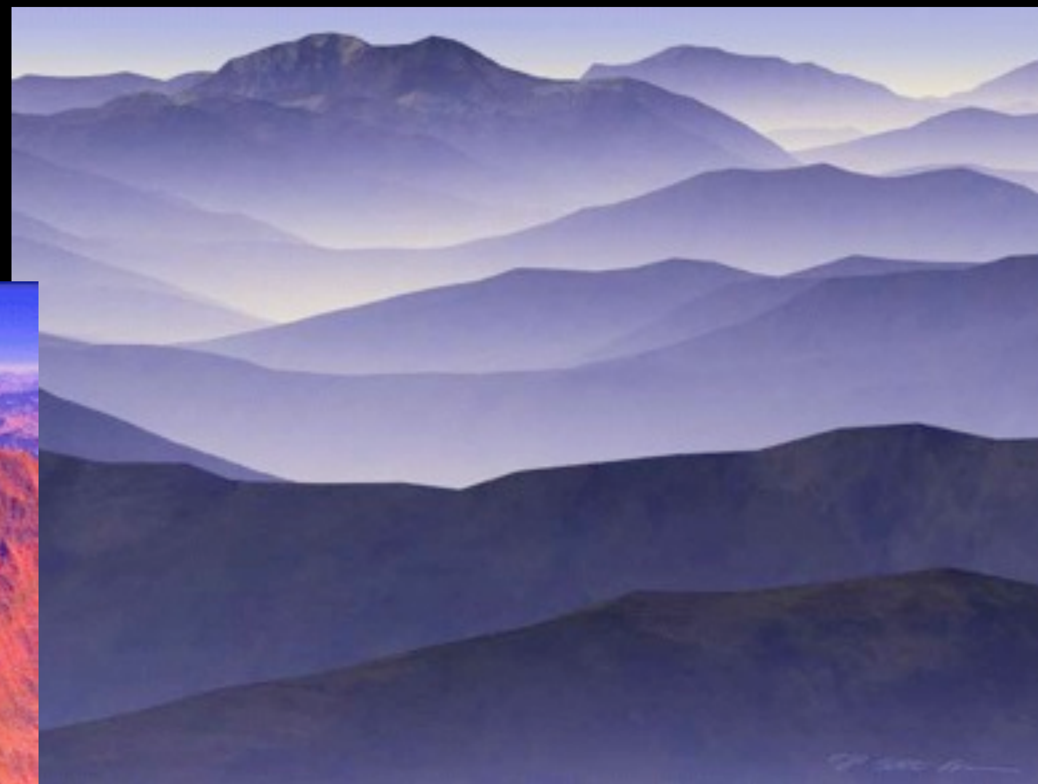

© Ken Musgrave

Applications: **Scene Synthesis**

- Height map for terrain

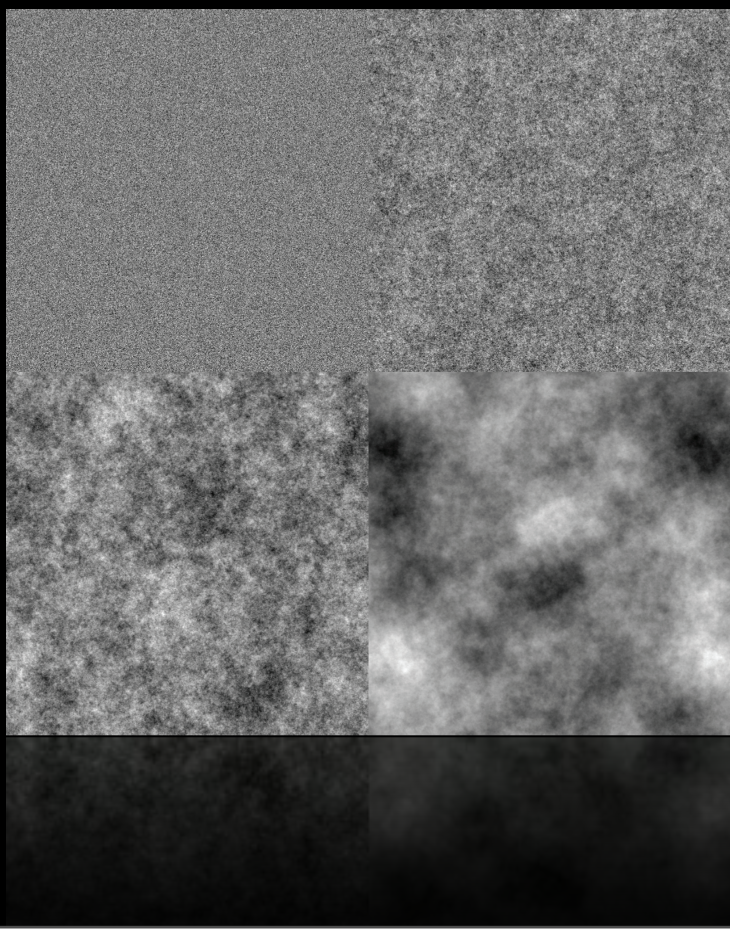

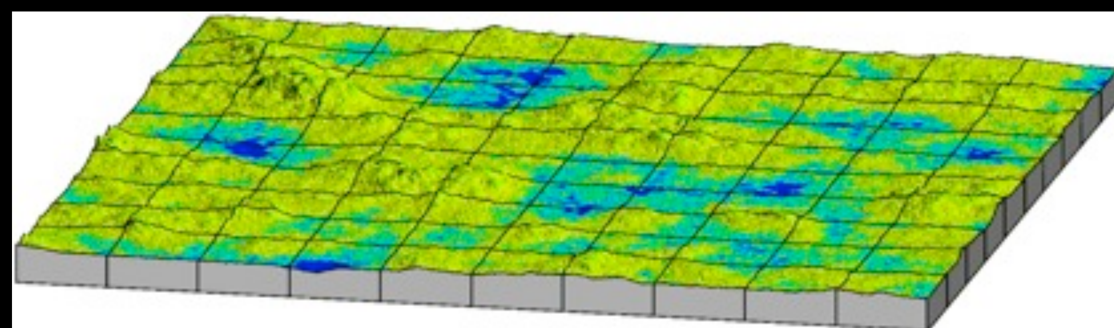

Deussen

Applications:
Scene Synthesis

- Plant modeling

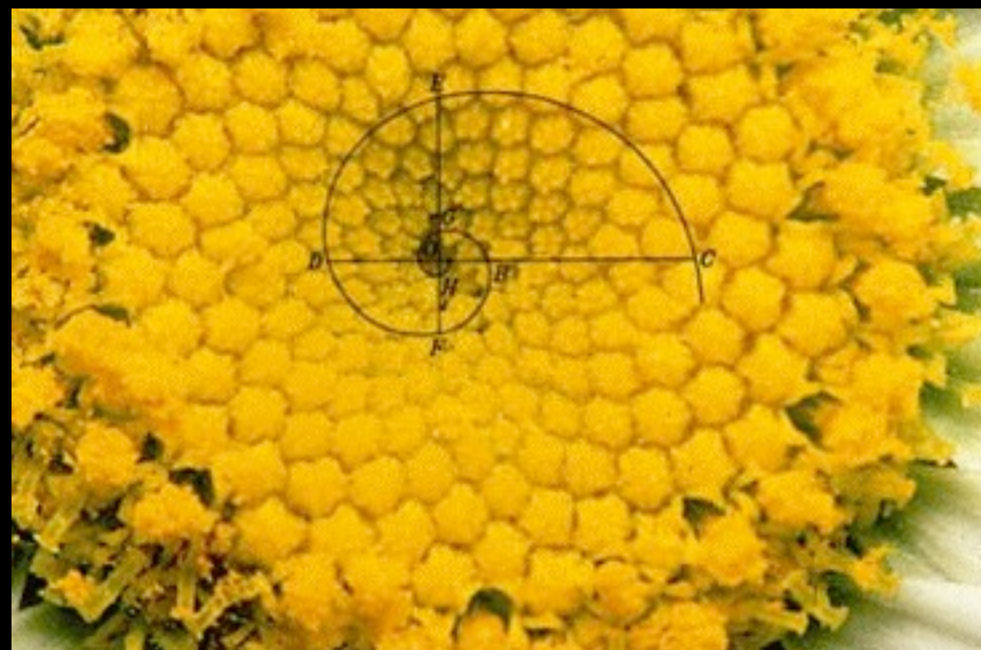

Weber

Applications:

Deblurring

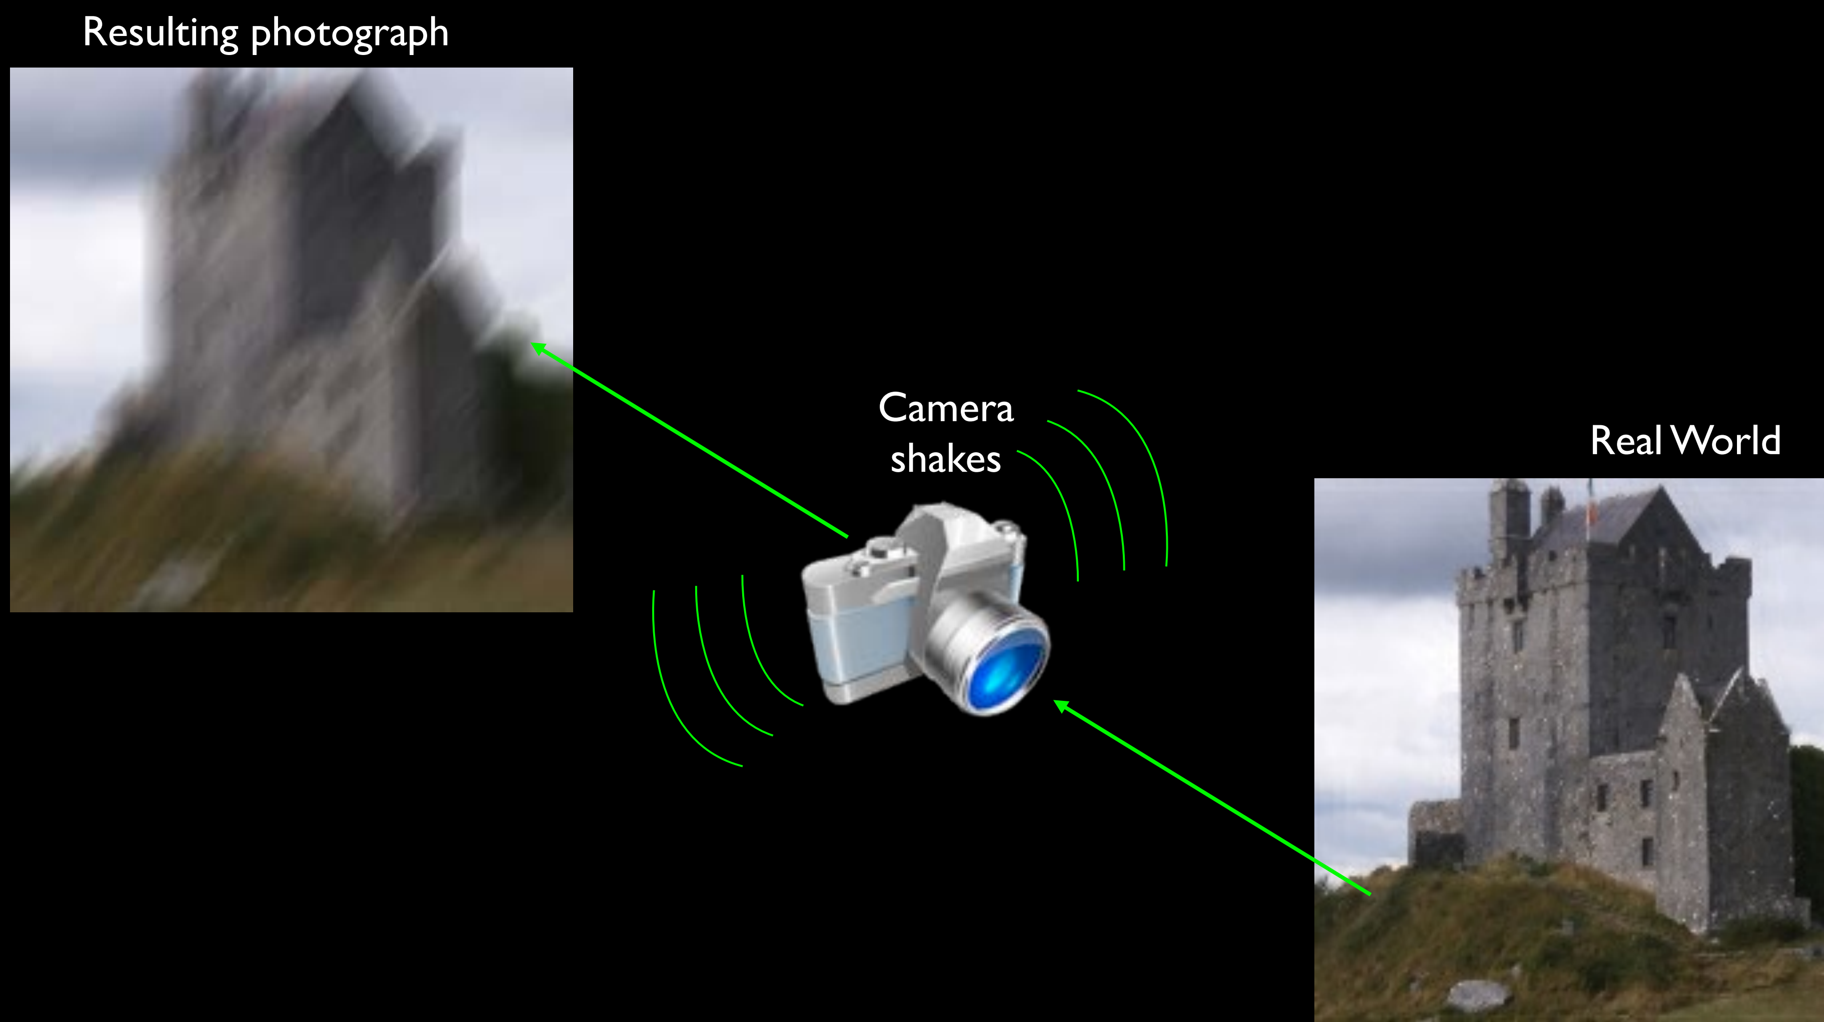

Applications:

Deblurring

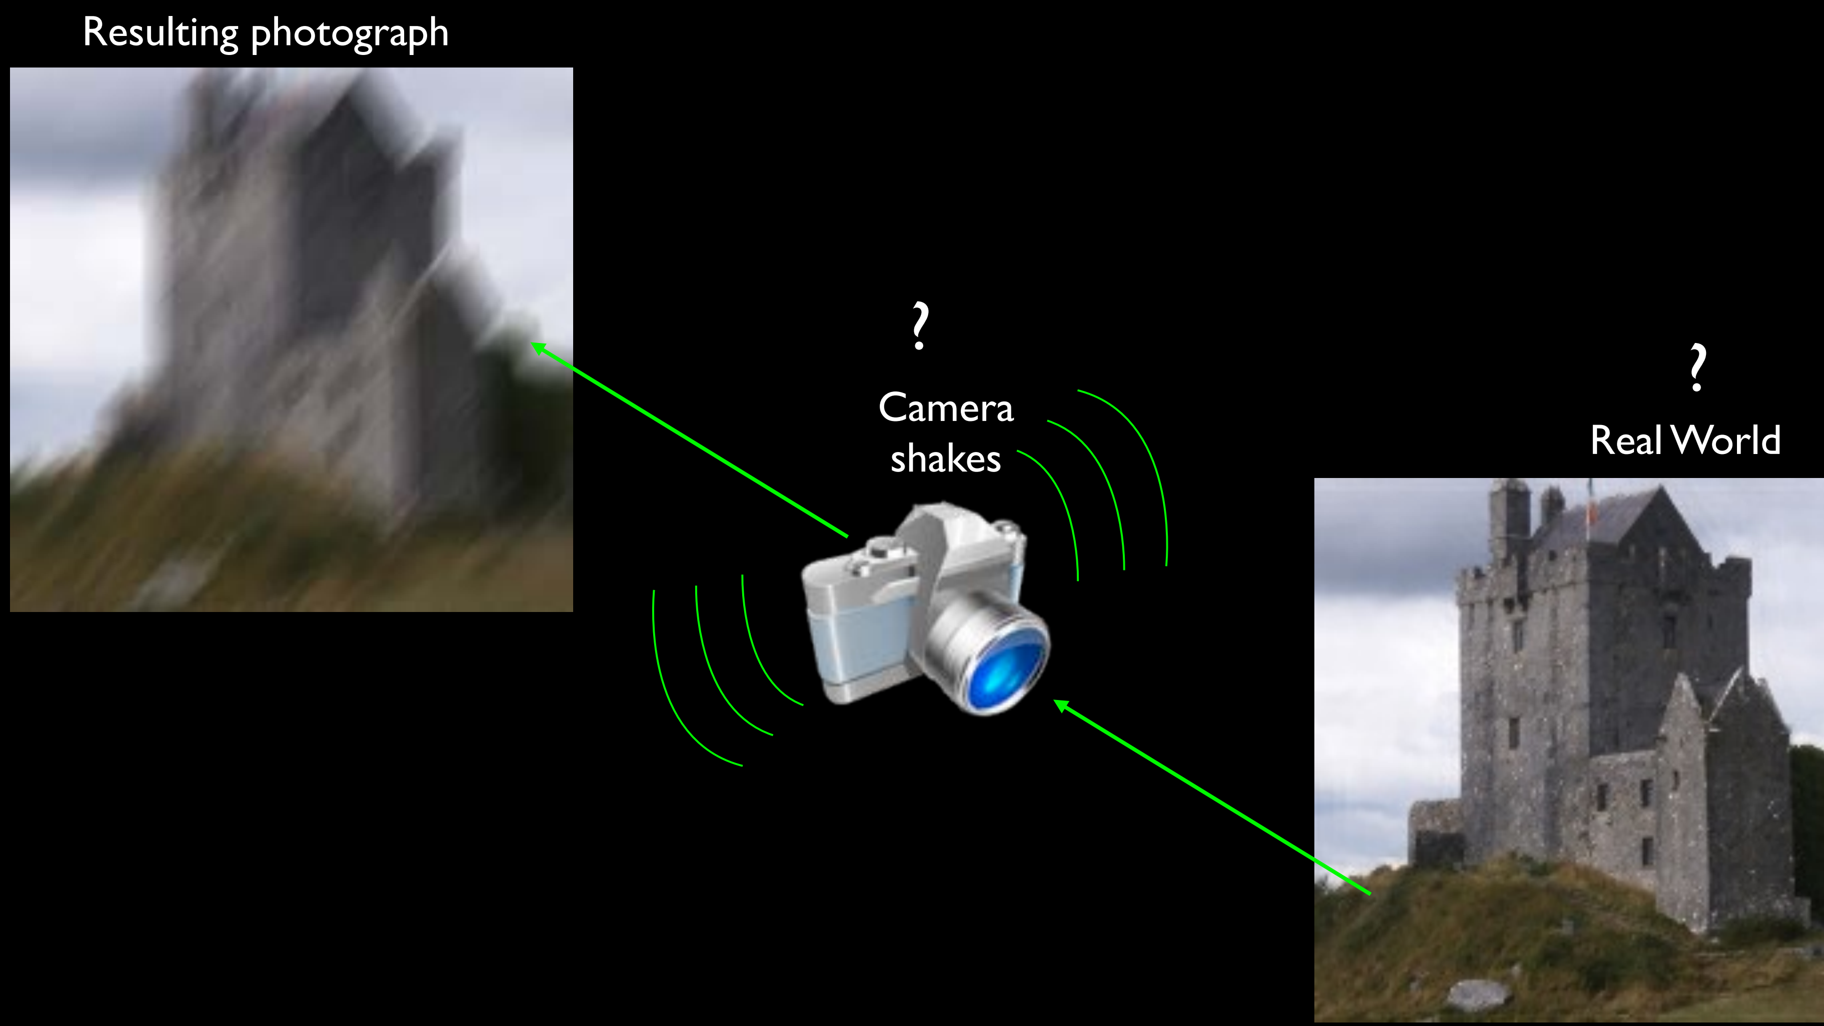

Applications:

Blind motion Deconvolution

Constraint: Real world image must follow power law
(Caron et al, 2002; Jalobeanu et al, 2002; Neelamani et al, 2004)

Constraint: Estimate Blur by optimizing to match real gradient distributions

Before

After

Applications:
Image Inpainting

- Remove an unwanted object from an image
- Fill in hole by copying from elsewhere in image

Match based on spectra (and other) information
(Hirani & Totsuka, 1996)

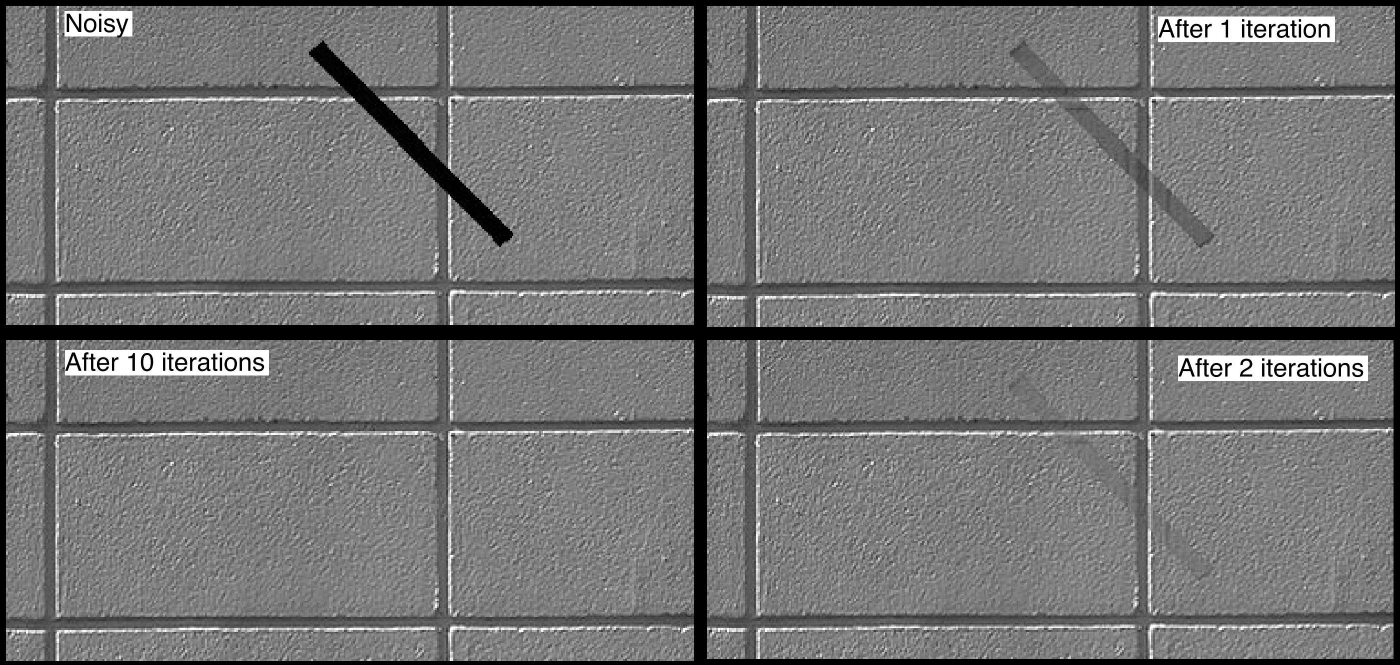

Wavelets

Erik Reinhard

Phase Structure

Wavelets

- Phase spectra are computed over entire images
- What about spatially localized analysis?
- Wavelets do this
- They are also selective to specific orientations and scales

Gabor Filters

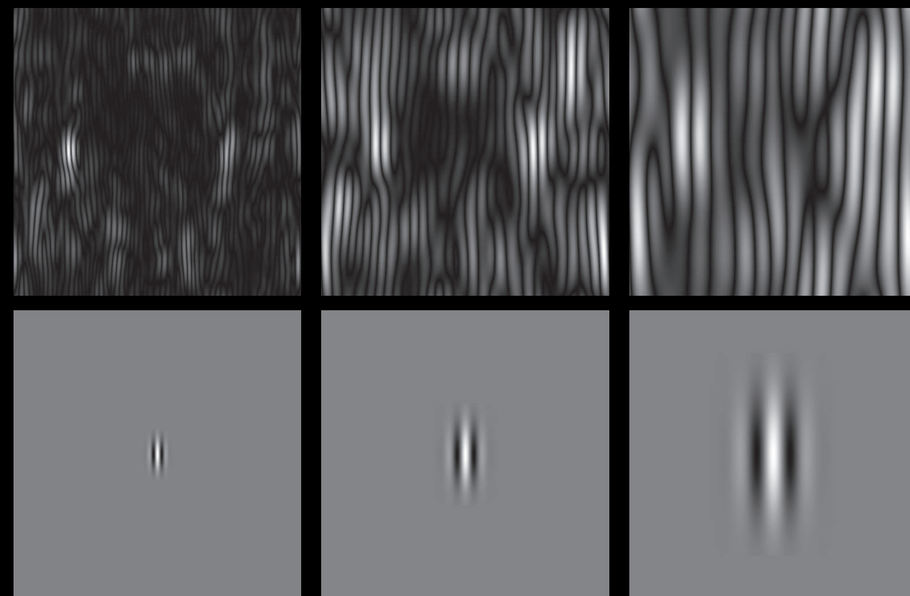

Sinusoids weighted by Gaussians

Haar Decomposition

Haar Decomposition

Haar Decomposition

Coefficient Histogram

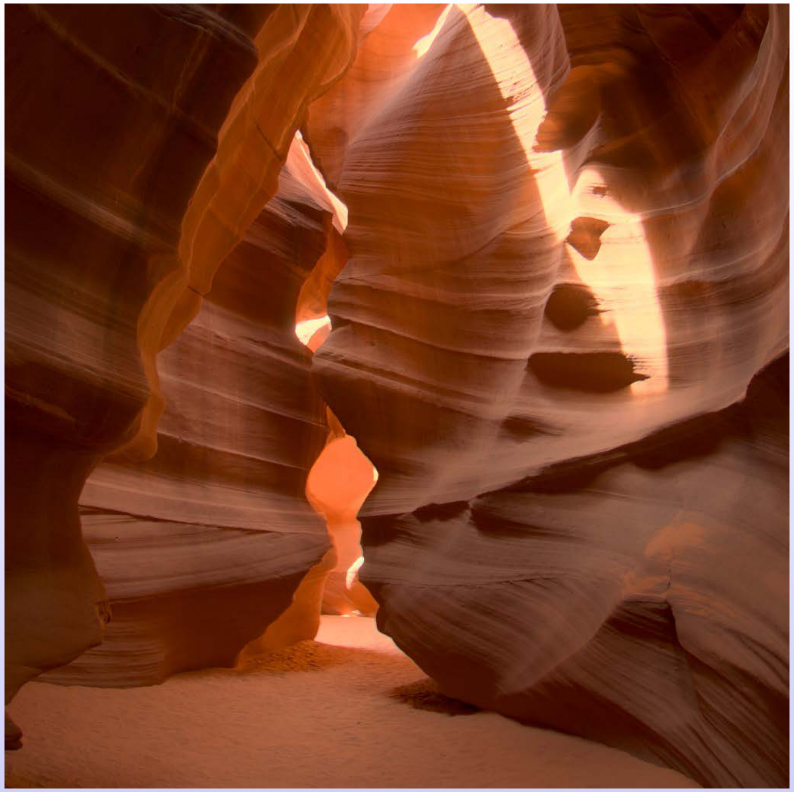

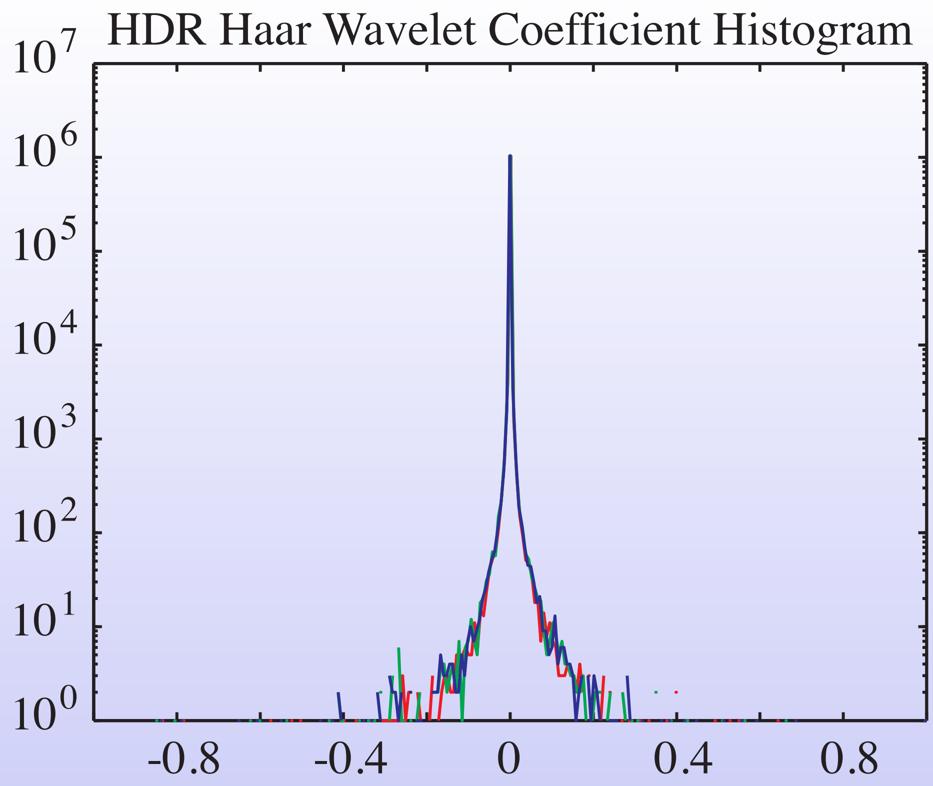

Wavelet Analysis

- Distributions of histograms of wavelet coefficients have high kurtosis, i.e. long tails
- Can be modeled with a Laplacian

Meaning of High Kurtosis

- Many natural image statistics end up showing high kurtosis
- This means that lots of values are small and some are large
- Effectively sparse coding

Sparse Coding

- In human vision, sparse coding is an important feature:
  - Variability of input is explained by fewer neurons
  - Metabolic efficiency
  - Minimizes wiring length
  - Increases capacity in associative memory

Wavelet Analysis

- Both phase and amplitude can be measured and correlated in a wavelet decomposition
- Surprising result: natural images are scale-invariant in both phase and amplitude

Complex Wavelet Amplitude

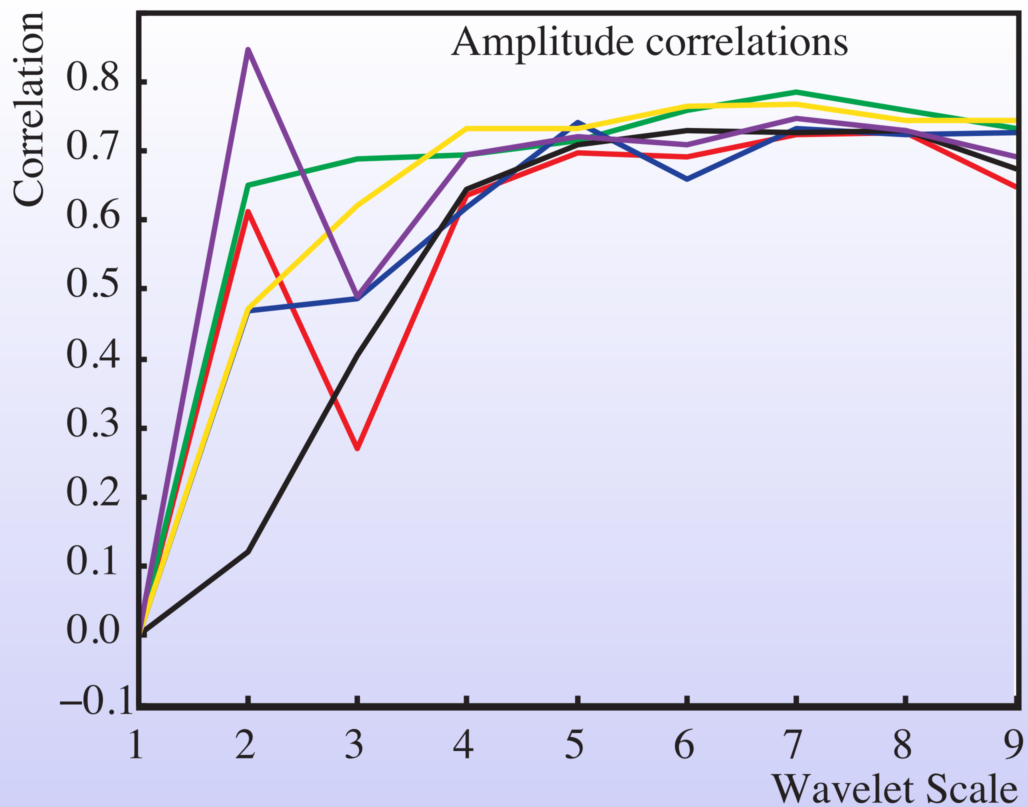

Complex Wavelet Amplitude

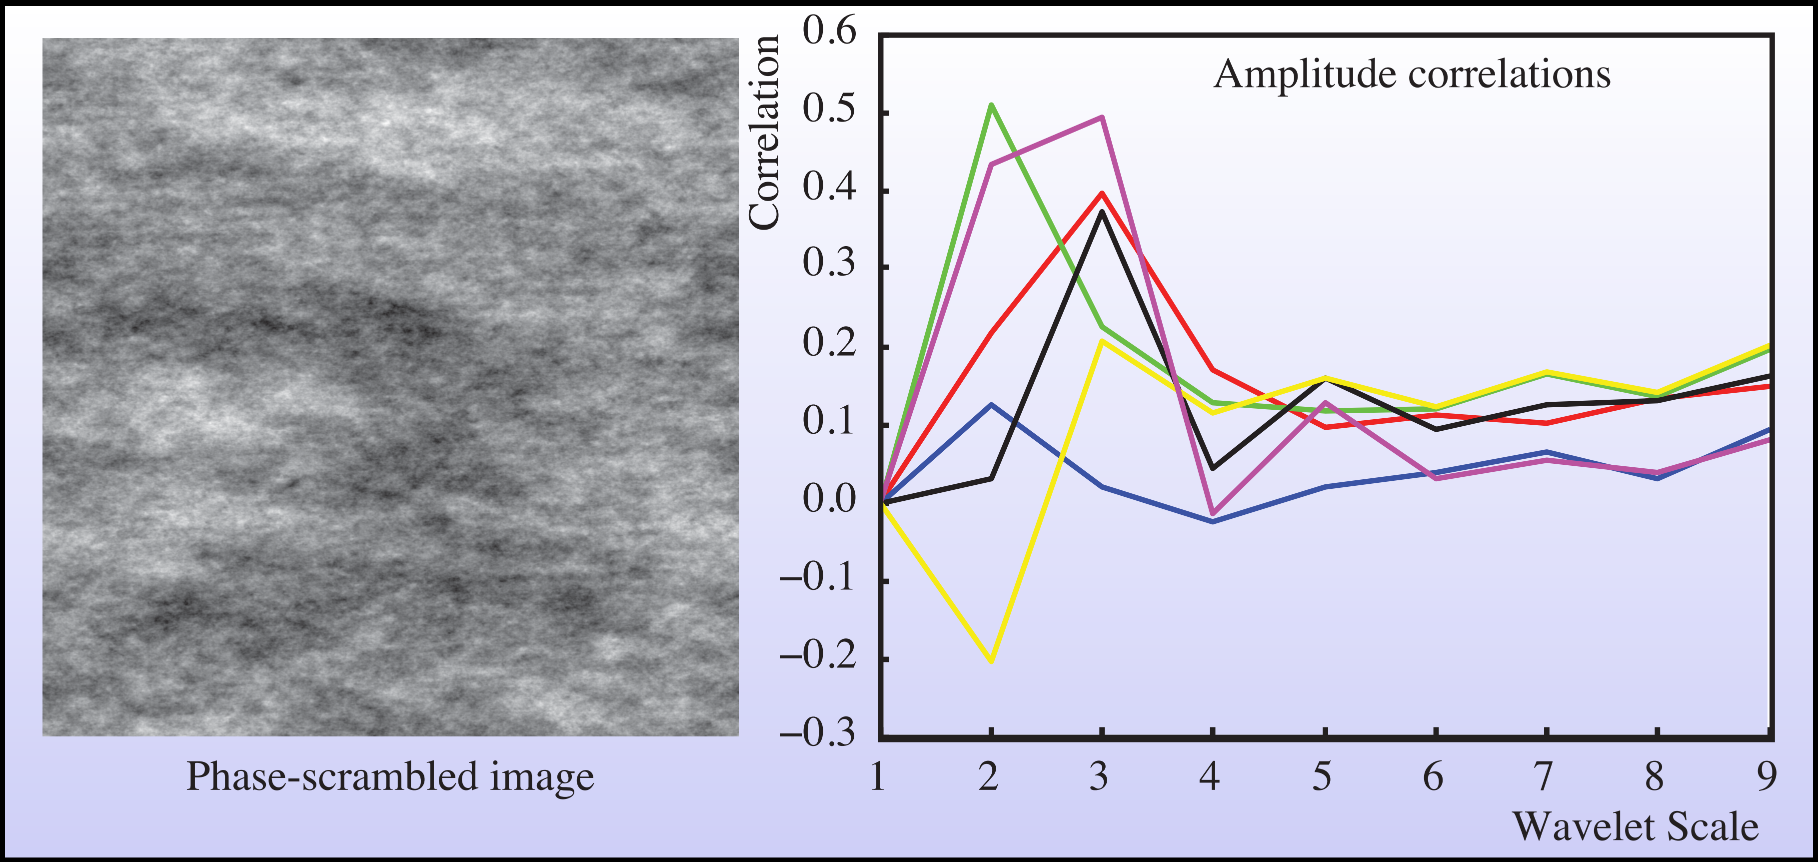

Applications of Wavelets

- Image denoising
- Image compression
- Object detection
- Image retrieval

Image Denoising

Face Detection

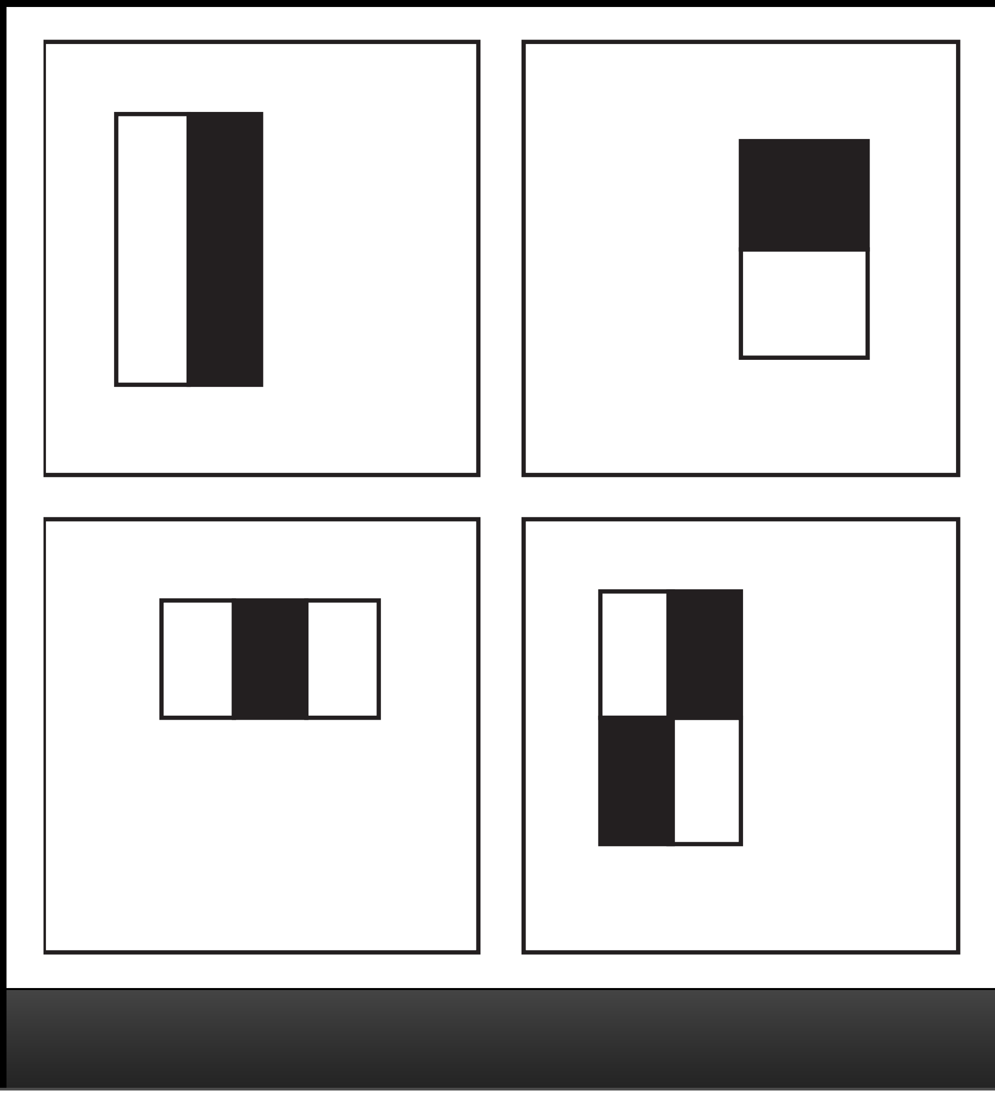

Viola & Jones use a small set of wavelet-like features to detect faces

Wavelet Reconstruction

Color Statistics

Erik Reinhard

Light Transduction

$$
\begin{aligned}
L_o(\lambda) &= L_e(\lambda) + \int_\Omega L_i(\lambda)\ f_r(\lambda)\ \cos(\Theta) d\omega \\
L &= \int_\lambda L_o(\lambda)\ \bar{l}(\lambda)\ d\lambda \\
M &= \int_\lambda L_o(\lambda)\ \bar{m}(\lambda)\ d\lambda \\
S &= \int_\lambda L_o(\lambda)\ \bar{s}(\lambda)\ d\lambda
\end{aligned}
$$

Implications

- Metamerism: different spectra integrate to the same cone responses, and are therefore perceived identically
- This allows us to build color displays, for instance
- Color statistics can be collected on tristimulus values, rather than color spectra

Color Constancy

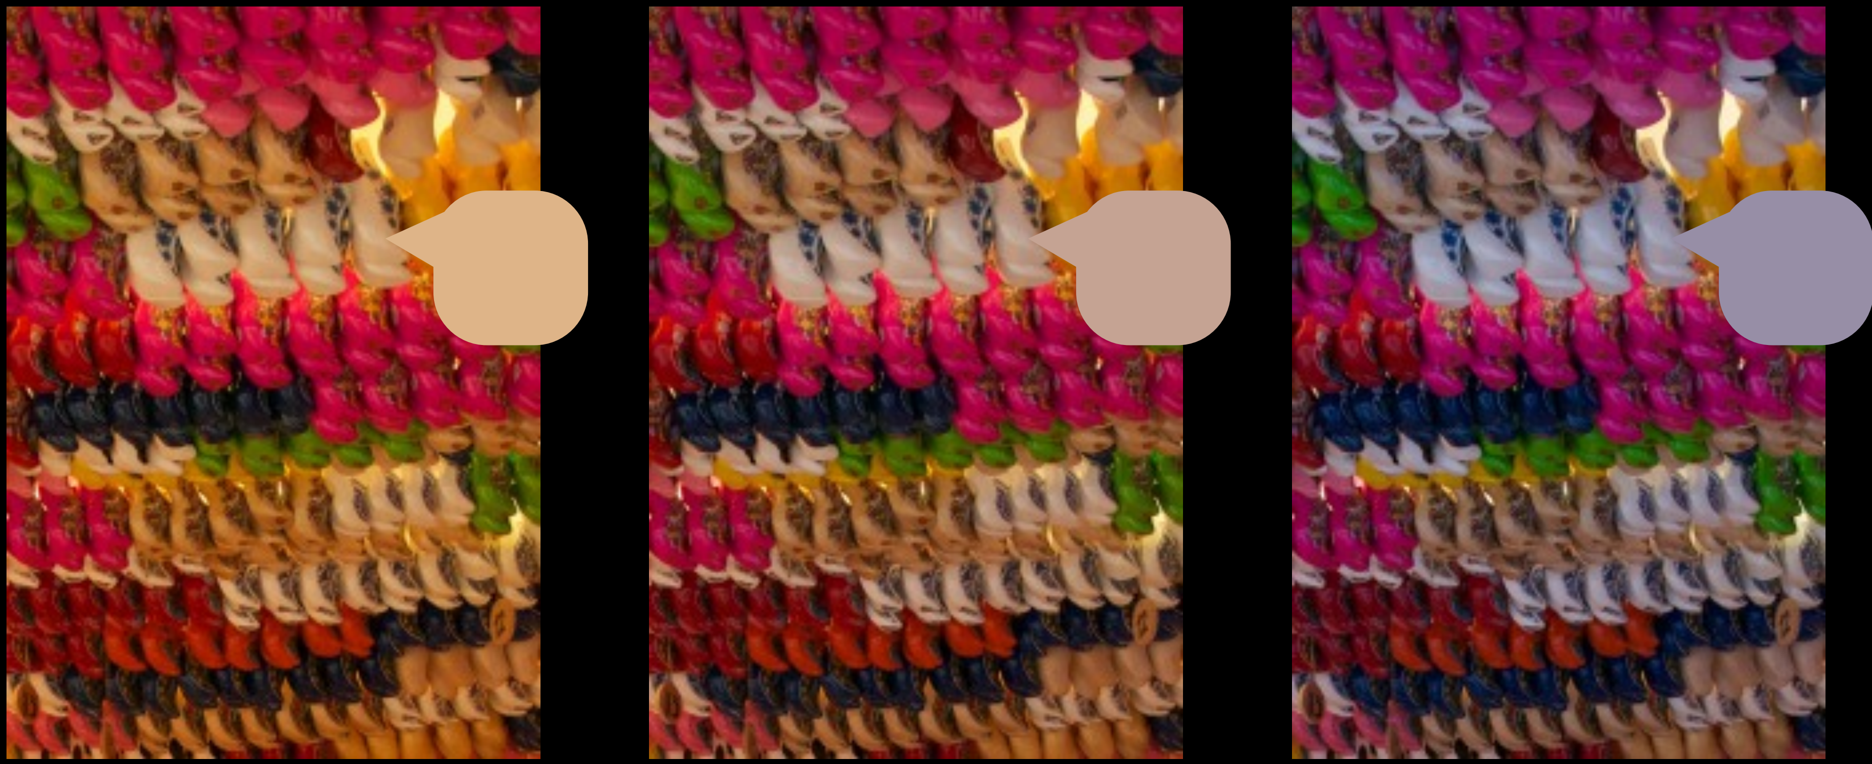

- Humans can discount the color of the illumination

Color Constancy

- Cannot be computed analytically from retinal input; it is an under-constrained problem
- Human vision makes statistical assumptions

Statistical Assumptions

- Grey world:
  - Spectrum of light sources usually off-white
  - Average BRDF of a scene often close to grey
  - Average color of an image yields estimate of dominant illuminant

Grey World

Grey World - Failure

Possible Fixes

- Exclude most saturated pixels from average
- Optionally: convert to CIELAB
  - Compute 2D histogram on a* and b* channels
  - Spread of histogram and distance to origin determine if color cast is likely due to illumination or reflectance

White Patch Algorithm

- Assume that lightest patches in the scene are neutral in color
- Their color therefore represents the illuminant

Grey-Edge Assumption

- The difference between two colored pixels tends to evaluate to grey

Algorithm Selection

- Different white balancing algorithms tend to work best on specific types of images
- Can therefore collect statistics on the image pixels and select an appropriate algorithm based on the outcome
- **Weibull distribution** is shown to be indicative

A Further Implication

Remember

- For grey values, in RGB (as well as LMS and similar color spaces) we have R=G=B
- If values average to grey, then in RGB-like color spaces strong correlations exist between channels

Statistical Decorrelation

Correlations in RGB/LMS

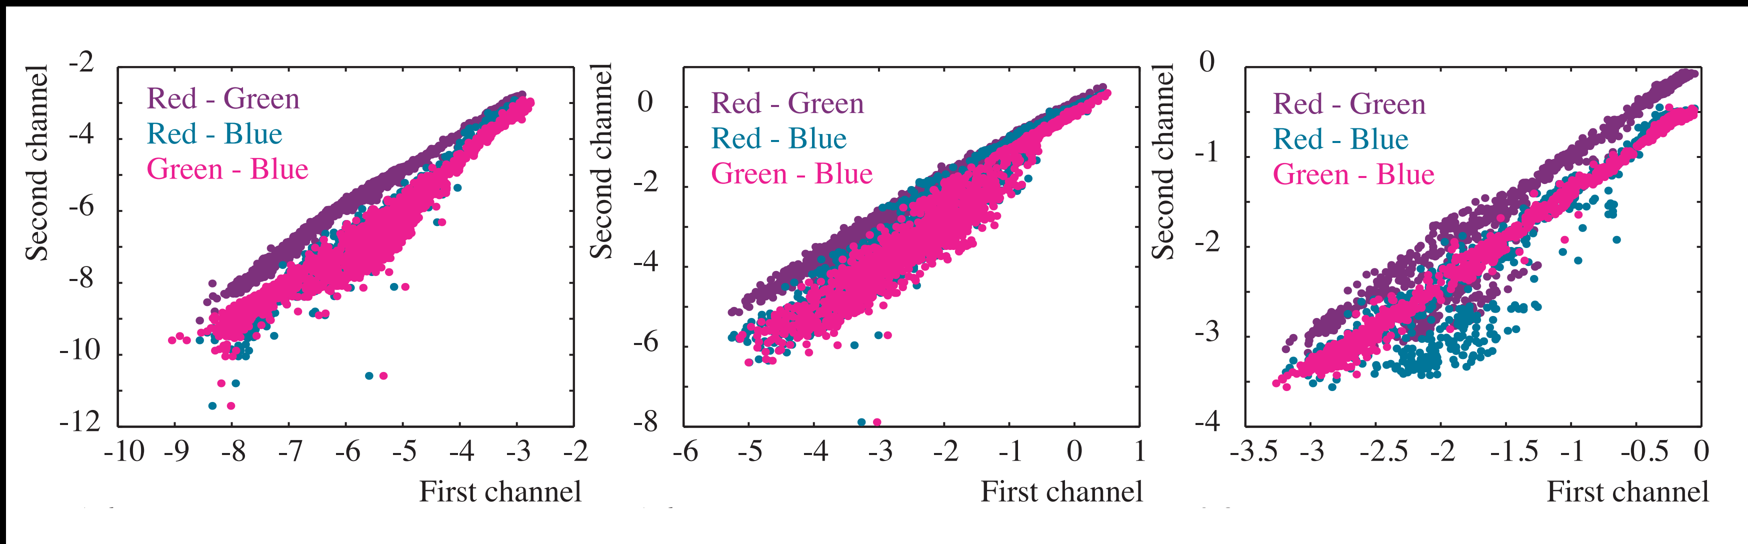

Color Opponent Space

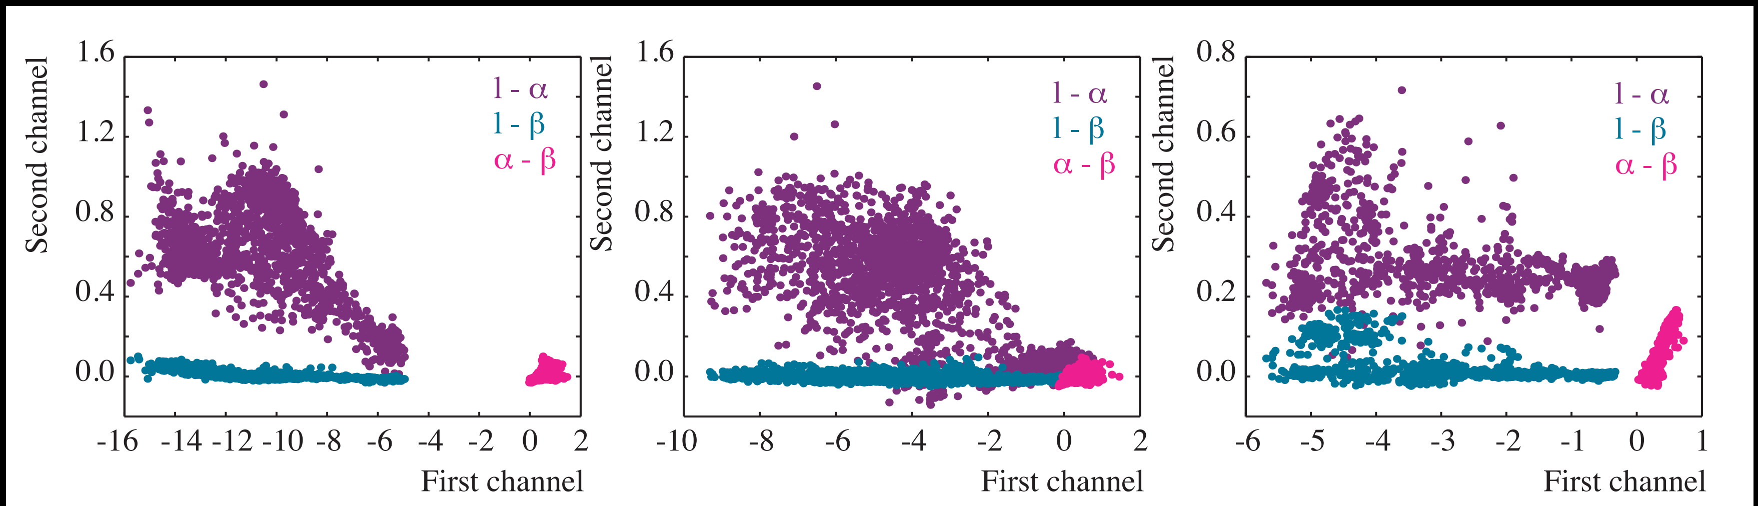

Statistical Decorrelation

- Lαβ color space
  - decorrelated - values of pixels in one channel do not predict the values in another
  - L - luminance
  - α,β - opponent channels

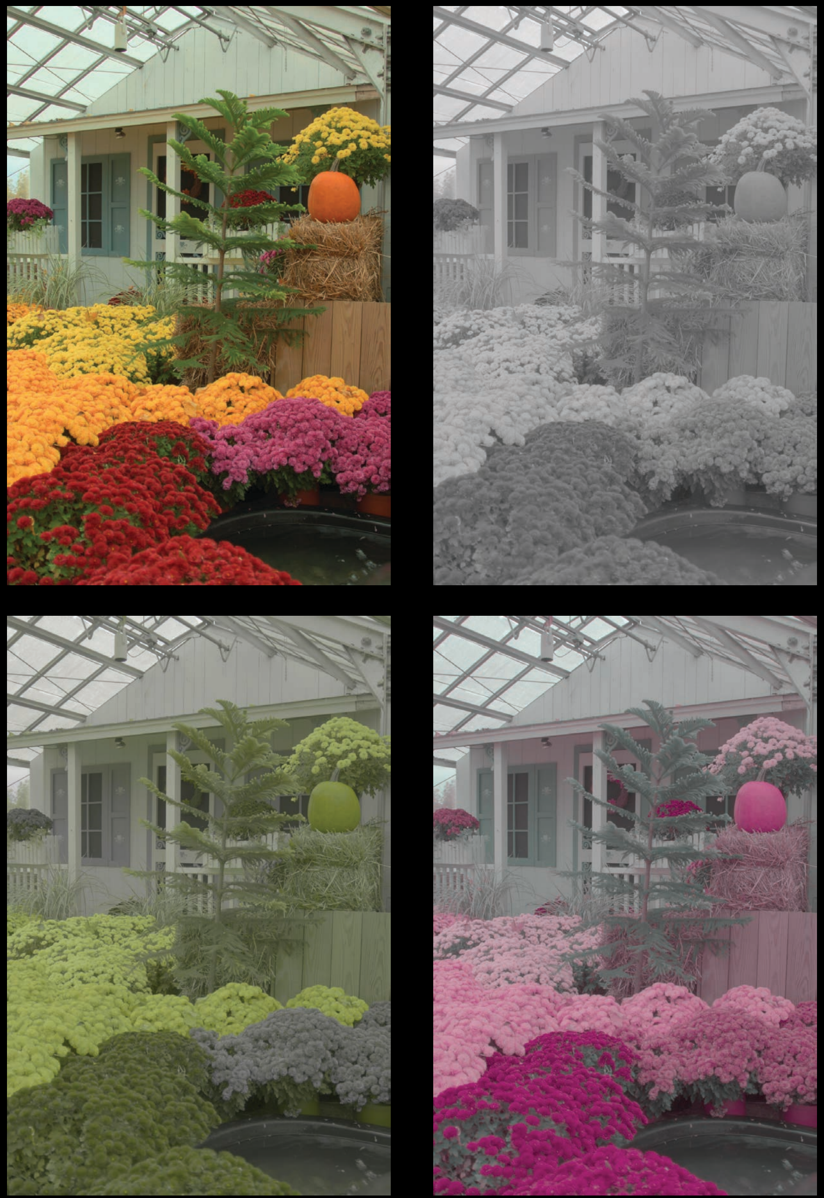

Histogram Matching

Reminder

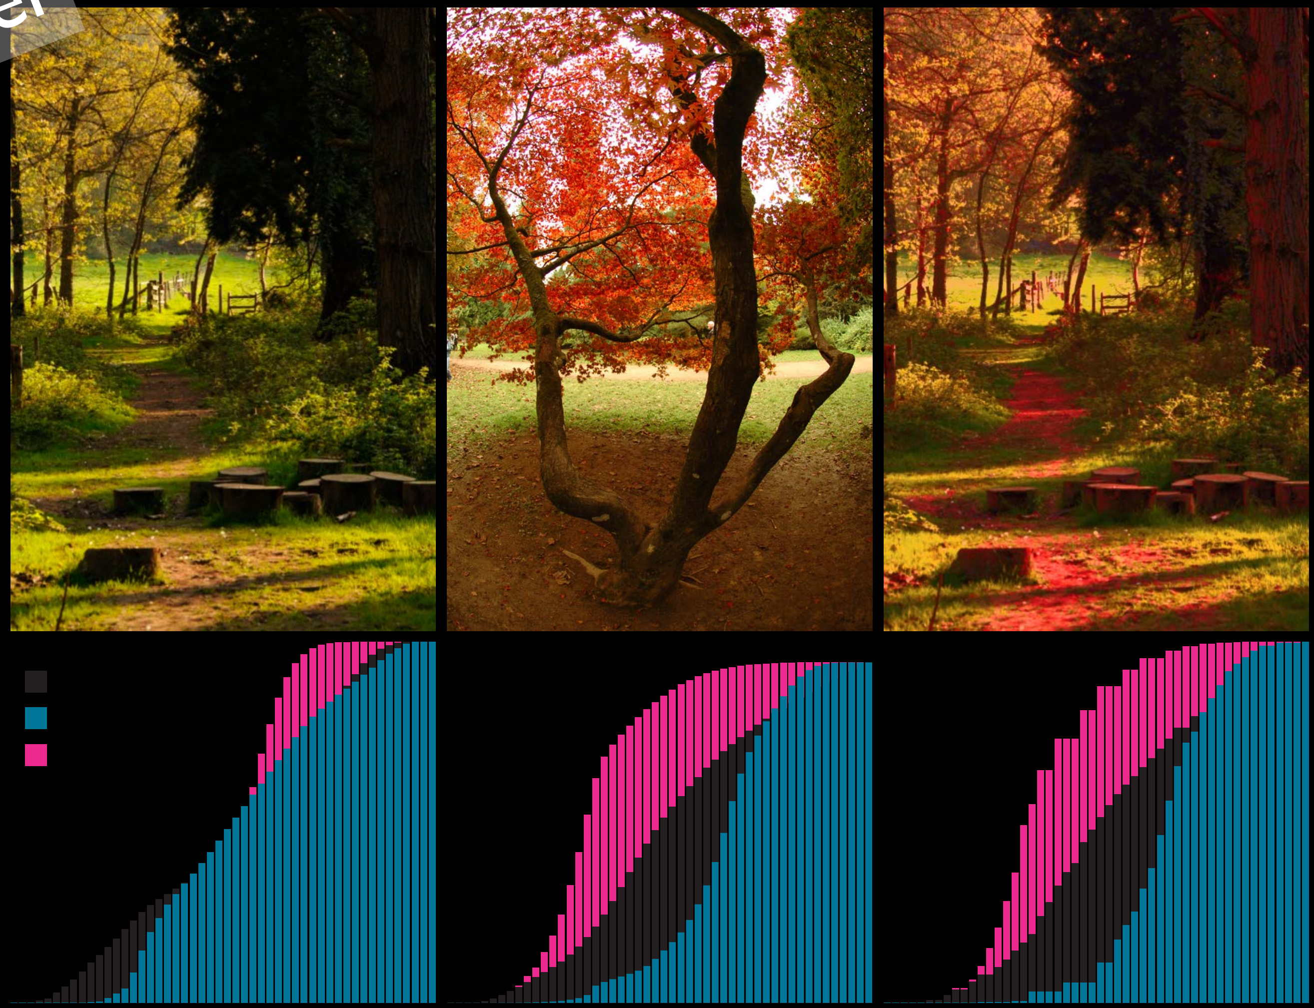

Source Target Result

Color Transfer

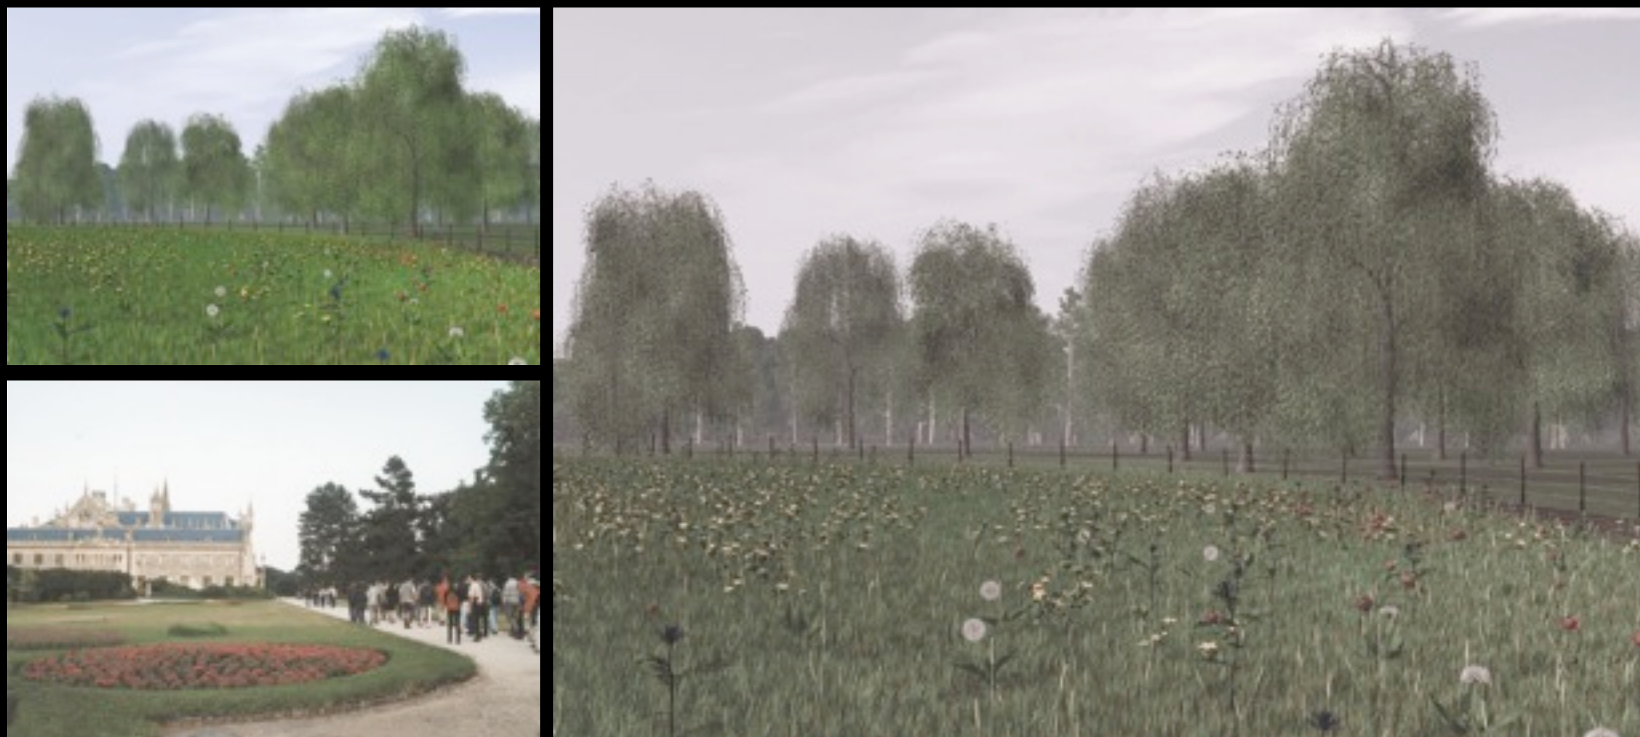

Color transfer between images
(Reinhard et al.2001)

Histogram Reshaping

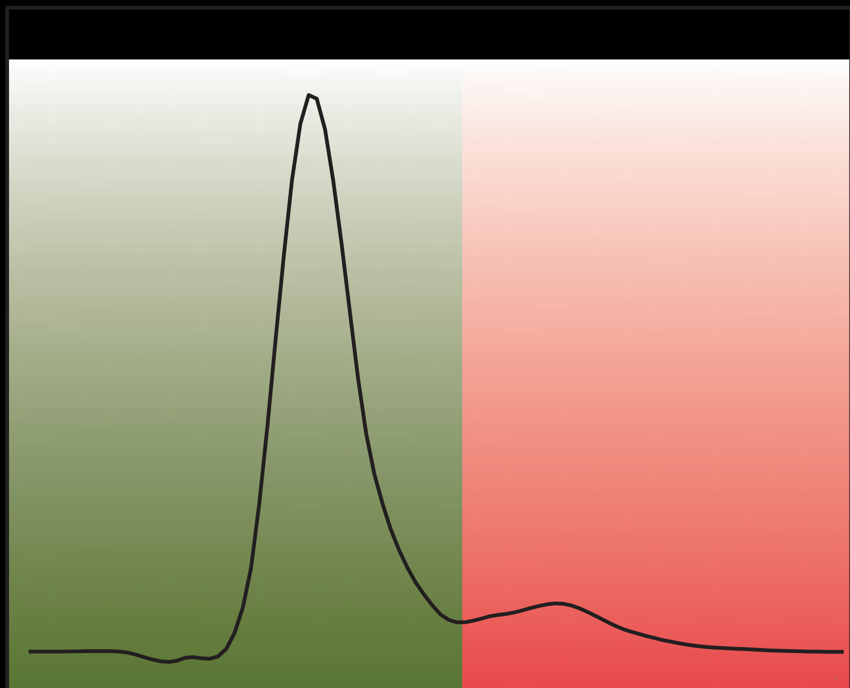

Histogram Reshaping

Histogram Reshaping

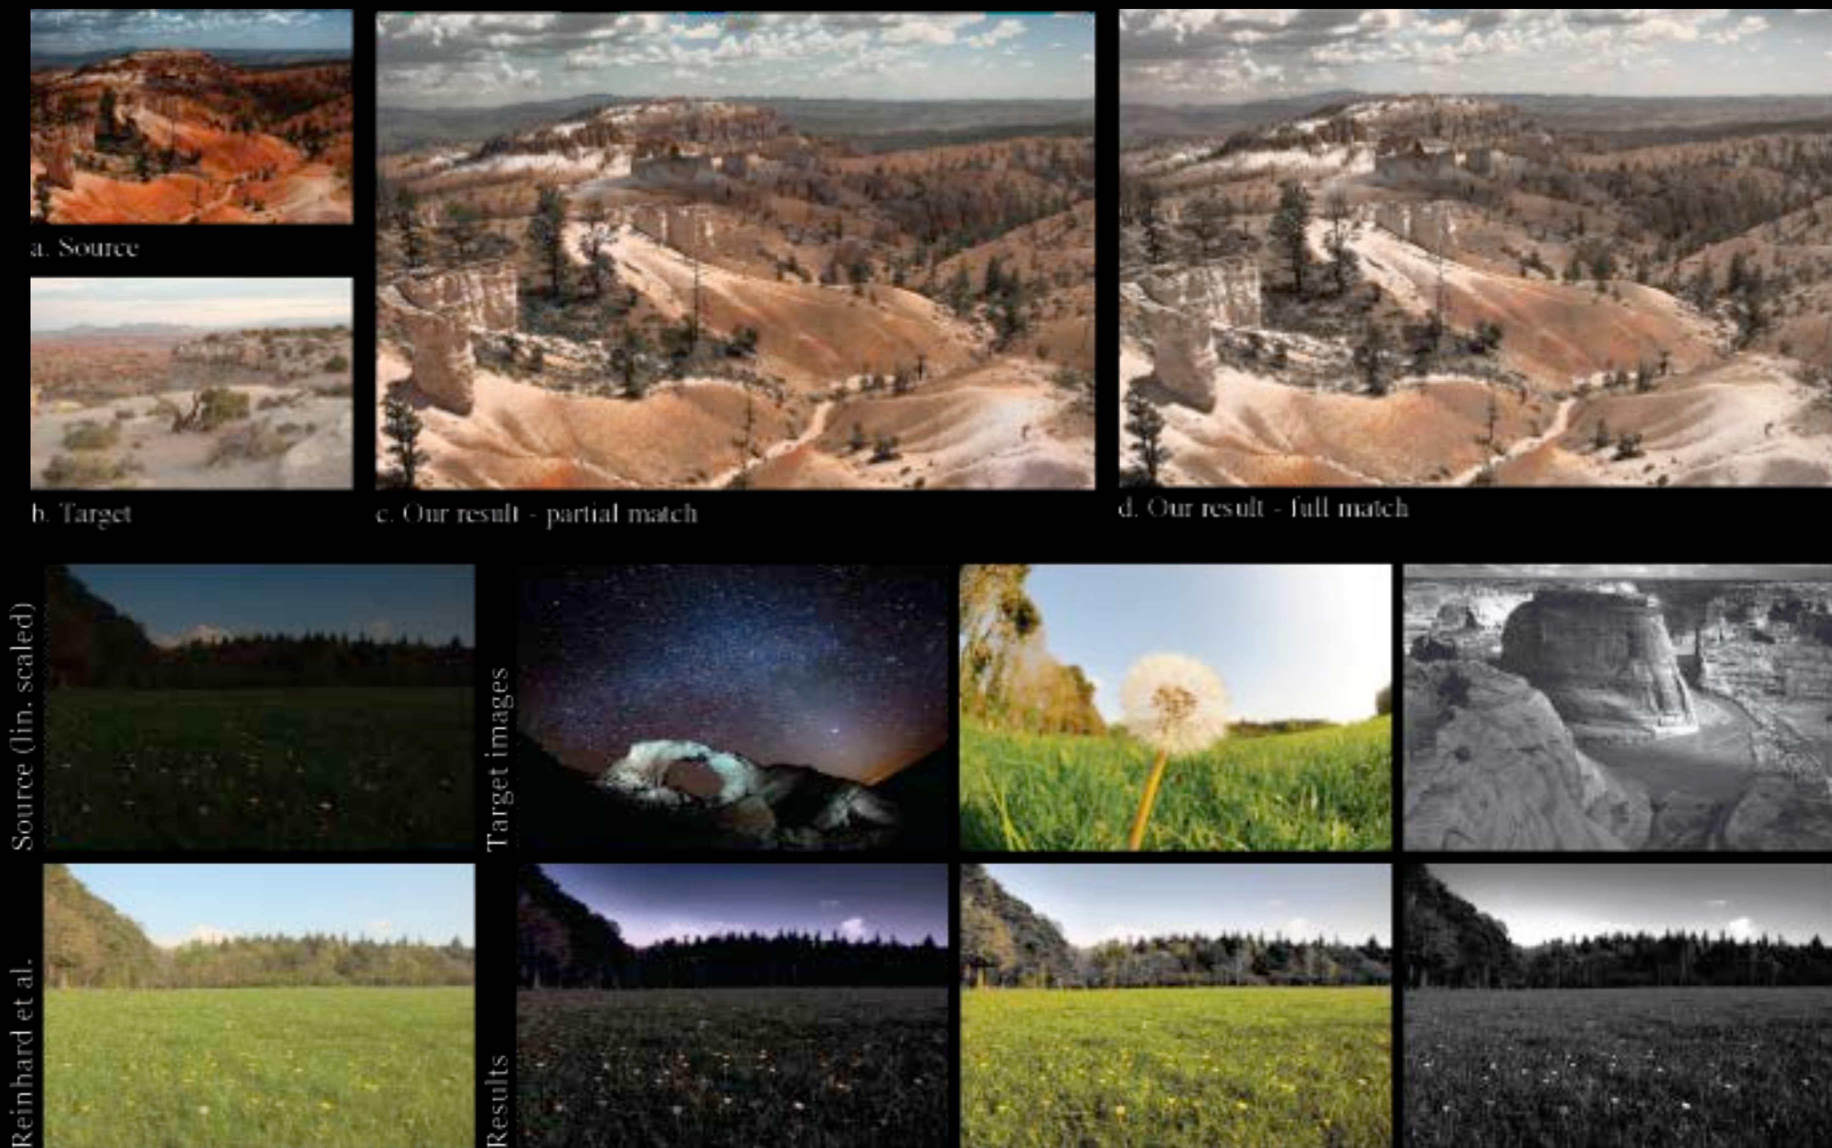

Conclusions

Statistics

- There are many ways to transform images, after which we can compute statistics
- When we transform images according to how we think the human visual system operates, we end up with highly kurtotic and sometimes independent representations
- Sparse coding is good for human vision, and probably good for solving engineering problems

Applications

- Many applications already known
  - Object detection
  - Compression
  - Deblurring
  - Inpainting
  - Color transfer
  - etc.

Applications

- Hopefully, as our knowledge of our environment increases, there will be many more to come
- Graphics, computer vision and image processing are prime areas of research that we think may benefit from natural image statistics

Questions?