High dynamic range and postprocessing for graphics

Advanced Graphics

Rafał Mantiuk

Computer Laboratory, University of Cambridge

Cornell Box: a rendering or photograph?

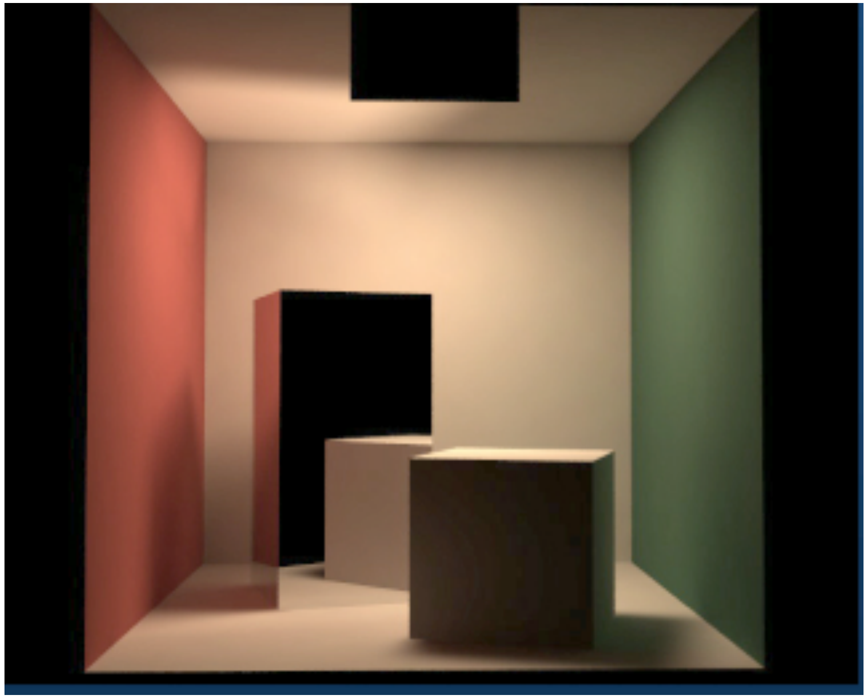

Rendering

Photograph

Real-world scenes are more challenging

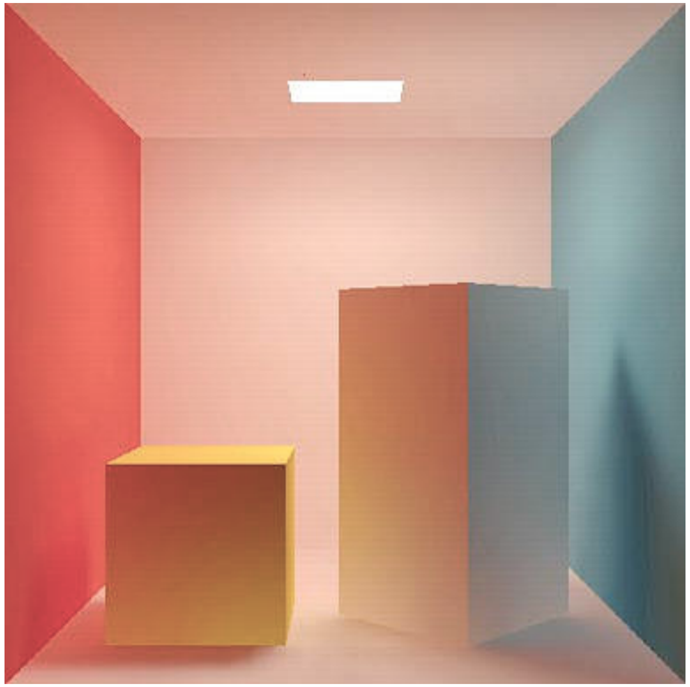

- The match could not be achieved if the light source in the top of the box was visible
- The display could not reproduce the right level of brightness

Glare

"Alan Wake" © Remedy Entertainment

Glare Illusion

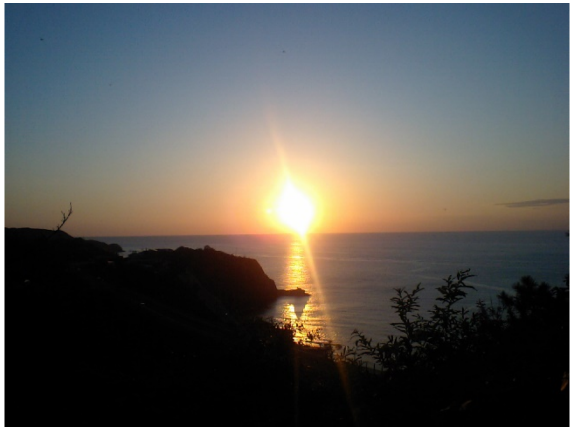

Photography

Painting

Computer Graphics
HDR rendering in games

Scattering of the light in the eye

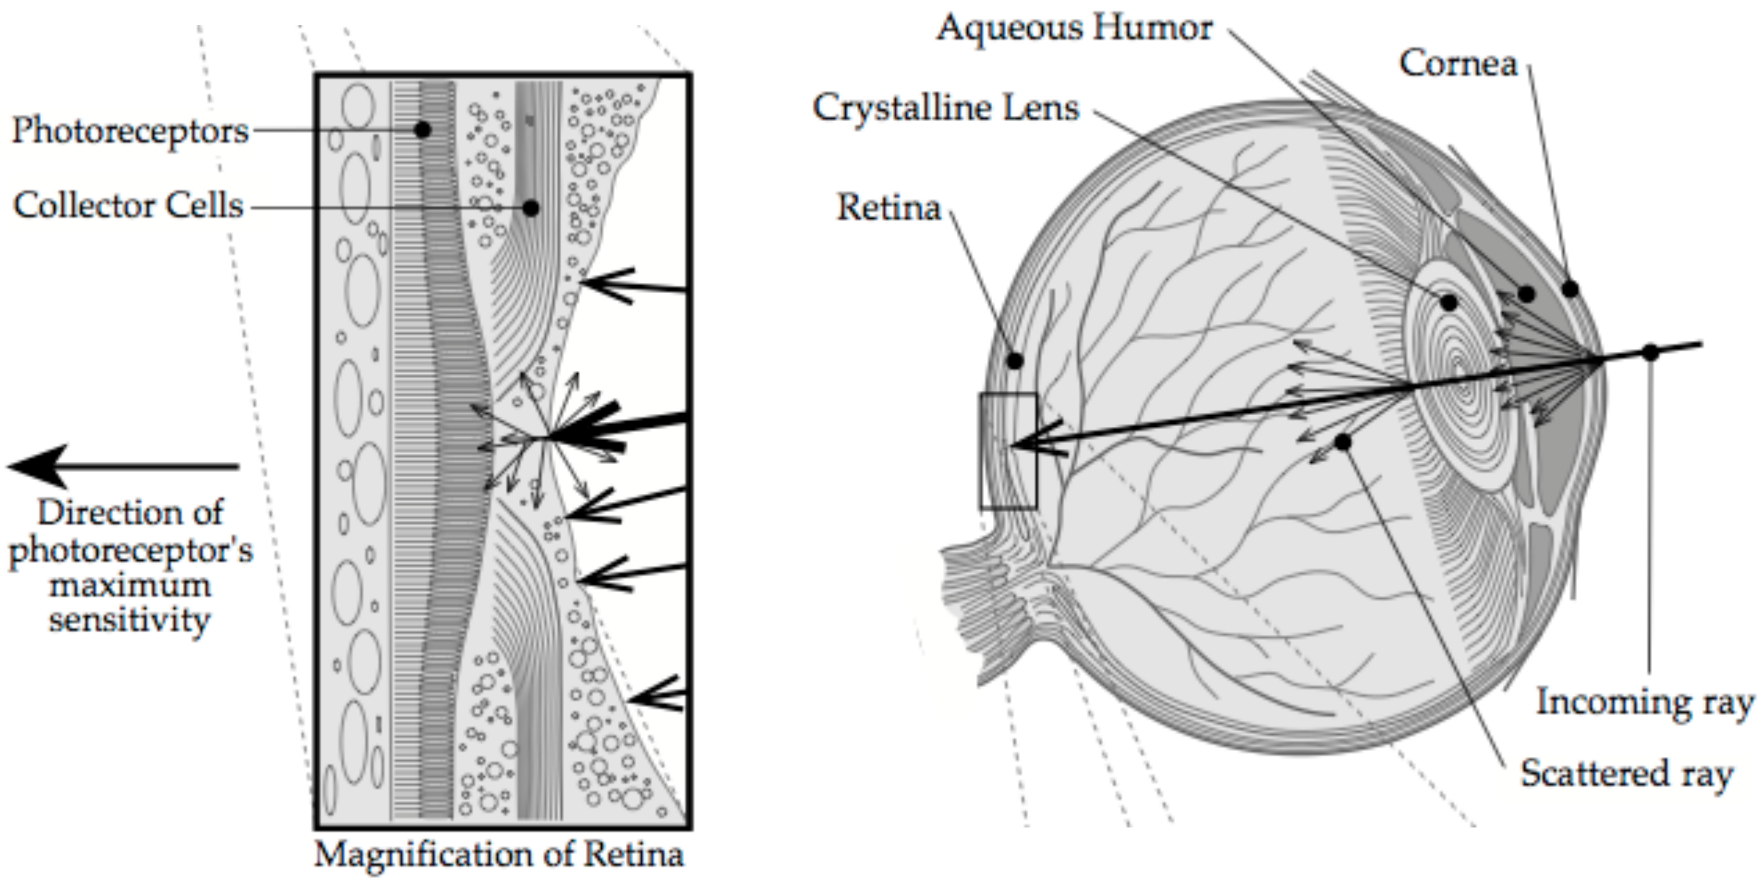

From: Sekuler, R., and Blake, R. Perception, second ed. McGraw- Hill, New York, 1990

Ciliary corona and lenticular halo

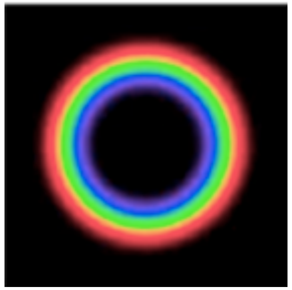

*

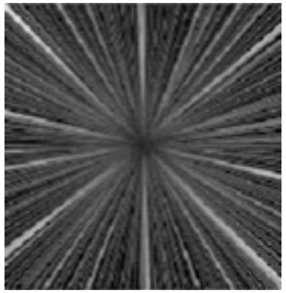

=

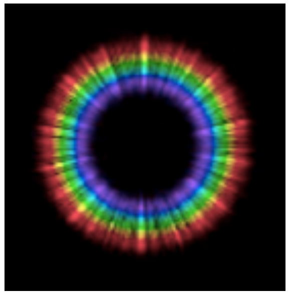 + 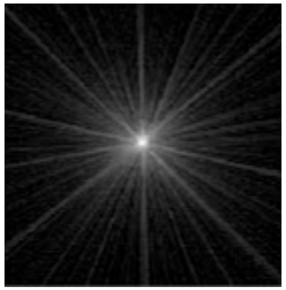 = 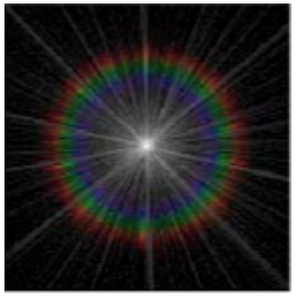

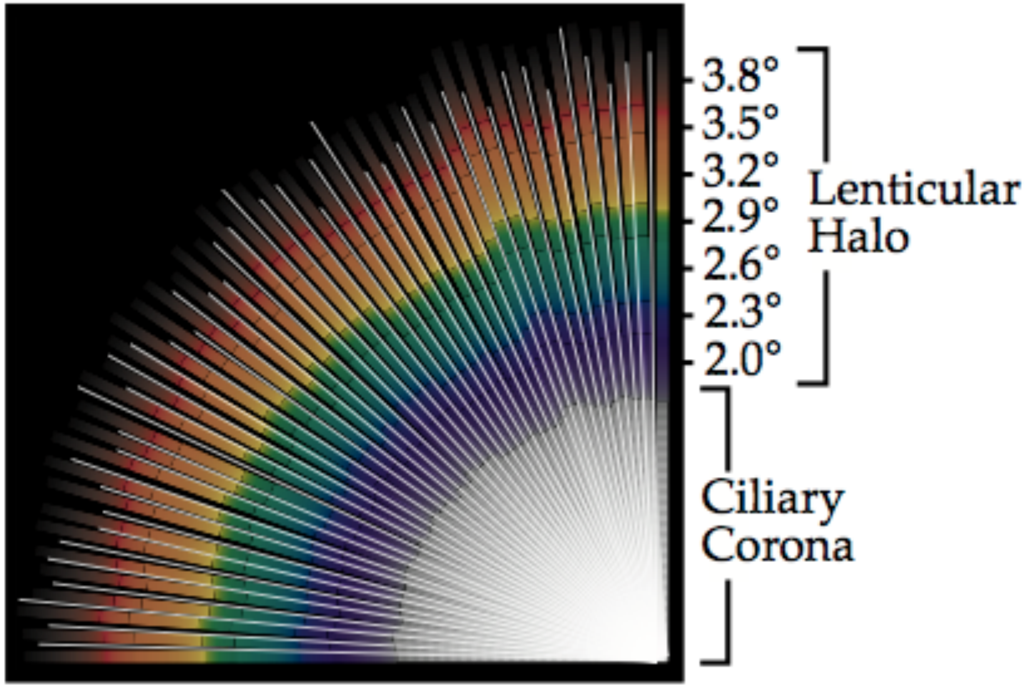

From: Spencer, G. et al. 1995. Proc. of SIGGRAPH. (1995)

Examples of simulated glare

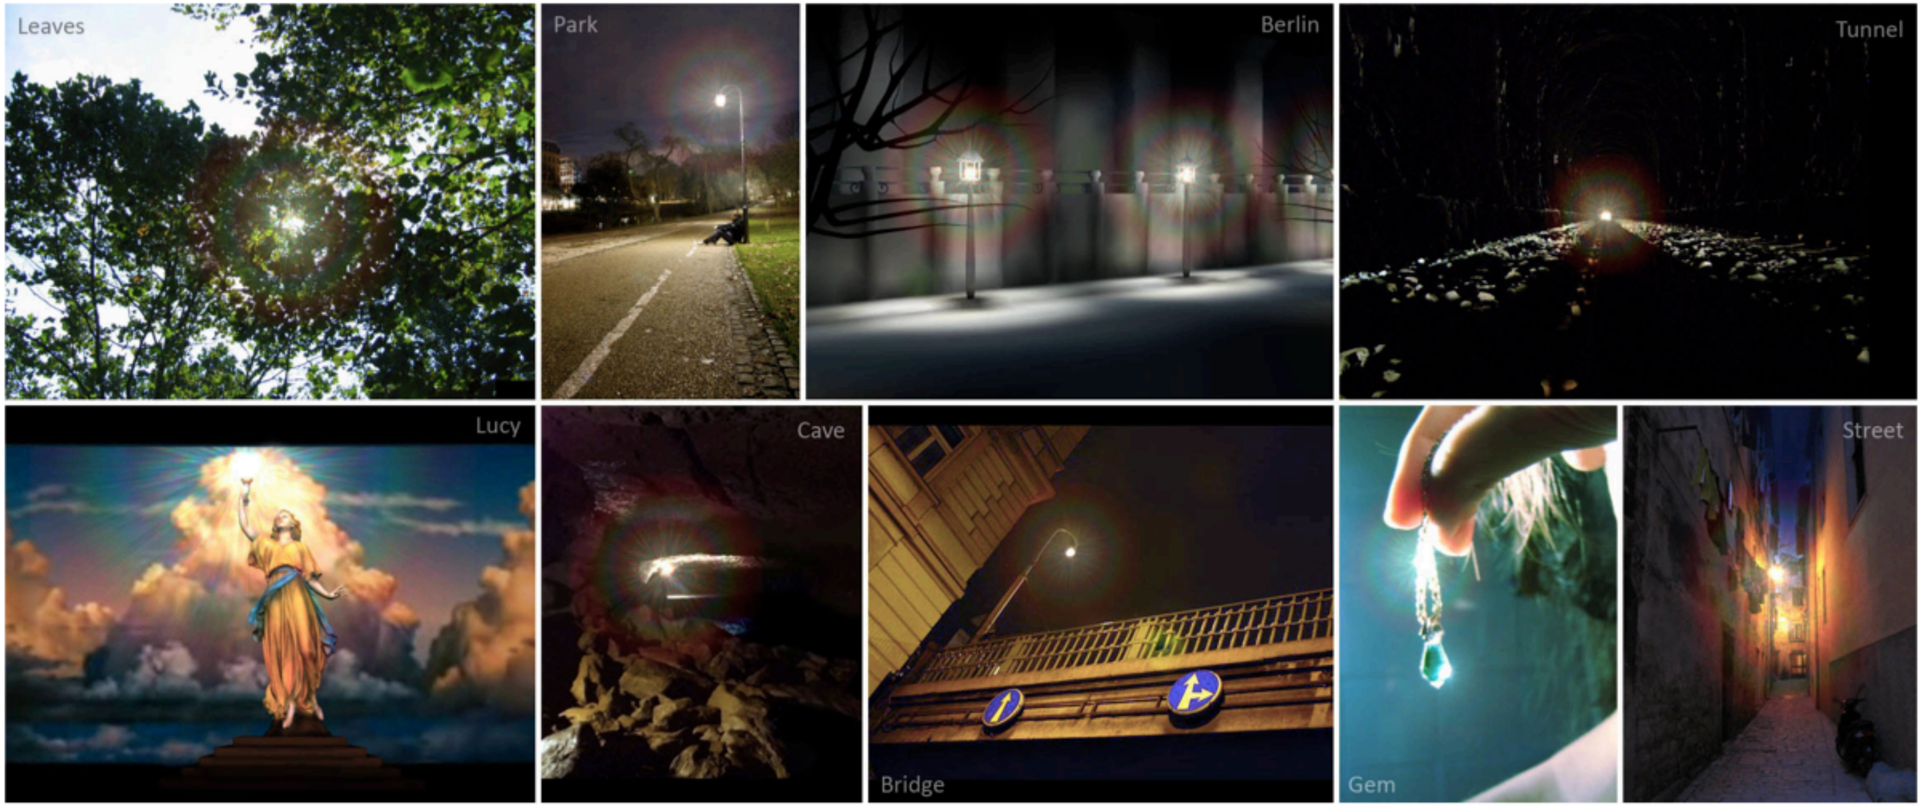

[From Ritschel et al, Eurographics 2009]

Temporal model of glare (low level)

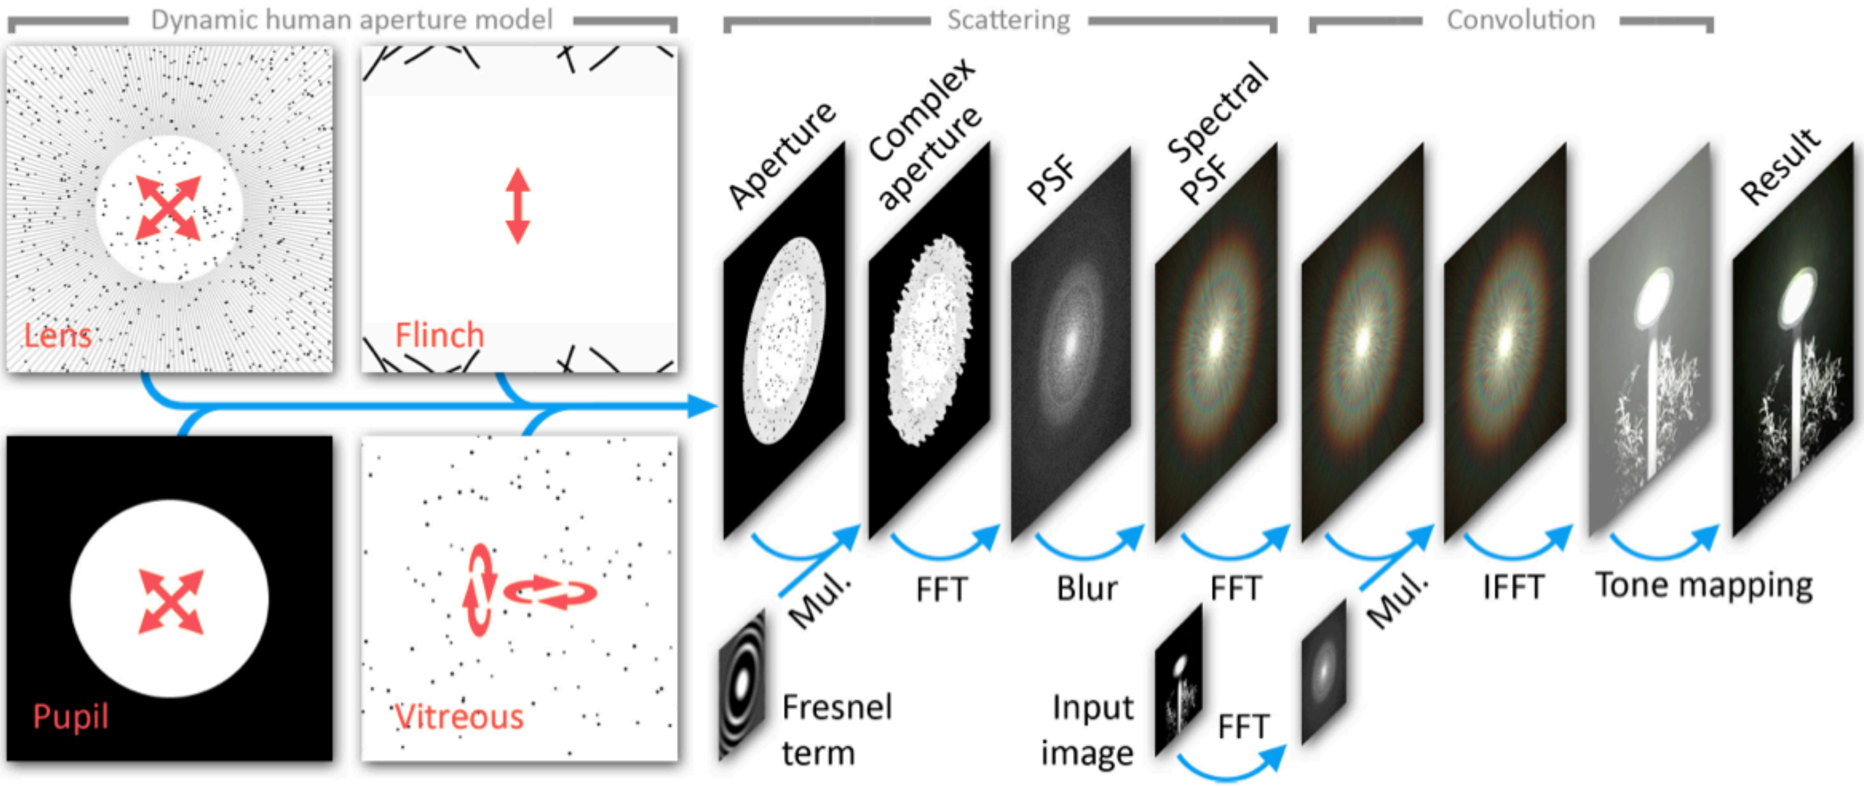

- The model assumes that glare is mostly caused by diffraction and scattering
- Can simulate temporal effects

[From Ritschel et al, Eurographics 2009]

Temporal glare

Point Spread Function of the eye

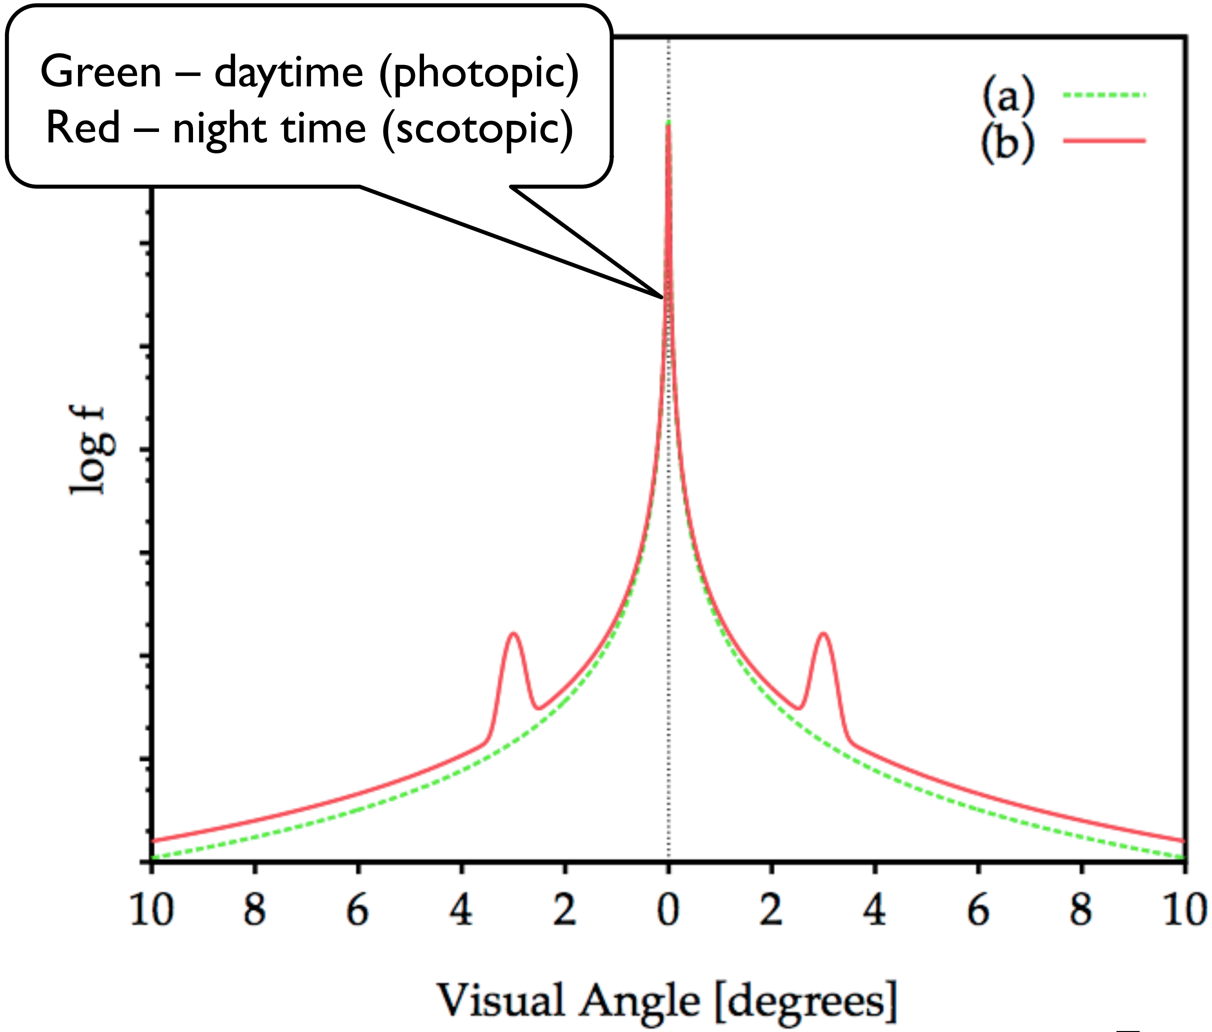

- What portion of the light is scattered towards a certain visual angle
- To simulate:
  - construct a digital filter
  - convolve the image with that filter

From: Spencer, G. et al. 1995. Proc. of SIGGRAPH. (1995)

The problem of double processing

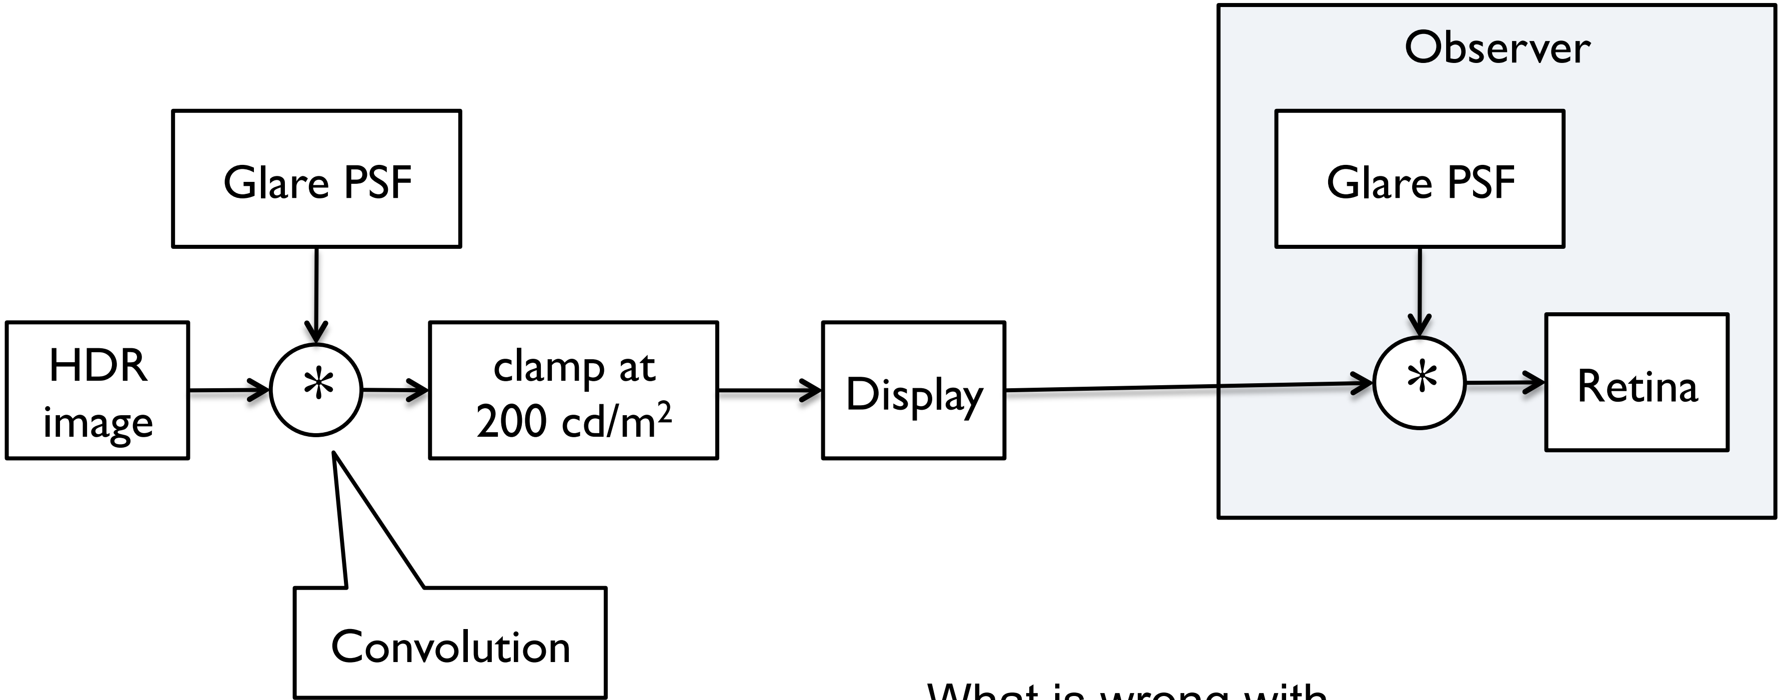

What is wrong with simulating glare this way?

The problem of double processing

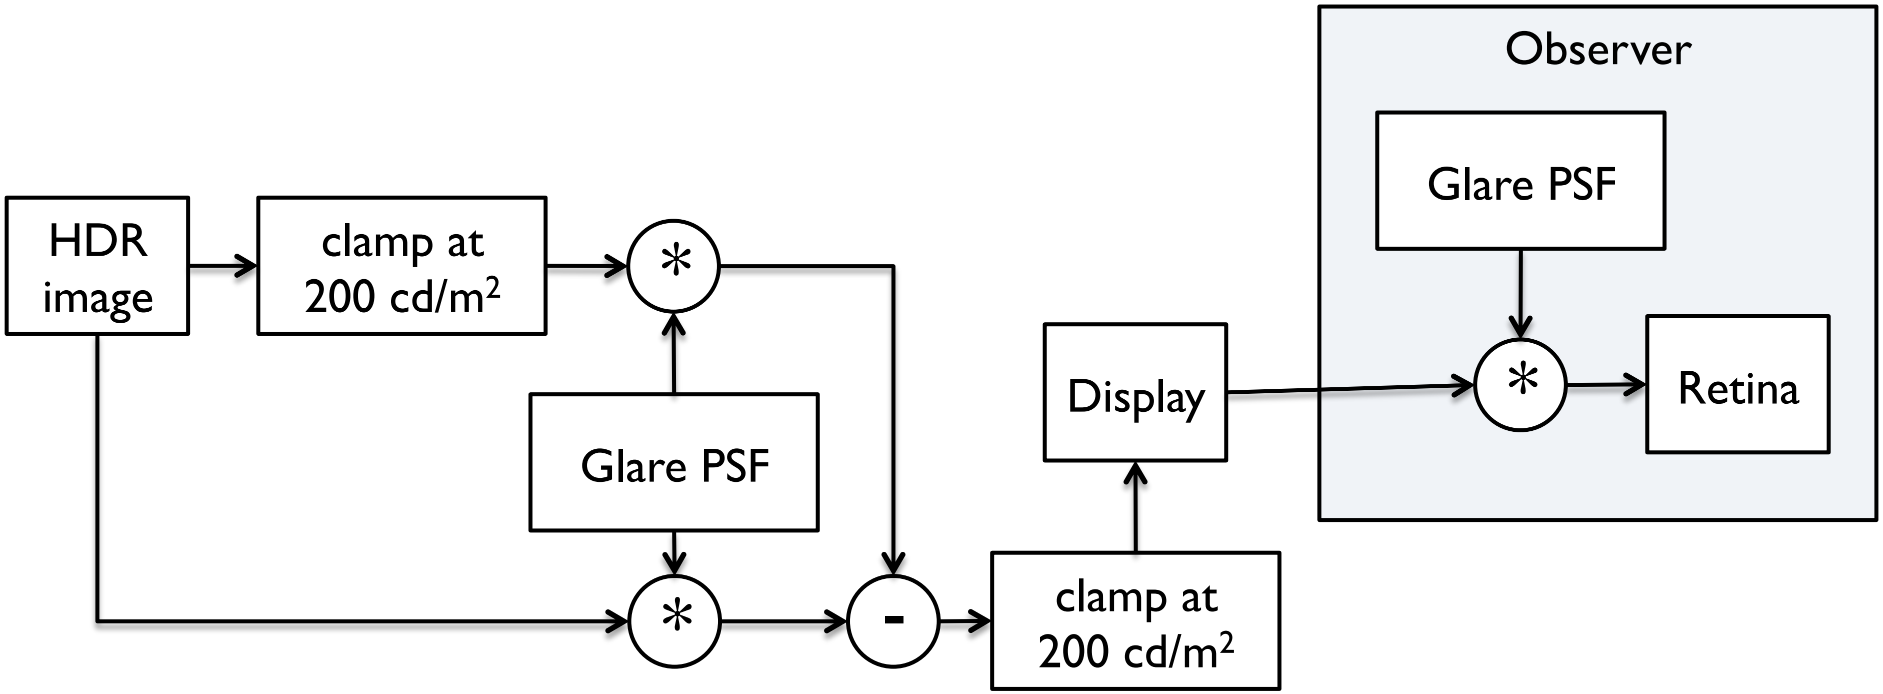

How does the diagram above avoid the problem of double processing?
Write down the operations as equations.
Can the processing be simplified?

PSF vs. OTF (Optical Transfer Function)

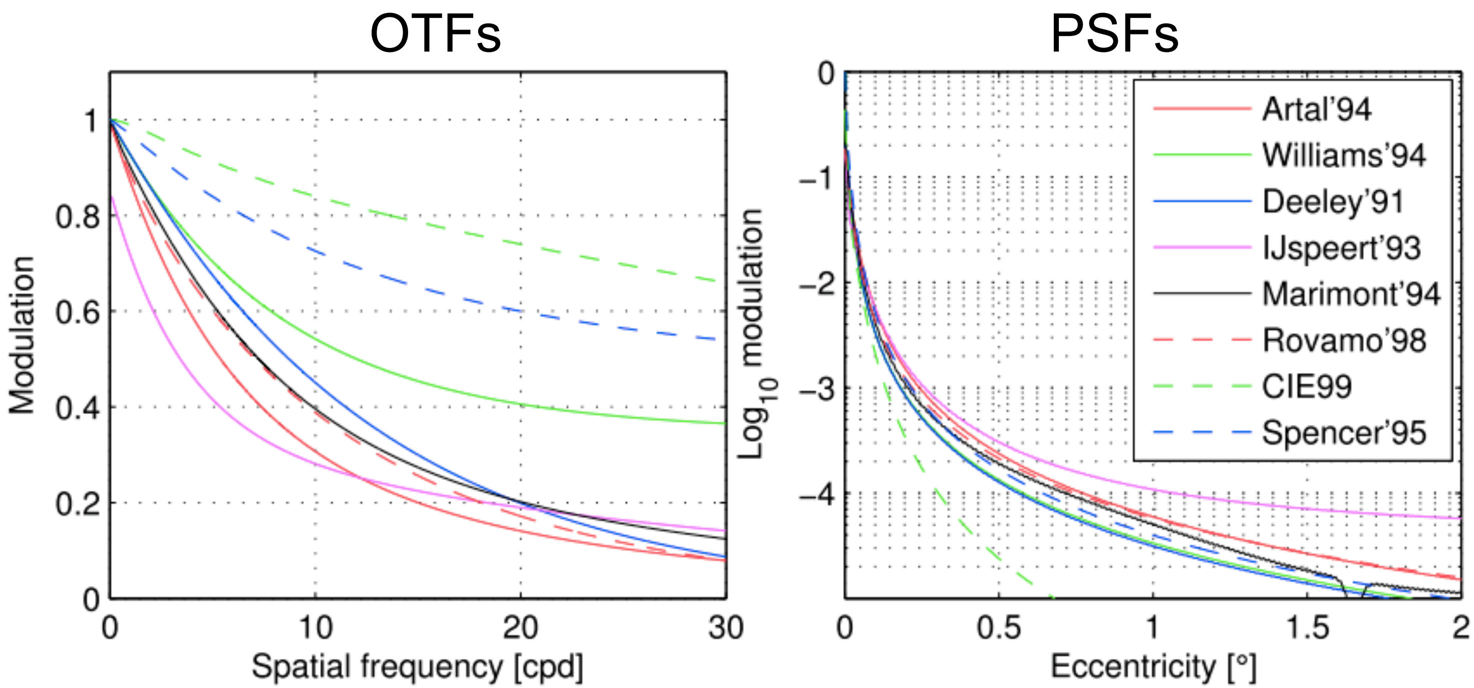

- An OTF is the Fourier transform of a PSF
- Convolution with larger kernels is faster in the Fourier domain

Glare (or bloom) in games

- Convolution with large, non-separable filters is too slow
- The effect is approximated by a combination of Gaussian filters
  - Each filter with different "sigma"
- The effect is meant to look good, not be be accurate model of light scattering
- Some games simulate camera rather than the eye

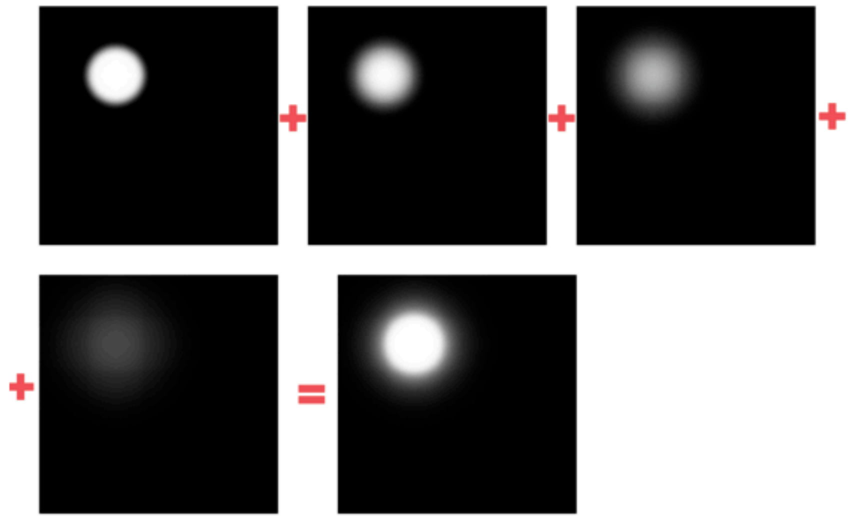

Does the exact shape of the PSF matter?

- The illusion of increased brightness works even if the PSF is very different from the PSF of the eye

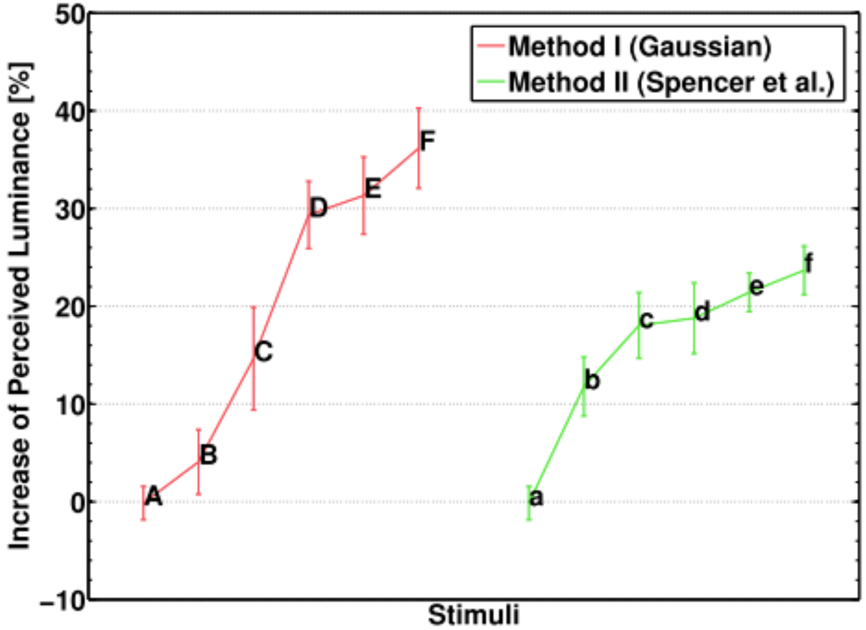

red - Gaussian

green - accurate

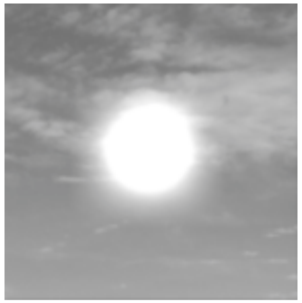

[Yoshida et al., APGV 2008]

Dynamic Range

Luminance

$$\frac{\max L}{\min L}$$

(for SNR>3)

Dynamic range (contrast)

- As ratio:

$$C = \frac{L_{\max}}{L_{\min}}$$

  - Usually written as 1000:1, etc.

- As "orders of magnitude" or log10 units:

$$C_{10} = \log_{10} \frac{L_{\max}}{L_{\min}}$$

- As stops:

$$C_2 = \log_2 \frac{L_{\max}}{L_{\min}}$$

One stop is doubling
of halving the amount of light

High dynamic range (HDR)

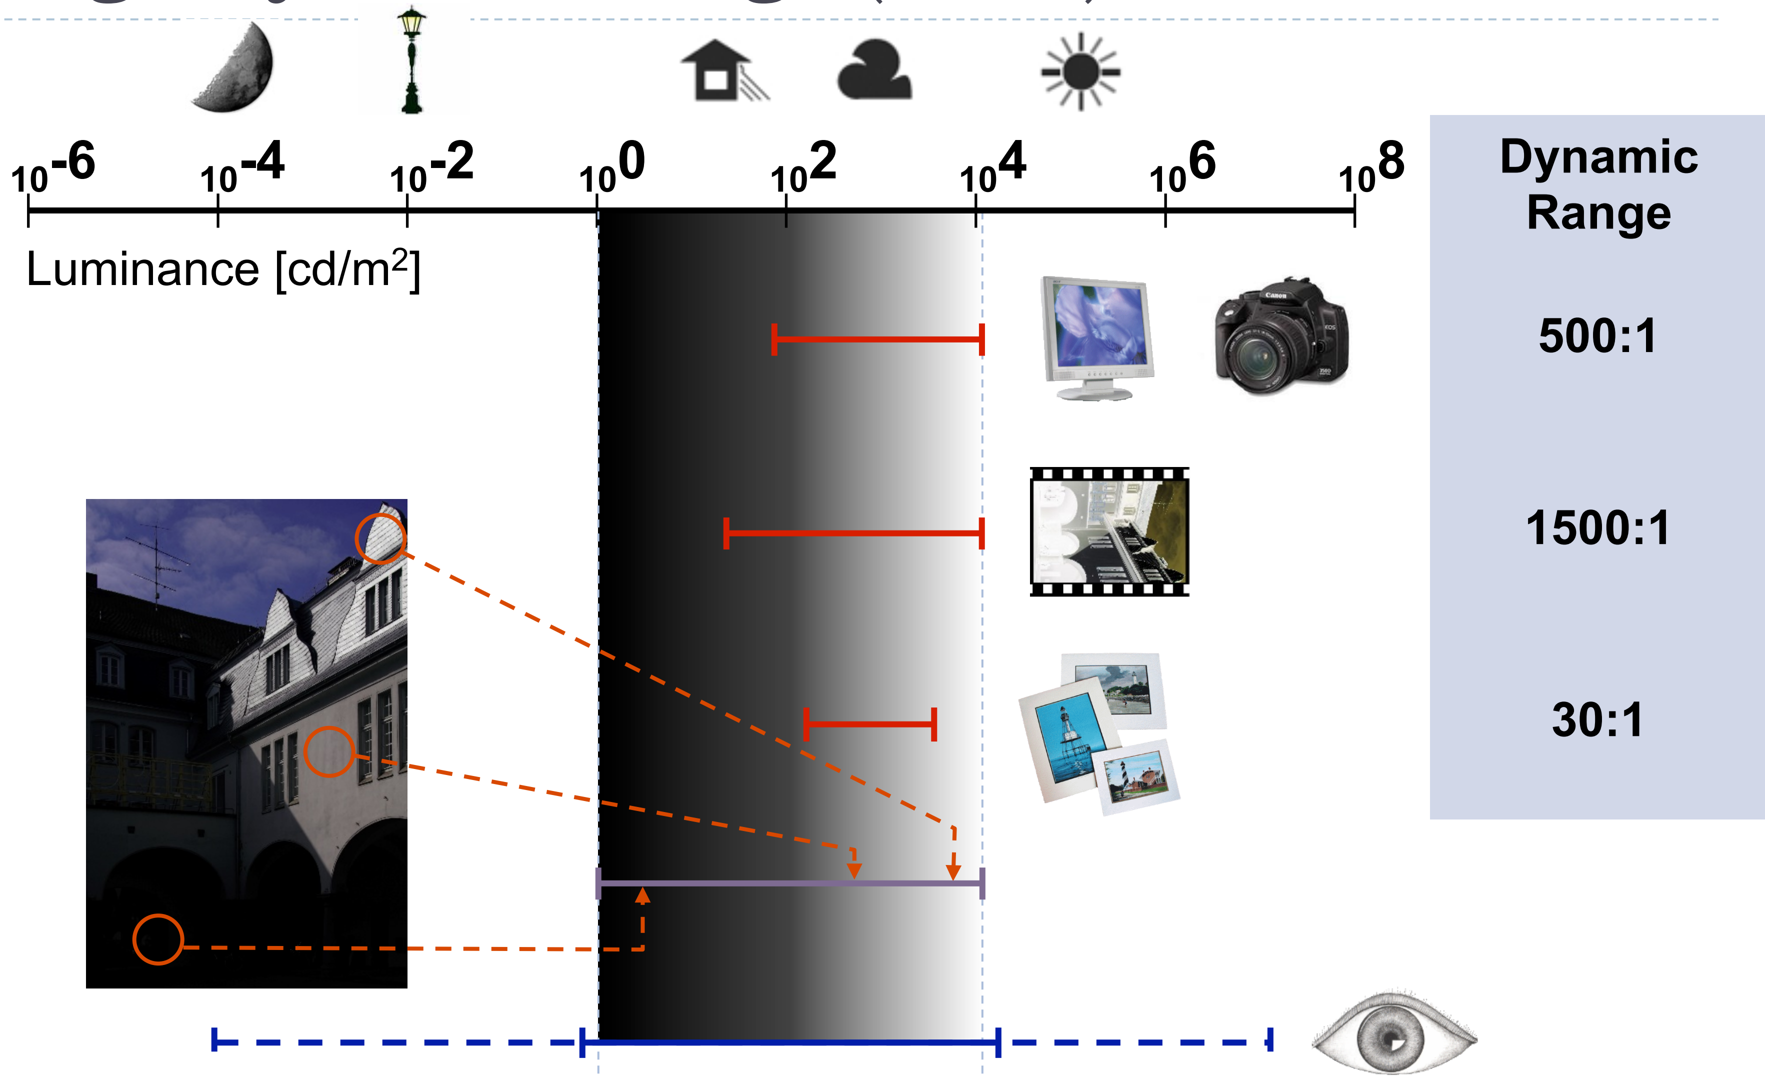

HDR is the „next big thing"

- HDR TV sets
  - LG, Sony, Sharp, Samsung & Panasonic
- HDR UltraHD broadcast
- Technicolor offers HDR color grading services
- Netflix & Amazon anounce HDR content streaming
- HDR experimental short films
- Better pixels instead of more pixels
- HDR is universarily accepted as „better"

technicolor

Visible colour gamut

- The eye can perceive more colours and brightness levels than
  - a display can produce
  - a JPEG file can store
- The premise of HDR:
  - Visual perception and not the technology should define accuracy and the range of colours

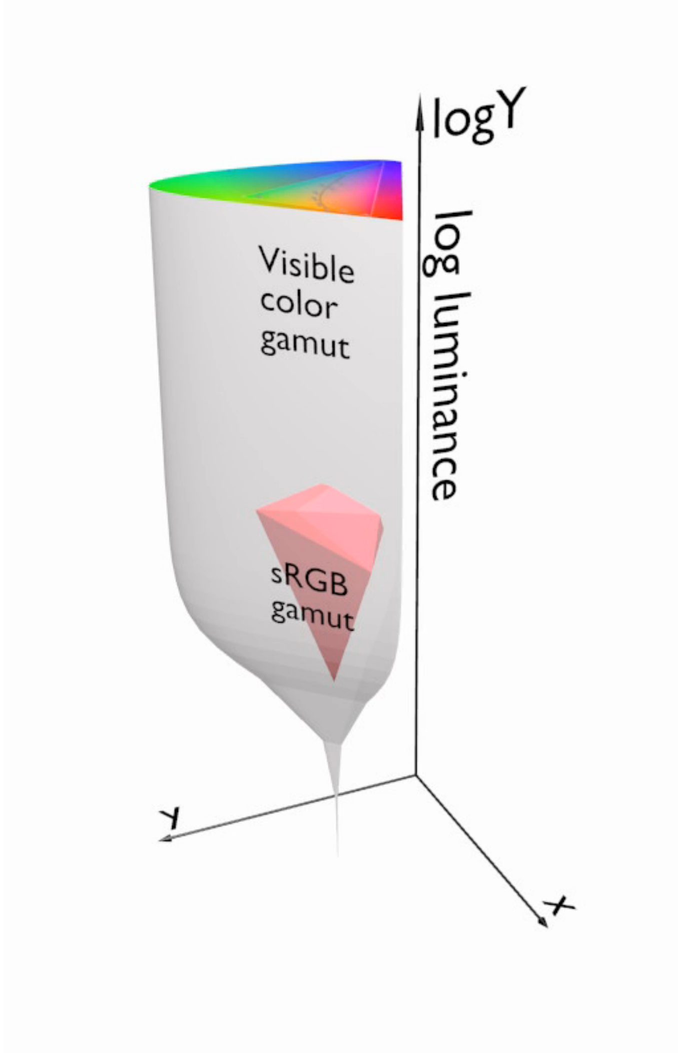

Why is HDR so important for graphics

- Physically accurate rendering
  - Image based lighting
  - Camera glare
  - Motion blur
- Simulation of perceptual phenomena
  - Adaptation
  - Glare in the eye (bloom)
  - Night / day-light vision

Image based lighting (IBL)

1. Capture an HDR image of a light probe

2. Create an illumination (cube) map

3. Use the illumination map as a source of light in the scene

The scene is surrounded by a cube map

Source: Image-based lighting, Paul Debevec, HDR Symposium 2009


Nick Bertke


C17
Nick Bertke
pogo767@gmail.com


Nick Bertke
pogo767@gmail.com

HDR rendering – motion blur

Tone-mapping problem

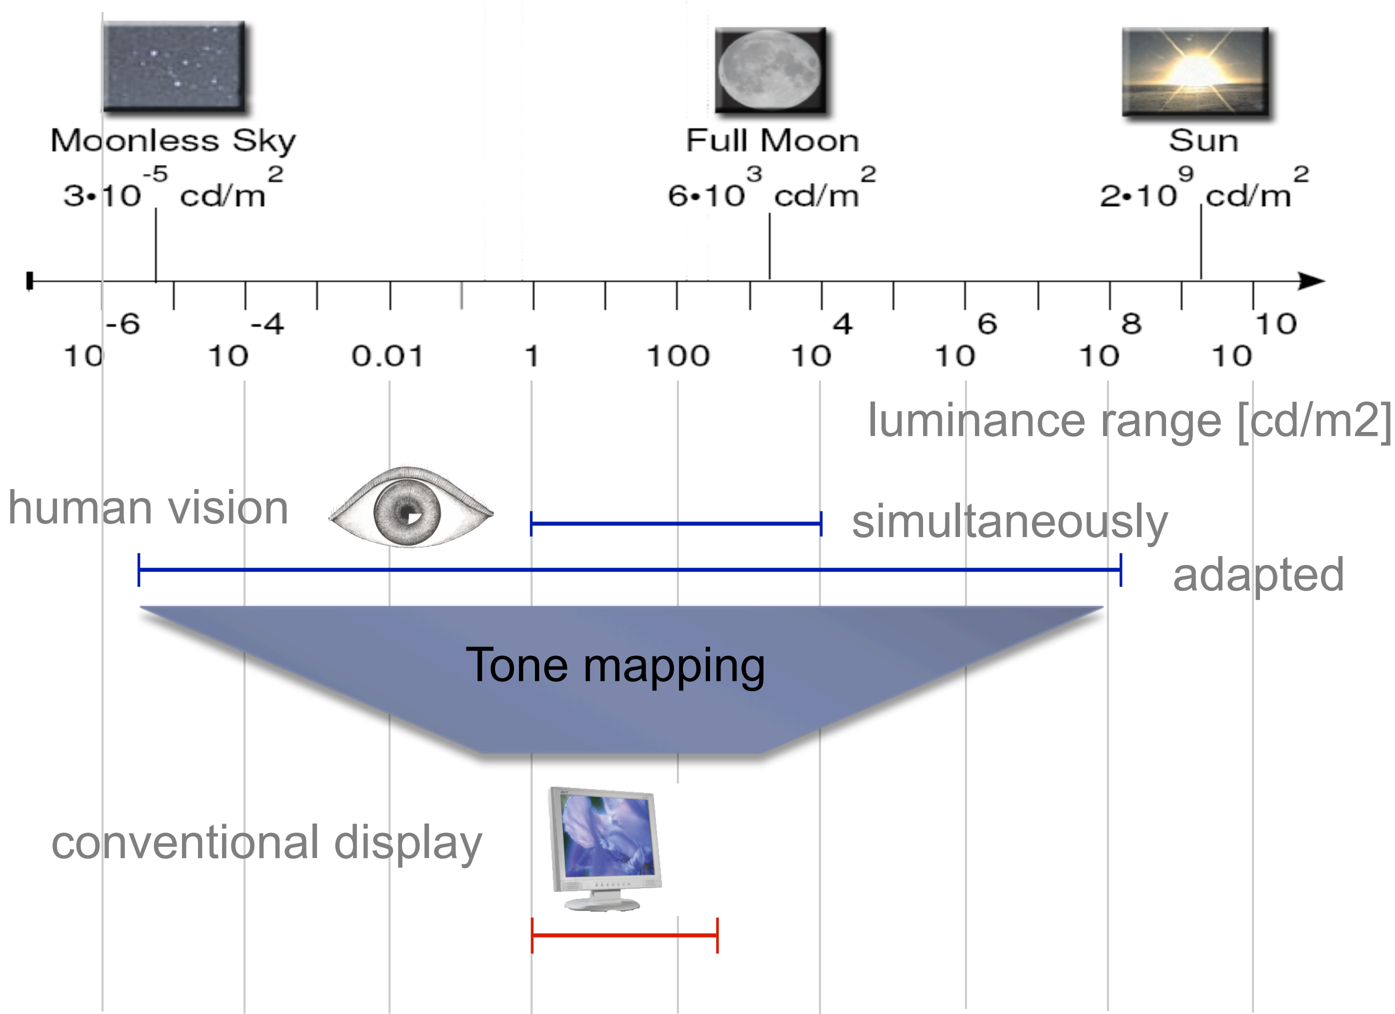

Three intents of tone-mapping

1. Scene reproduction operator
2. Visual system simulator
3. Best subjective quality

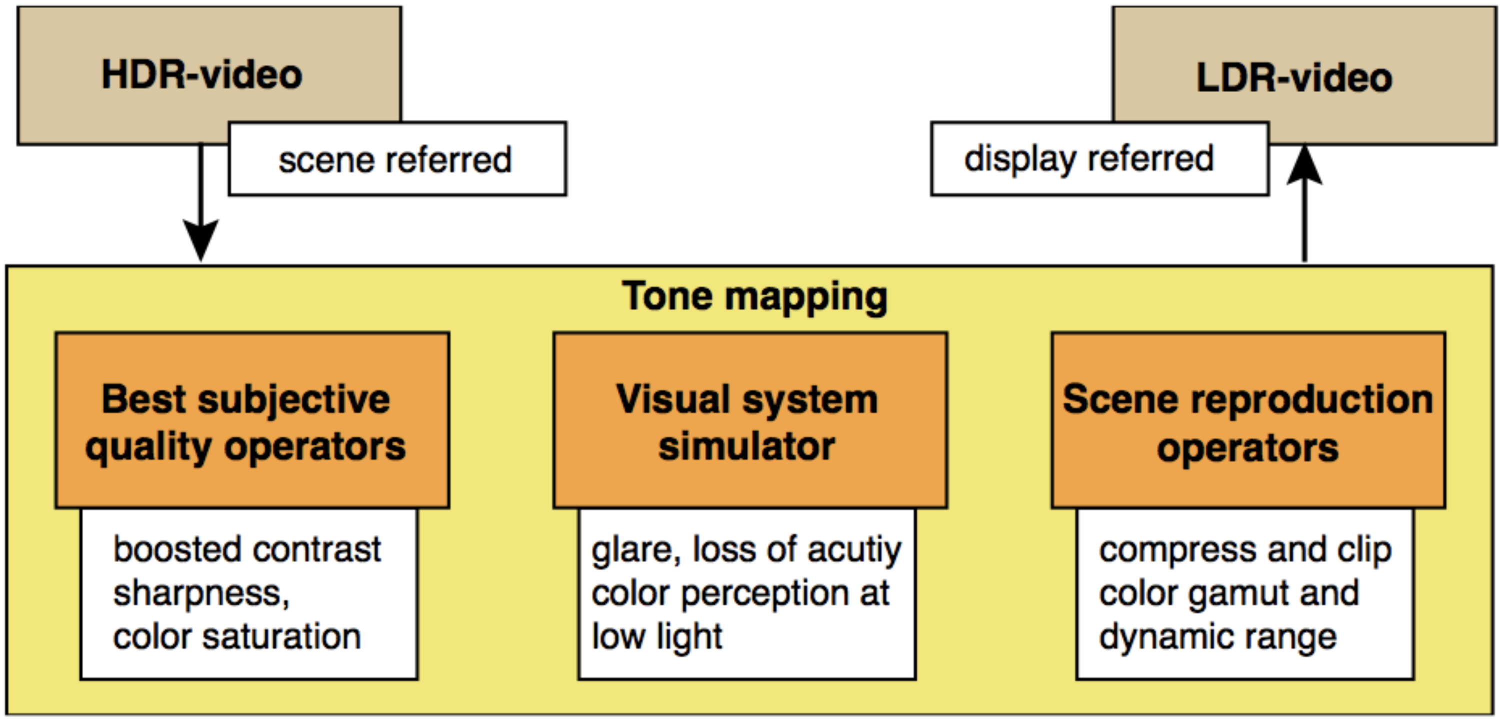

Intent #1: Best subjective quality

- Tools
  - Photoshop
  - Lightroom
  - Photomatix
- Techniques
  - Color-grading
- Often artistic intent

Intent #2: Visual system simulator

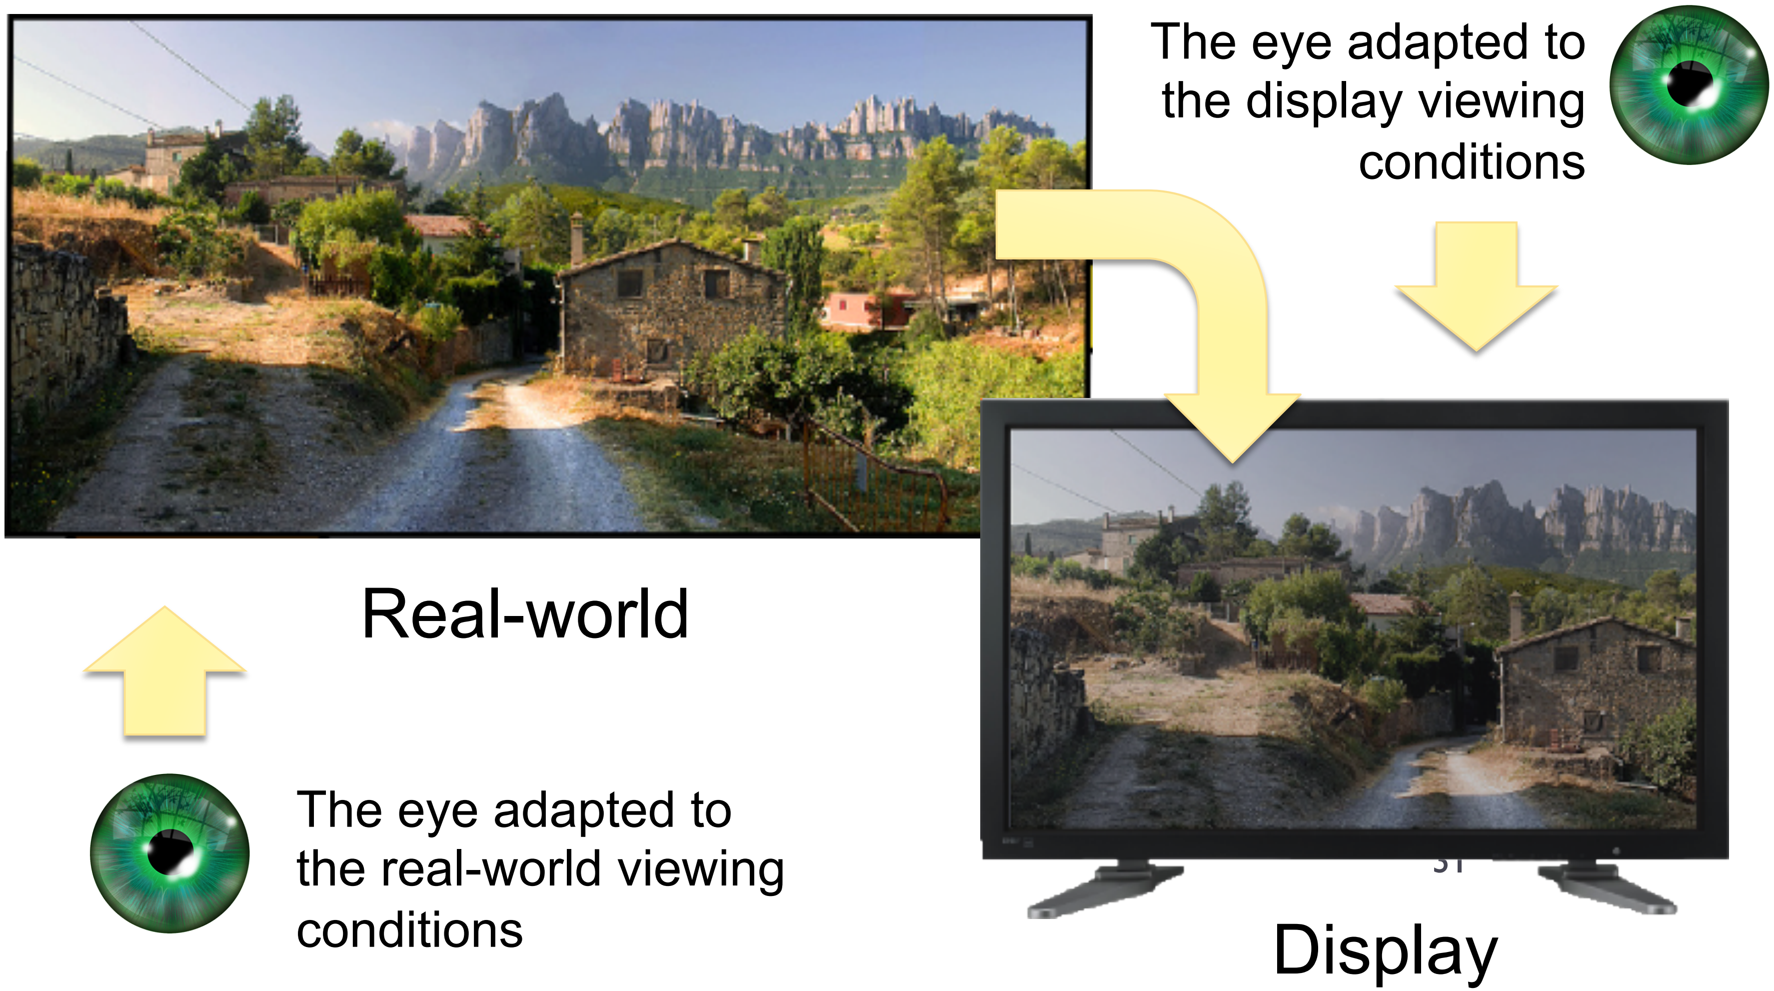

Goal: match color appearance

Visual system simulator - example

- Simulation of glare

Age-adaptive night vision

Video 4

Rivoli

Simulation of age-adaptive night vision

Intent #3: Scene reproduction problem

Real-world

Display

Goal: map colors to a restricted color space

Mapping problem

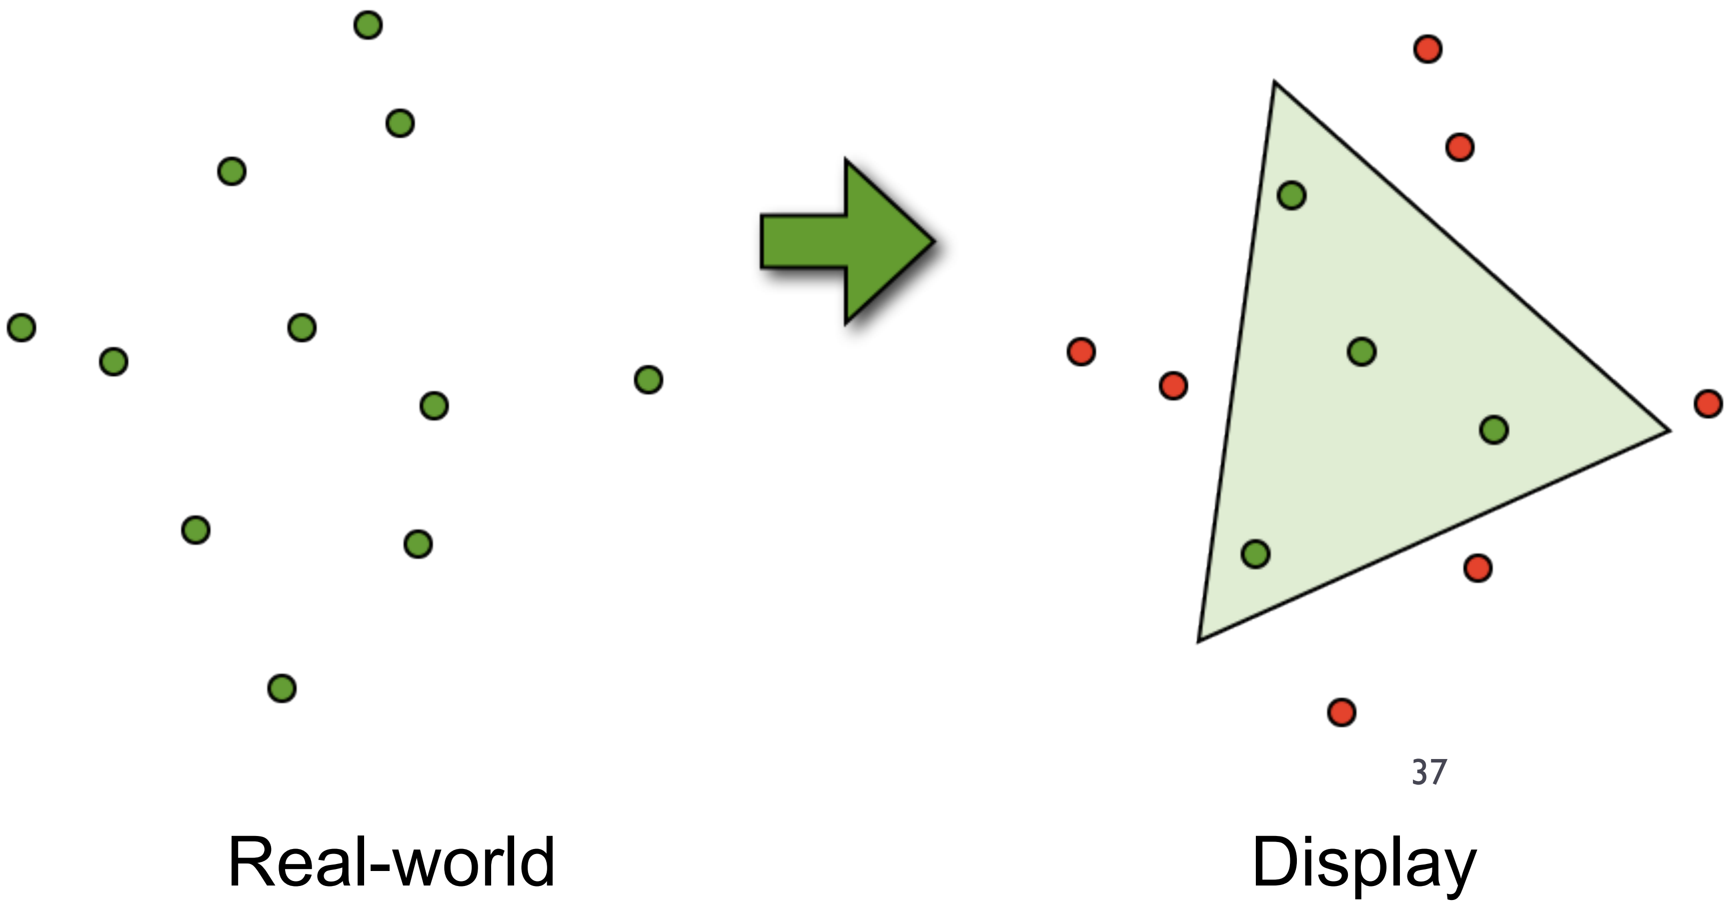

Luminance

- Luminance – perceived brightness of light, adjusted for the sensitivity of the visual system to wavelengths

Luminance

$$L_V = \int_0^\infty L(\lambda) \cdot V(\lambda) d\lambda$$

Light spectrum (radiance)

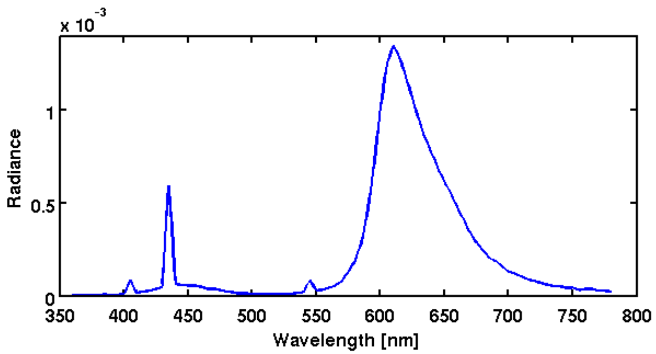

Luminous efficiency function (weighting)

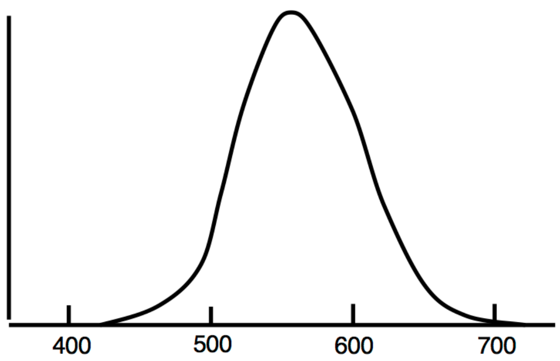

Luminance and Luma

- Luminance
  - Photometric quantity defined by the spectral luminous efficiency function
- $L \approx 0.2126\ R + 0.7152\ G + 0.0722\ B$
- Units: cd/m^2

- Luma
  - Gray-scale value computed from LDR (gamma corrected) image
- $Y = 0.2126\ R' + 0.7152\ G' + 0.0722\ B'$
- Unitless

(Forward) Display model

- GOG: Gain-Gamma-Offset

Luminance

Peak luminance

Gamma

Display black level

Screen reflections

$$L = \underbrace{(L_{peak} - L_{black})}_{\text{Gain}} V^{\gamma} + \underbrace{L_{black} + L_{refl}}_{\text{Offset}}$$

Pixel value 0-1

Reflectance factor (0.01)

$$L_{refl} = \frac{k}{\pi} E_{amb}$$

Ambient illumination (in lux)

Inverse display model

Symbols are the same as for the forward display model

$$V = \left(\frac{L - L_{black} - L_{refl}}{L_{peak} - L_{black}} \right)^{(1/\gamma)}$$

Note: This display model does not address any colour issues. The same equation is applied to red, green and blue color channels. The assumption is that the display primaries are the same as for the sRGB color space.

Two ways to do tone-mapping

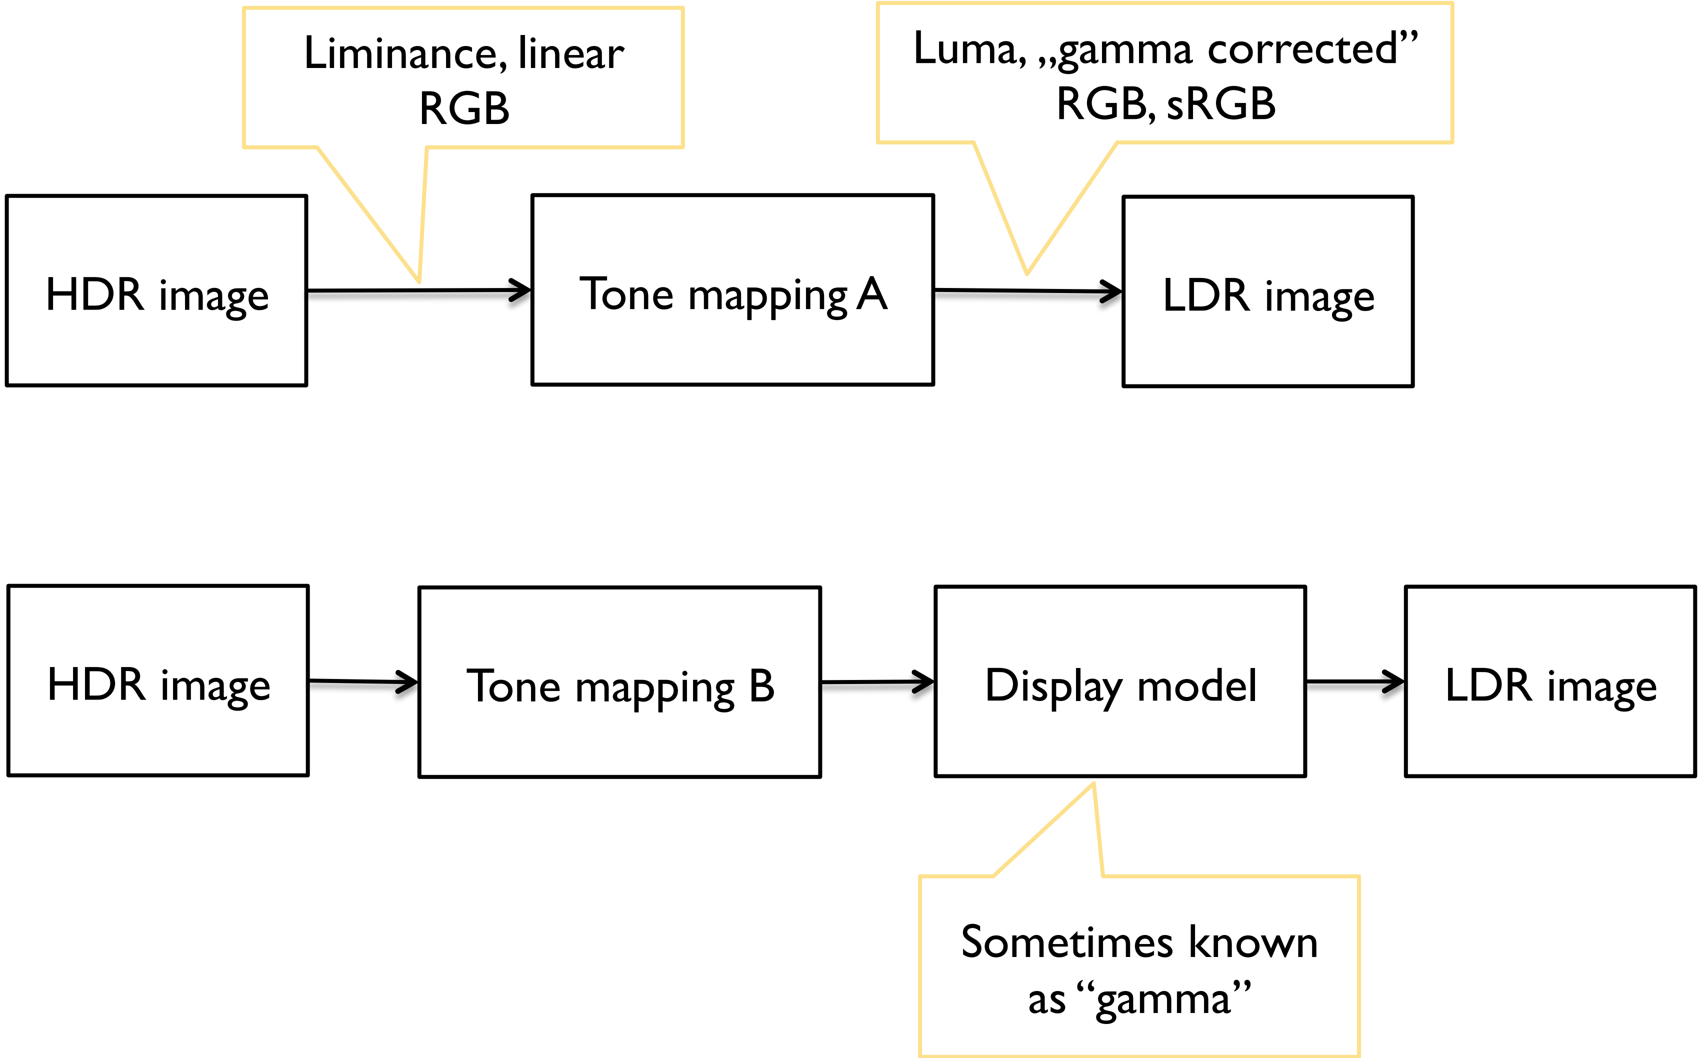

Arithmetic of HDR images

- How does the basic arithmetic operations
  - Addition
  - Multiplication
  - Power function
- affect the appearance of an HDR image
- We work in the luminance space (NOT luma)

Multiplication – brightness change

Resulting luminance

Input luminance

$$T(L_p) = B \cdot L_p$$

Brightness change parameter

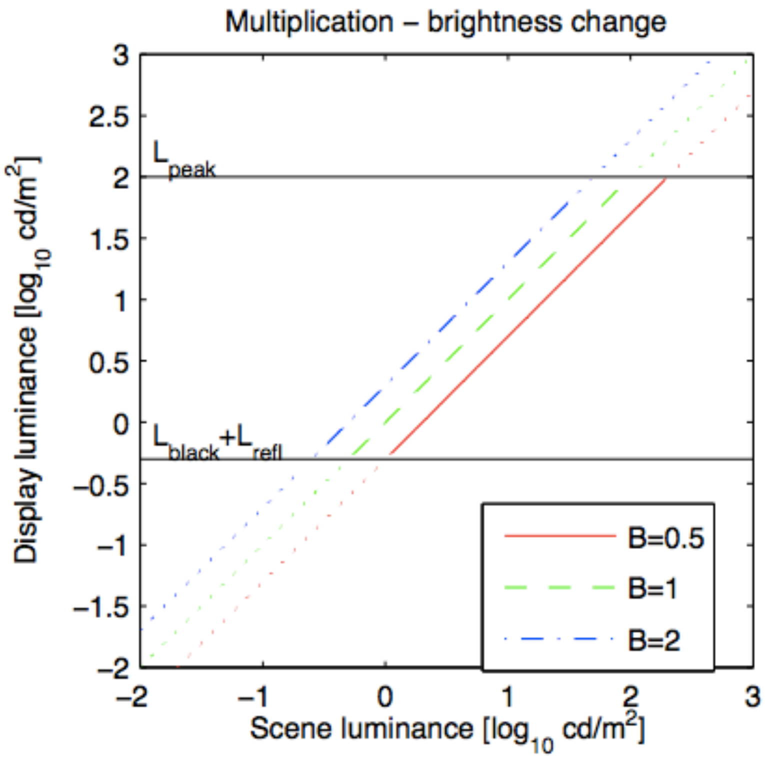

- Multiplication makes the image brighter or darker
- It does not change the dynamic range!

Power function – contrast change

$$T(L_p) = \left(\frac{L_p}{L_{white}}\right)^c$$

Contrast change (gamma)

Luminance of white

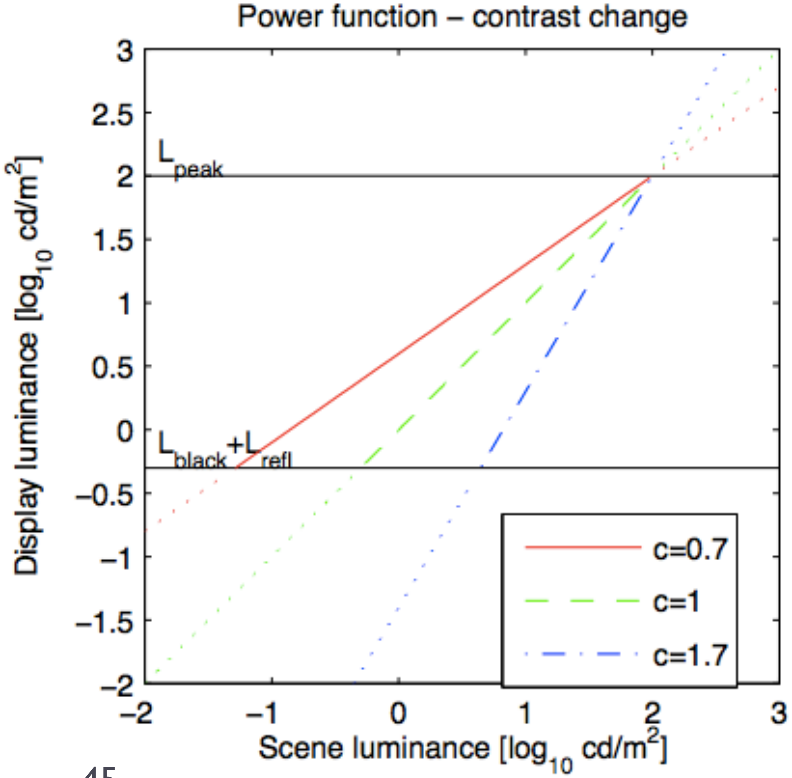

- Power function stretches or shrinks image dynamic range
- It is usually performed relative to reference white
- Apparent brightness changes is the side effect of pushing tones towards or away from the white point

Addition – black level

$$T(L_p) = L_p + F$$

Black level (flare, fog)

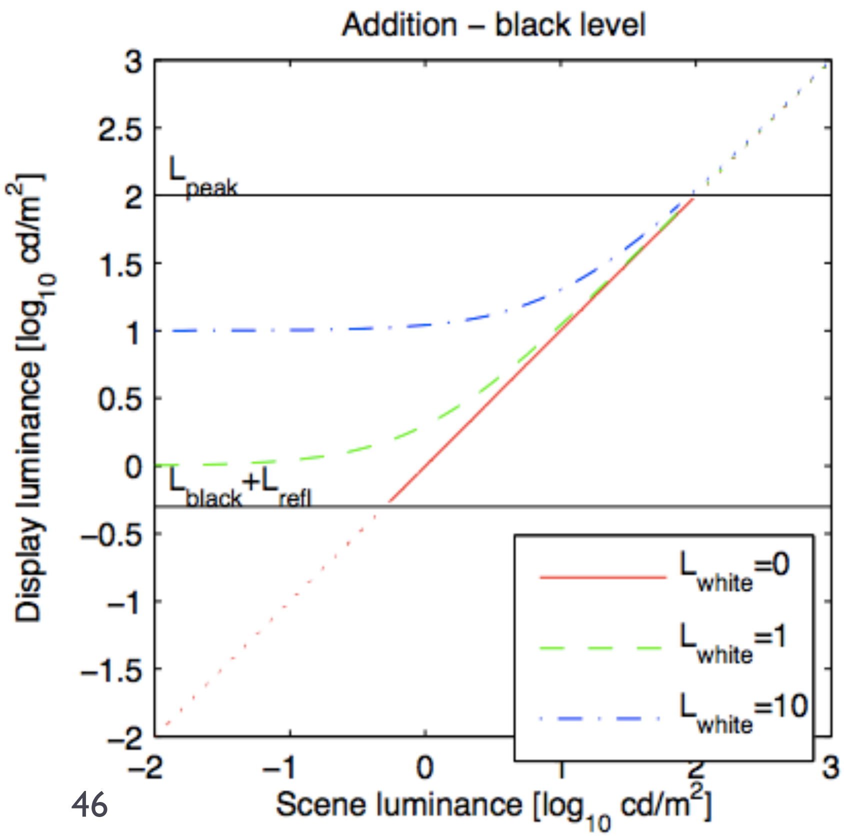

- Addition elevates black level, adds fog to an image
- It does NOT make the overall image brighter
- It reduces dynamic range

Outline

- What is tone-mapping?
- The perception of HDR scenes
- Major approaches to tone-mapping
  - Illumination & reflectance separation
  - Forward visual model
  - Forward & inverse visual model
  - Constraint mapping problem

Outline

- What is tone-mapping?
- The perception of HDR scenes
- Major approaches to tone-mapping
  - Illumination & reflectance separation
  - Forward visual model
  - Forward & inverse visual model
  - Constraint mapping problem

Illumination & reflectance separation

Input

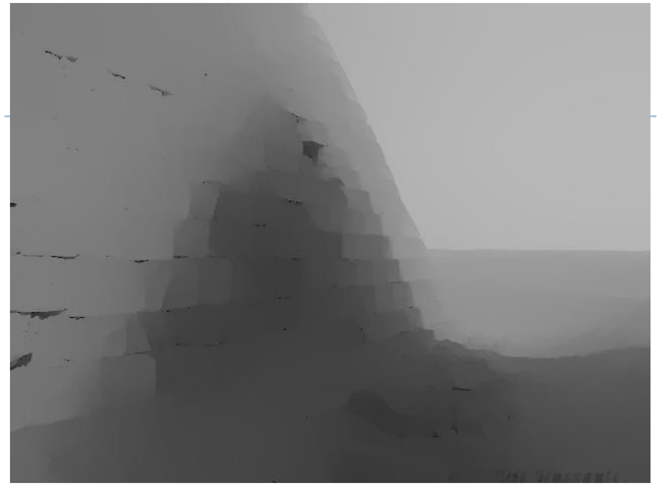

Illumination

Reflectance

Illumination and reflectance

Illumination

- White ≈ 90%
- Black ≈ 3%
- Dynamic range < 100:1
- Reflectance critical for object & shape detection

Reflectance

- Sun ≈ 10^9 cd/m^2
- Lowest perceivable luminance ≈ 10^{-6} cd/m^2
- Dynamic range 10,000:1 or more
- Visual system partially discounts illumination

Reflectance & Illumination TMO

- Hypothesis: *Distortions in reflectance are more apparent than the distortions in illumination*
- Tone mapping could preserve reflectance but compress illumination

Illumination

Tone-mapped image

$$L_d = R \cdot T(I)$$

Reflectance

Tone-mapping

- for example:

$$L_d = R \cdot (I / L_{white})^c \cdot L_{white}$$

How to separate the two?

- (Incoming) illumination – slowly changing
  - except very abrupt transitions on shadow boundaries
- Reflectance – low contrast and
  high frequency variations

Gaussian filter

$$f(x) = \frac{1}{2\pi\sigma_s} e^{\frac{-x^2}{2\sigma_s^2}}$$

- First order approximation

- Blurs sharp boundaries
- Causes halos

Tone mapping result

Bilateral filter

$$I_p \approx \frac{1}{k_s} \sum_{t \in \Omega} f(p-t)\, g(L_p - L_t)\, L_p$$

- Better preserves sharp edges

- Still some blurring on the edges
- Reflectance is not perfectly separated from illumination near edges

Tone mapping result

[Durand & Dorsey, SIGGRAPH 2002]

WLS filter

- Weighted-least-squares optimization

Make reconstructed image *u* possibly close to input *g*

Smooth out the image by making partial derivatives close to 0

$$\sum_p \left((u_p - g_p)^2 + \lambda \left(a_{x,p}(g) \left(\frac{\partial u}{\partial x} \right)_p^2 + a_{y,p}(g) \left(\frac{\partial u}{\partial y} \right)_p^2 \right) \right) \quad \text{-> min}$$

Spatially varying smoothing – less smoothing near the edges

- [Farbman et al., SIGGRAPH 2008]

WLS filter

- Stronger smoothing and still distinct edges

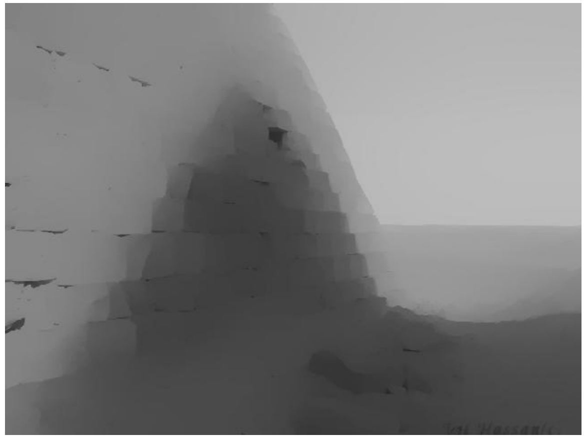

- Can produce stronger effects with fewer artifacts

Tone mapping result

Retinex

- Retinex algorithm was initially intended to separate reflectance from illumination [Land 1964]
  - There are many variations of Retinex, but the general principle is to eliminate from an image small gradients, which are attributed to the illumination

1 step: compute gradients in log domain

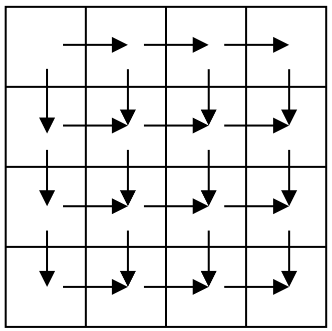

2nd step: set to 0 gradients less than the threshold

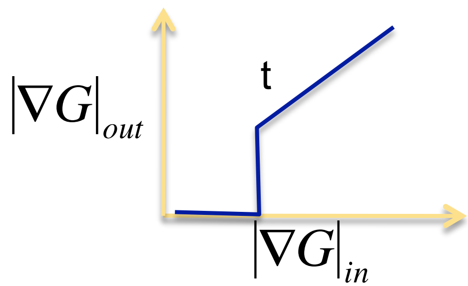

3rd step: reconstruct an image from the vector field

$$\nabla^2 I = \operatorname{div} G$$

For example by solving the Poisson equation

Retinex examples

From: http://dragon.larc.nasa.gov/retinex/757/

From:http://www.ipol.im/pub/algo/lmps_retinex_poisson_equation/#ref_1

Gradient domain HDR compression

[Fattal et al., SIGGRAPH 2002]

- Similarly to Retinex, it operates on log-gradients
- But the function amplifies small contrast instead of removing it
  - Contrast compression achieved by global contrast reduction
    - Enhance reflectance, then compress everything

Retinex

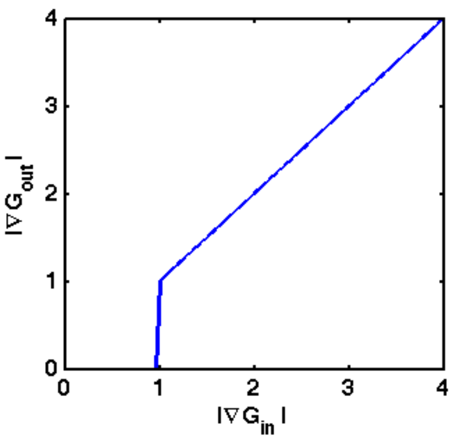

Gradient domain

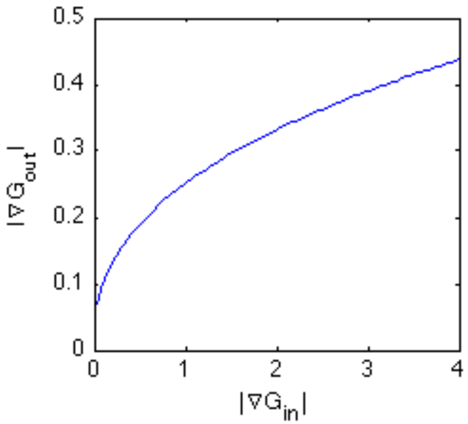

Tone mapping in photography

- Dodging and burning
  - Darken on brighten image parts by occluding photographic paper during exposure
  - Ansel Adams, *The print*, 1995
  - Photoshop tool

- Essentially – attenuate low-pass frequencies that contain scene illumination

Automatic dodging and burning

- Reinhard et al., *Photographic tone reproduction for digital images.* SIGGRAPH 2002
- Choose dodging an burning kernel size adaptively
  - depending on the response of the center-surround filter
  - thus avoid halo artifacts

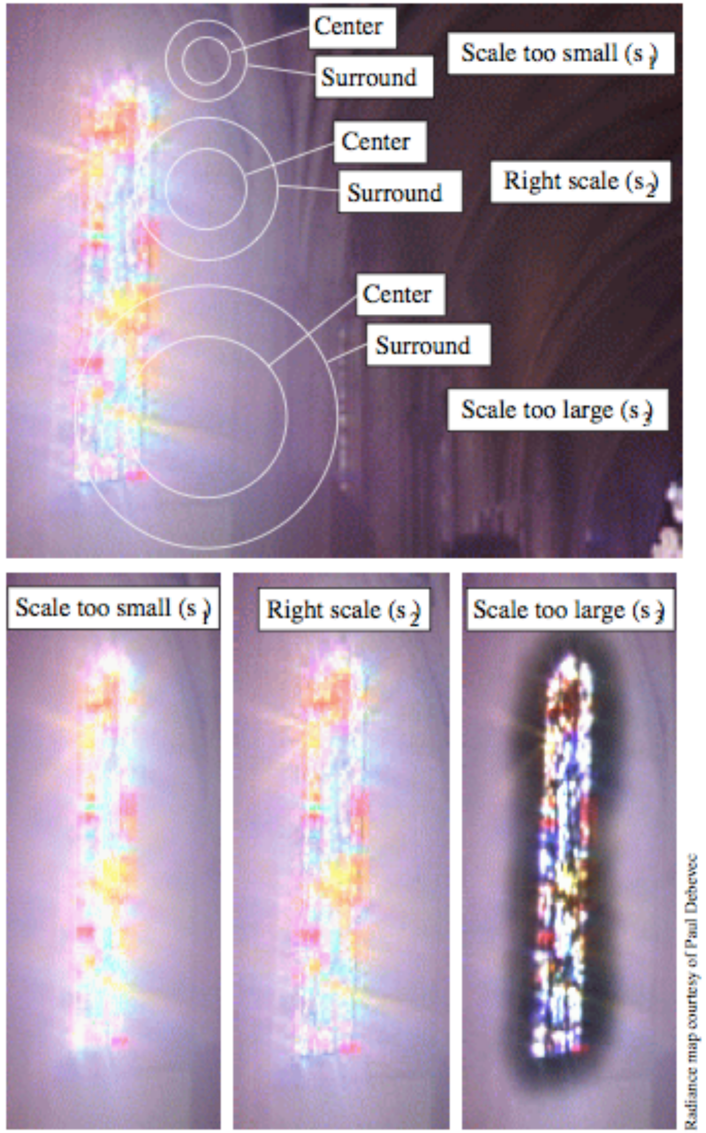

Outline

- What is tone-mapping?
- The perception of HDR scenes
- Major approaches to tone-mapping
  - Illumination & reflectance separation
  - Forward visual model
  - Forward & inverse visual model
  - Constraint mapping problem

Forward visual model

- Mimic the processing in the human visual system

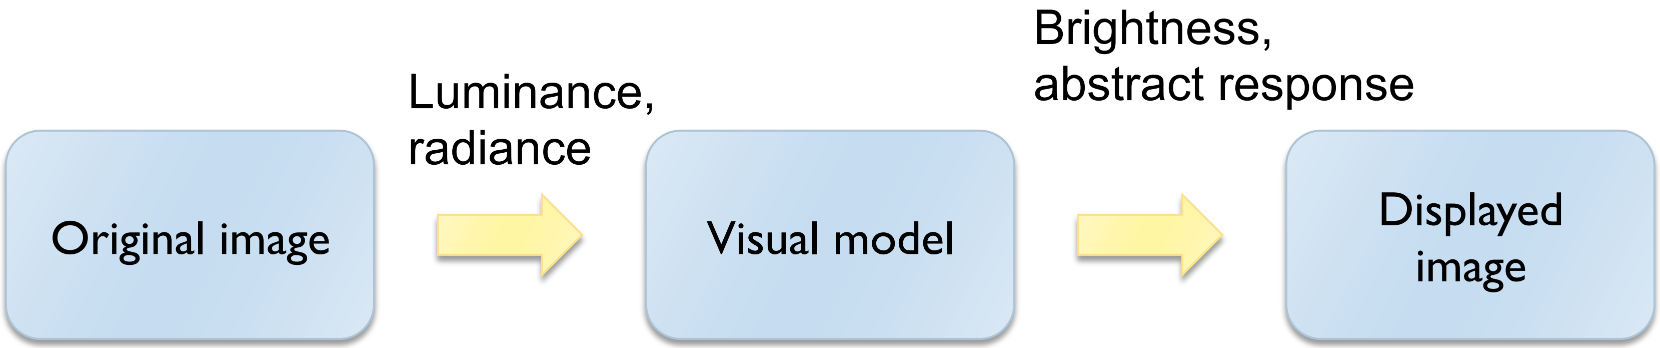

- Assumption: what is displayed is brightness or abstract response of the visual system
- Problem: double processing

Forward visual model: Retinex

- Remove illumination component from an image
  - Because the visual system also discounts illuminant
- Display 'reflectance' image on the screen

- Assumption:
  - The abstract 'reflectance' contains most important visual information
  - Illumination is a distraction for object recognition and scene understanding

Photoreceptor response

- Dynamic range reduction inspired by photoreceptor physiology
  - [Reinhard & Devlin ‘05]

$$V = \frac{I}{I + \sigma(I_a)} V_{max}$$

$$\sigma(I_a) = (f I_a)^m.$$

- From gamma to sigmoidal response:

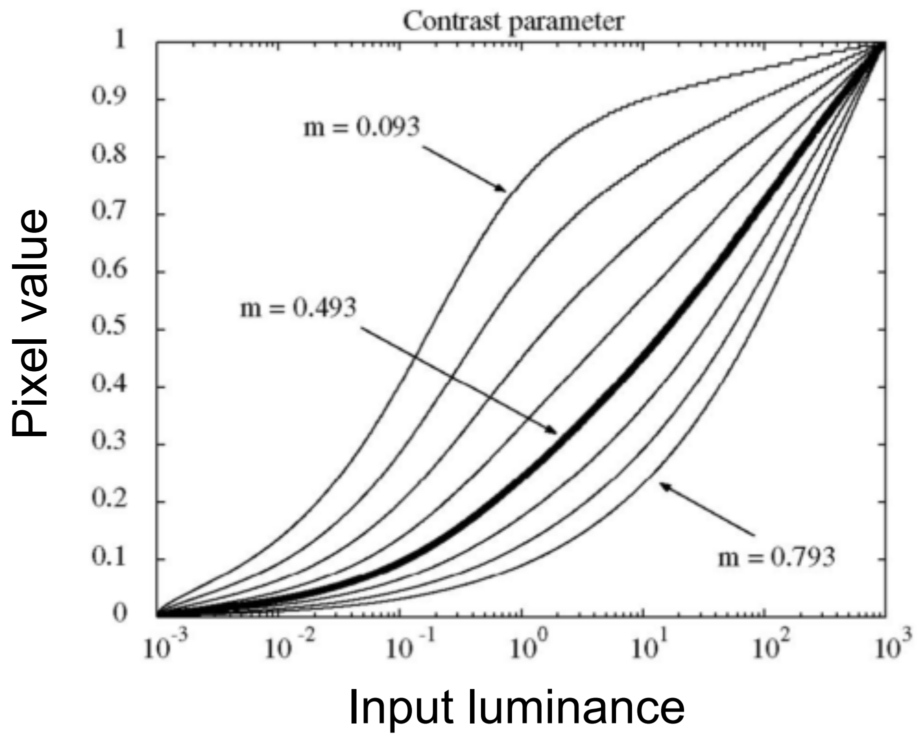

Task 7: Photoreceptor TMO

- Implement photoreceptor TMO

Output pixel value (RGB)

Input radiance (RGB)

Maximum response (set it to 1)

$$V = \frac{I}{I + \sigma(I_a)} V_{max}$$

$$\sigma(I_a) = (f I_a)^m$$

Parameter (set it to 1 and experiment with different values)

Semi-saturation constant

Global/local adaptation luminance (make it local and equal geoman(I))

Parameter (between 0 and 1)

Results: photoreceptor TMO

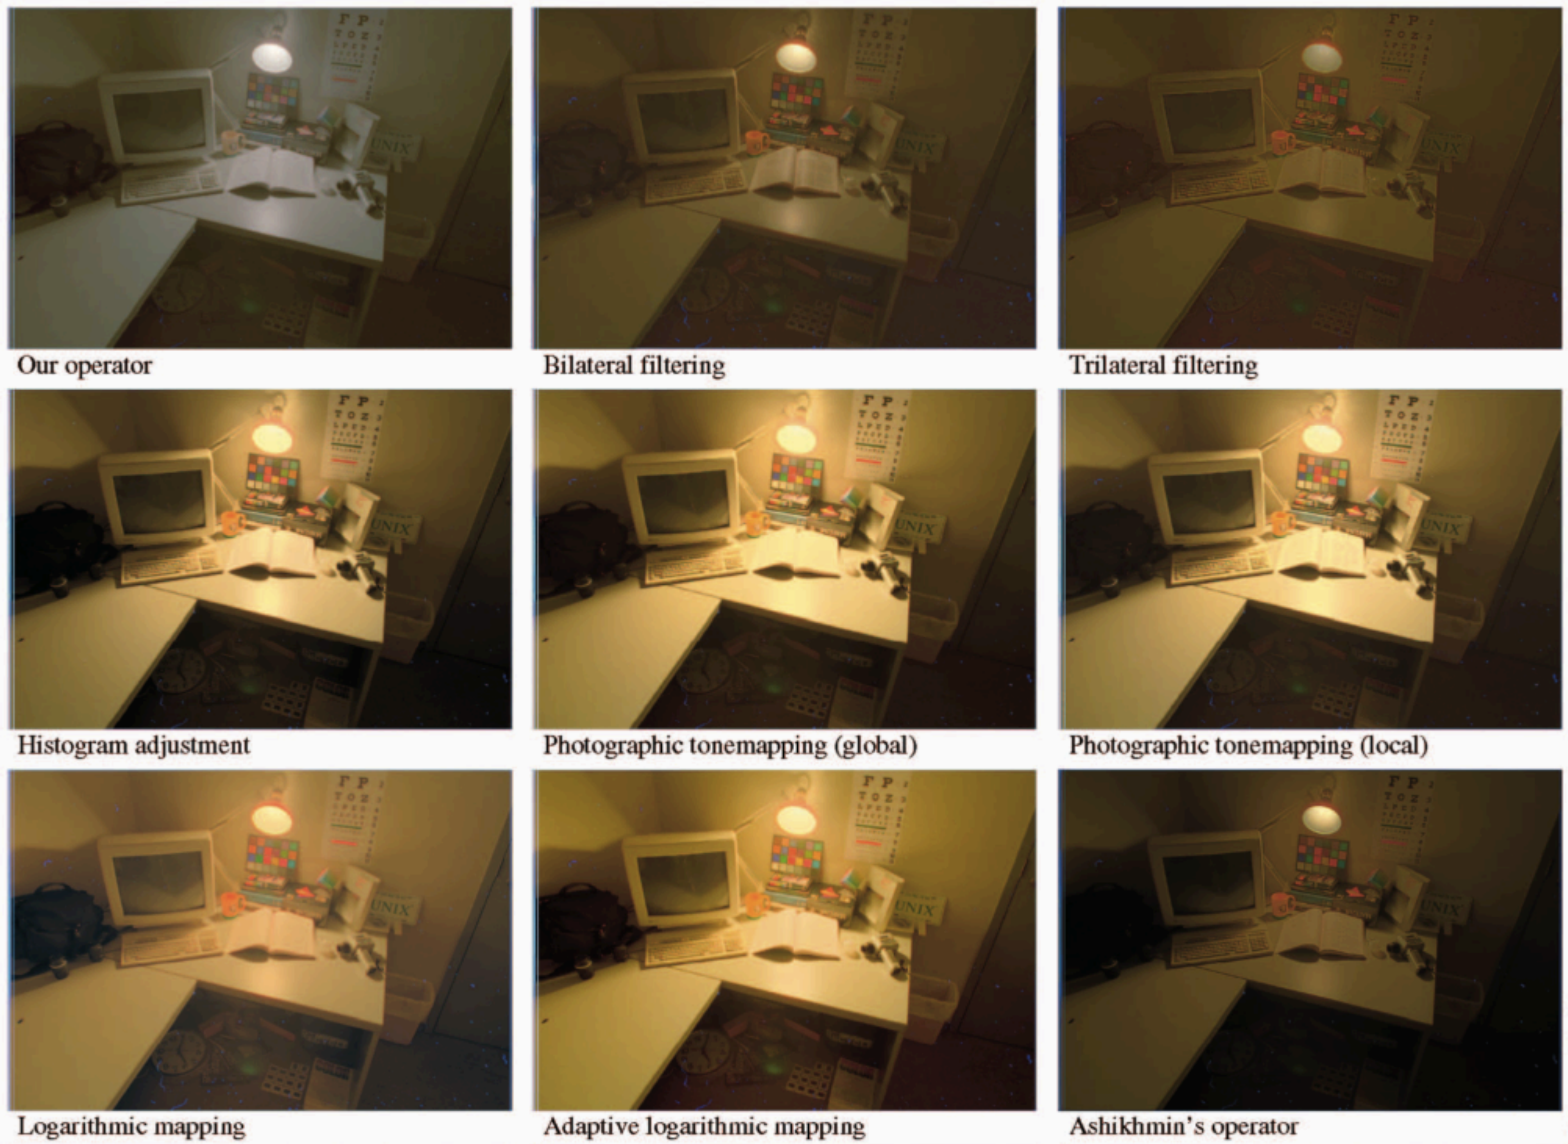

Photoreceptor models

- Naka-Rushton equation:

$$\frac{R}{R_{max}} = \frac{Y^n}{Y^n + \sigma^n}$$

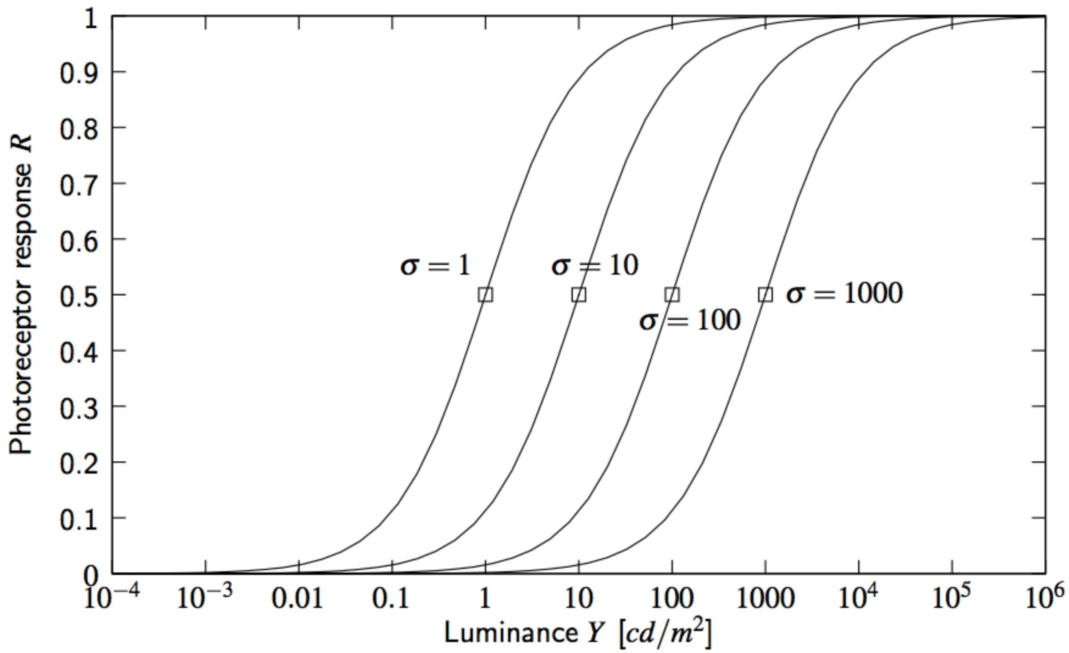

Experiment:

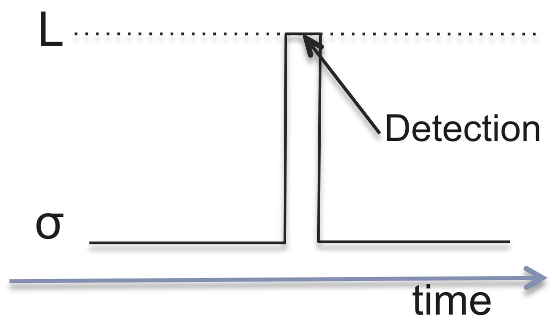

- Response of the photoreceptor to a short flicker of light - less applicable to viewing static images

Sigmoidal tone-curves

- Very common in digital cameras
  - Mimic the response of analog film
  - Analog film has been engineered for many years to produce optimum tone-reproduction (given that he tone curve must not change)

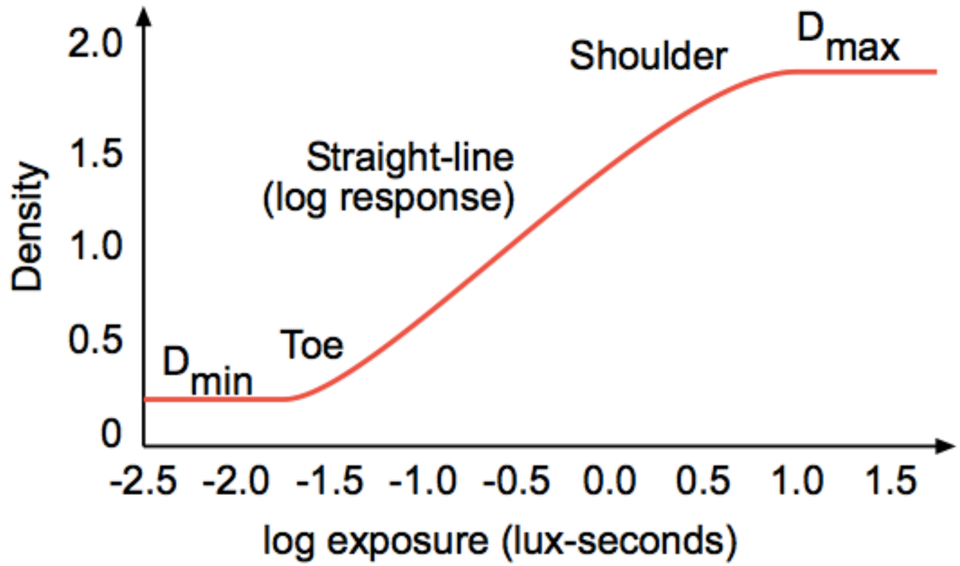

- In practice - the most commonly used tone-mapping!

[Speculative] Why do sigmoidal tone-curves work?

- Because they mimic photoreceptor response
  - Unlikely, because photoreceptor response to steady light is not sigmoidal
- Because they preserve contrast in mid-tones, which usually contains skin color
  - We are very sensitive to variation in skin color
- Because an image on average has Gaussian distribution of log-luminance
  - S-shape function is the result of histogram equalization of an image with a Gaussian-shape histogram

Outline

- What is tone-mapping?
- The perception of HDR scenes
- Major approaches to tone-mapping
  - Illumination & reflectance separation
  - Forward visual model
  - Forward & inverse visual model
  - Constraint mapping problem

Forward and inverse visual model

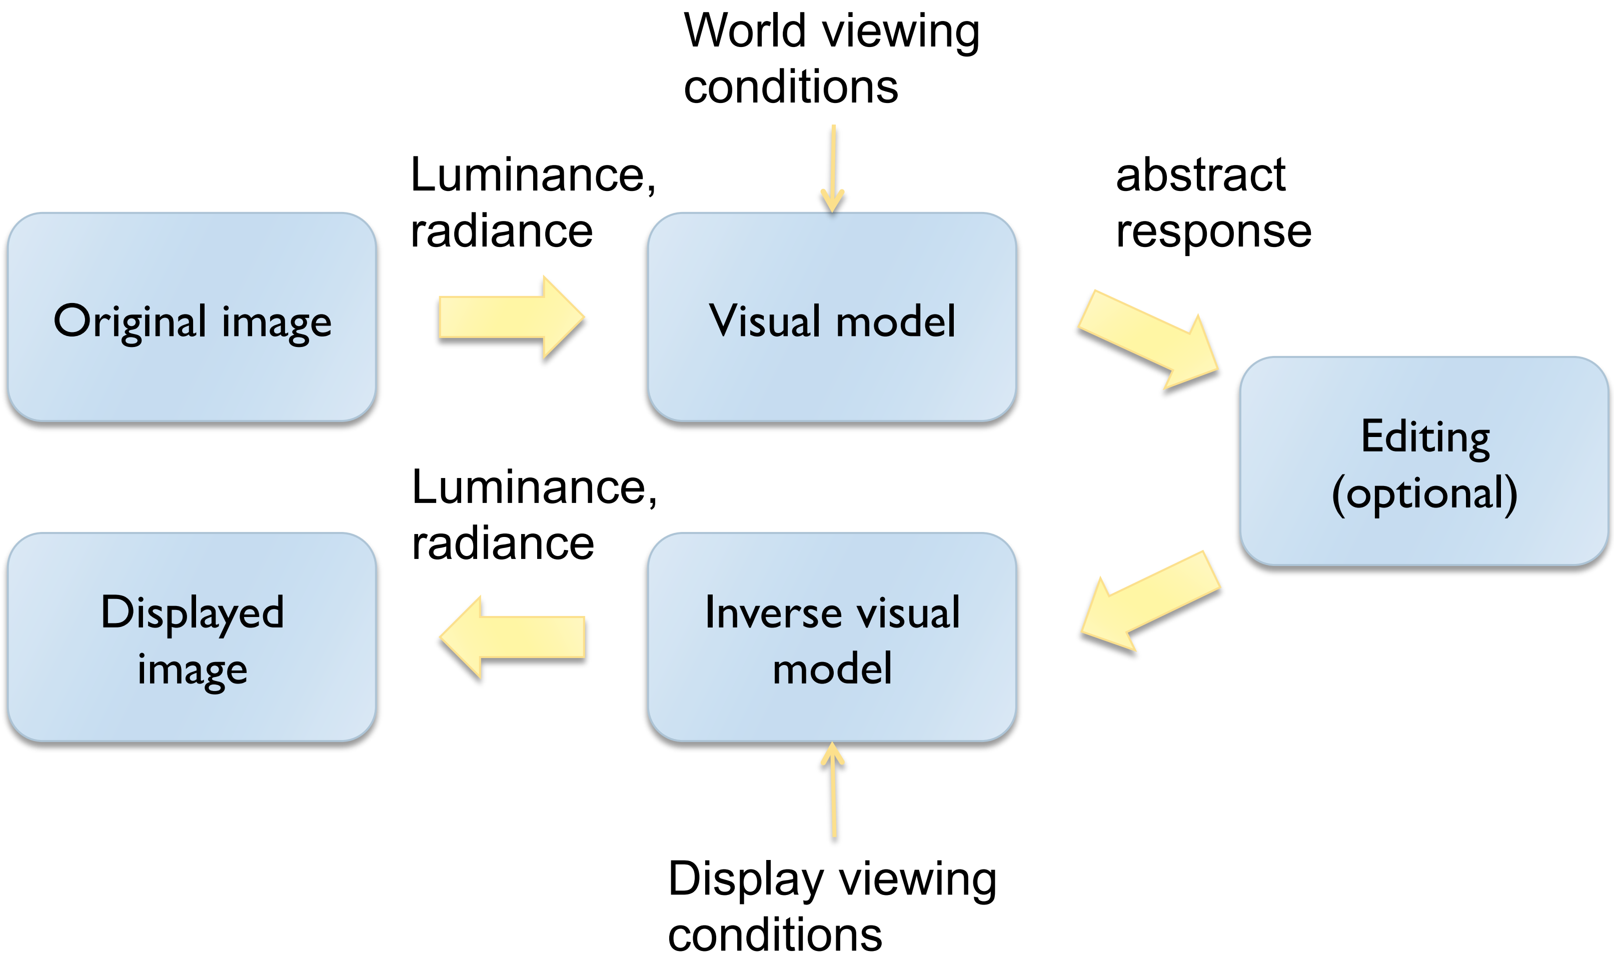

Multi-scale model

- Multi-scale model of adaptation and spatial vision and color appearance
  - [Pattanaik et al. '98]
- Combines
  - psychophysical threshold and superthreshold visual models
  - light & dark adaptation models
  - Hunt's color appearance model
- One of the most sophisticated visual models

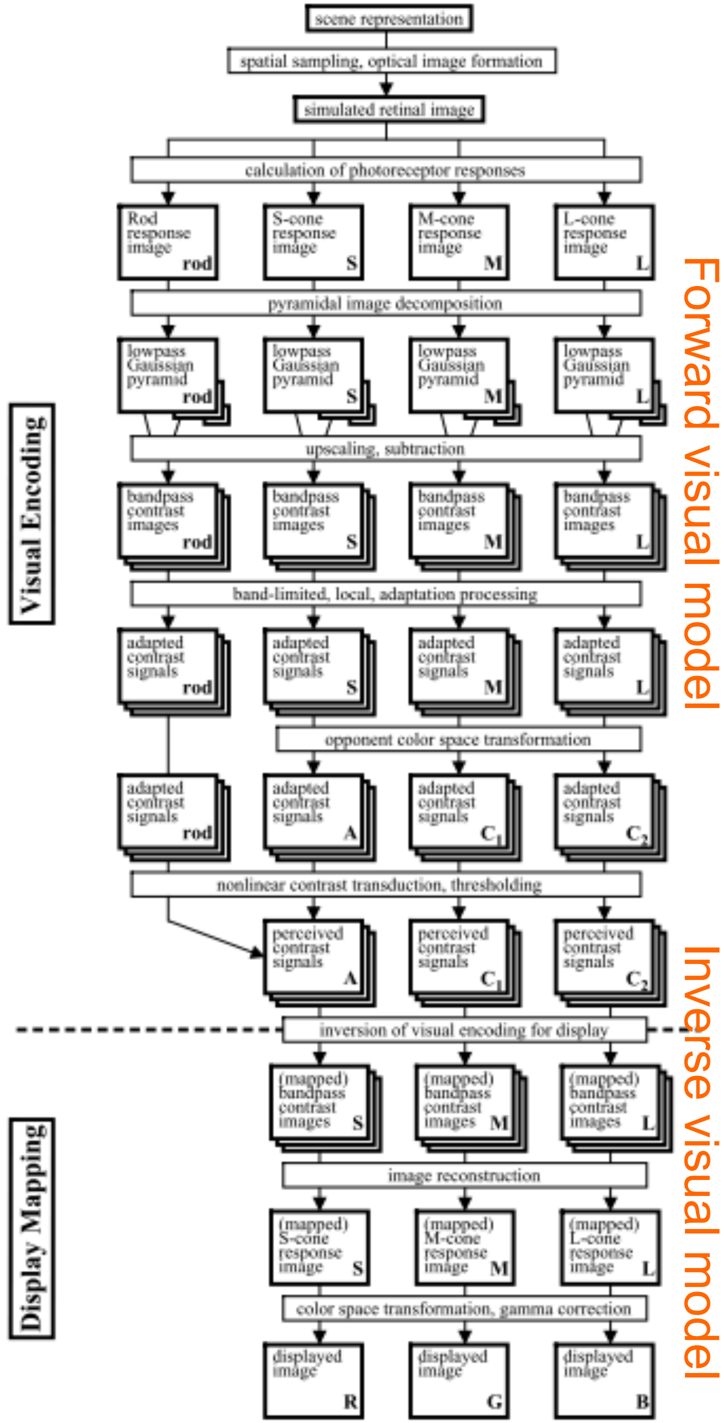

Results: multi-scale model …

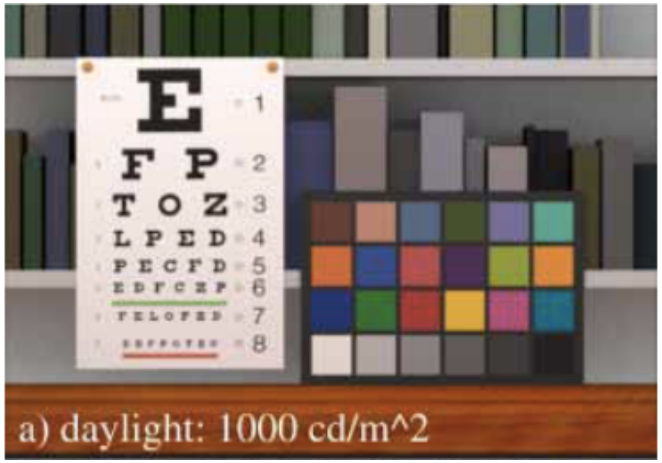

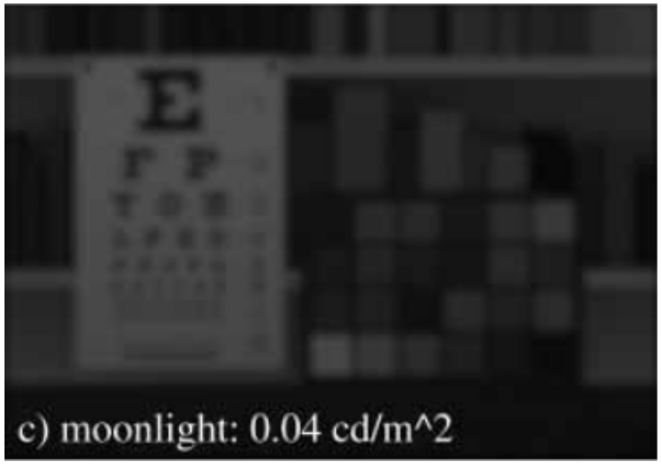

Forward and inverse visual model

- Advantages of F&I visual models
  - Can render images for different viewing conditions
    - Different state of chromatic or luminance adaptation
  - Physically plausible
    - output in the units of luminance or radiance
- Shortcomings F&I visual models
  - Assume that a standard display can reproduce the impression of viewing much brighter or darker scenes
  - Cannot ensure that the resulting image is within the dynamic range of the display
    - Not necessary meant to reduce the dynamic range
  - Visual models are difficult to invert

Outline

- What is tone-mapping?
- The perception of HDR scenes
- Major approaches to tone-mapping
  - Illumination & reflectance separation
  - Forward visual model
  - Forward & inverse visual model
  - Constraint mapping problem

Constraint mapping problem

- Goal: to restrict the range of values while reducing inflicted damage

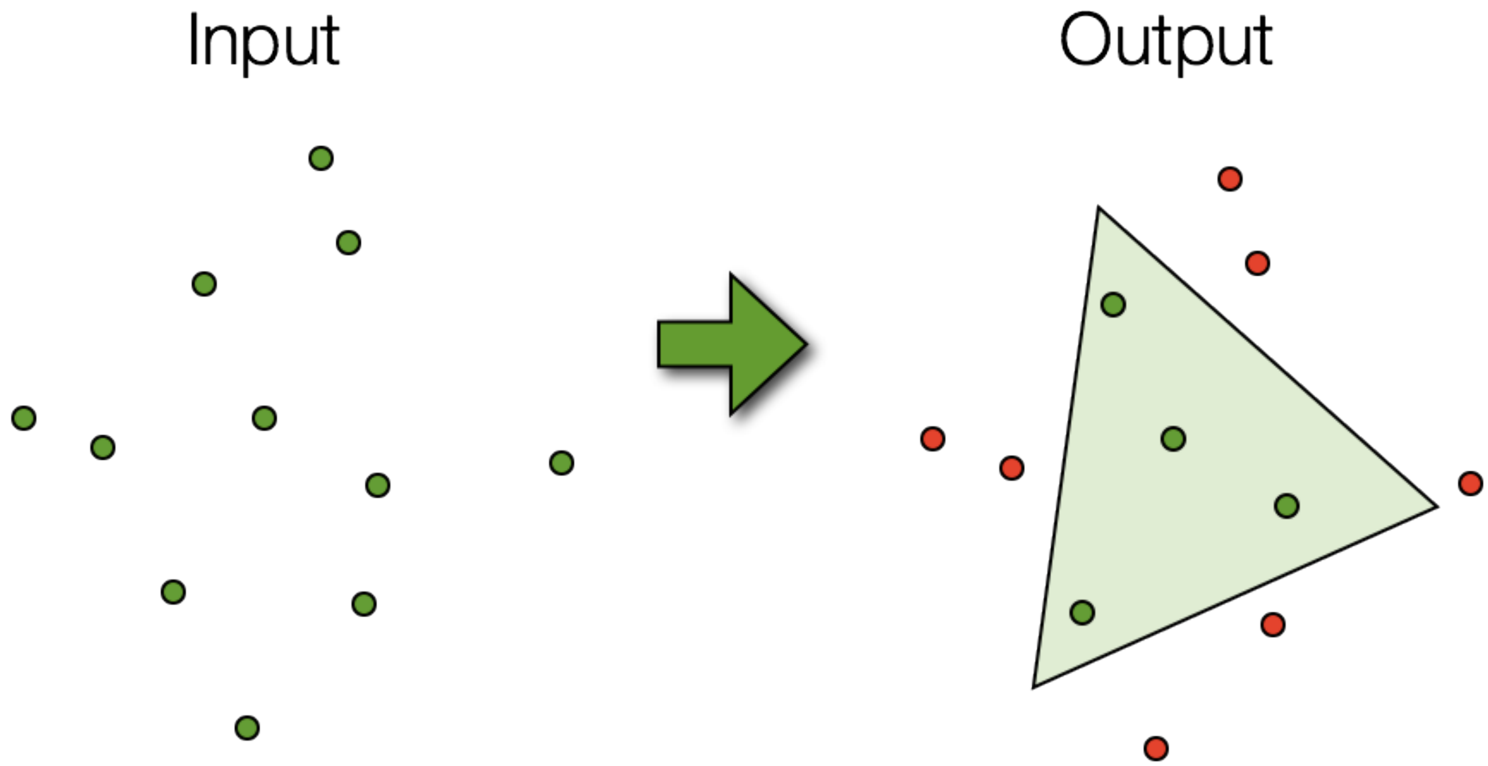

Global tone mapping operator

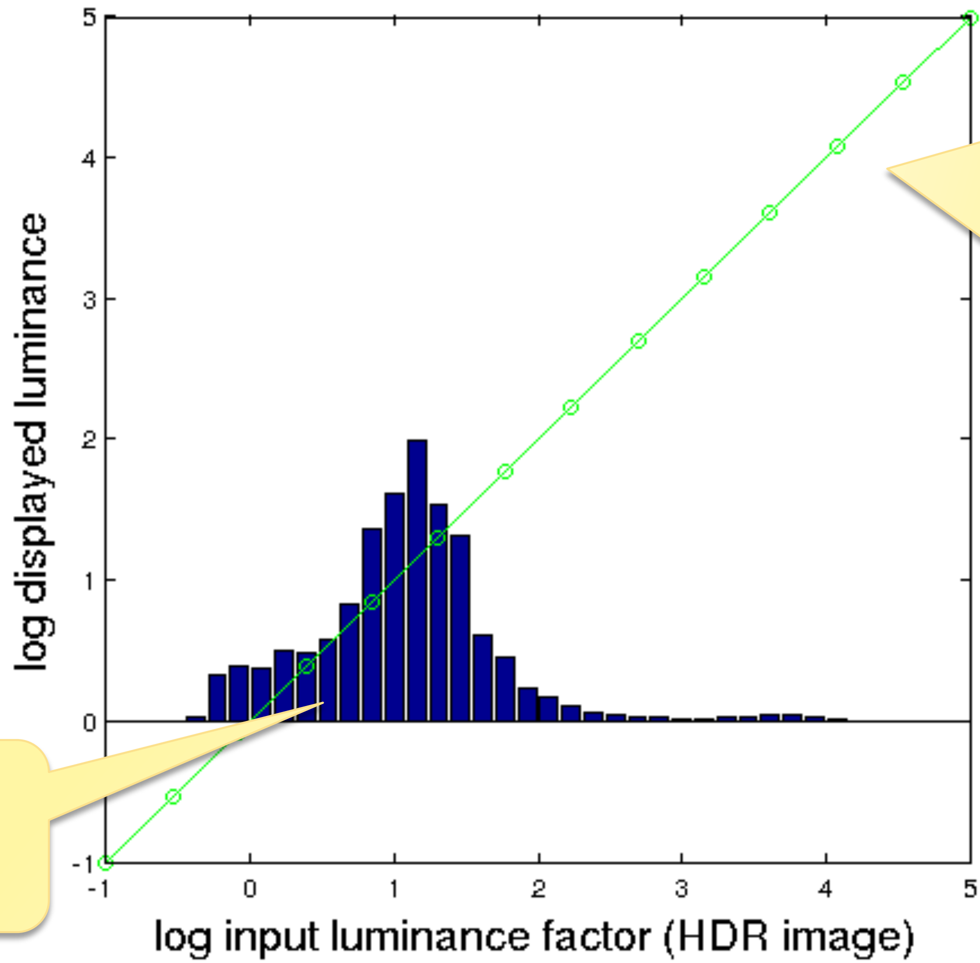

Image histogram

Best tone-mapping is the one which does not do anything, i.e. slope of the tone-mapping curves is equal to 1.

Display limitations

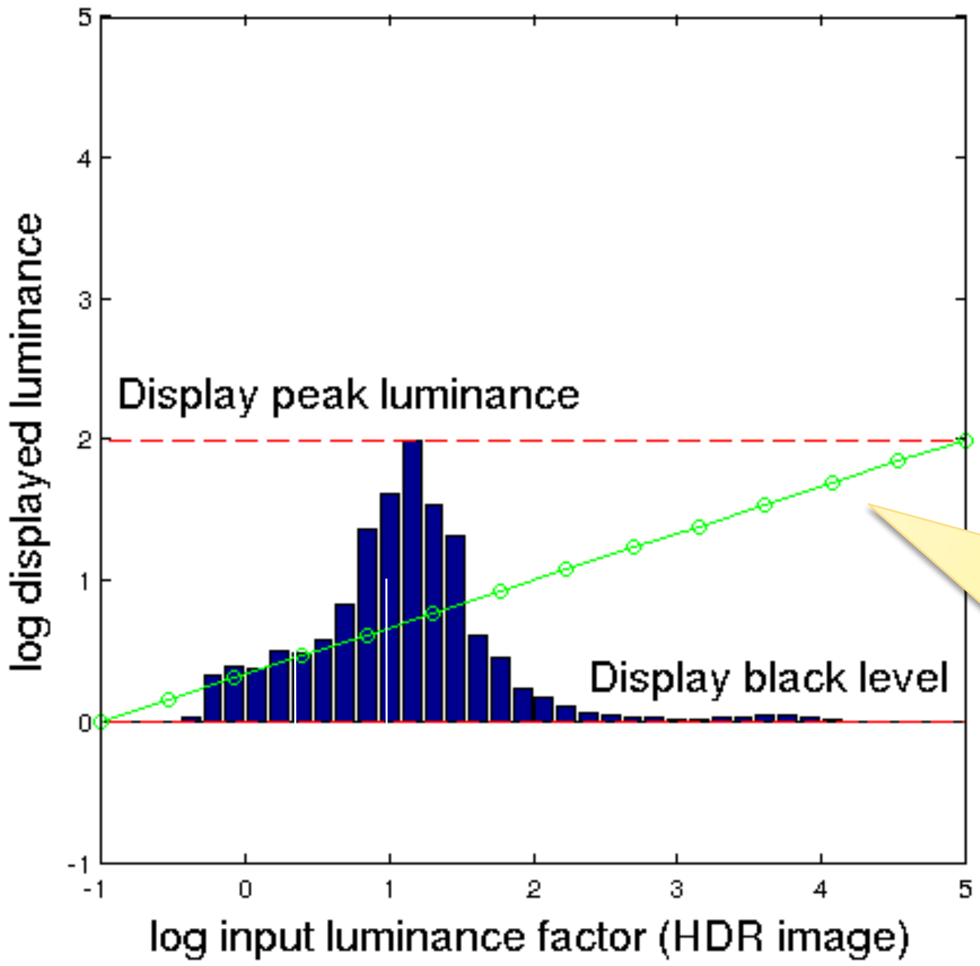

But in practice contrast (slope) must be limited due to display limitations.

Tone mapping

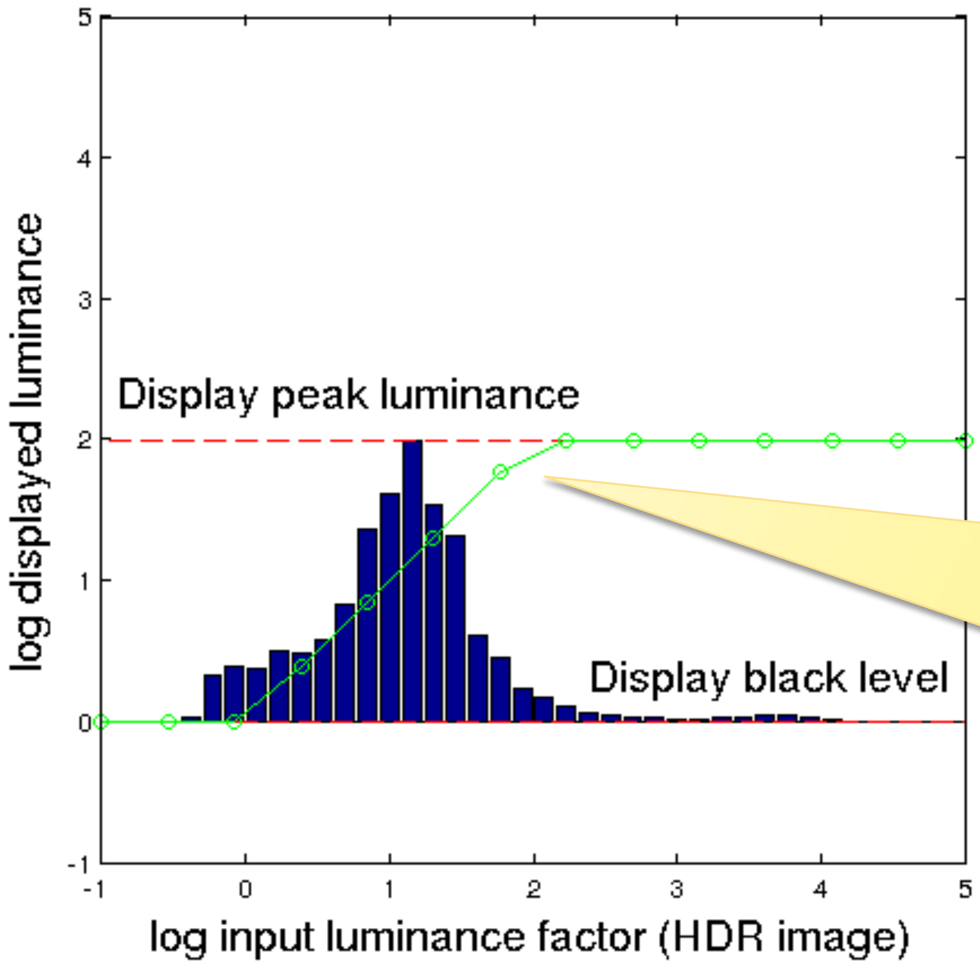

Global tone-mapping is a compromise between clipping and contrast compression.

Histogram equalization

- 1. From histogram compute cumulative distribution funct.

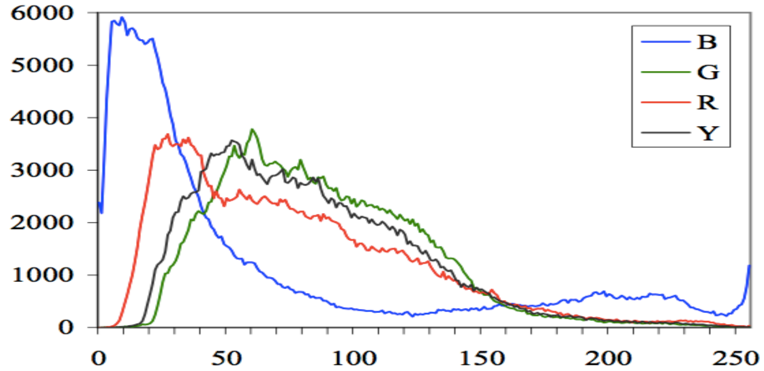

$$c(I) = \frac{1}{N}\sum_{i=0}^{I} h(i) = c(I-1) + \frac{1}{N}h(I)$$

- 2. Use that function to assign new pixel values

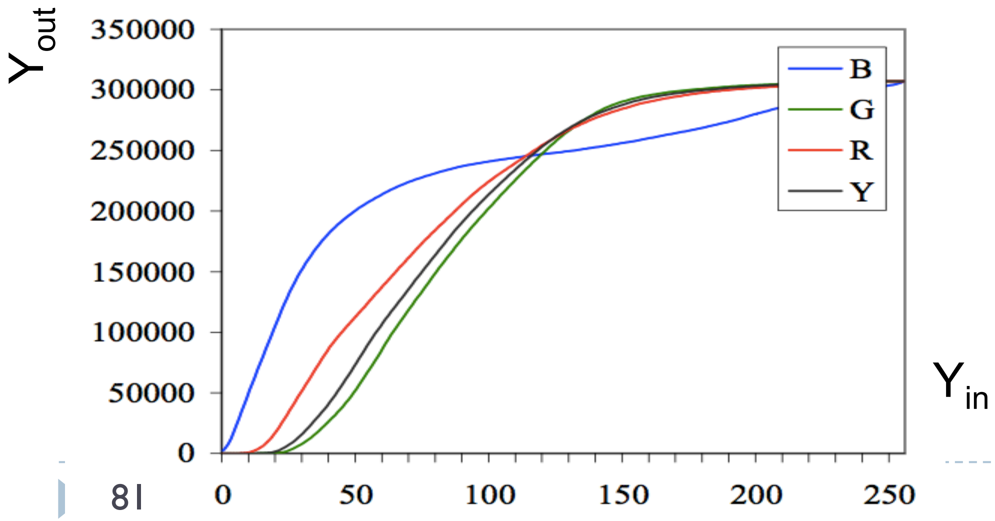

$$Y_{out} = c(Y_{in})$$

Histogram equalization

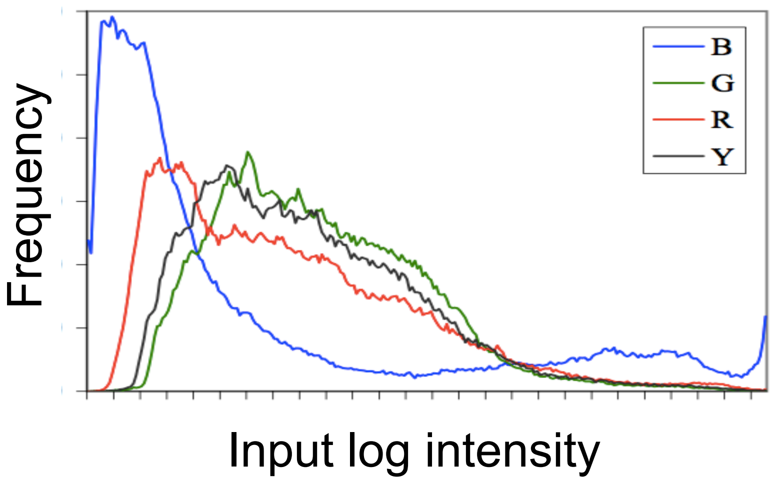

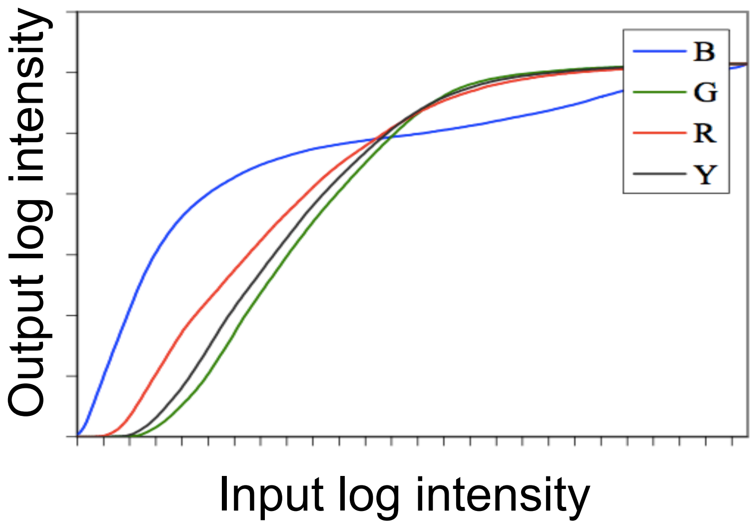

- Steepest slope for strongly represented bins
  - Enhance contrast, if many pixels
  - Reduce contrast, if few pixels
- HE distributes contrast distortions relative to the “importance” of a brightness level

Histogram adjustment with a linear ceiling

- [Larson et al. 1997, IEEE TVCG]

Linear mapping

Histogram equalization

Histogram equalization with ceiling

Histogram adjustment with a linear ceiling

- Truncate the bins that exceed the ceiling
- Recompute the ceiling based on the truncated histogram
- Repeat until converges

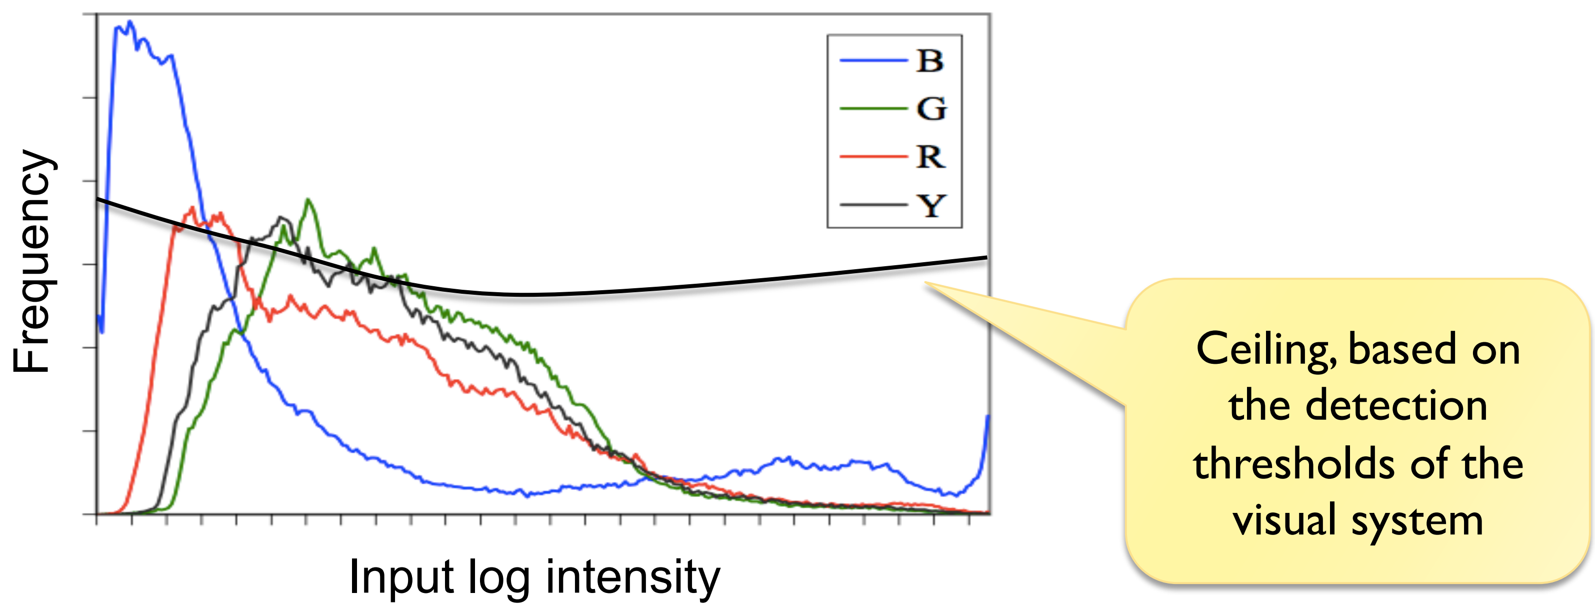

Task 8: Histogram equalization

- Implement histogram equalization of HDR images
- Operate on luminance only, then add color
- Operate in the $\log_{10}$ domain
- Rescale the result 0-2 (log domain) or 1-100 (linear domain)
- Use inverse display model to map from resulting luminance to pixel values
- Optional: implement capping histogram values, so that the maximum slope <= 1

Display adaptive tone-mapping

Goal: Minimize the visual difference between the input and displayed images

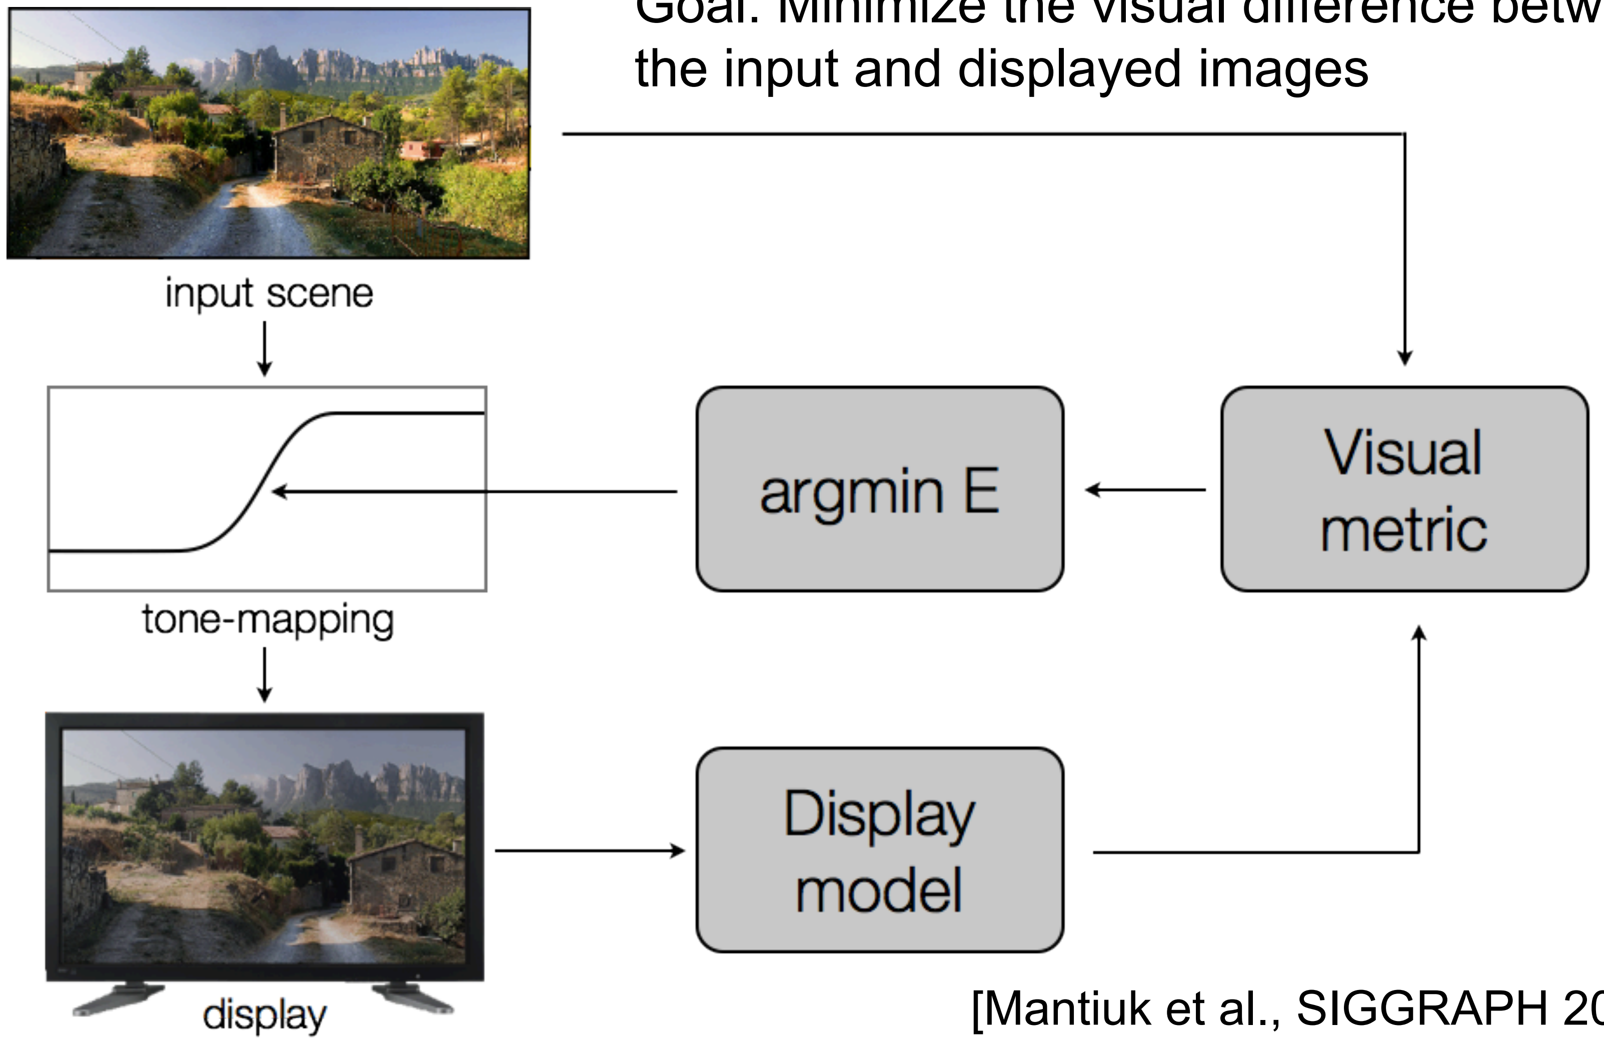

[Mantiuk et al., SIGGRAPH 2008]

Ambient illumination compensation

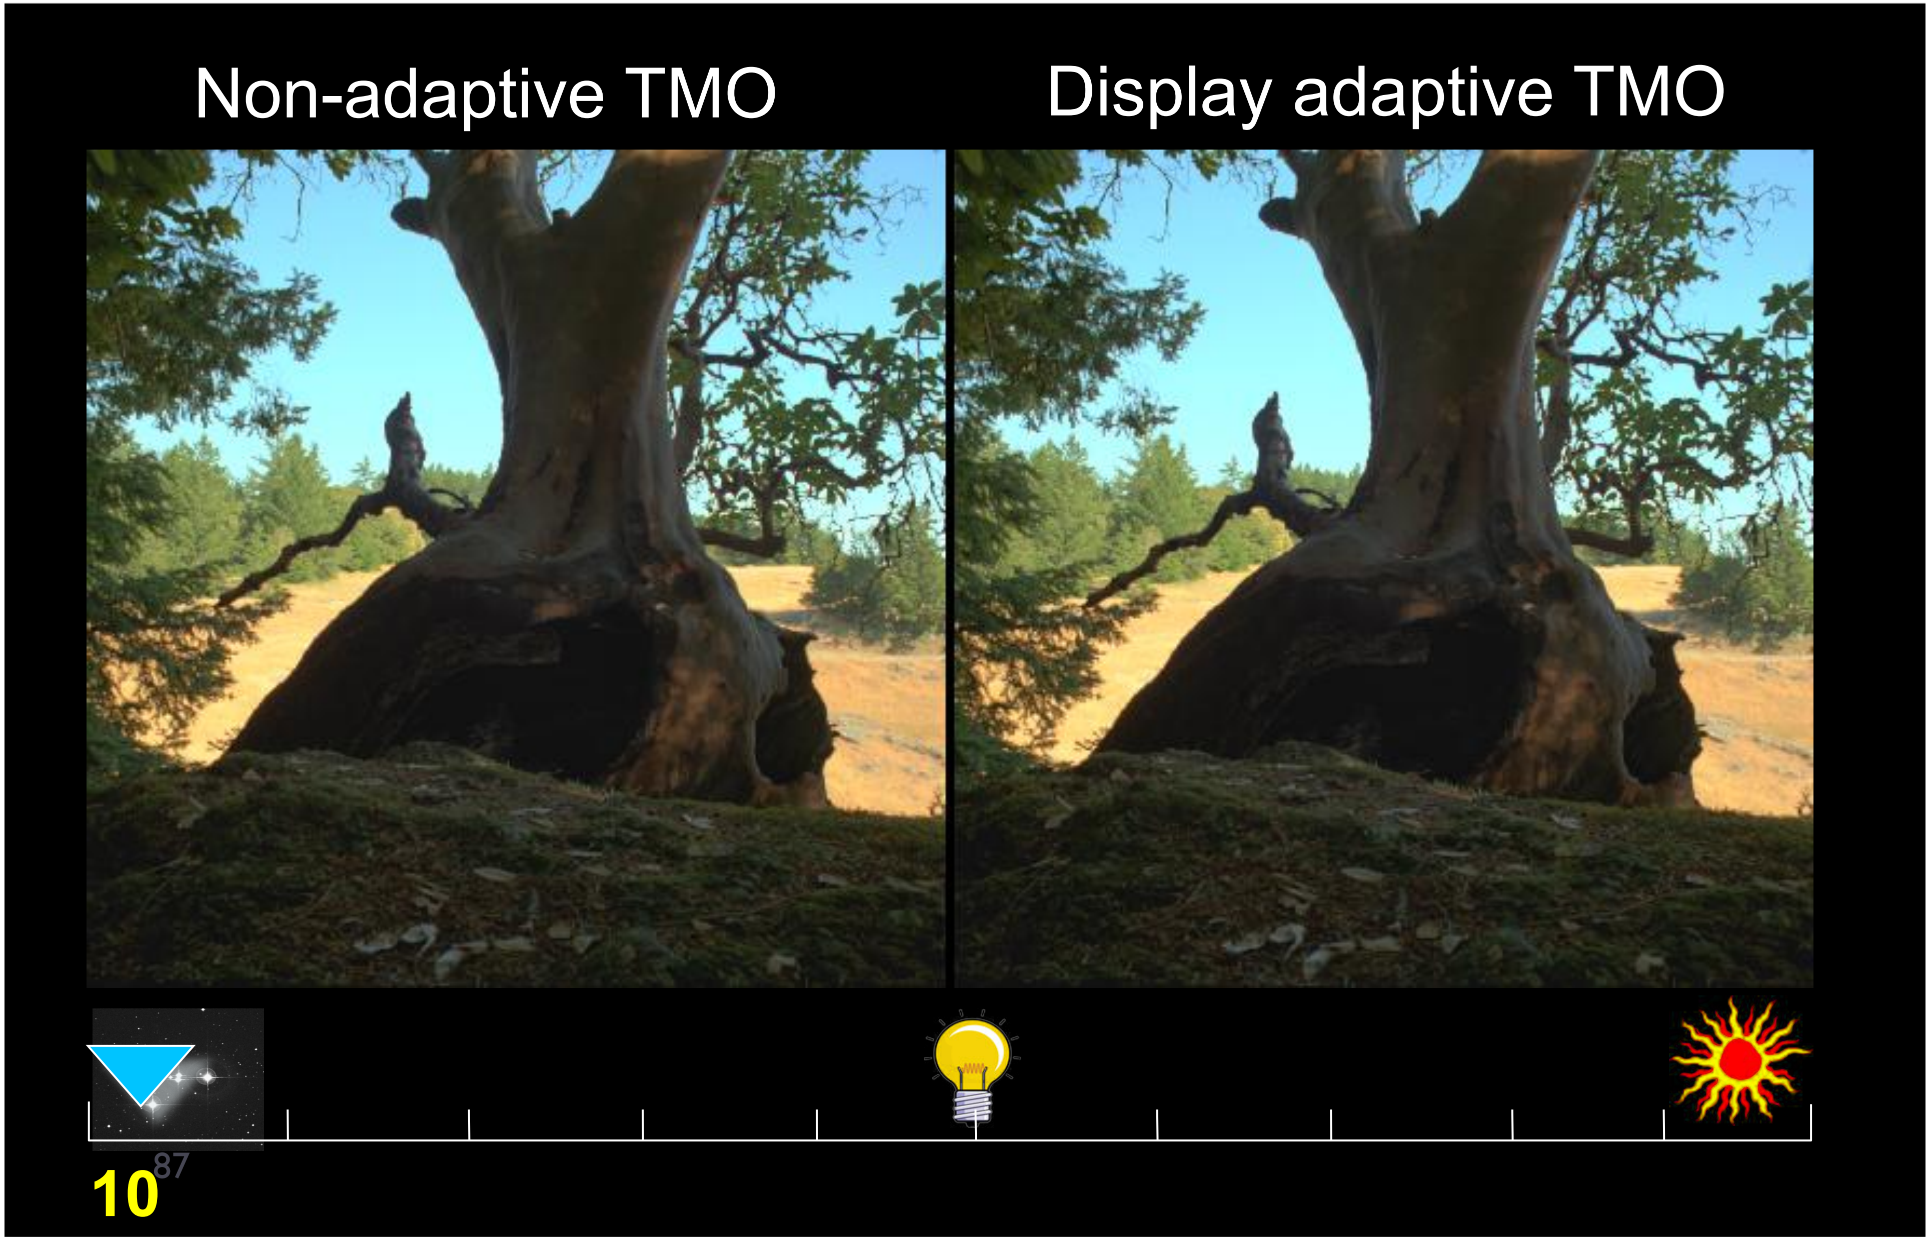

Ambient illumination compensation

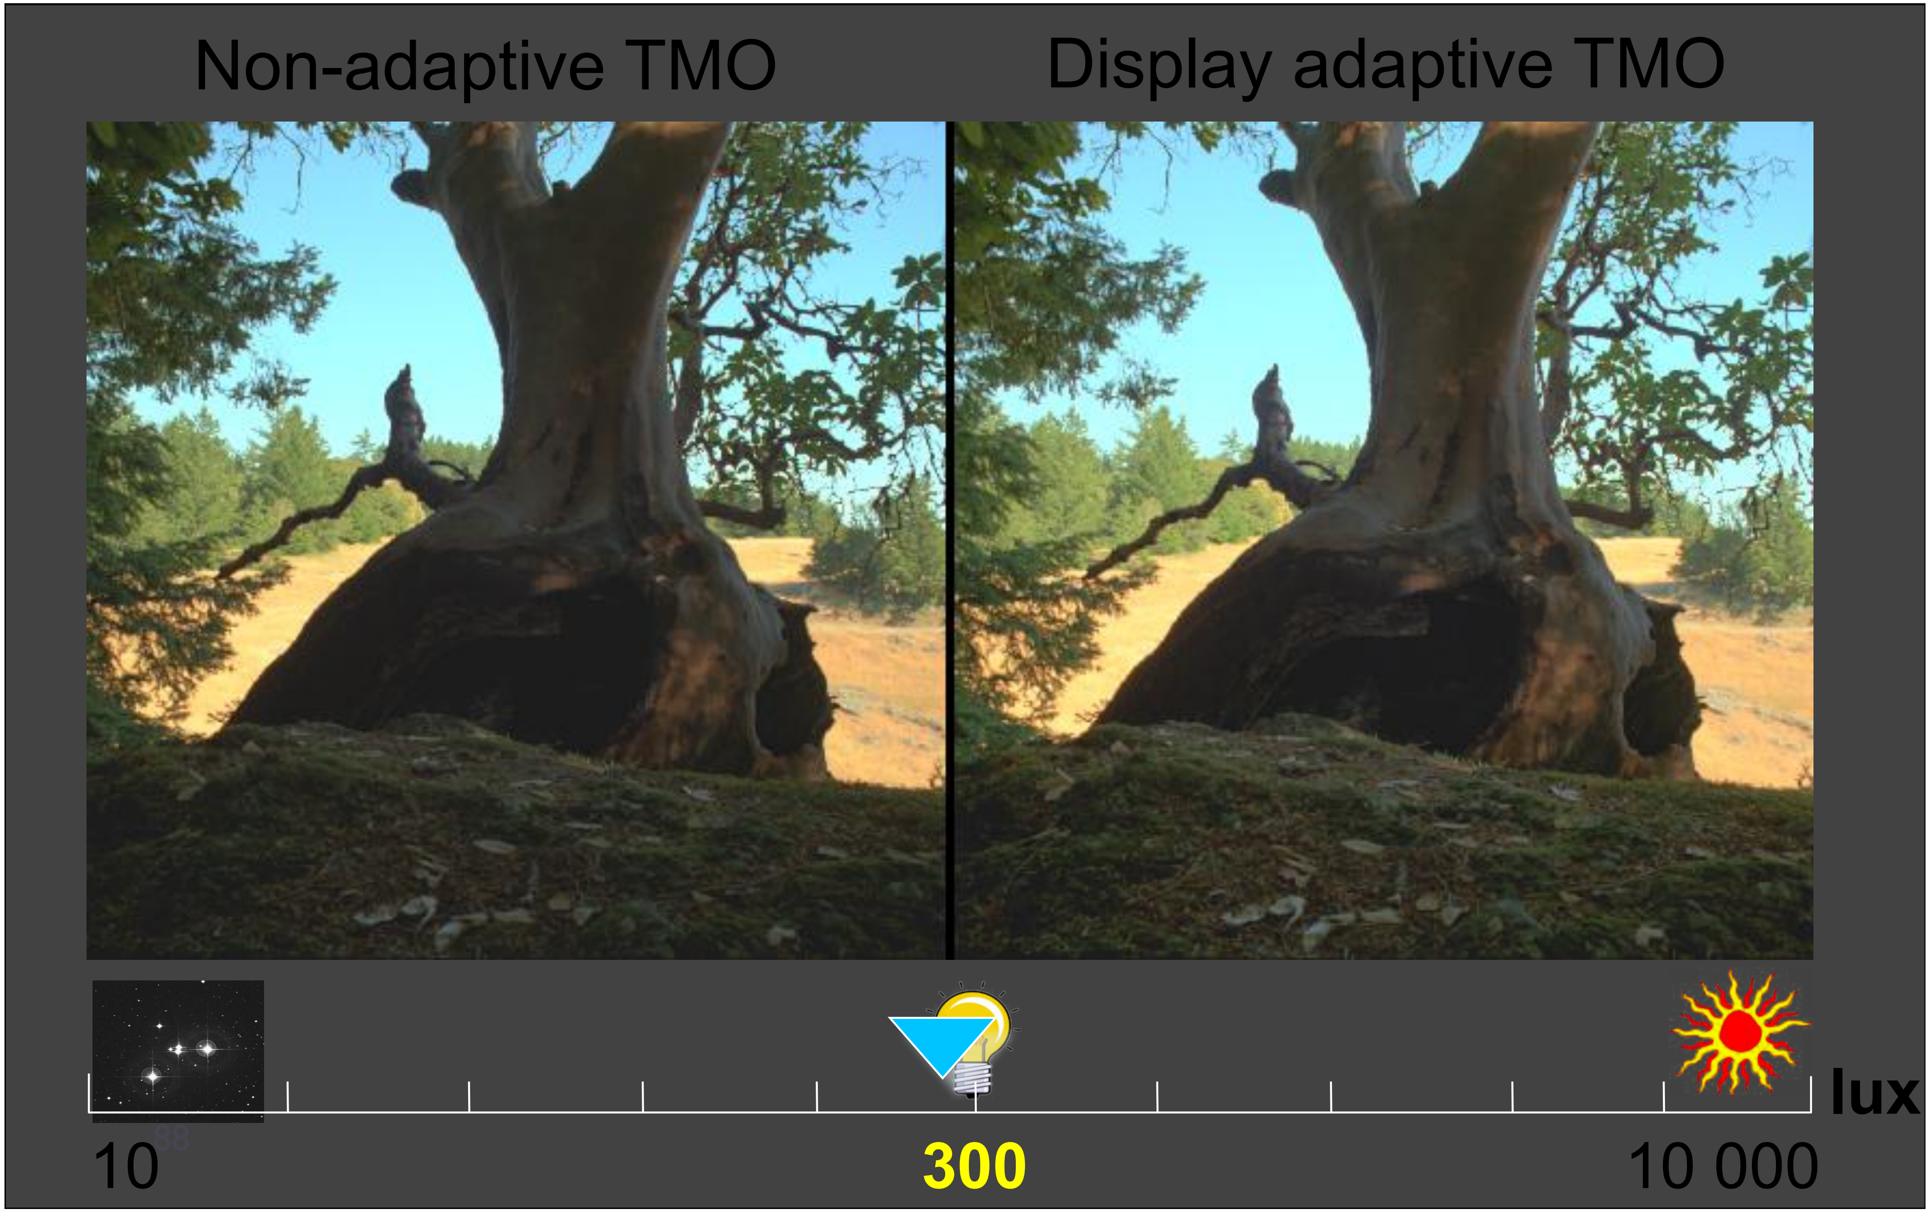

Ambient illumination compensation

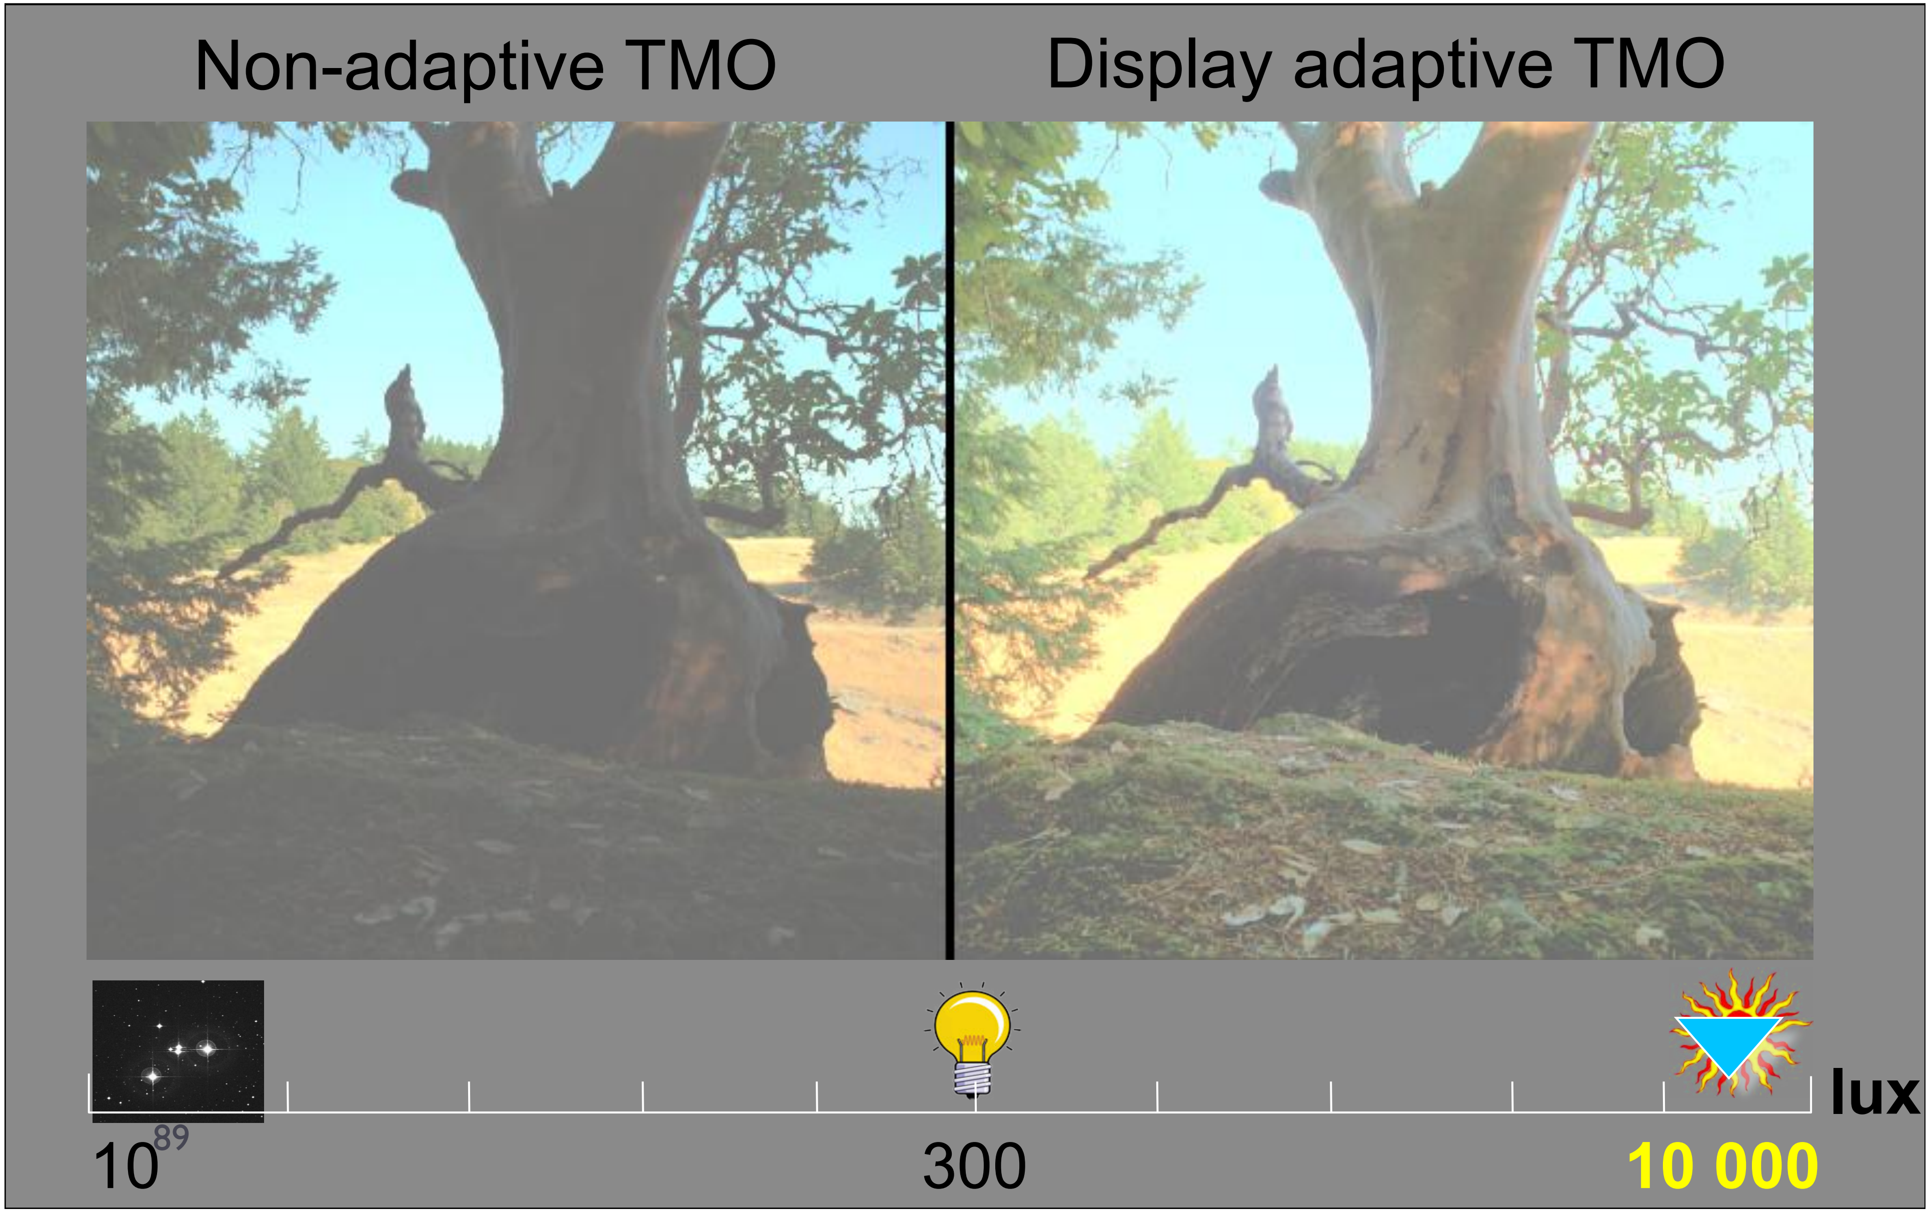

Results: display contrast

ePaper | standard LCD | HDR display

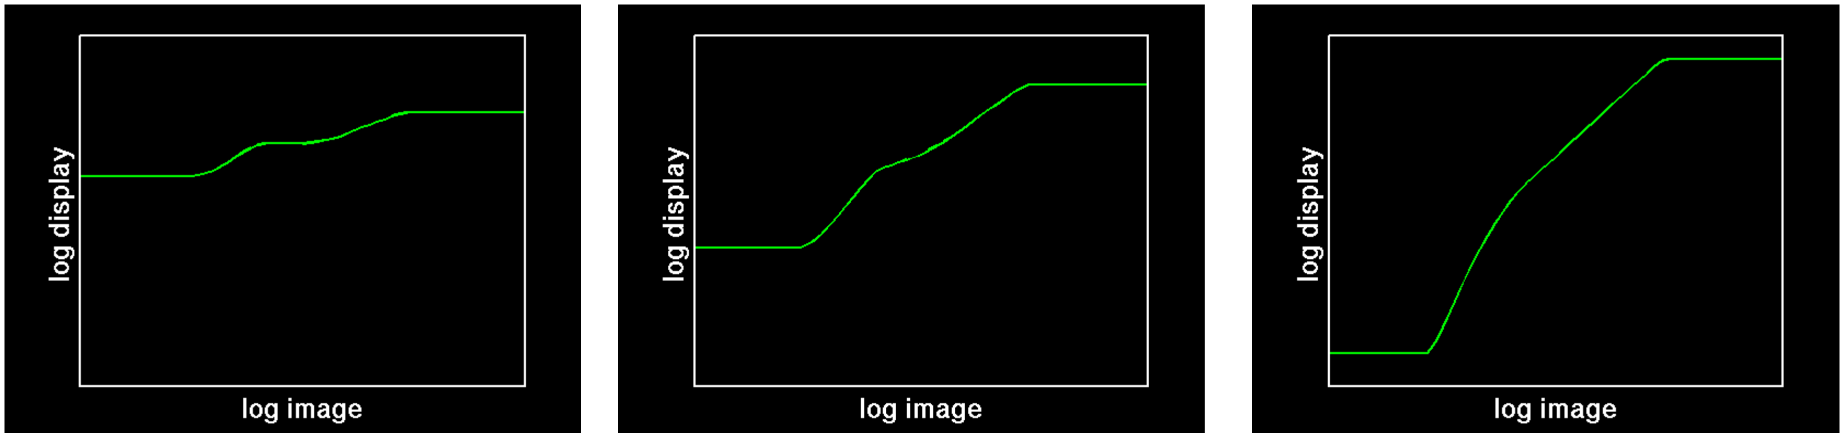

References

- Overview of HDR imaging & tone-mapping
  - http://www.cl.cam.ac.uk/~rkm38/hdri_book.html
- Papers on the simulation of glare

  - Spencer, G. et al. 1995. Physically-Based Glare Effects for Digital Images. Proceedings of SIGGRAPH. (1995), 325–334
  - Ritschel, T. et al. 2009. Temporal Glare: Real-Time Dynamic Simulation of the Scattering in the Human Eye. Computer Graphics Forum. 28, 2 (Apr. 2009), 183–192