Perception in tone-mapping

Rafał Mantiuk
Computer Laboratory, University of Cambridge

Cornell Box: a rendering or photograph?

Rendering

Photograph

Real-world scenes are more challenging

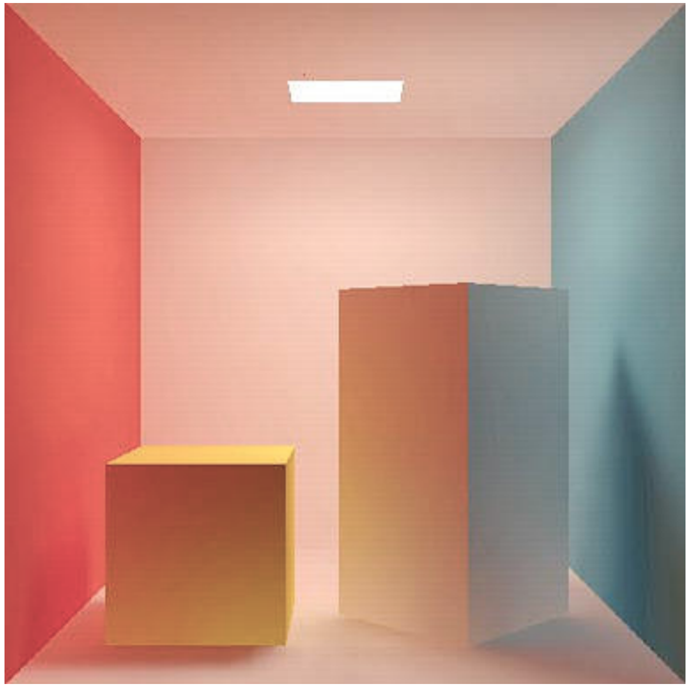

- The match could not be achieved if the light source in the top of the box was visible
- The display could not reproduce the right level of brightness

Outline

- Fundamentals
  - Intents of tone-mapping
  - Display model, luminance and luma
  - Four approaches to tone-mapping
- Perceptual issues
  - Color
  - Glare
  - Visual metrics-driven TMO
  - Light and dark adaptation

Intents of tone mapping

Three intents of tone-mapping

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

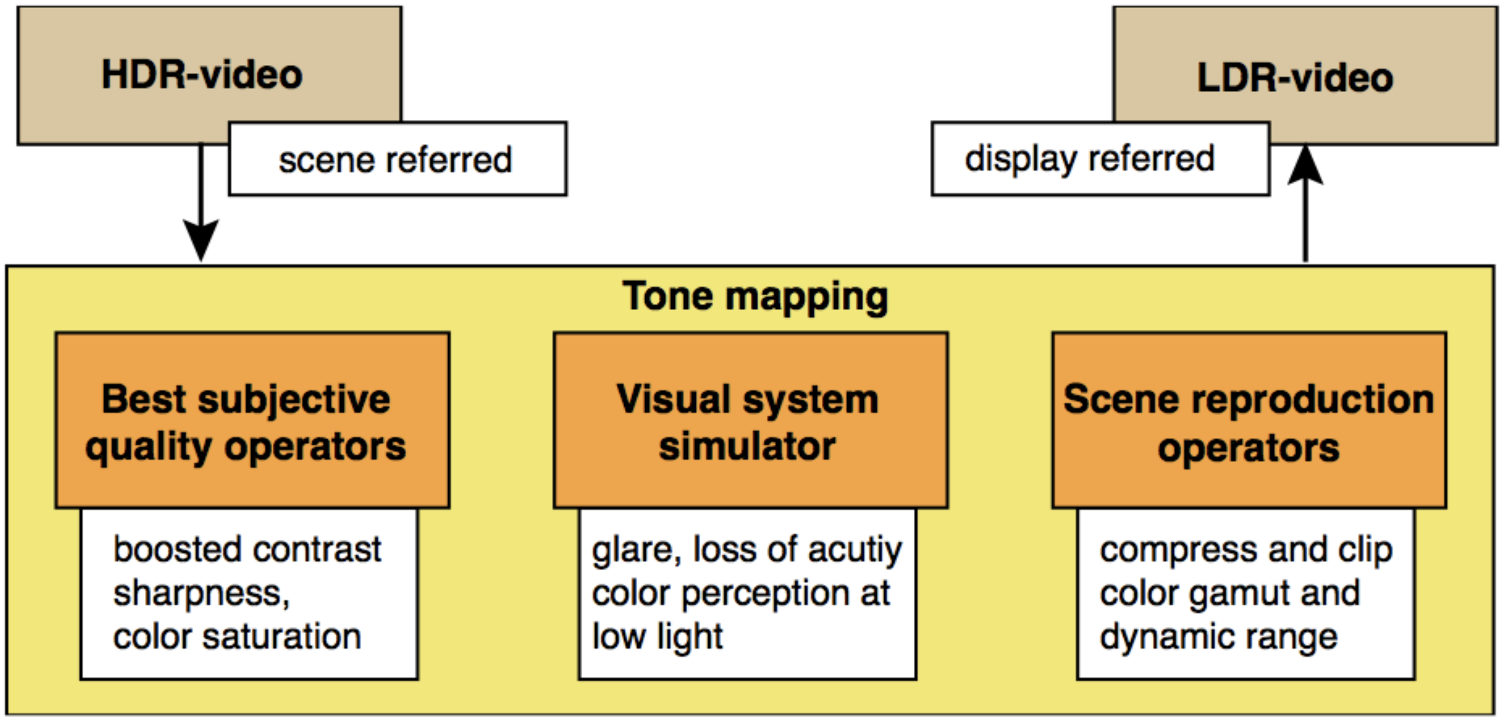

Intent #1: Best subjective quality

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

Intent #2: Visual system simulator

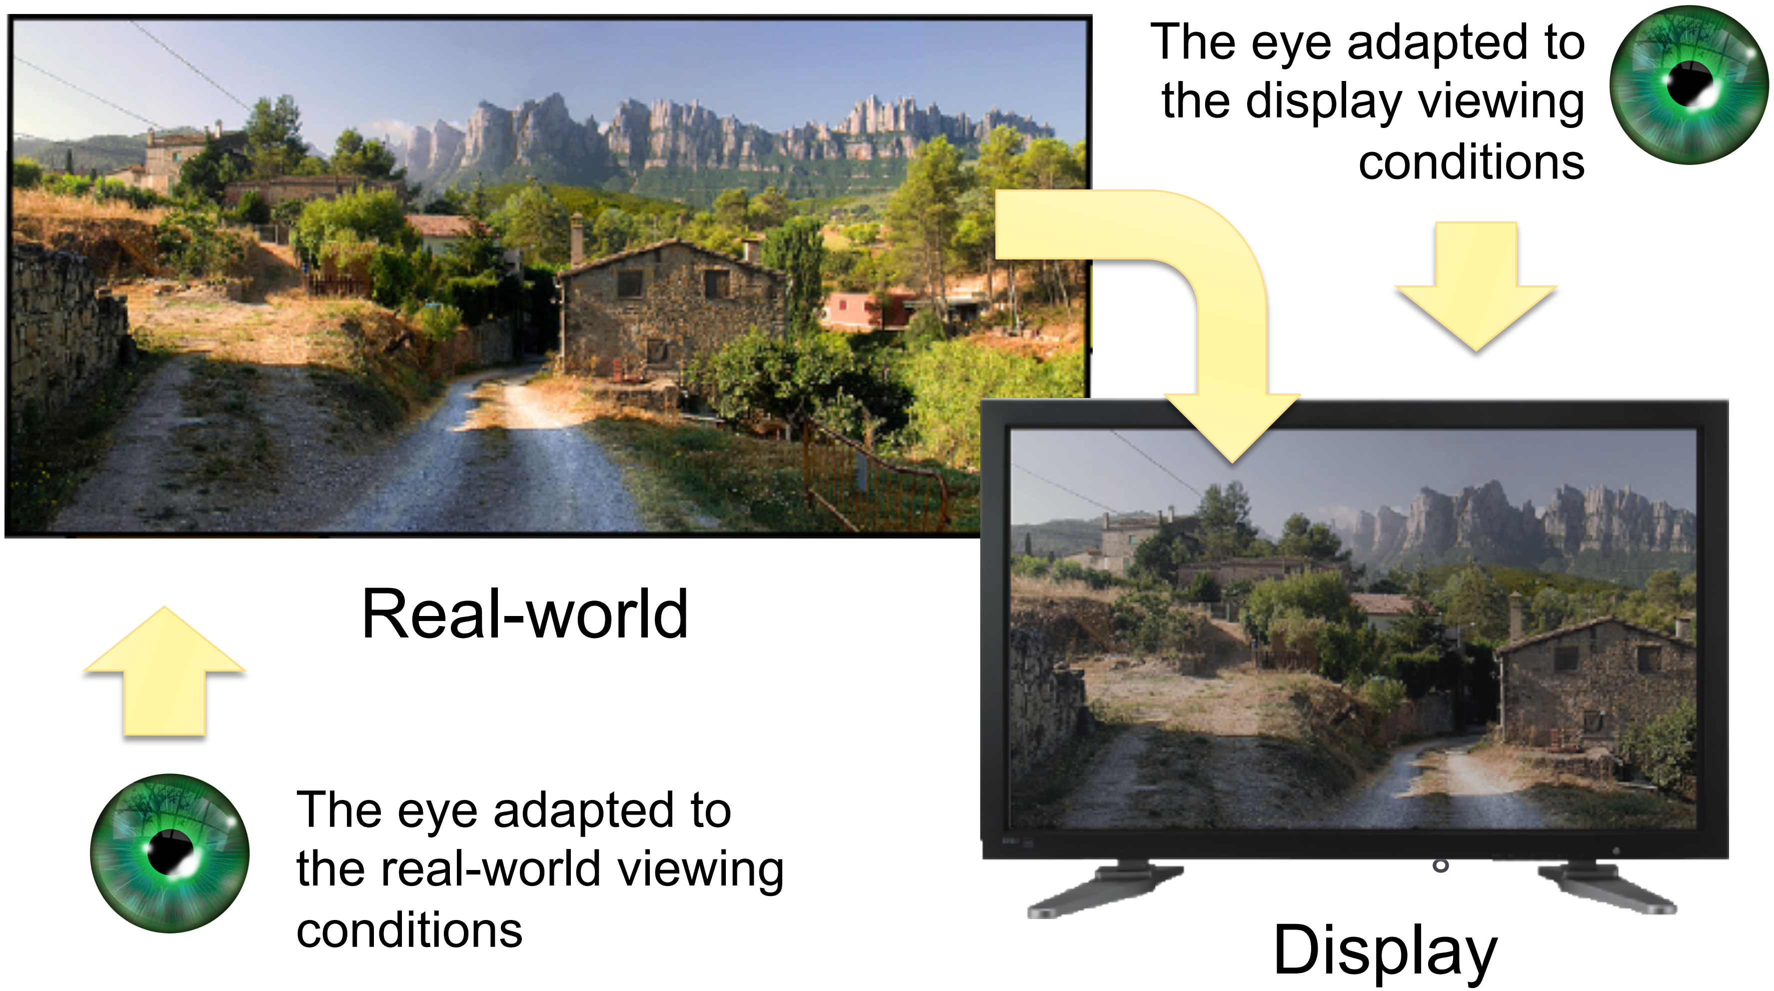

Goal: match color appearance

Intent #3: Scene reproduction problem

Real-world

Display

Goal: map colors to a restricted color space

Mapping problem

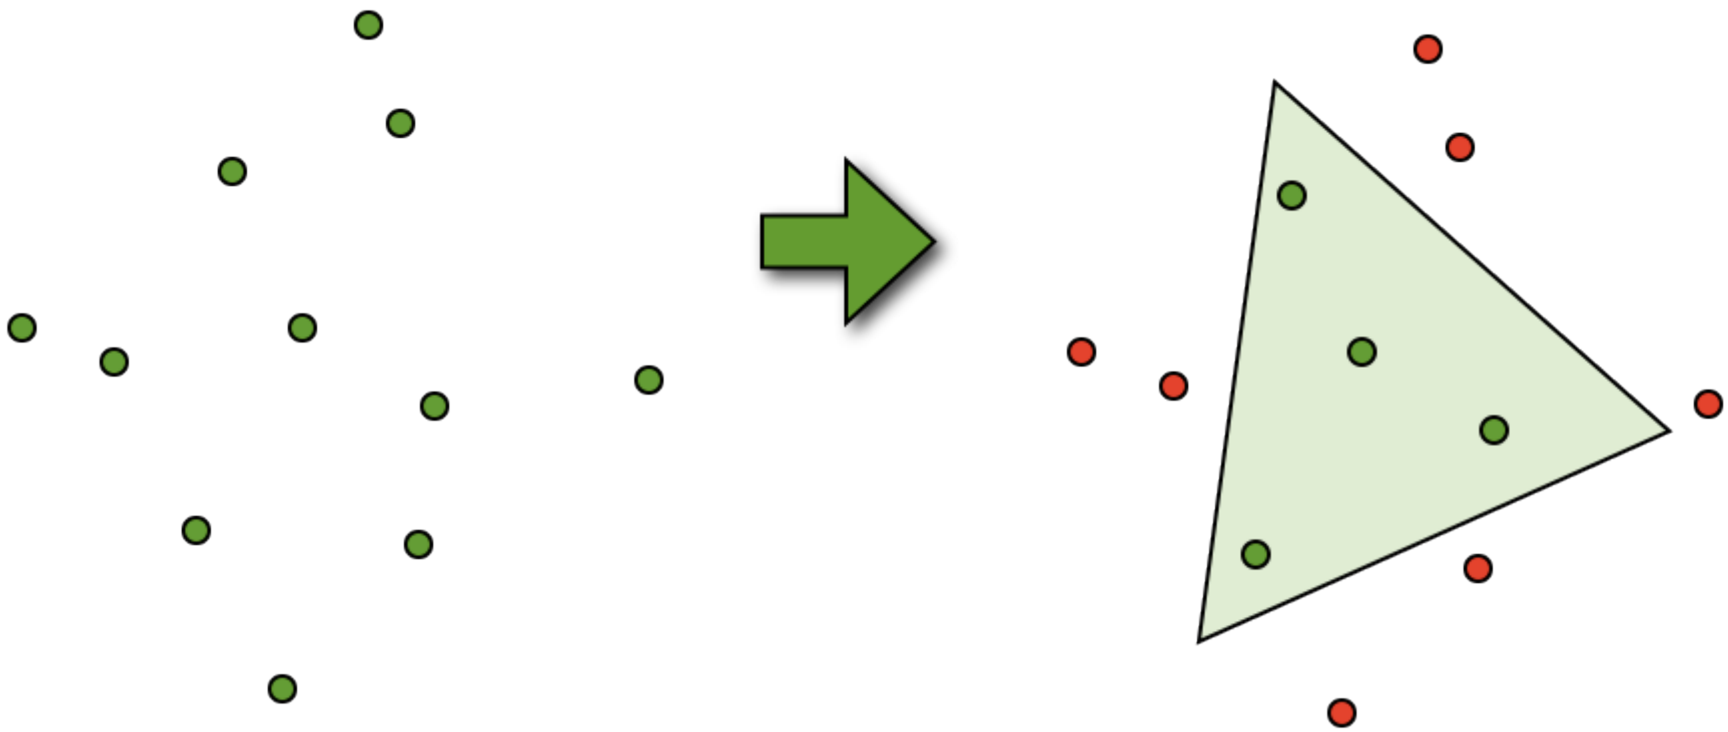

Real-world

Display

Display adaptive tone-mapping

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

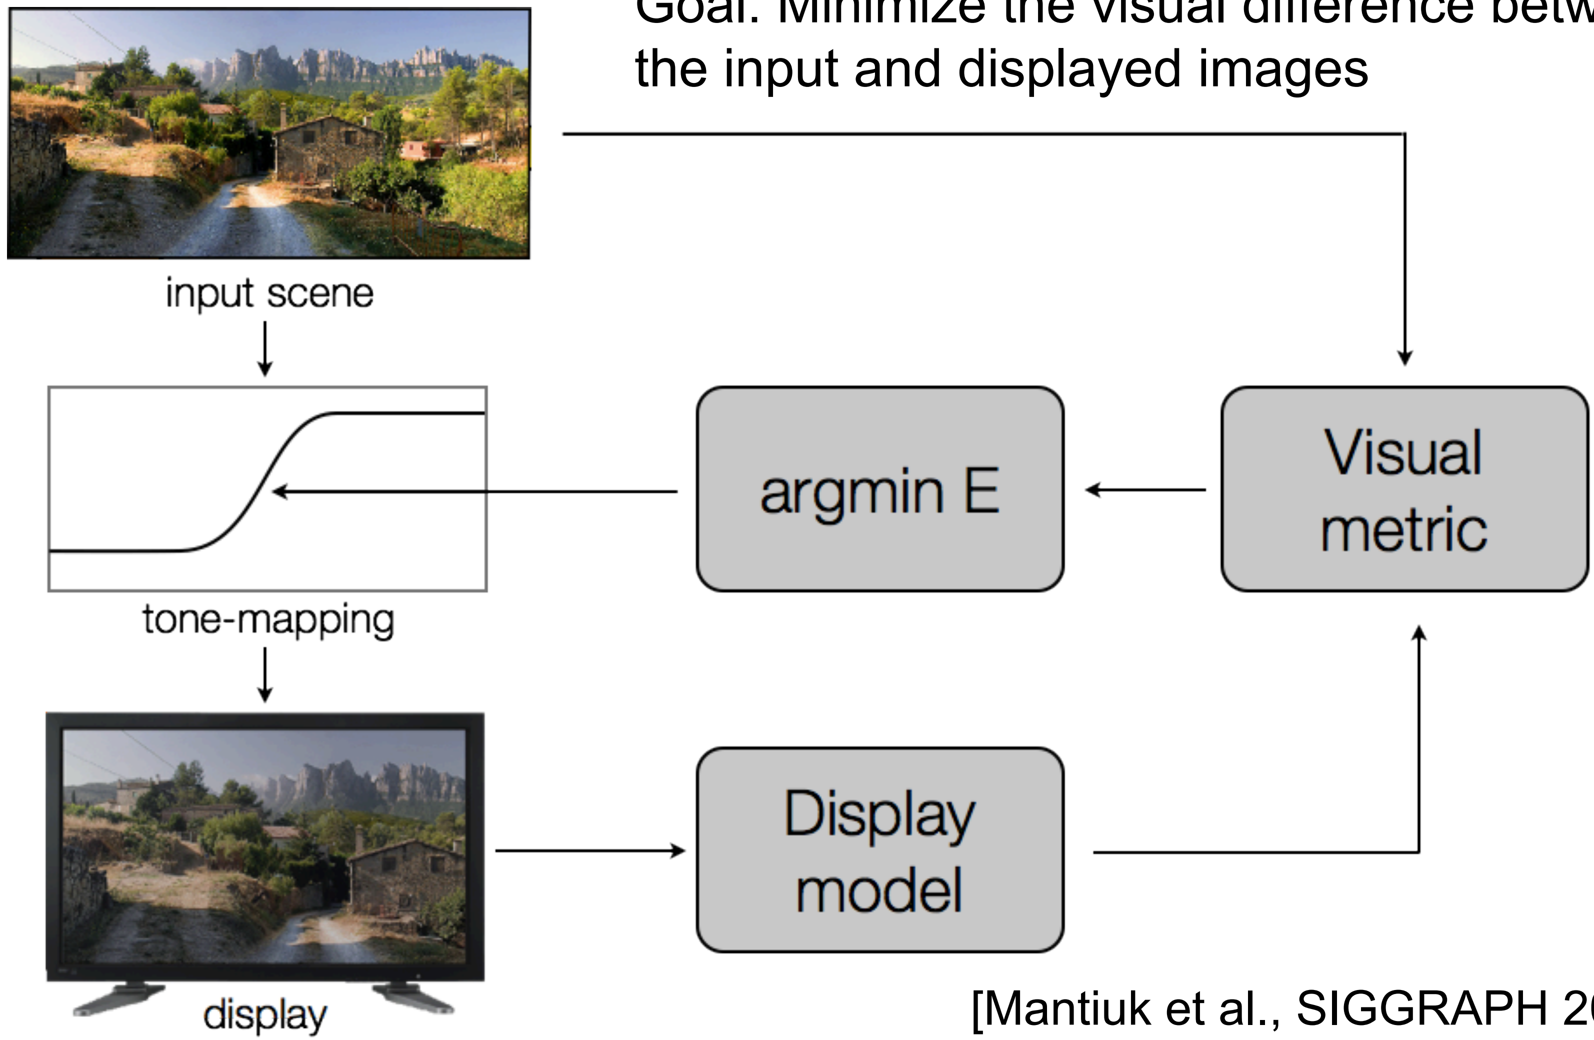

[Mantiuk et al., SIGGRAPH 2008]

Display model, luminance and luma

Forward and inverse display model

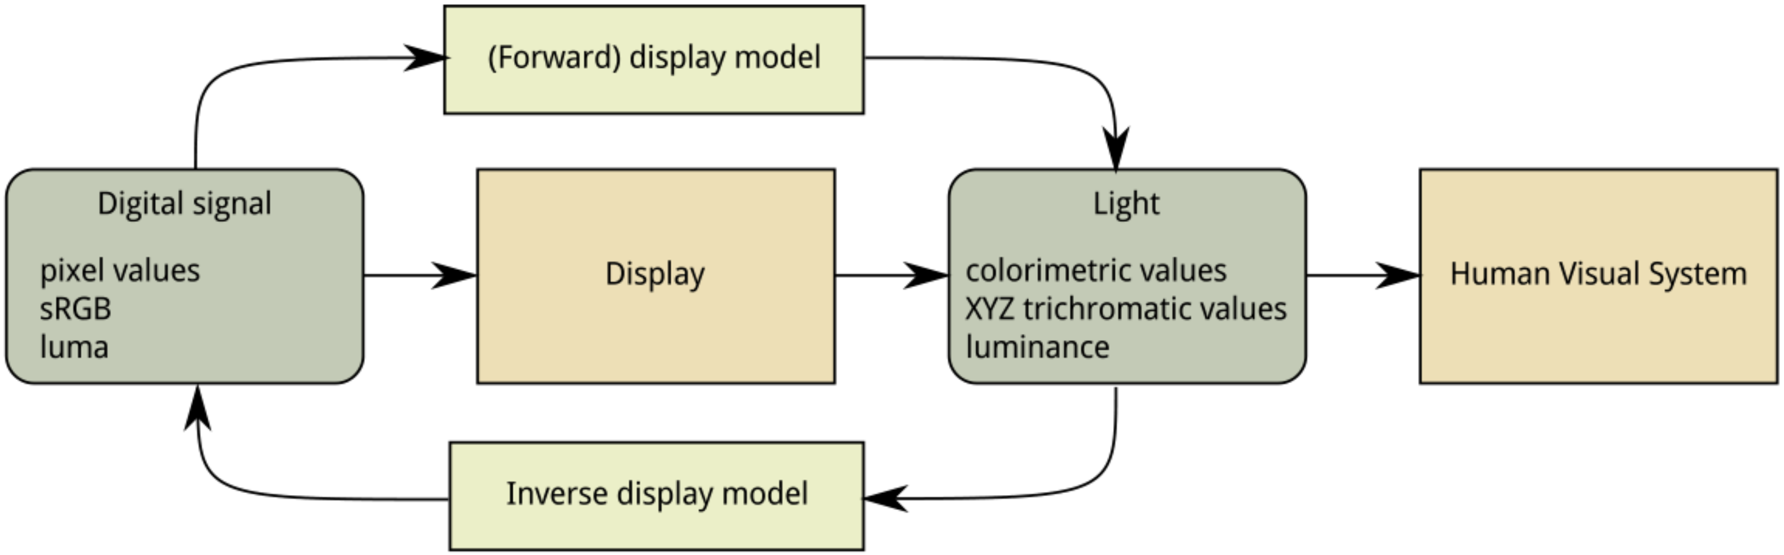

Luminance

- 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)

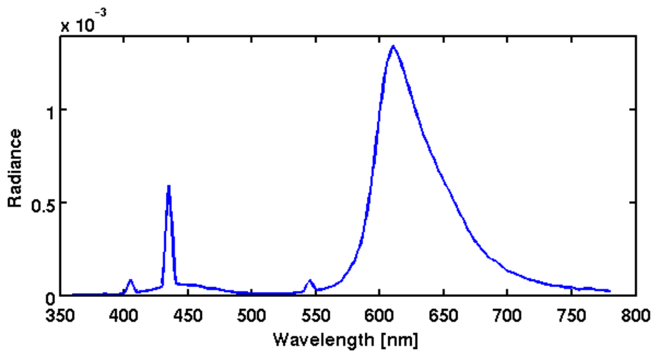

Luminous efficiency function (weighting)

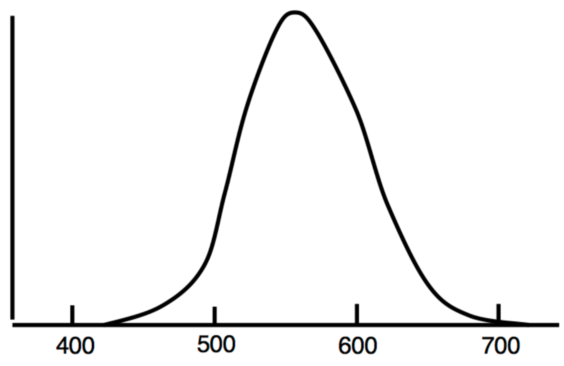

Luminance and Luma

- 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

Two ways to do tone-mapping

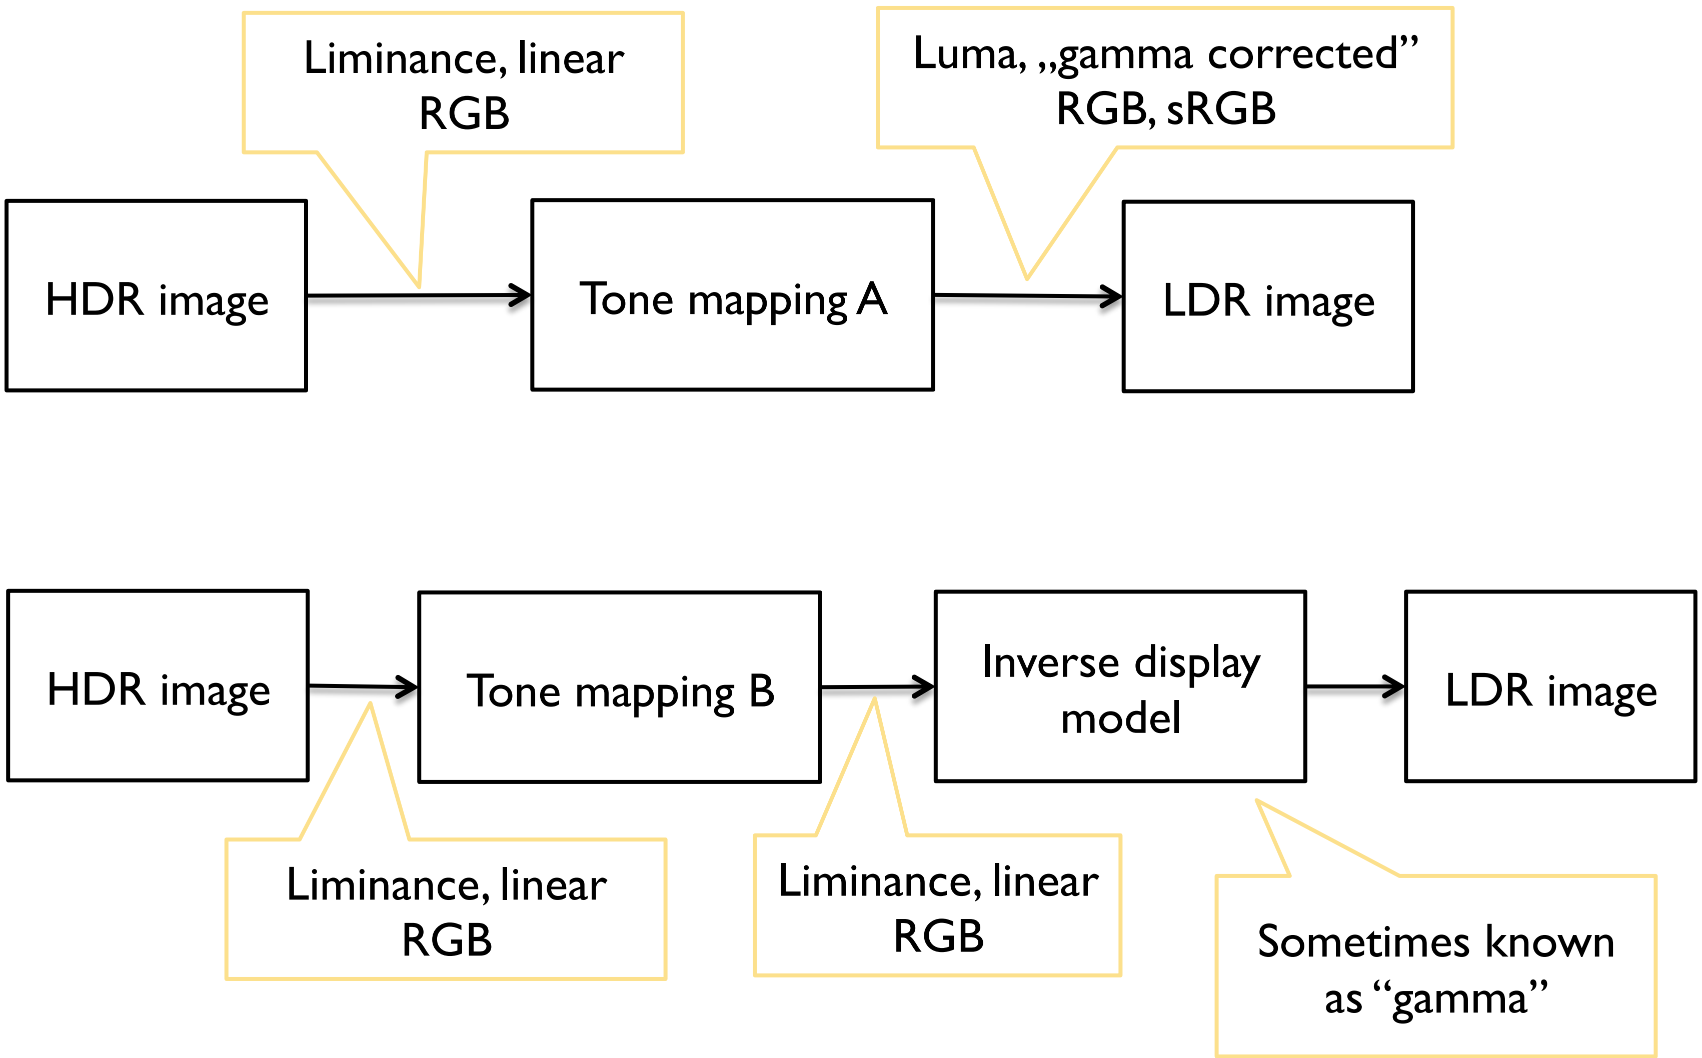

(Forward) Display model

- GOG: Gain-Gamma-Offset

Luminance

Peak luminance

Gamma

Display black level

Screen reflections

$$L = (L_{peak} - L_{black}) V^{\gamma} + L_{black} + L_{refl}$$

Gain

Pixel value 0-1

Offset

Reflectance factor (0.01)

$$L_{refl} = \frac{k}{2\pi} E_{amb}$$

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.

Arithmetic of HDR images

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

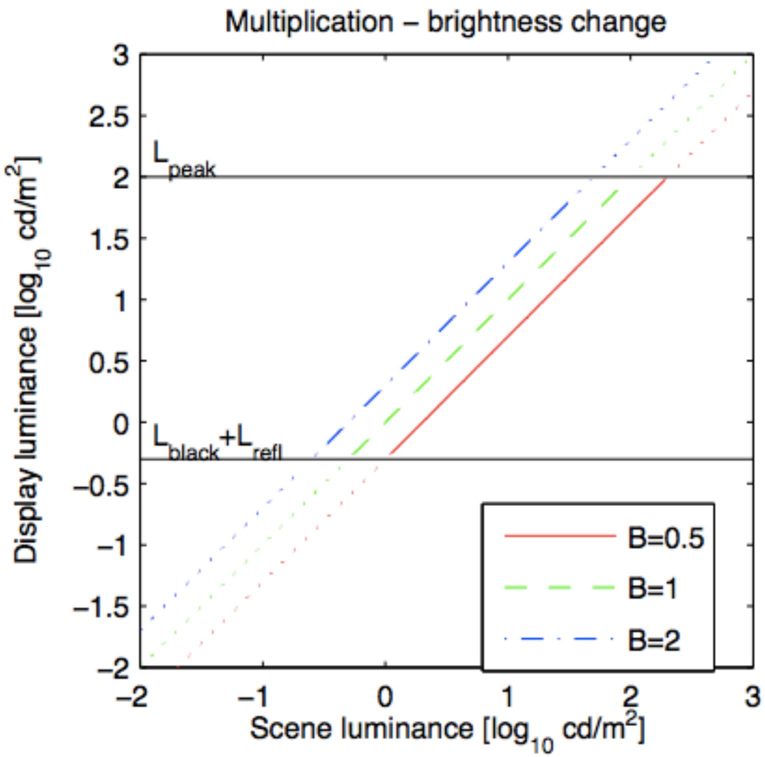

- Multiplication makes the image brighter or darker
- It does not change the dynamic range!

Power function – contrast change

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

Contrast change (gamma)

Luminance of white

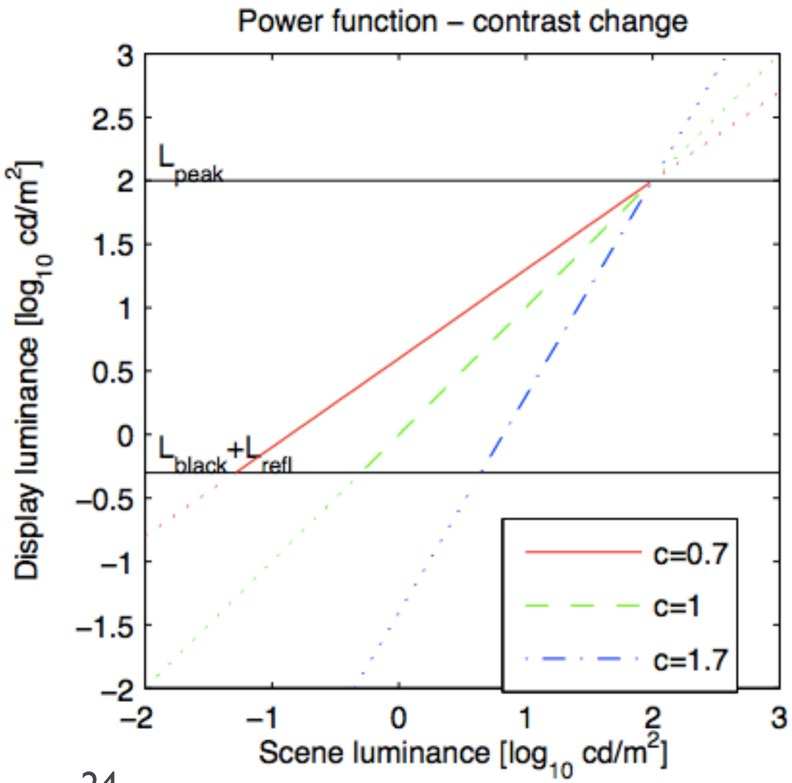

- Power function stretches or shrinks image dynamic range
- It is usually performed relative to reference white
- Apparent brightness changes is the side effect of pushing tones towards or away from the white point

Addition – black level

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

Black level (flare, fog)

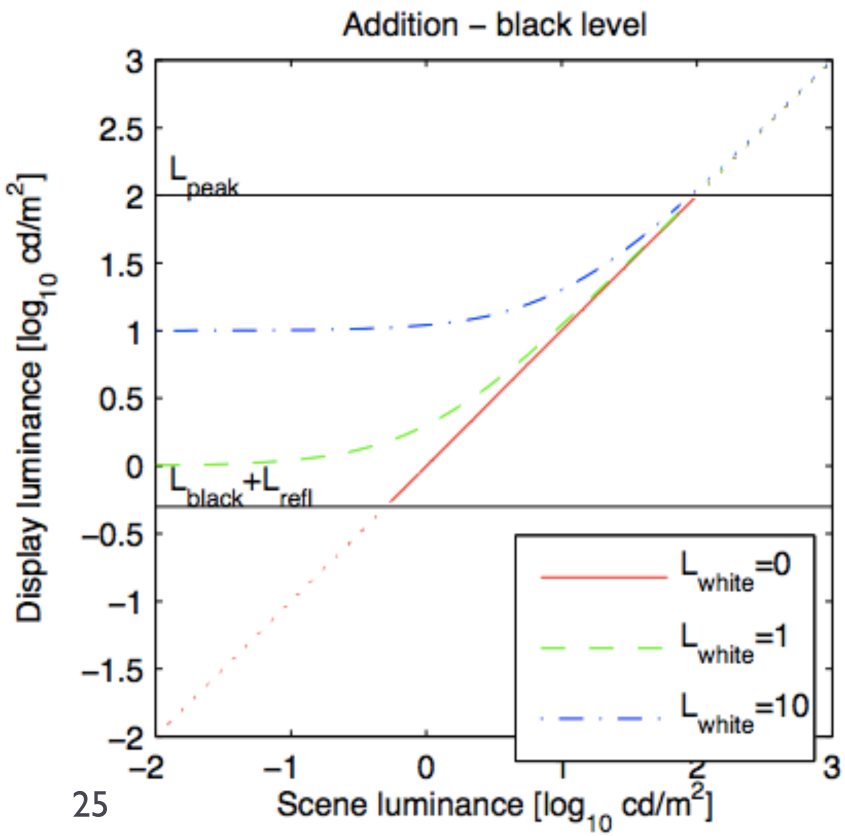

- Addition elevates black level, adds fog to an image
- It does NOT make the overall image brighter
- It reduces dynamic range

4 approaches to tone mapping

4 (major) approaches to tone mapping

- Illumination & reflectance separation
- Forward visual model
- Forward & inverse visual model
- Constrained mapping problem

Illumination & reflectance separation

Input

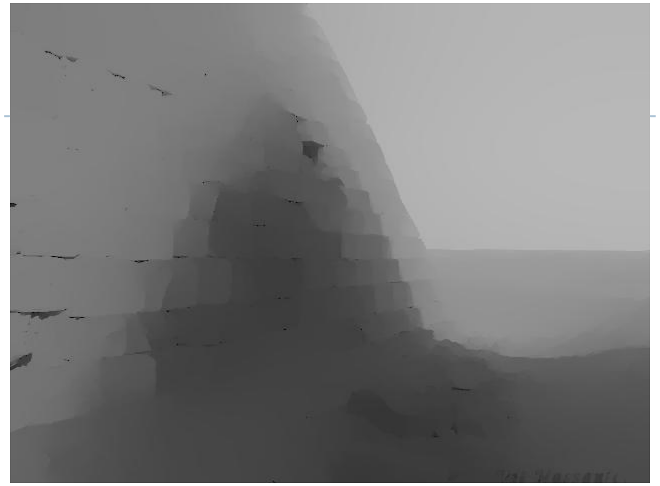

Illumination

Reflectance

$$Y = I \cdot R$$

Image

Illumination

Reflectance

Illumination and reflectance

Reflectance

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

Illumination

- 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

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

Tone-mapped image

Reflectance

Tone-mapping

Illumination

- 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

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

Spatially varying smoothing – less smoothing near the edges

- [Farbman et al., SIGGRAPH 2008]

WLS filter

- Stronger smoothing and still distinct edges

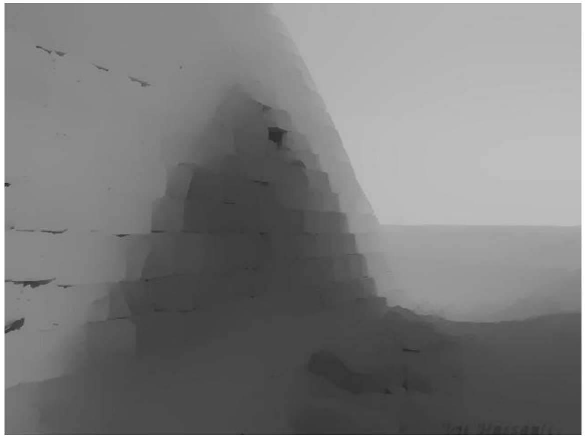

Tone mapping result

- Can produce stronger effects with fewer artifacts

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

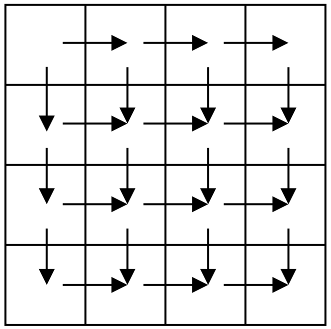

2nd step: set to 0 gradients less than the threshold

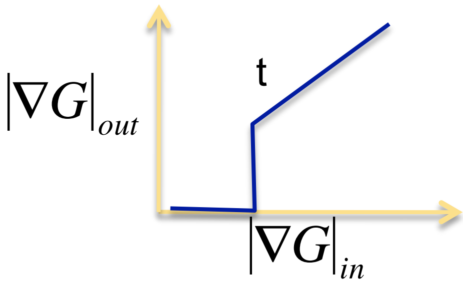

3rd step: reconstruct an image from the vector field

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

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

Retinex

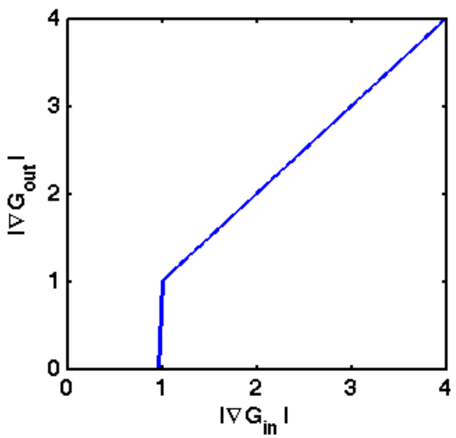

Gradient domain

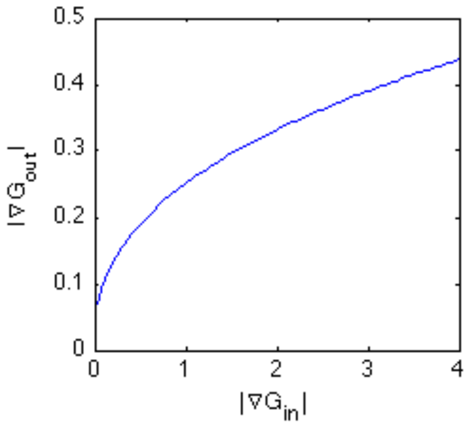

- Contrast compression achieved by global contrast reduction
  - Enhance reflectance, then compress everything

4 (major) approaches to tone mapping

- Illumination & reflectance separation
- Forward visual model
- Forward & inverse visual model
- Constrained mapping problem

Forward visual model

- Mimic the processing in the human visual system

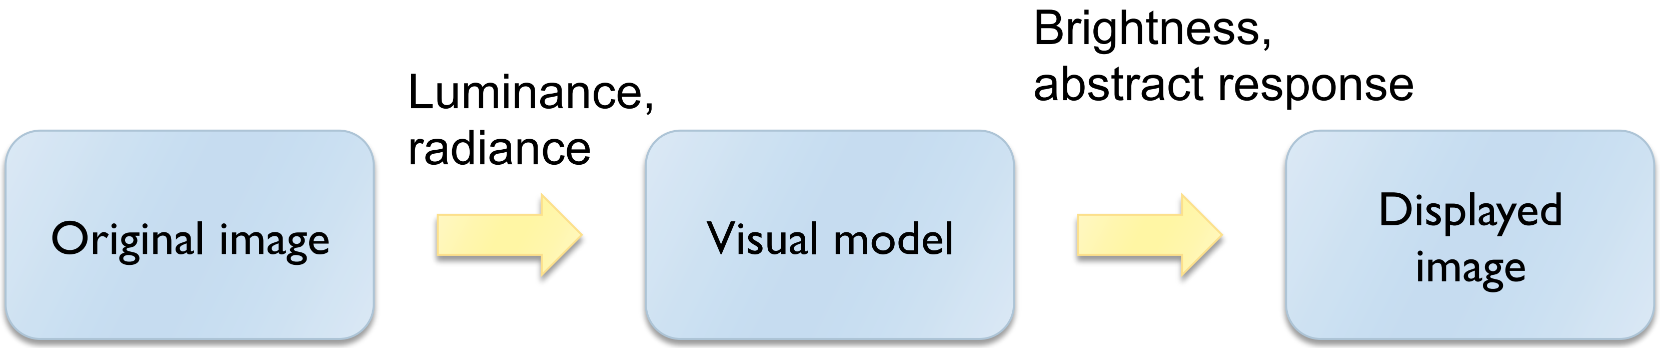

- Assumption: what is displayed is brightness or abstract response of the visual system
- Problem: double processing

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:

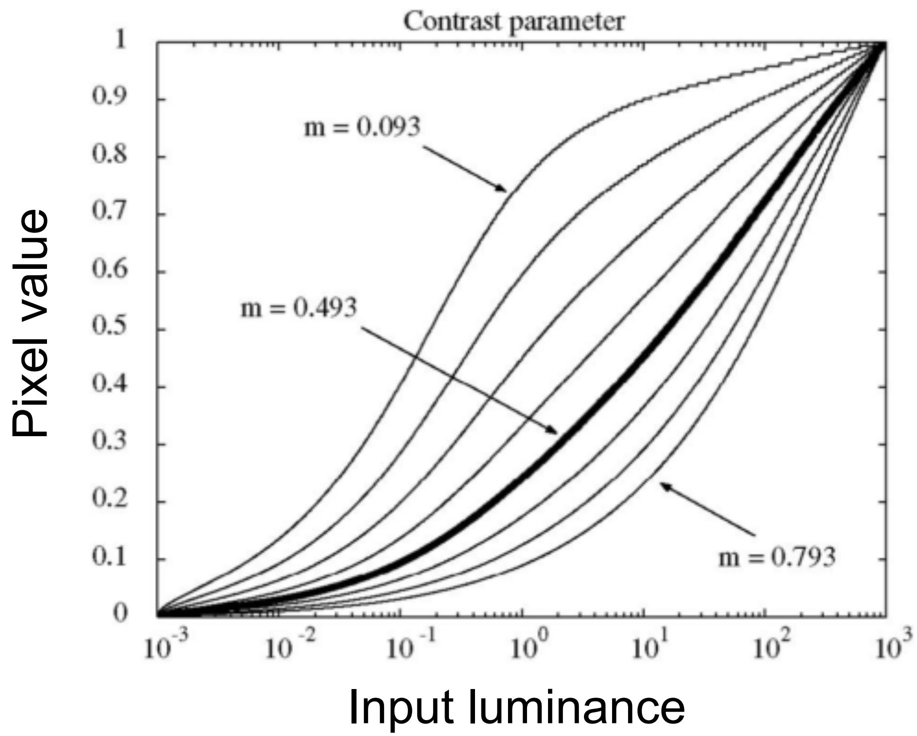

Photoreceptor TMO

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

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

Output pixel value (RGB)

Input radiance (RGB)

Maximum response (set it to 1)

Semi-saturation constant

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

Global/local adaptation luminance

Parameter (between 0 and 1)

Results: photoreceptor TMO

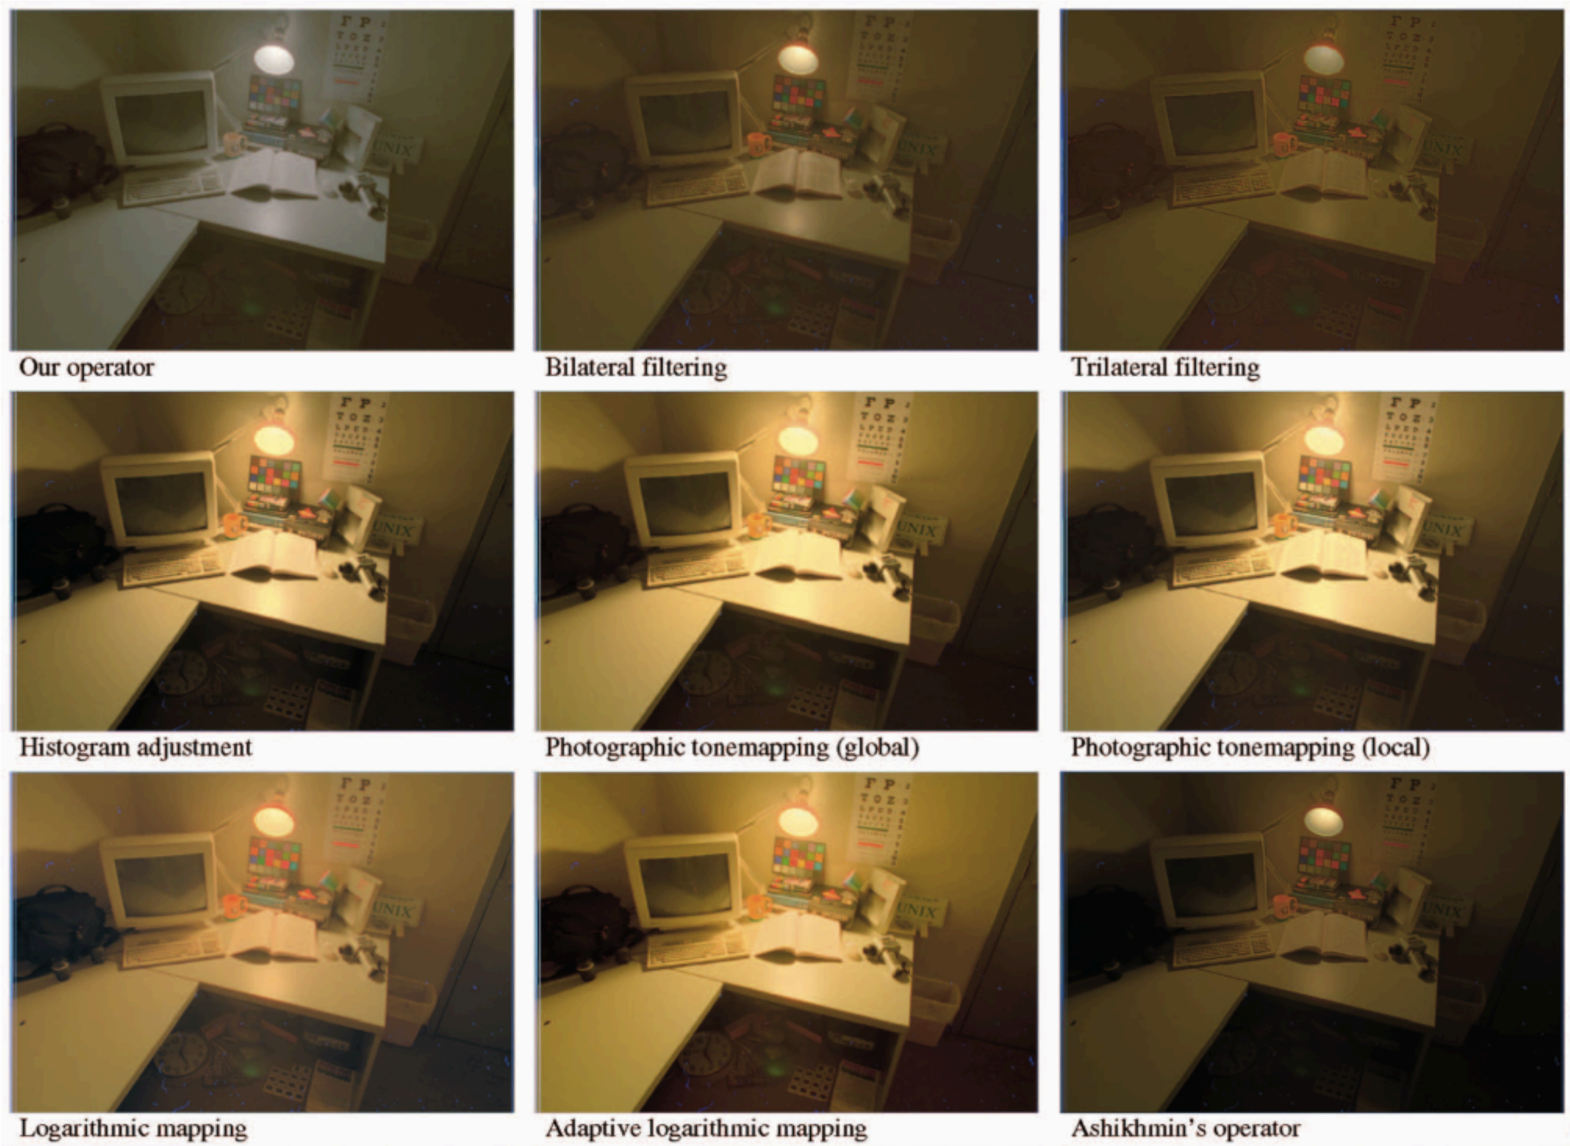

4 (major) approaches to tone mapping

- Illumination & reflectance separation
- Forward visual model
- Forward & inverse visual model
- Constrained mapping problem

Forward and inverse visual model

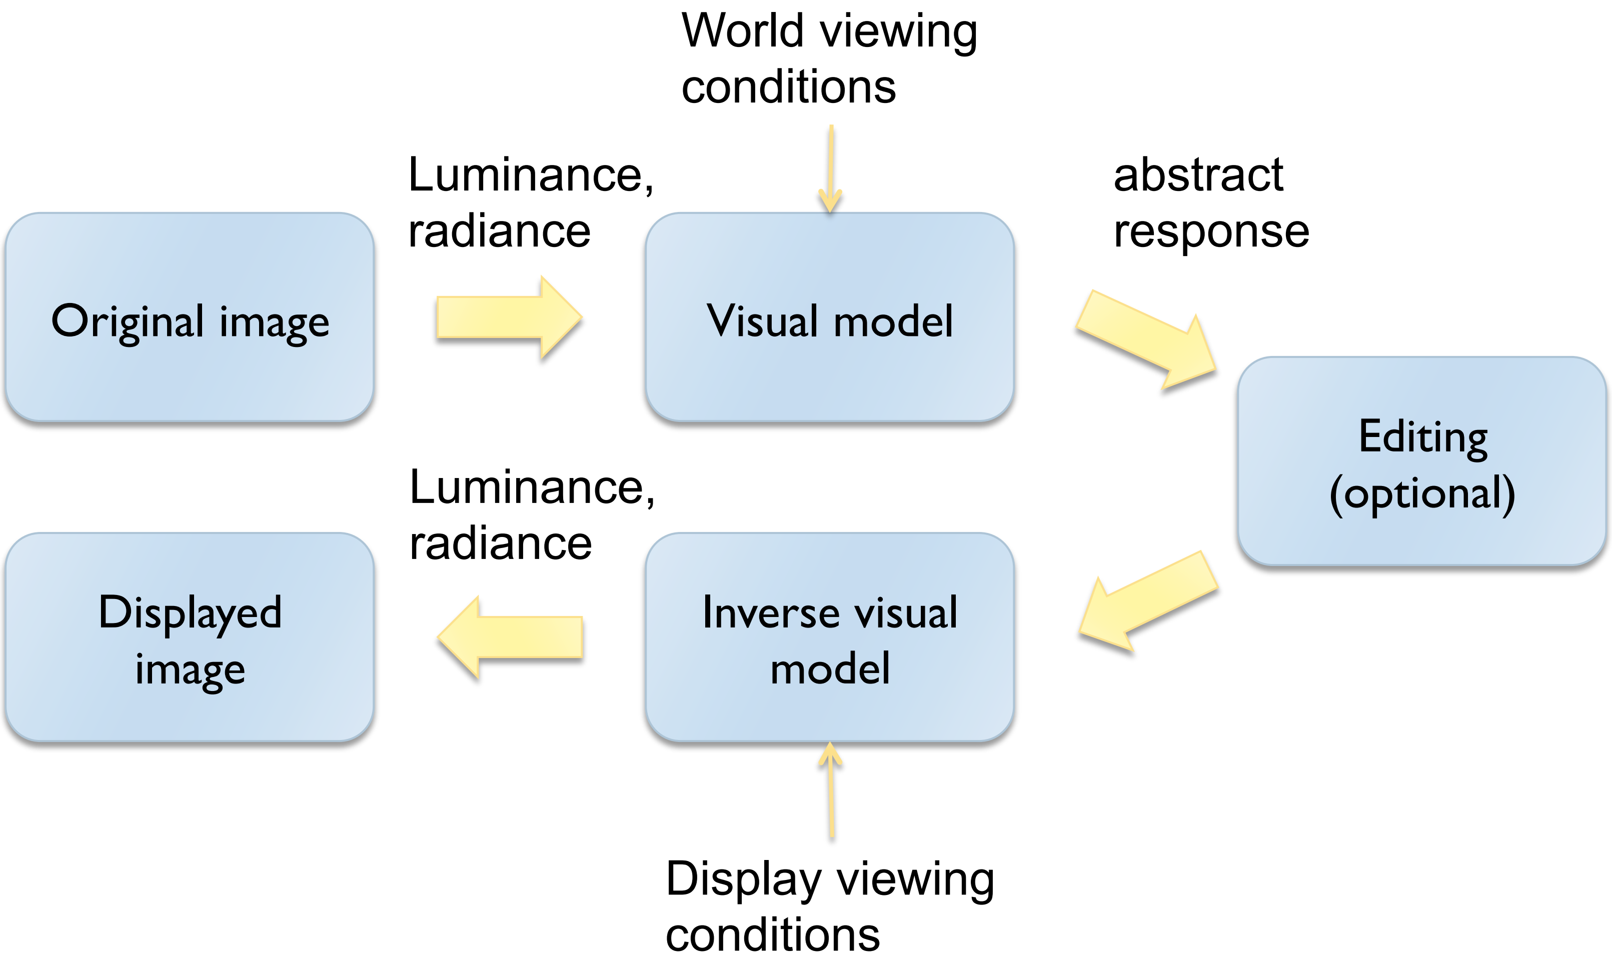

Multi-scale model

- 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

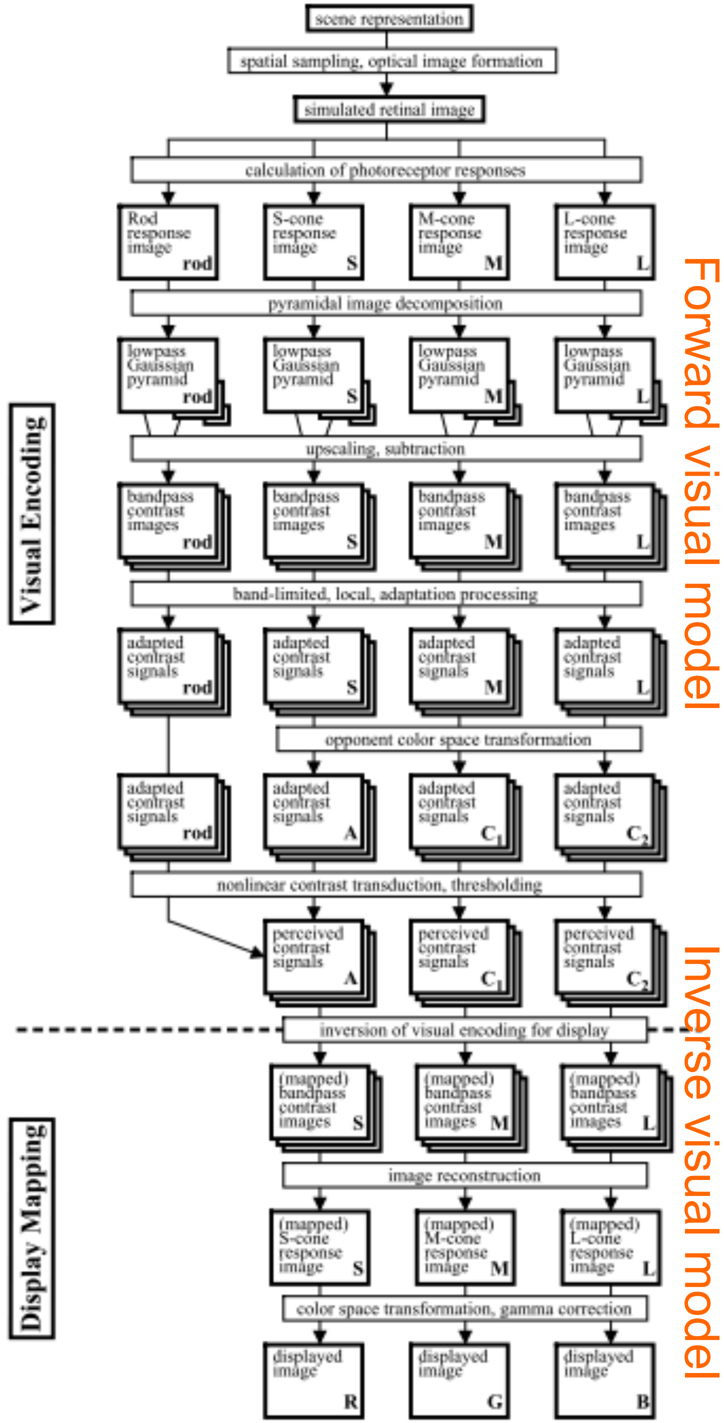

Forward and inverse visual model

- 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

4 (major) approaches to tone mapping

- Illumination & reflectance separation
- Forward visual model
- Forward & inverse visual model
- Constrained mapping problem

Global tone mapping operator

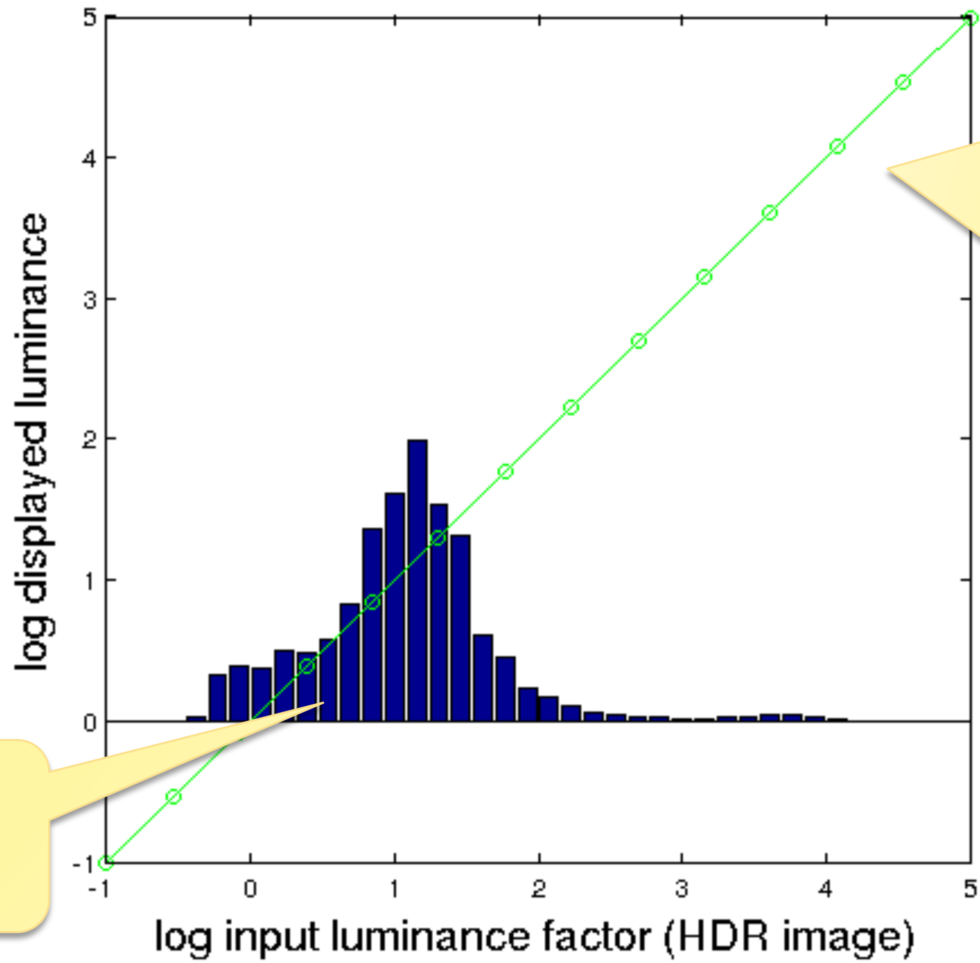

Best tone-mapping is the one which does not do anything, i.e. slope of the tone-mapping curves is equal to 1.

Image histogram

Display limitations

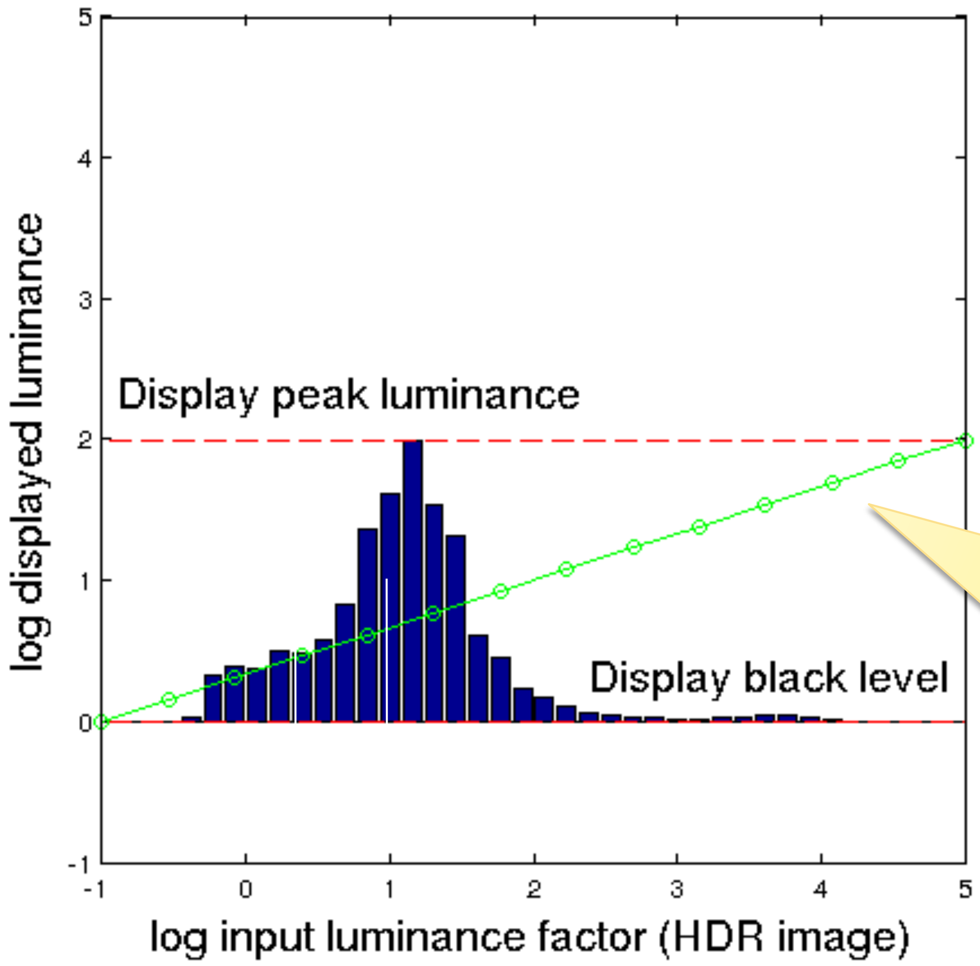

But in practice contrast (slope) must be limited due to display limitations.

Tone mapping

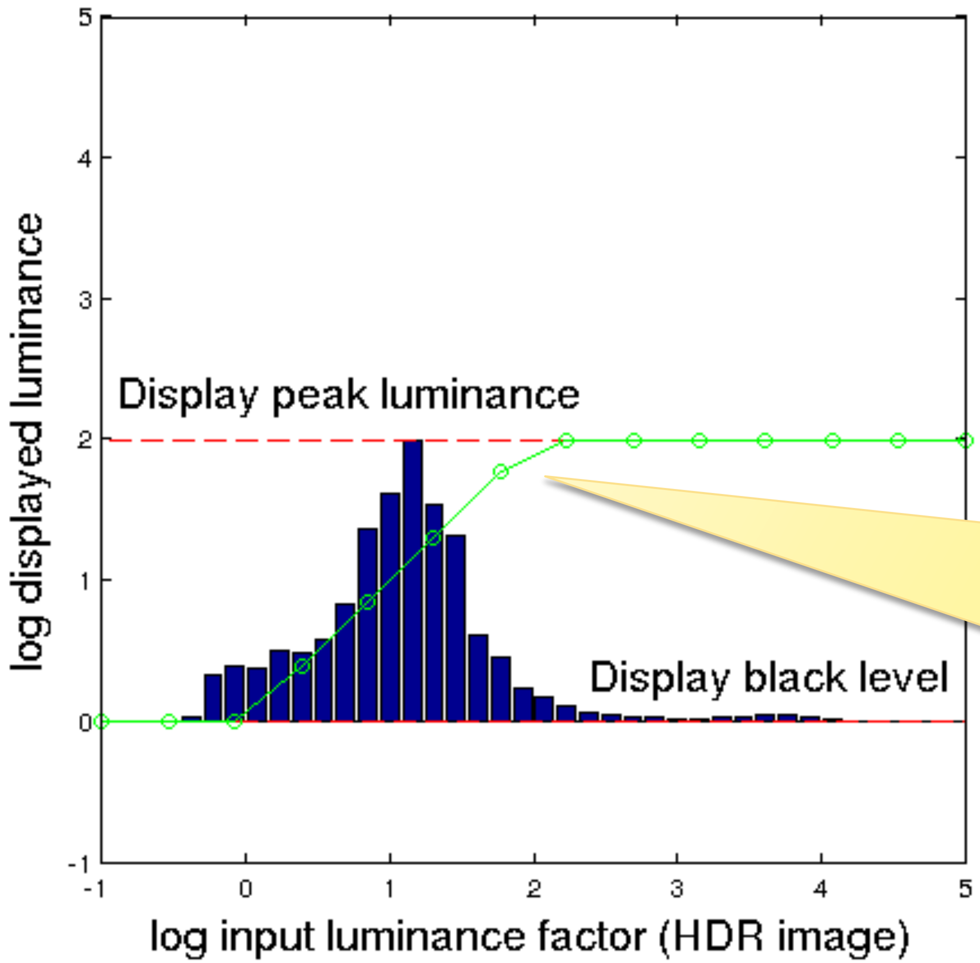

Global tone-mapping is a compromise between clipping and contrast compression.

Constrained mapping + visual models

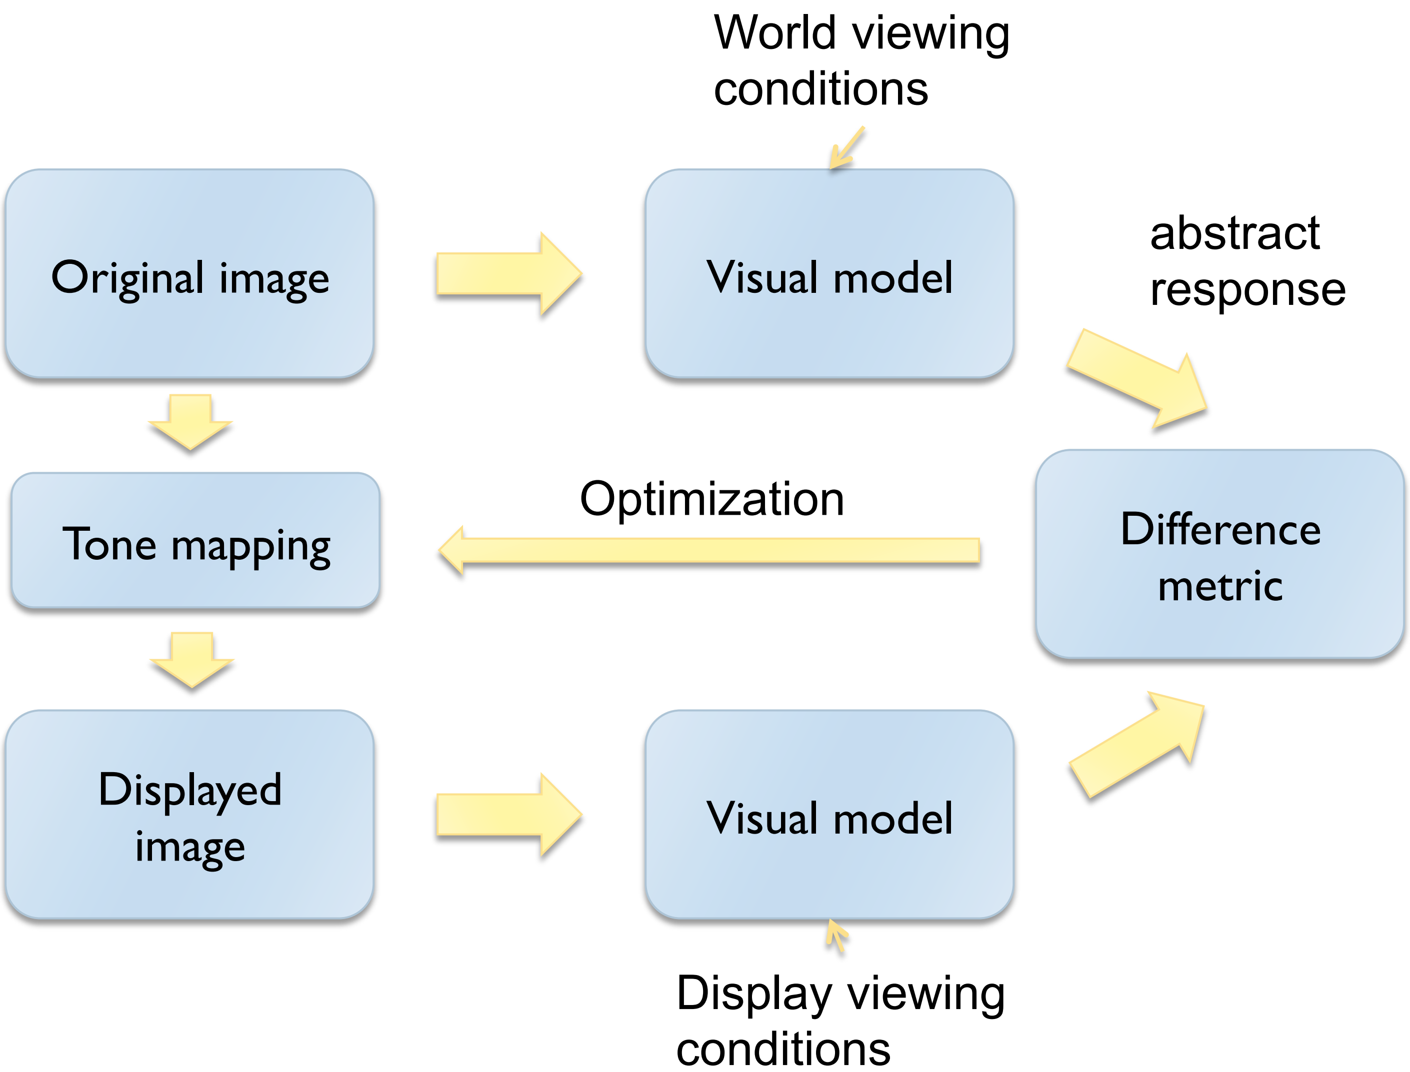

4 (major) approaches to tone mapping

- Illumination & reflectance separation
- Forward visual model
- Forward & inverse visual model
- Constrained mapping problem

Color

Motivation

To preserve the appearance of color after tone mapping

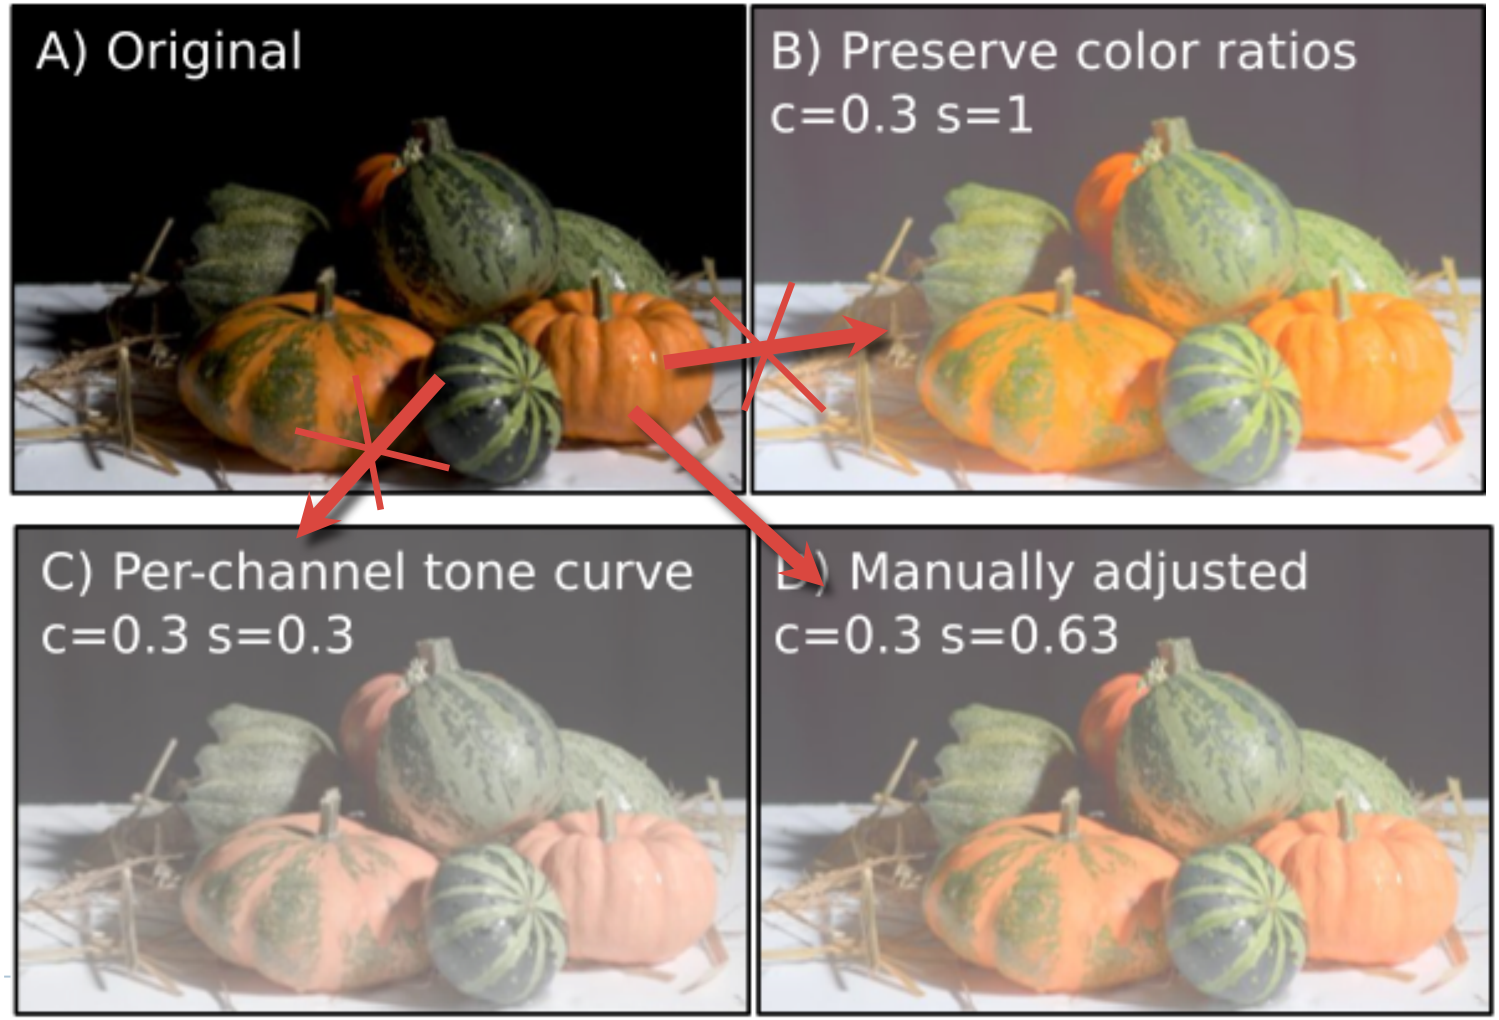

Color processing

- Existing methods
  - Apply the same processing to all color channels
    - works for simple (global) TMOs
  - **Operate on luminance, transfer color from the original HDR image**
  - Model color appearance

Color processing

- Transfer color from the original image

$$C_{out} = \left(\frac{C_{in}}{L_{in}} \right)^{s} \cdot L_{out}$$

Output color channel

Saturation parameter

Resulting luminance

- The heuristic from Fattal et al. 2002
  - works reasonably well in practice
- Difficulty:
  - How to select value ‘s’

Color correction for tone-mapping

Empirical formula for saturation correction
Estimated from experimental data

[Mantiuk et al., CGF 2009]

Global contrast compression (gamma)

$$s(c) = \frac{(1 + k_1)\, c^{k_2}}{1 + k_1\, c^{k_2}}$$

k1=2.3892, k2=0.8552

$$C_{out} = \left(\frac{C_{in}}{L_{in}} \right)^{s} L_{out}$$

What is wrong with the color-correction formula?

$$C_{out} = \left(\frac{C_{in}}{L_{in}} \right)^s \cdot L_{out}$$

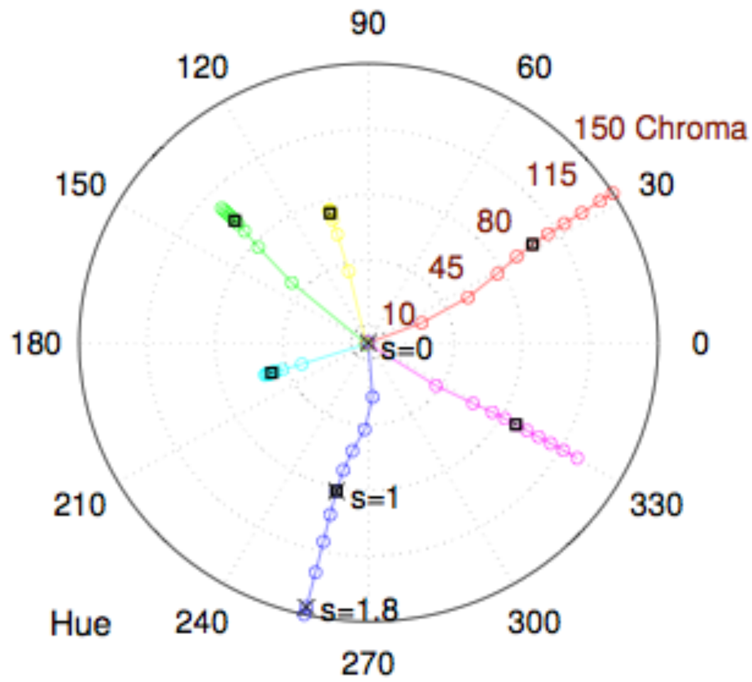

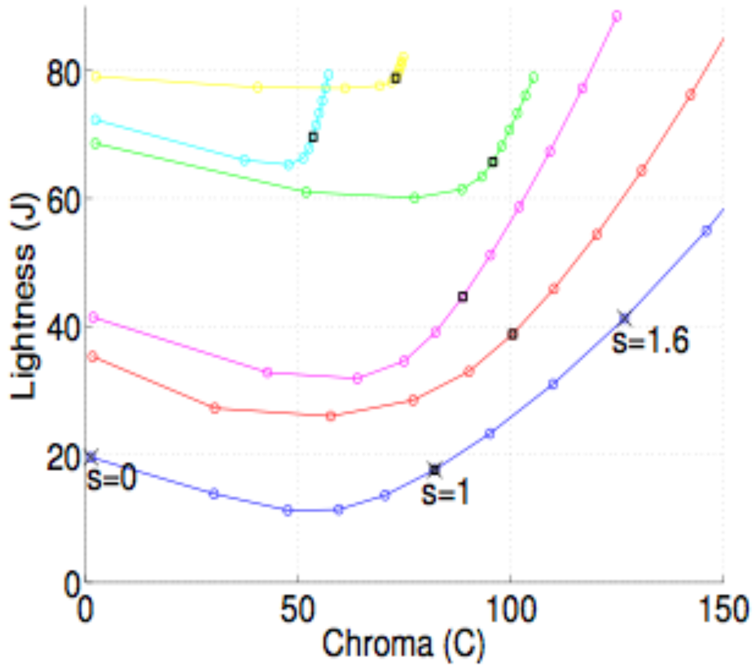

Visualization of CIECAM02 appearance correlates (hue=f(chroma), lightness=f(chroma)) for 6 sample colors.

Glare / blooming

Glare / blooming

From: Half Life 2: Lost Coast preview

Glare illusion

"Alan Wake" © Remedy Entertainment

Glare illusion

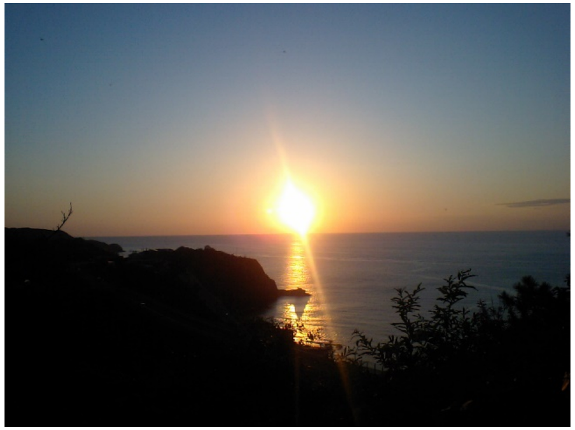

Photography

Painting

Computer Graphics
HDR rendering in games

Scattering of the light in the eye

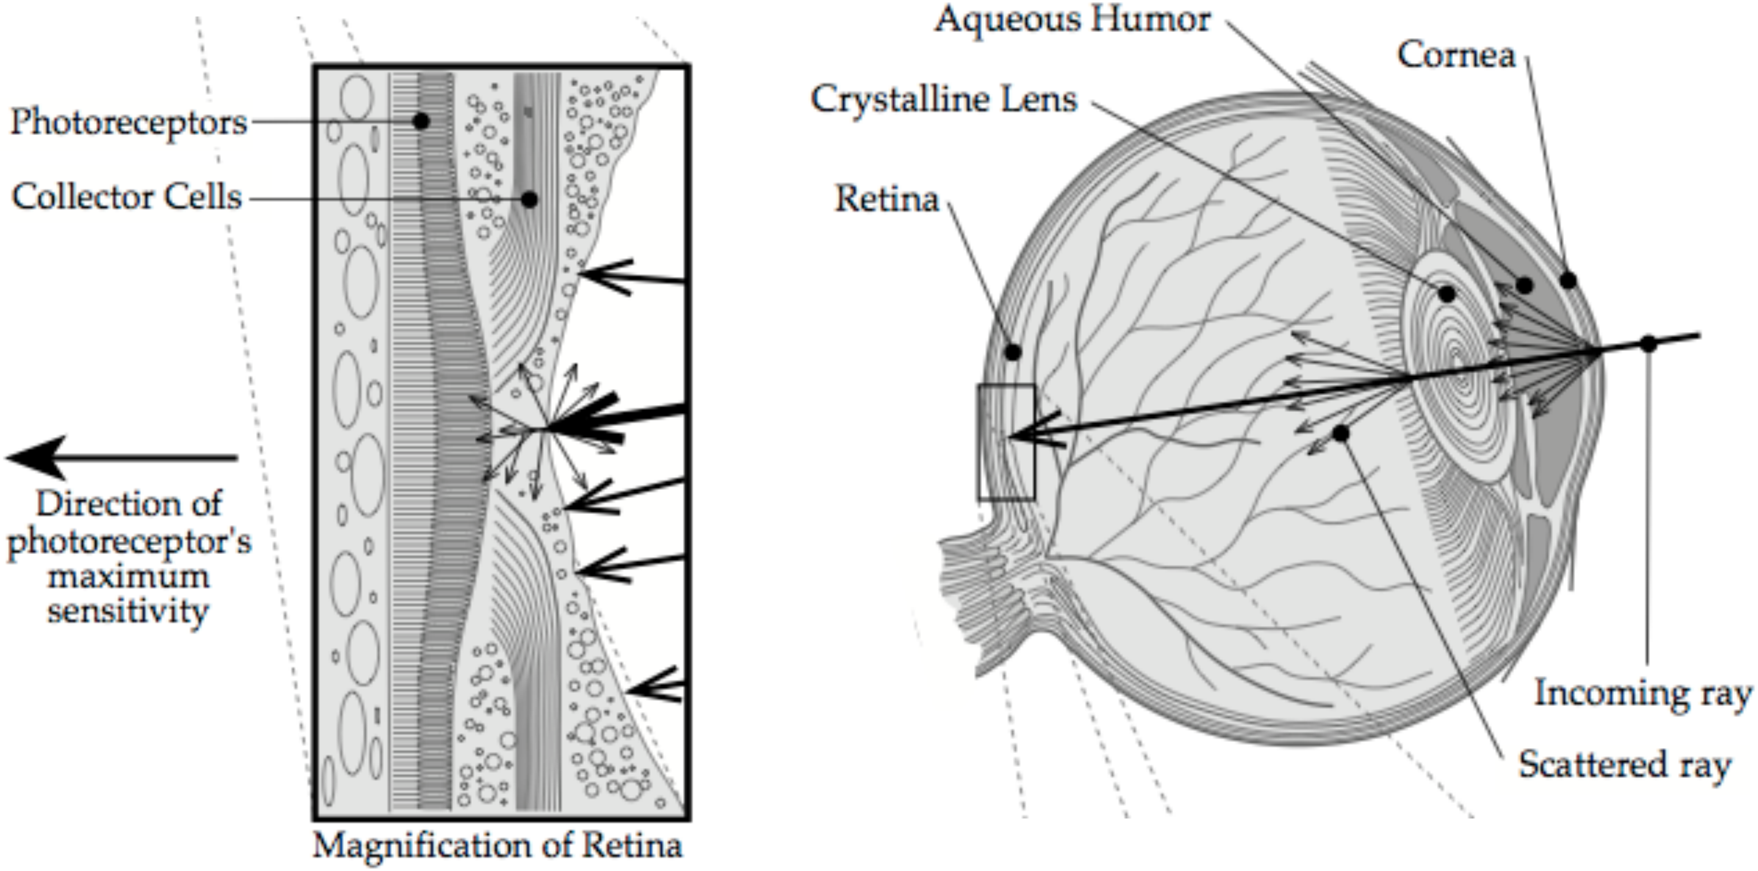

From: Sekuler, R., and Blake, R. Perception, second ed. McGraw- Hill, New York, 1990

Ciliary corona and lenticular halo

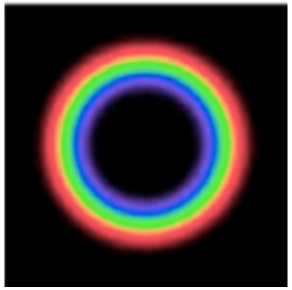

*

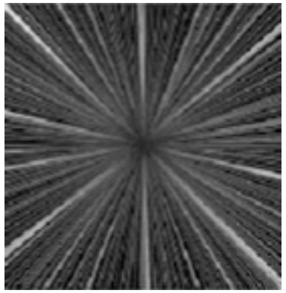

=

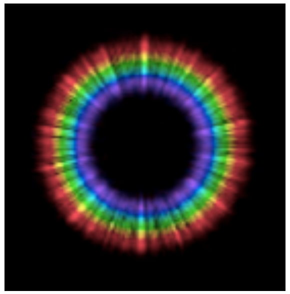

+

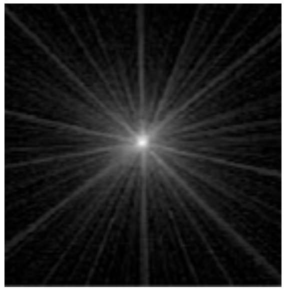

=

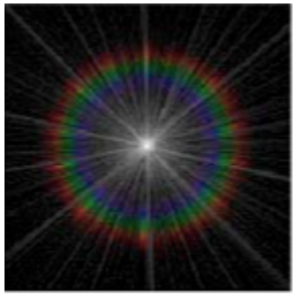

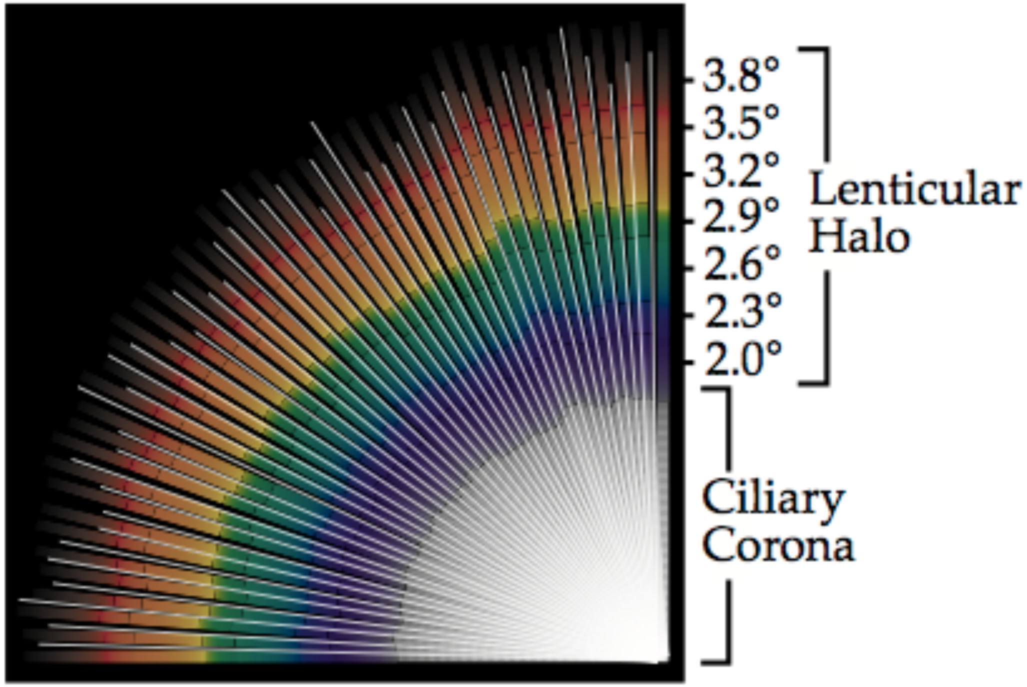

From: Spencer, G. et al. 1995. Proc. of SIGGRAPH. (1995)

Temporal model of glare (low level)

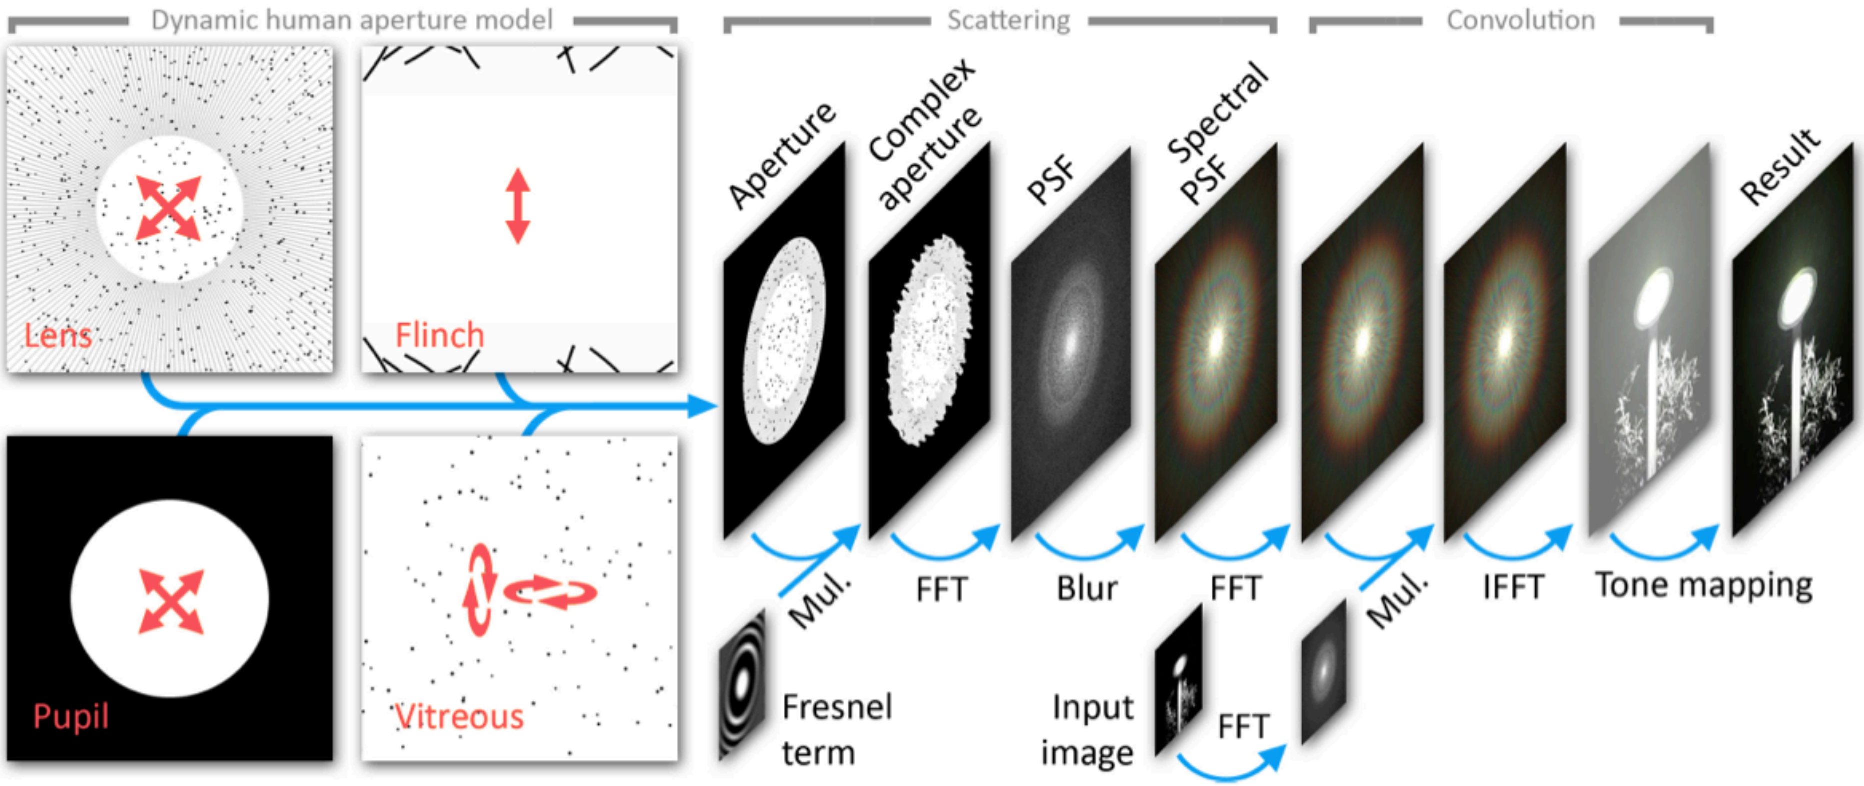

- The model assumes that glare is mostly caused by diffraction and scattering
- Can simulate temporal effects

[From Ritschel et al, Eurographics 2009]

Examples of simulated glare

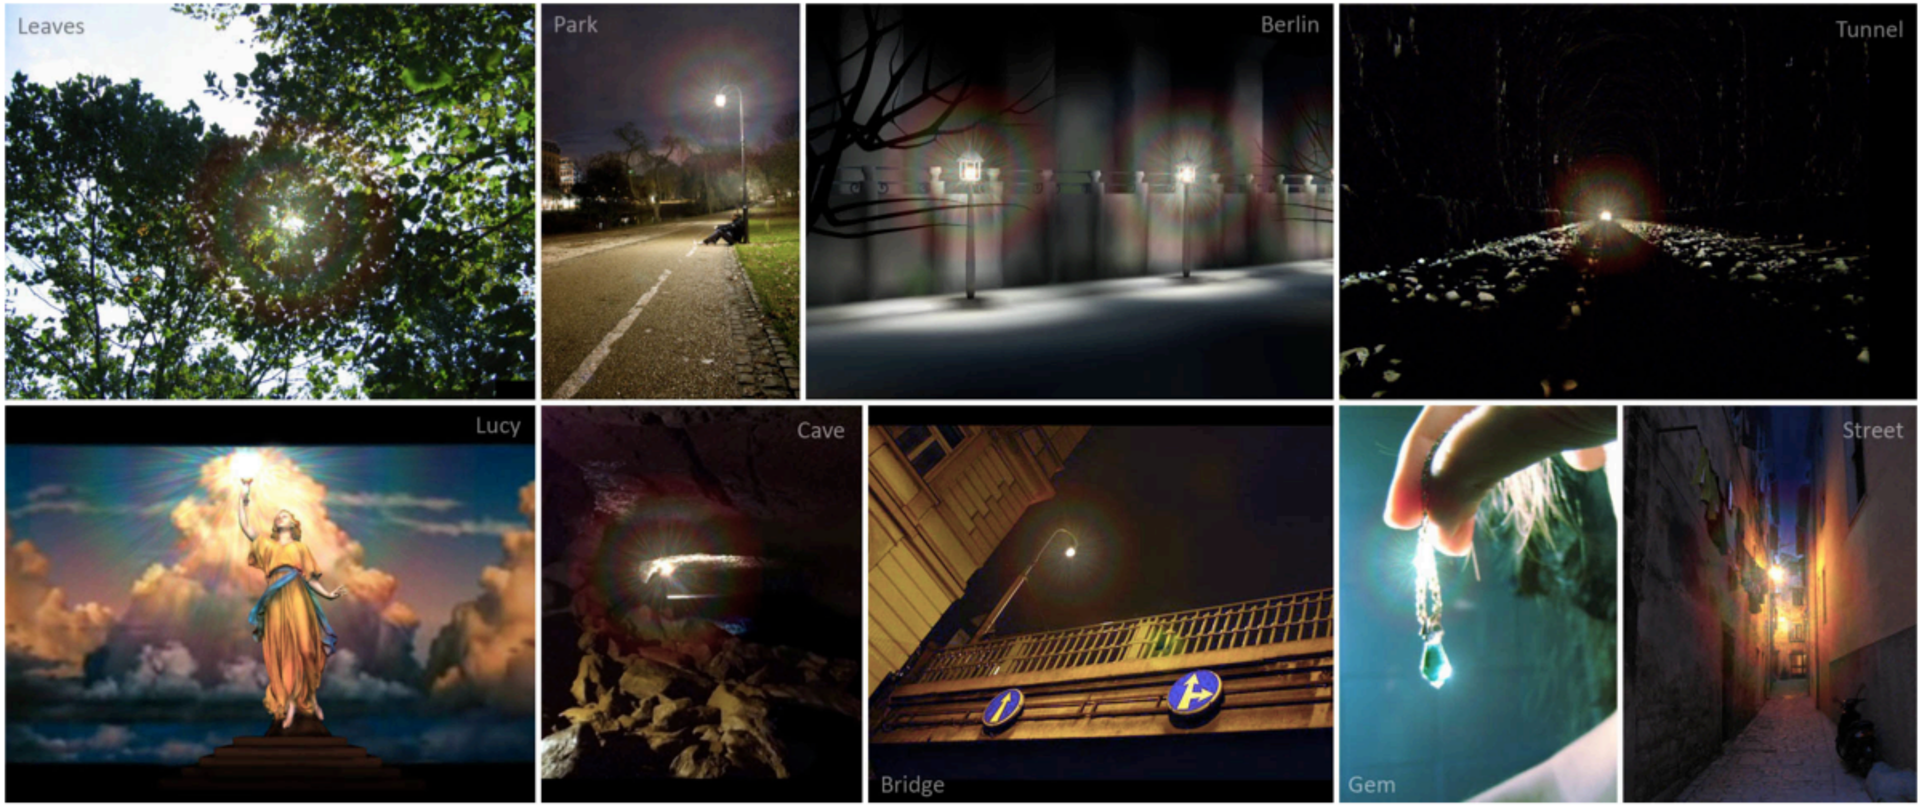

[From Ritschel et al, Eurographics 2009]

Temporal glare

[From Ritschel et al, Eurographics 2009]

Point Spread Function of the eye

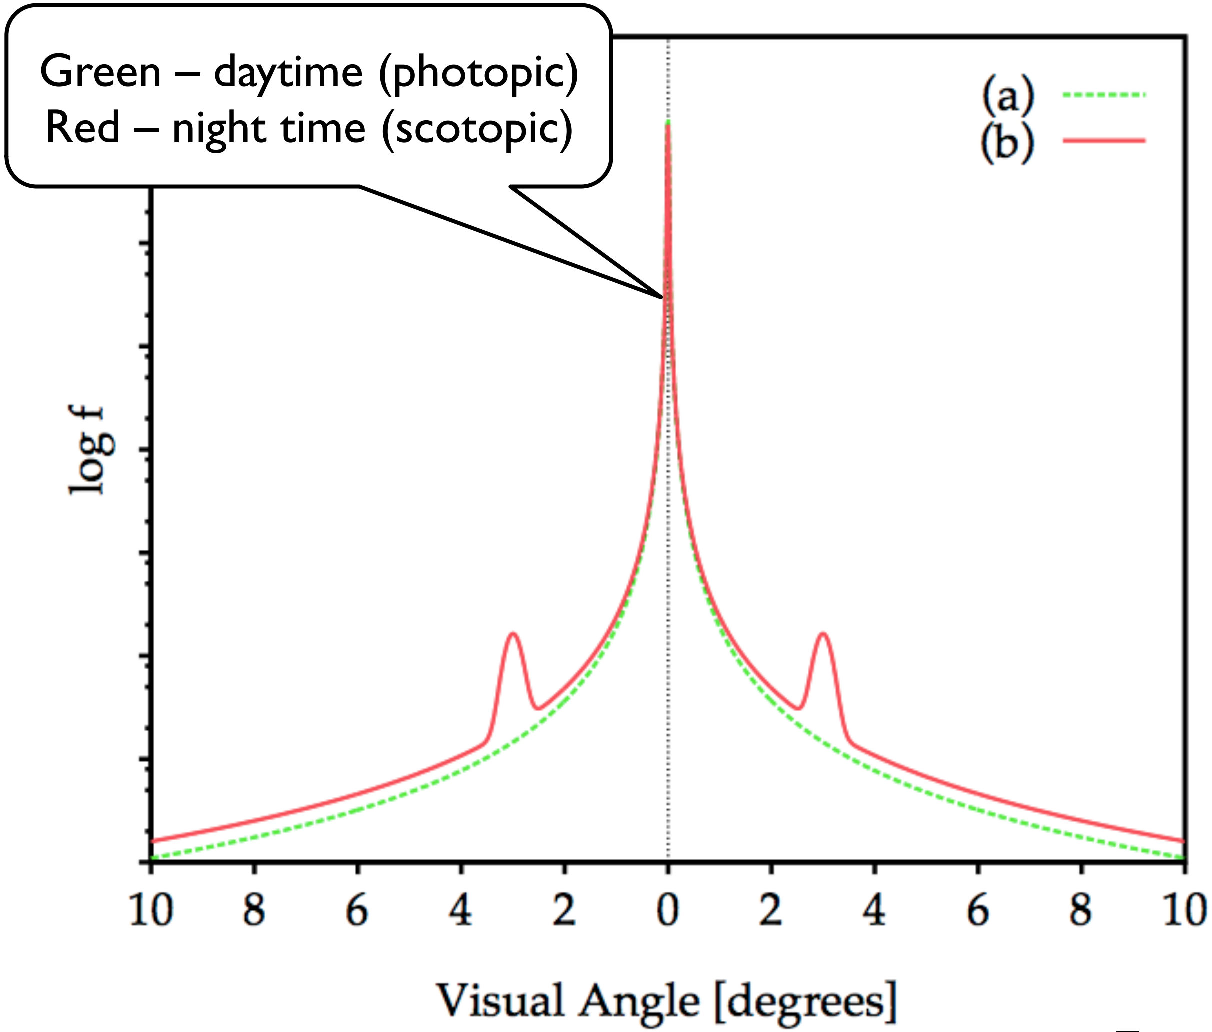

- What portion of the light is scattered towards a certain visual angle
- To simulate:
  - construct a digital filter
  - convolve the image with that filter

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

The problem of double processing

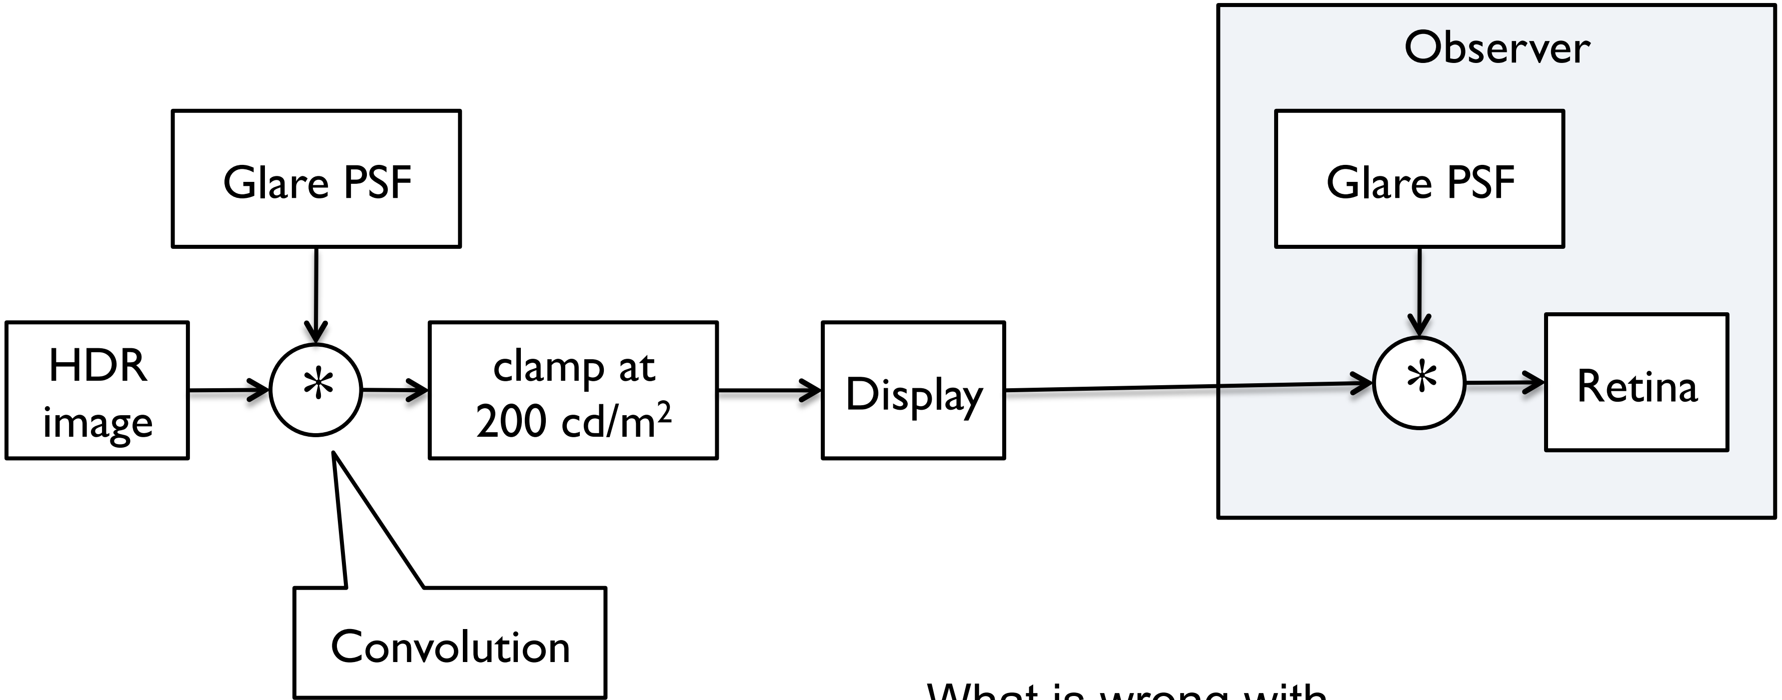

What is wrong with simulating glare this way?

The problem of double processing

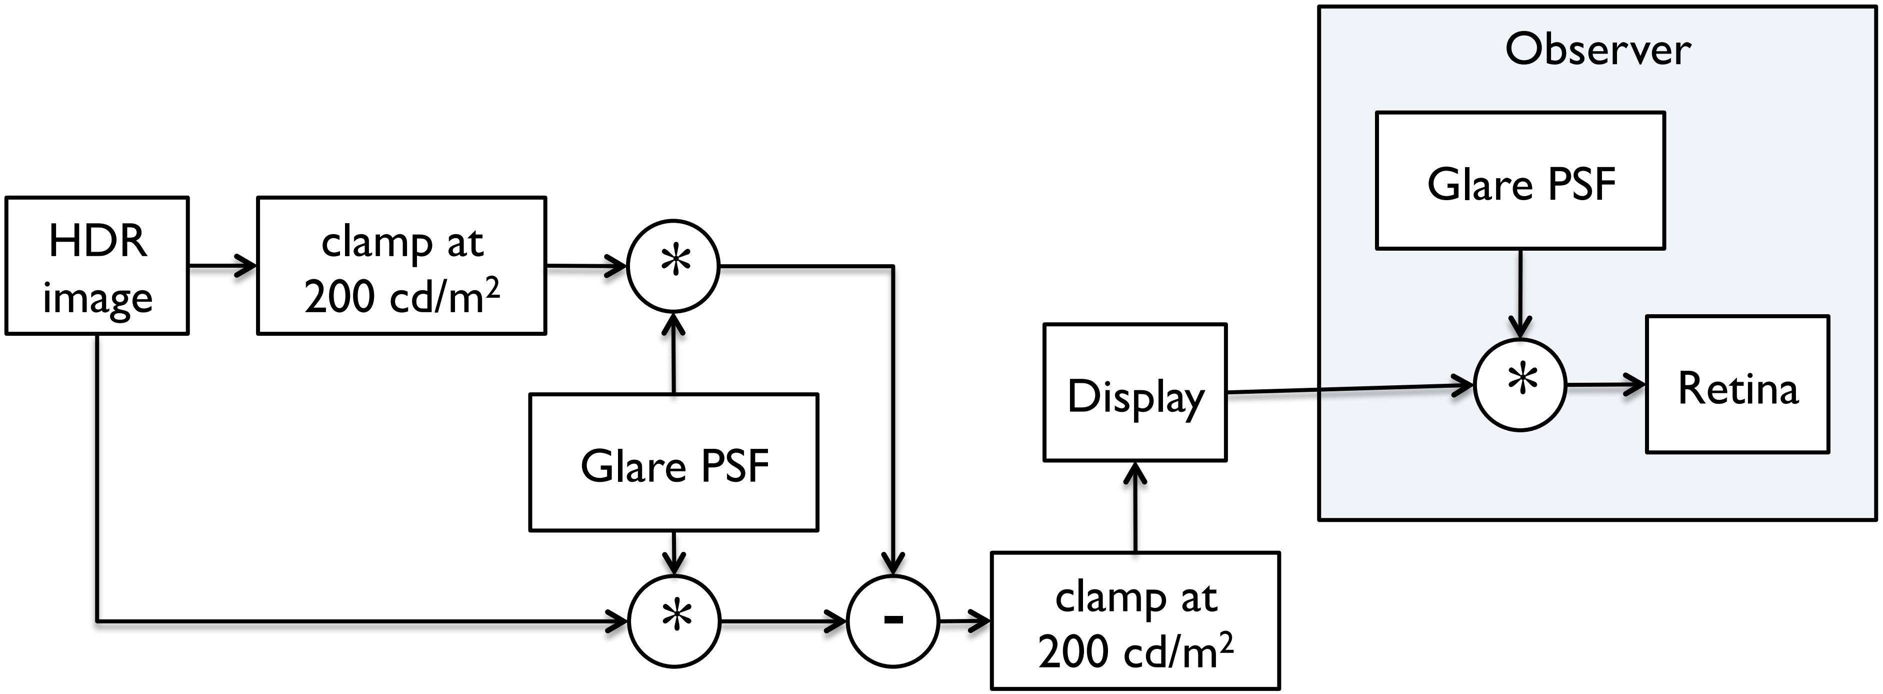

How does the diagram above avoid the problem of double processing?
Write down the operations as equations.
Can the processing be simplified?

PSF vs. OTF (Optical Transfer Function)

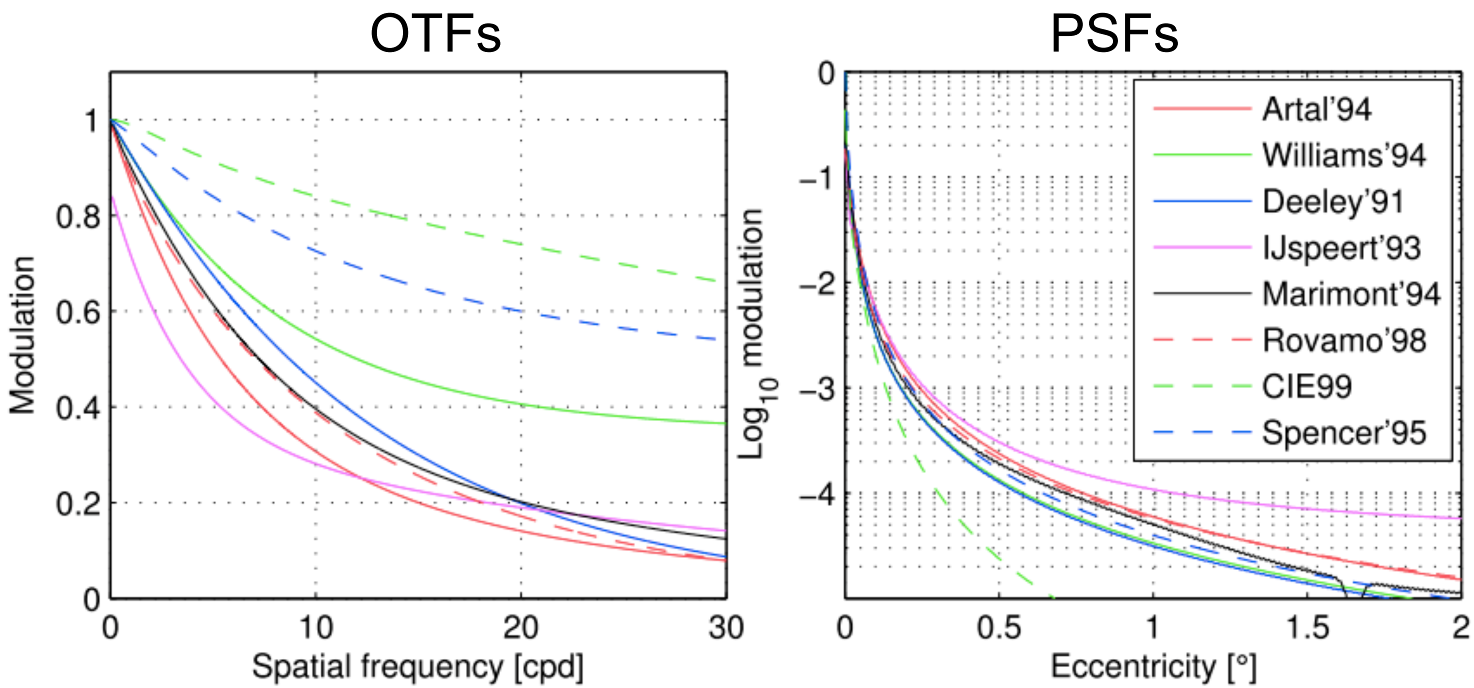

- An OTF is the Fourier transform of a PSF
- Convolution with larger kernels is faster in the Fourier domain

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

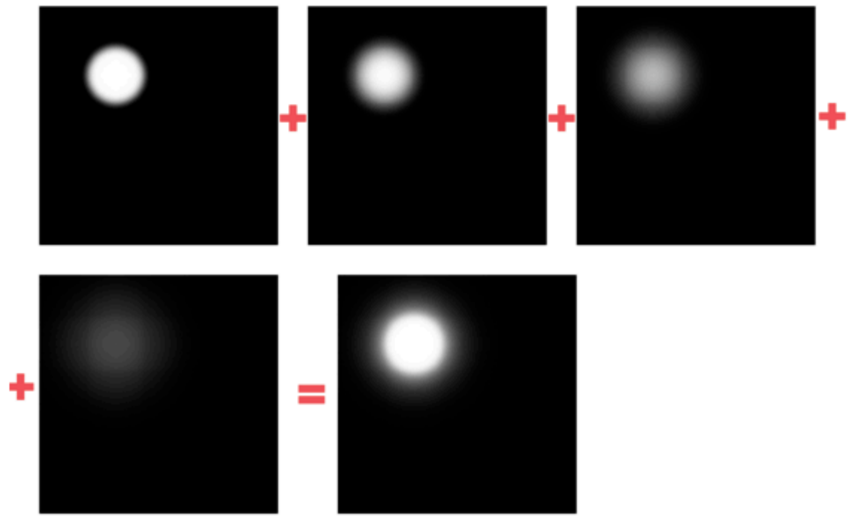

Does the exact shape of the PSF matter?

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

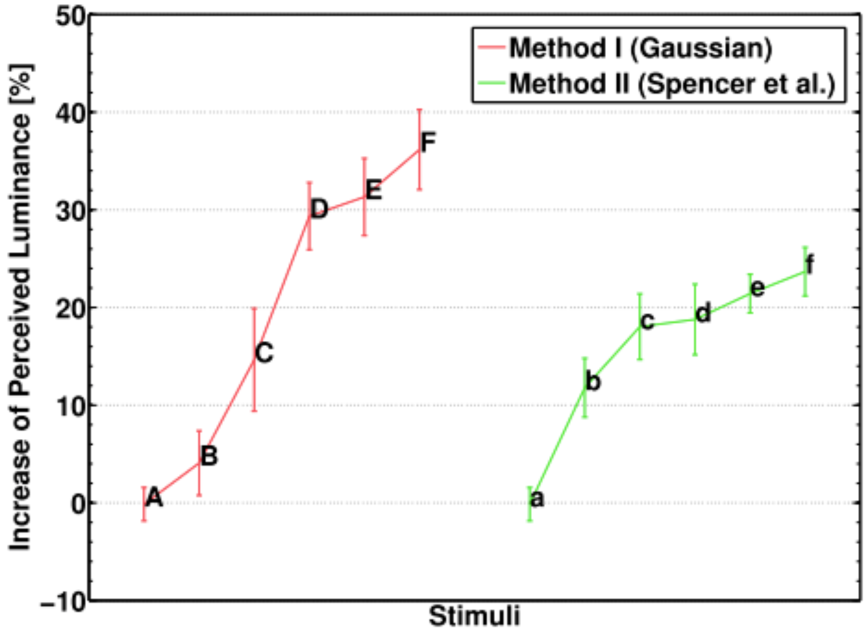

red - Gaussian

green - accurate

[Yoshida et al., APGV 2008]

Visual metric-driven TMOs

Display adaptive tone-mapping

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

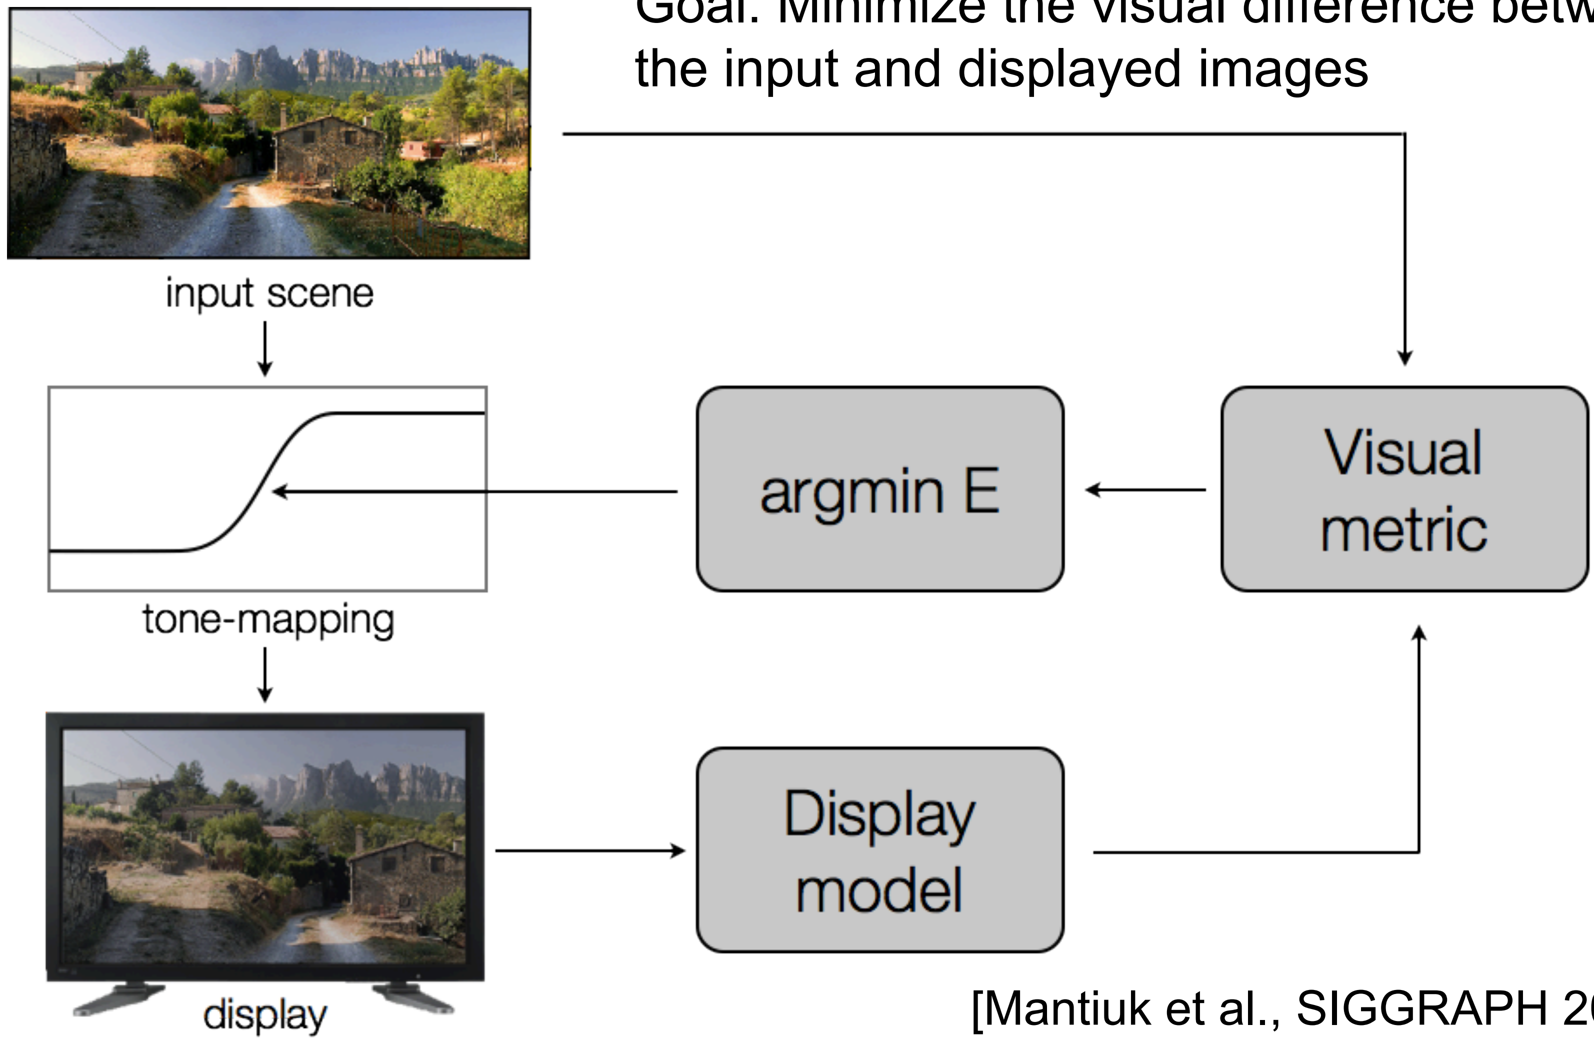

[Mantiuk et al., SIGGRAPH 2008]

Visual metric

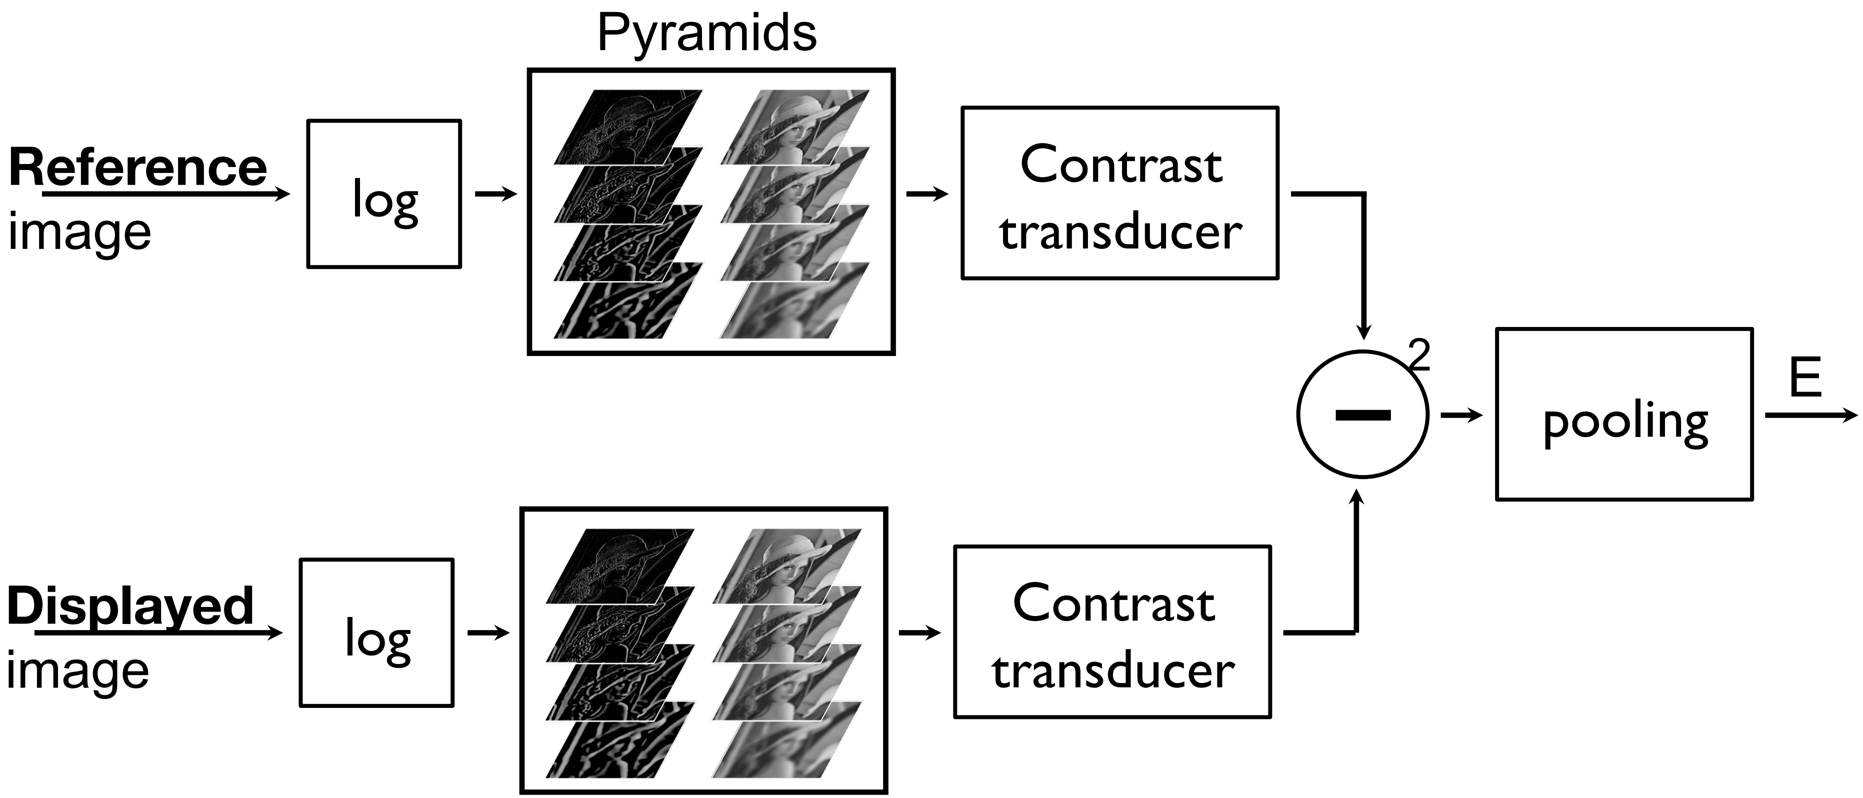

Visual metric

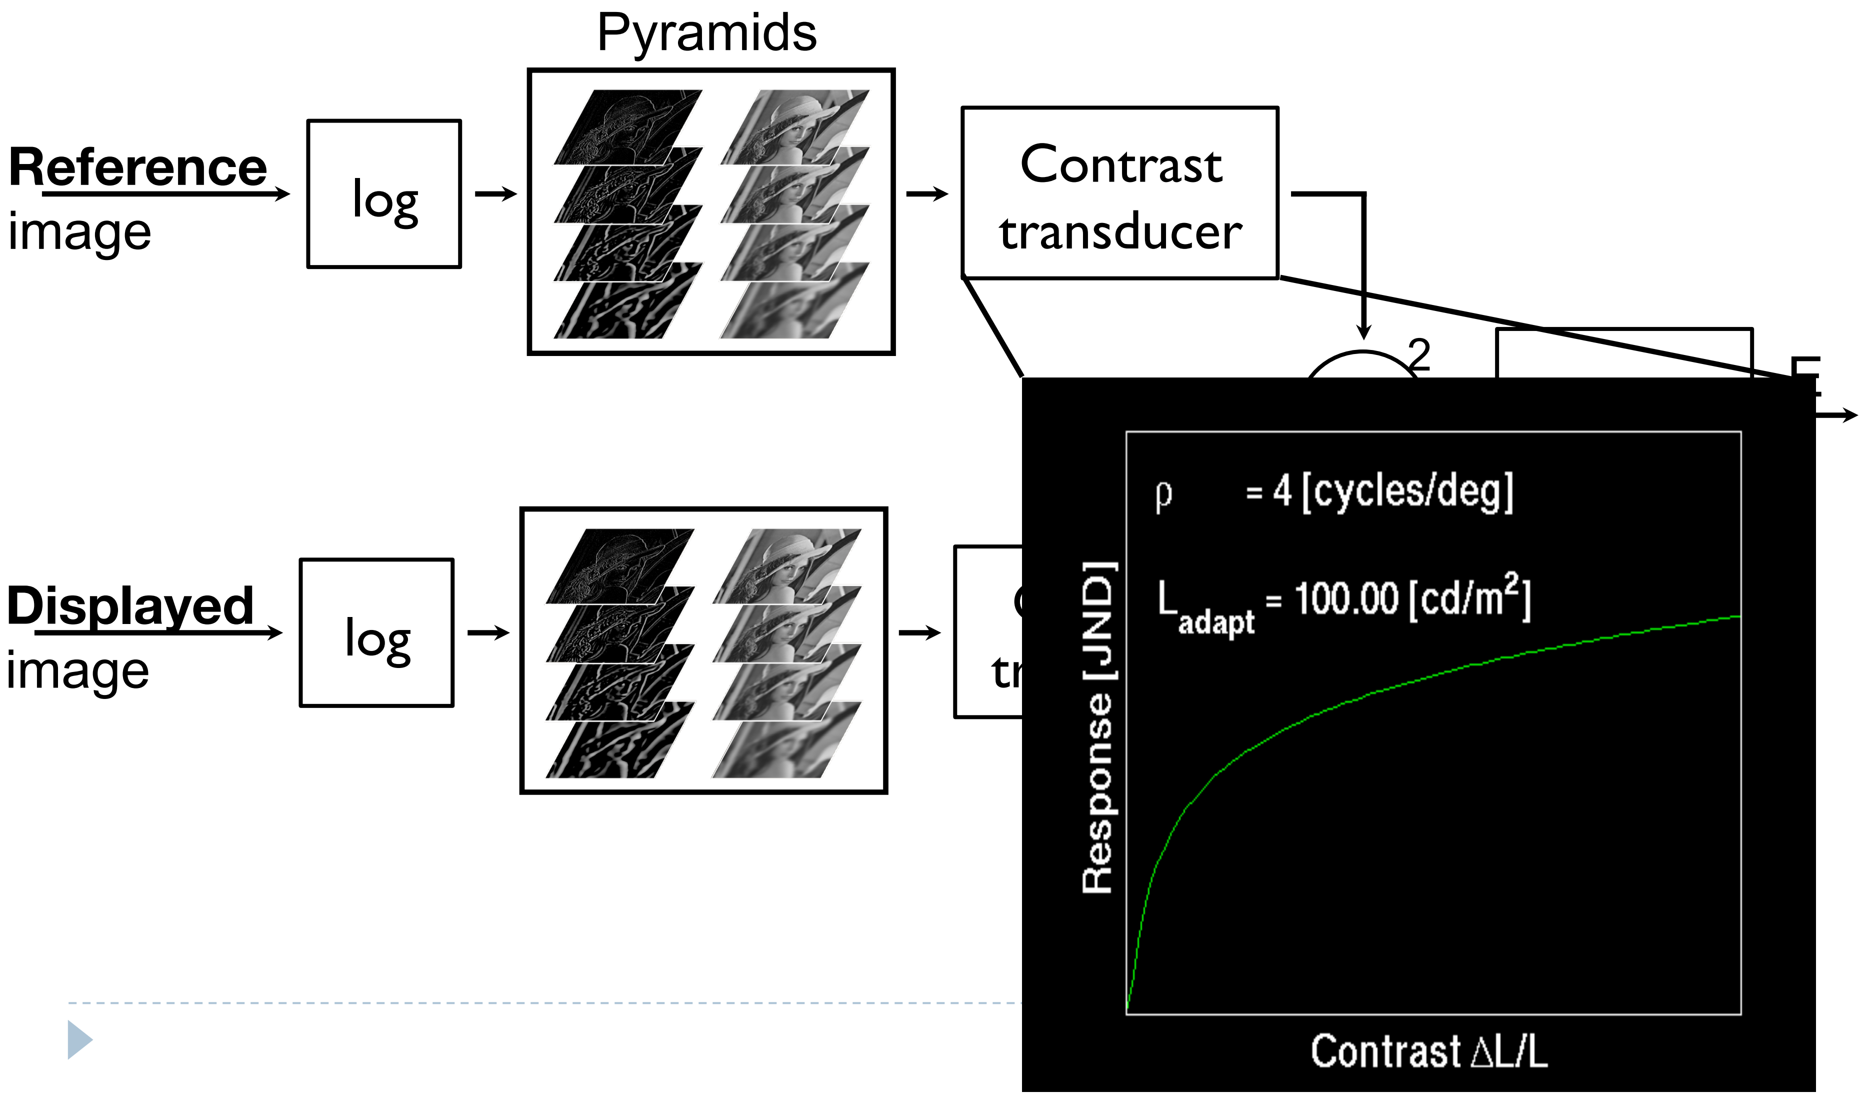

To accelerate: create a statistical model of an image

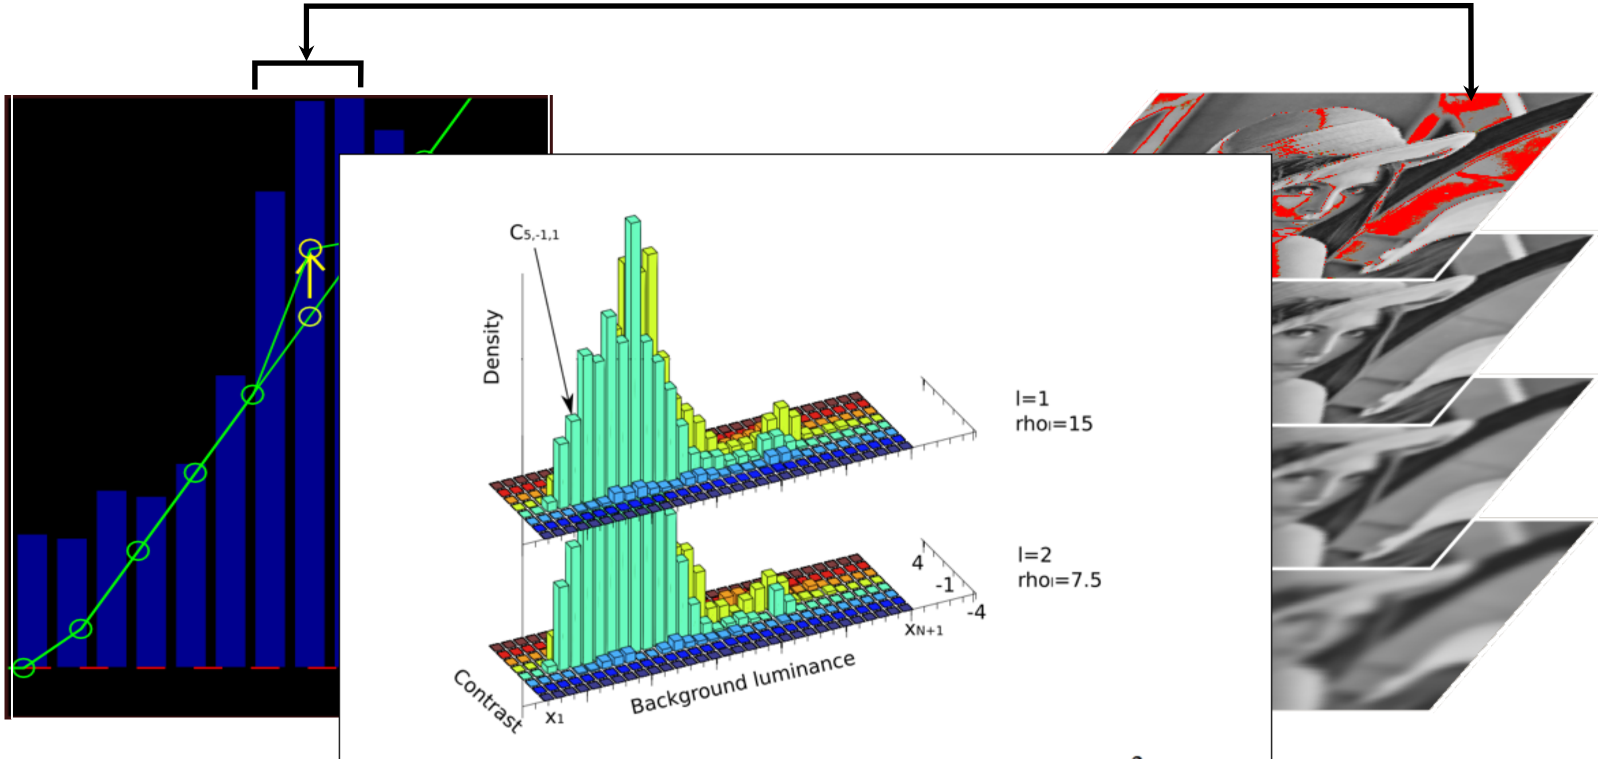

$$\underset{d_1,..,d_{N-1}}{\arg\min} = \sum_{l} \sum_{i=1}^{N-1} \sum_{\substack{m=-M \\ m\neq 0}}^{M} \left[T\left(\sum_{k\in\phi} d_k, S_d\right) - T\left(e \sum_{k\in\phi} \delta, S_r\right) \right]^2 \cdot c_{i,m,l}$$

Quadratic programming problem - 30-40 variables

Ambient illumination compensation

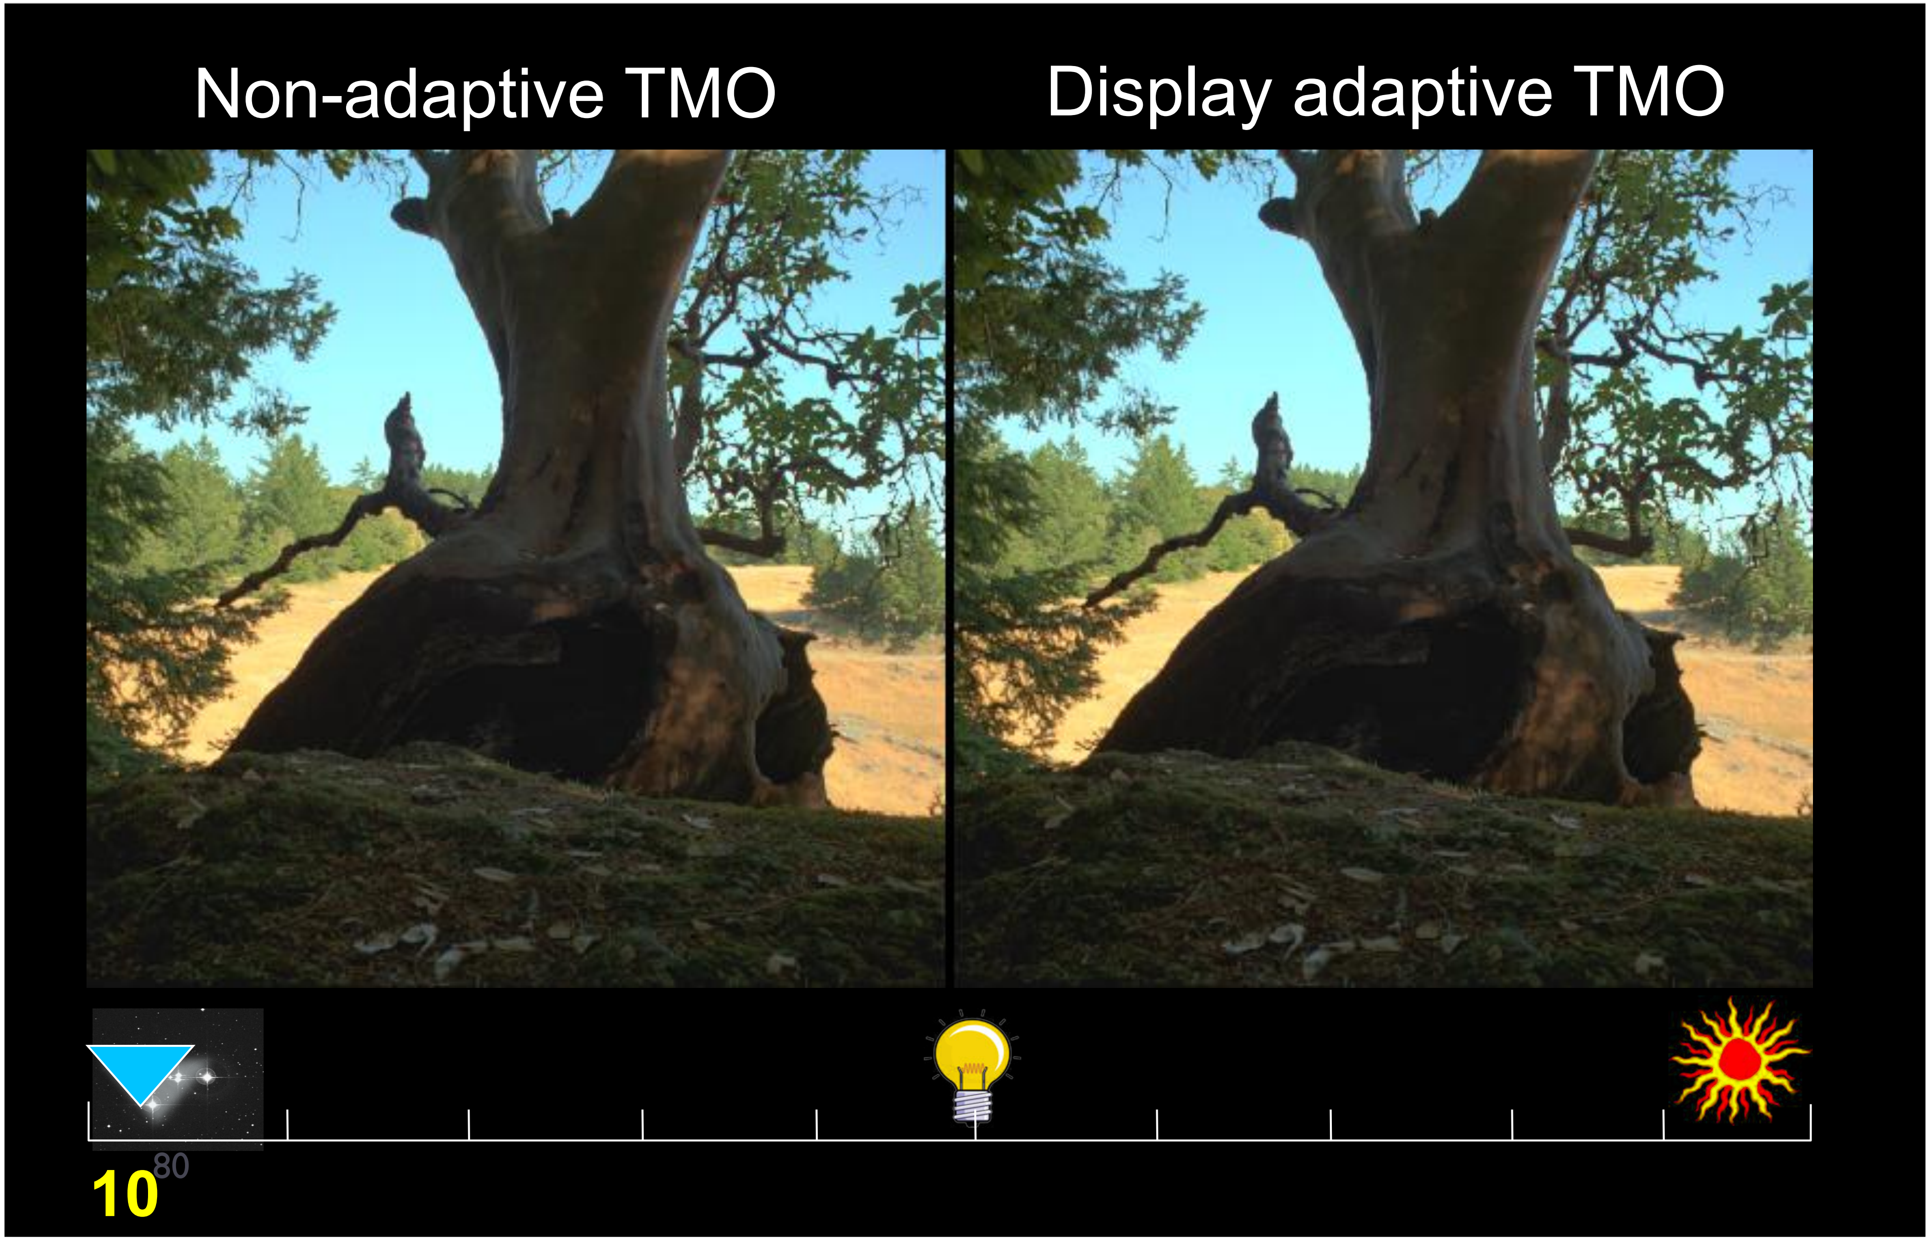

Ambient illumination compensation

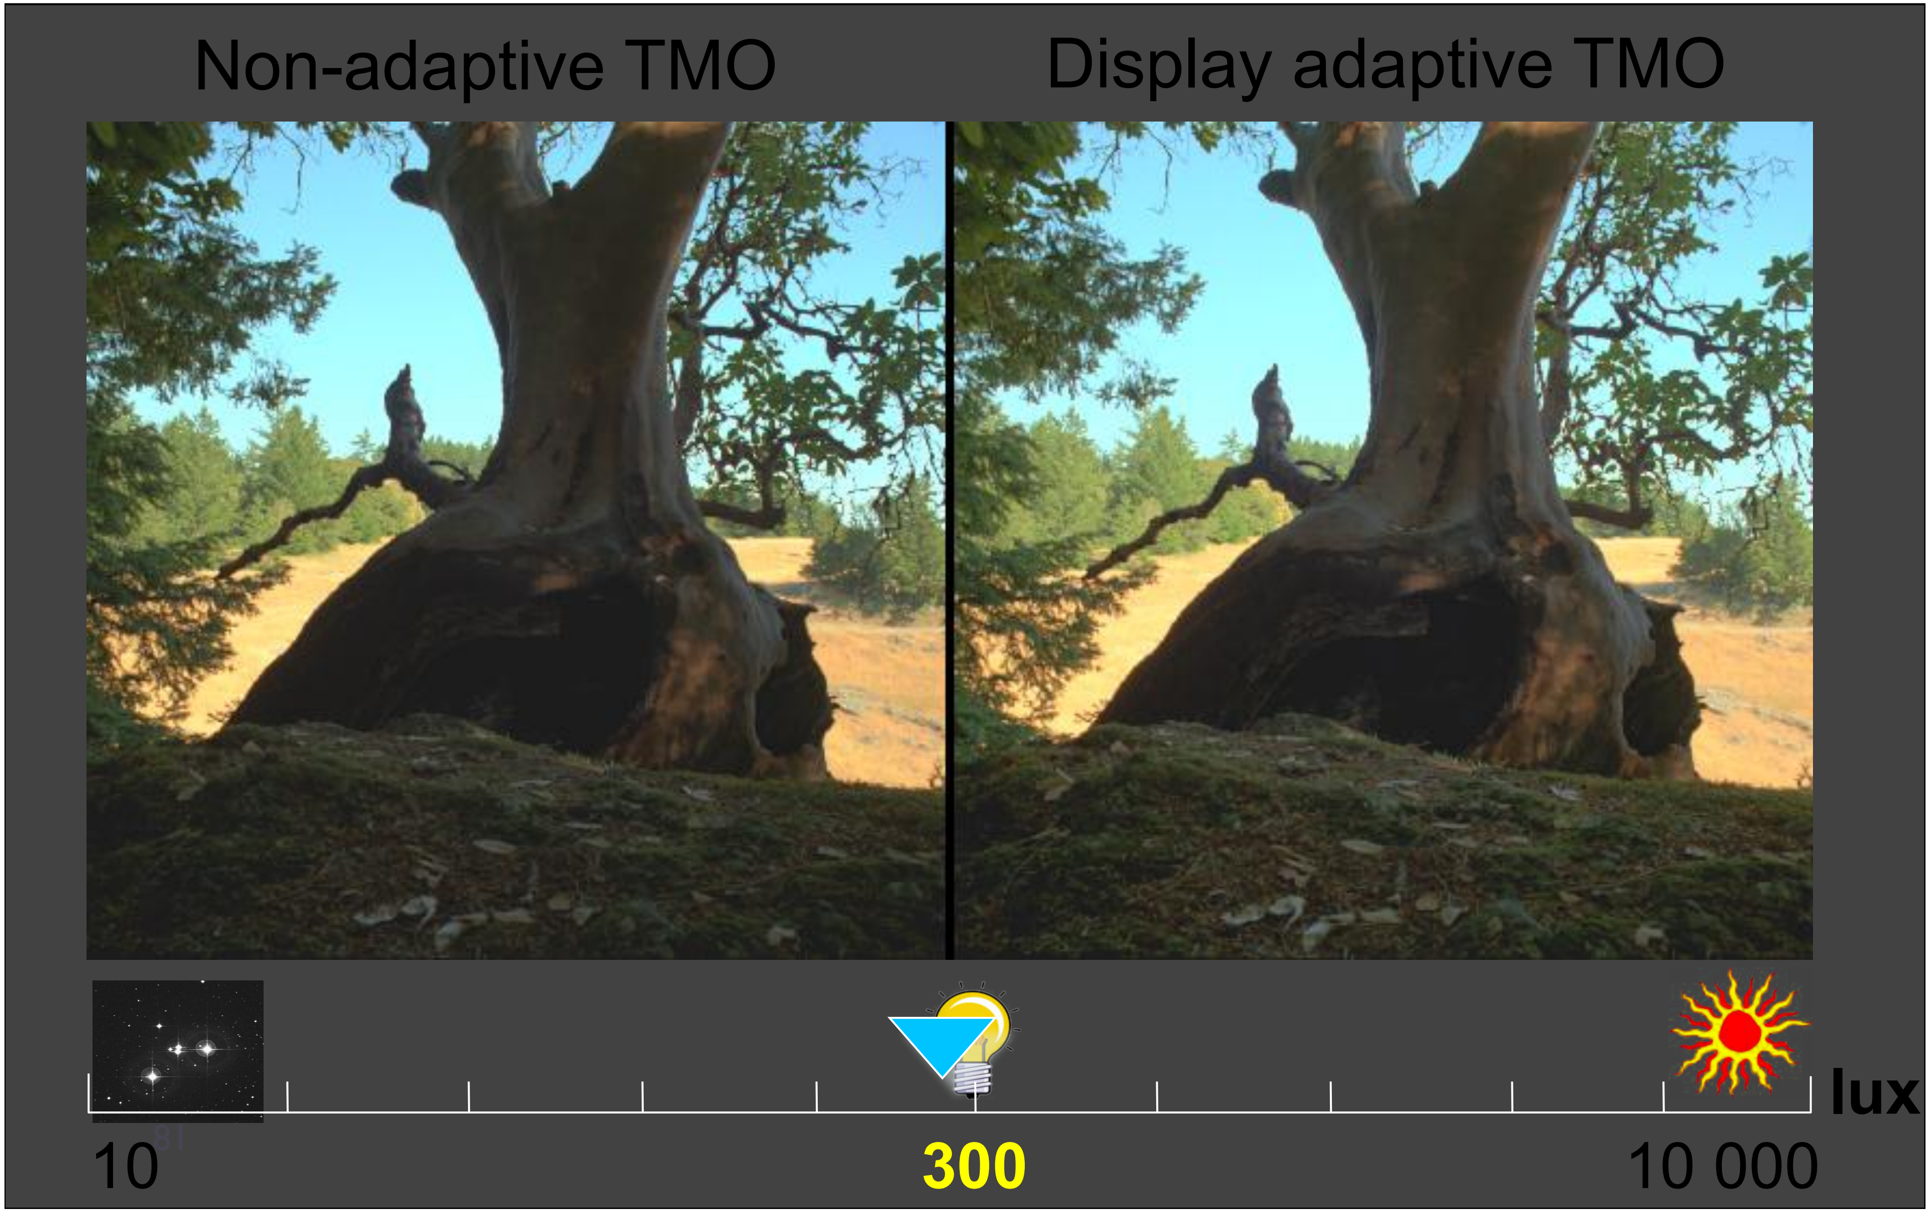

Ambient illumination compensation

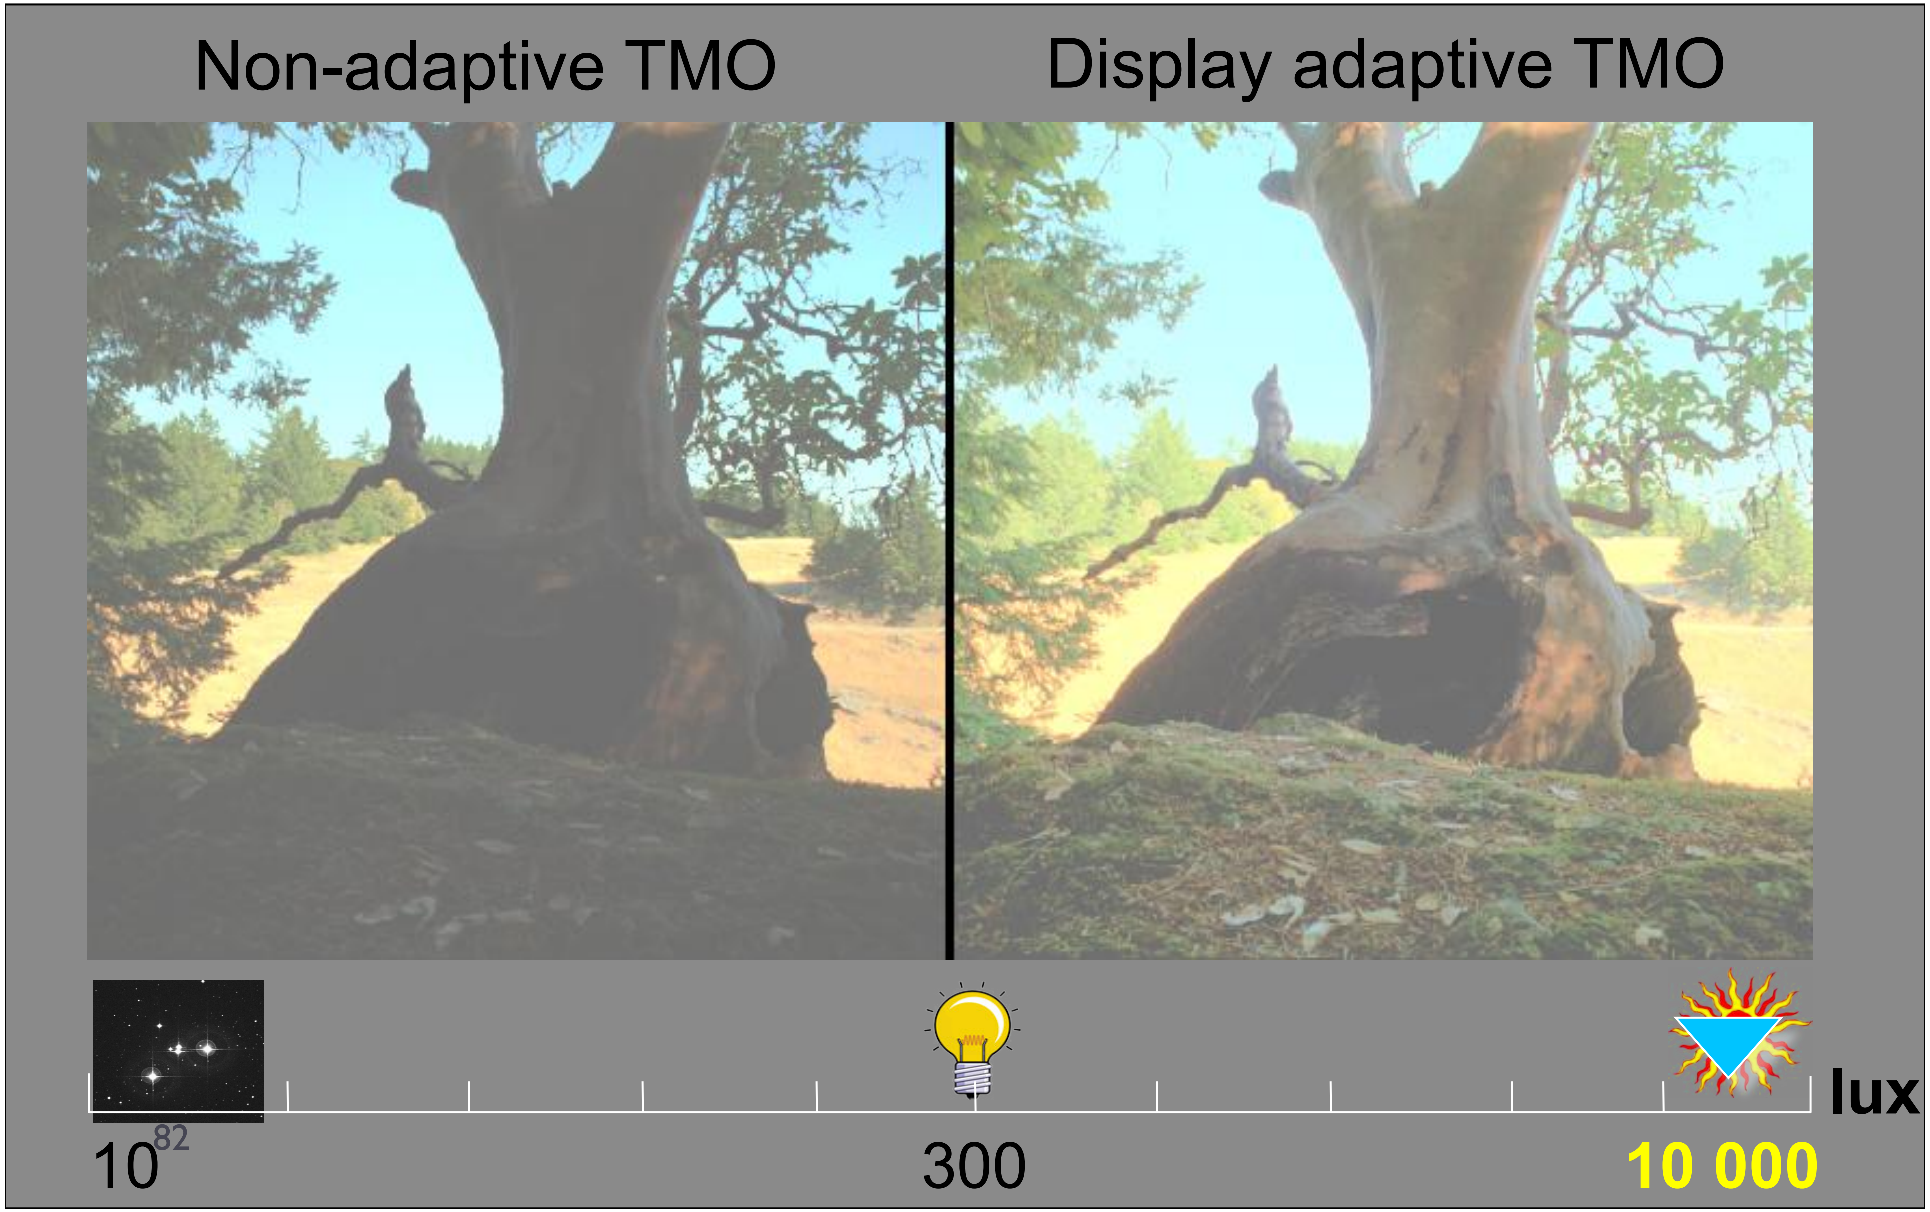

Results: display contrast

ePaper　　standard LCD　　HDR display

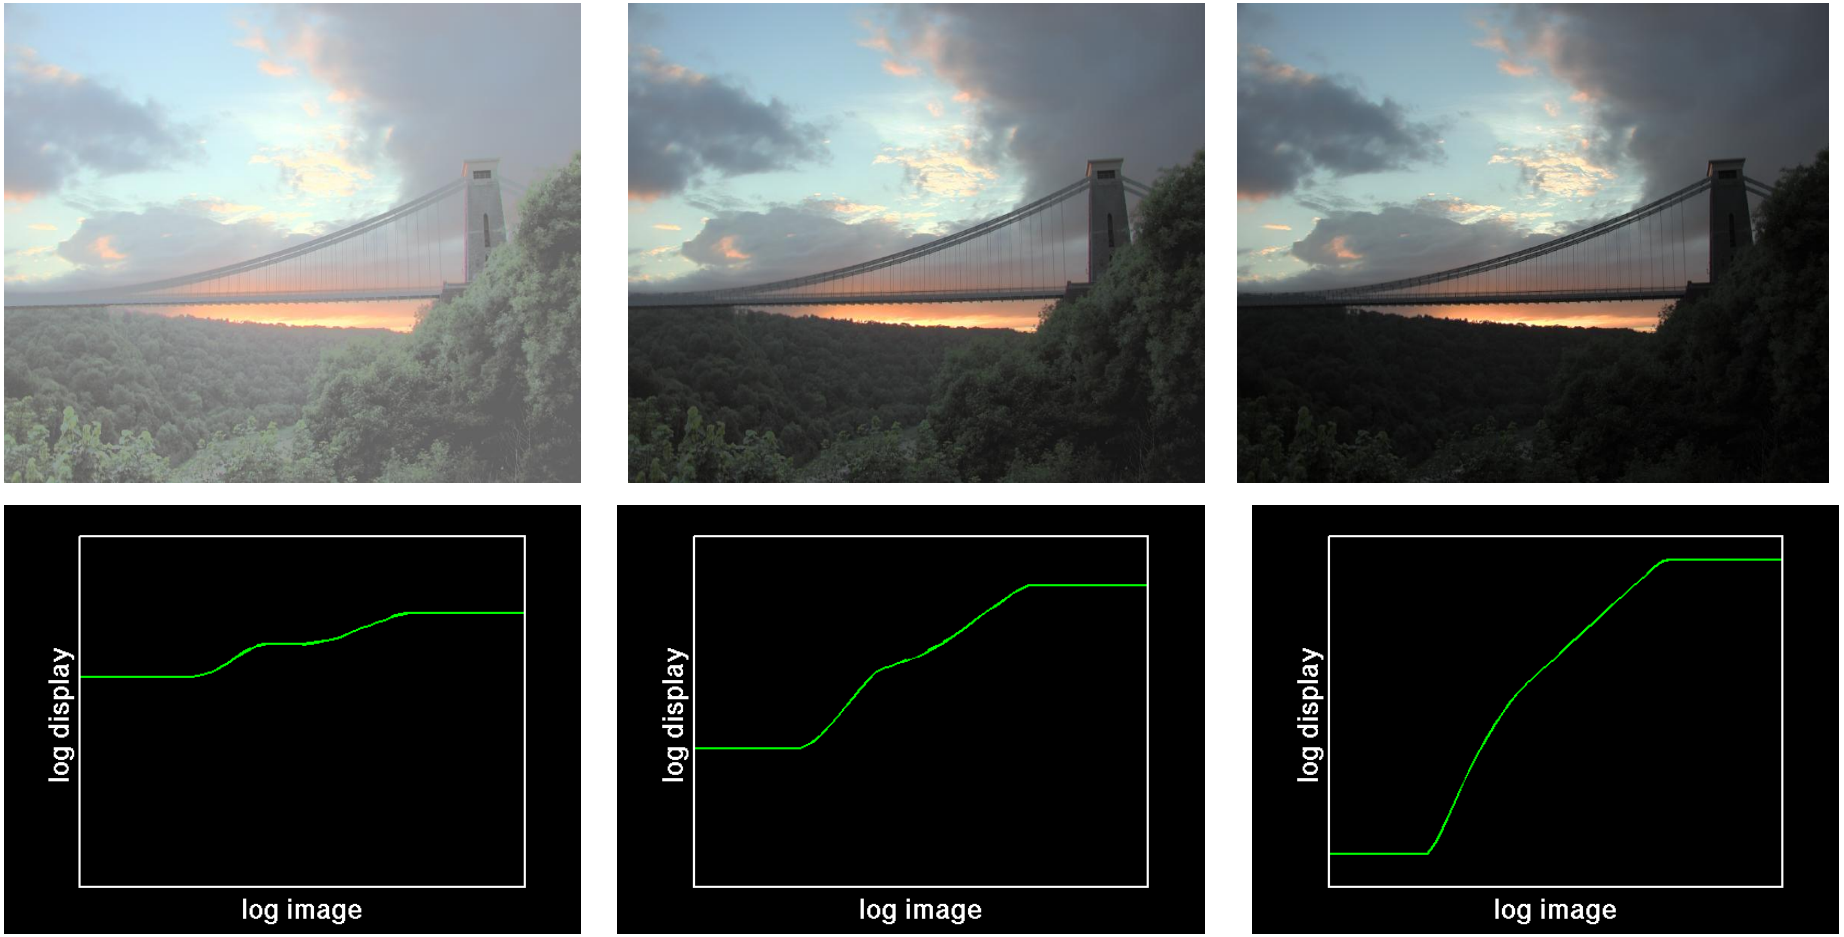

Light and dark adaptation

Temporal adaptation

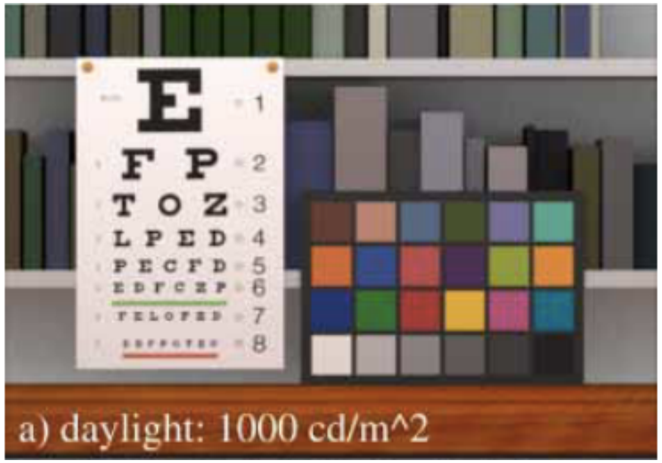

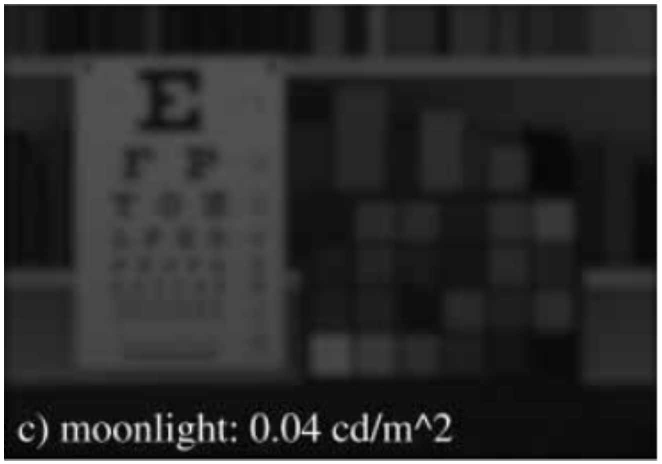

[Pattanaik et al. '98]

Time-course of temporal adaptation

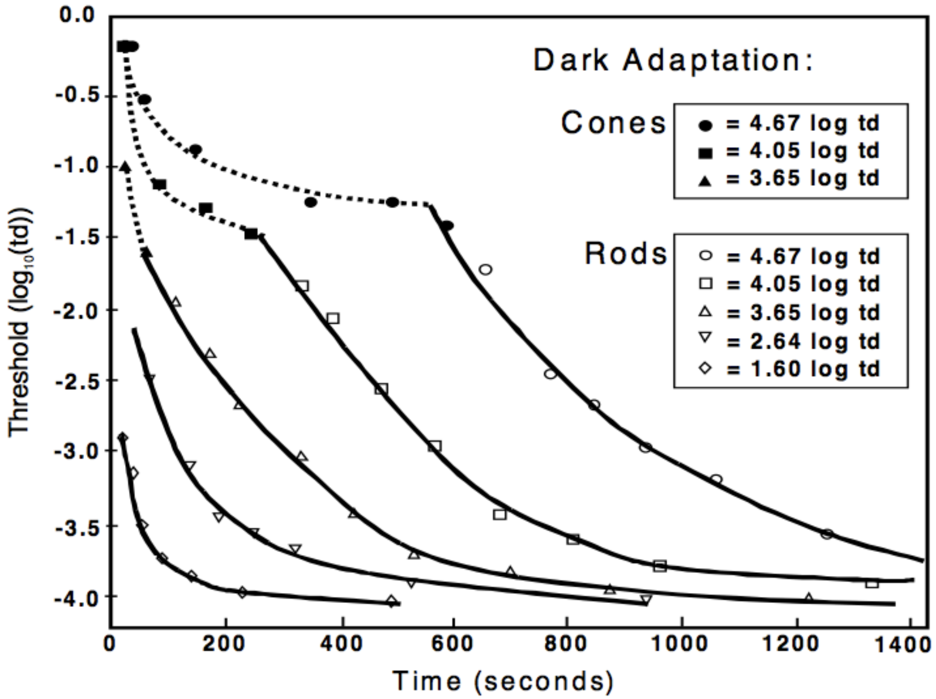

Bright -> Dark

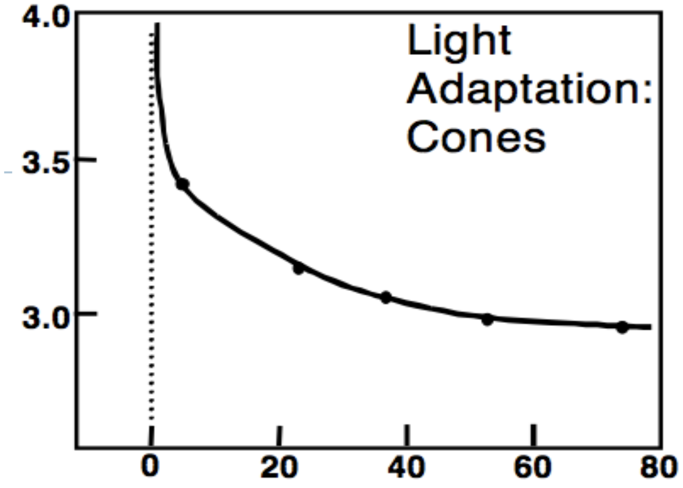

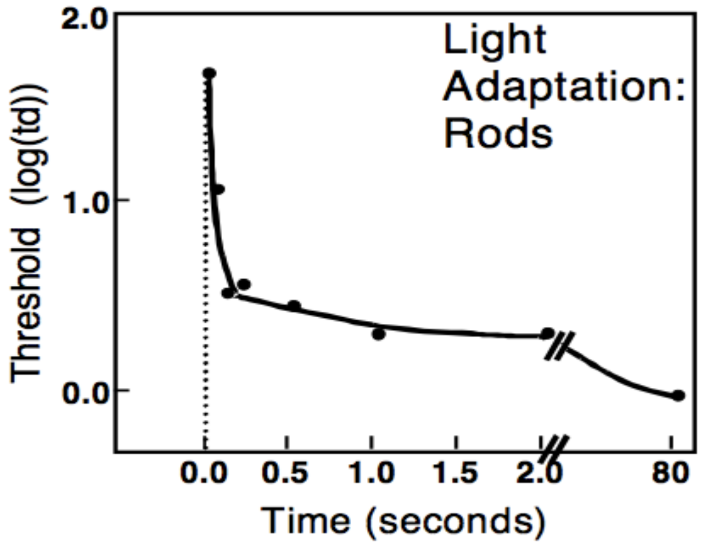

Dark -> Bright

In tone mapping: Photoreceptor models

- Naka-Rushton equation:

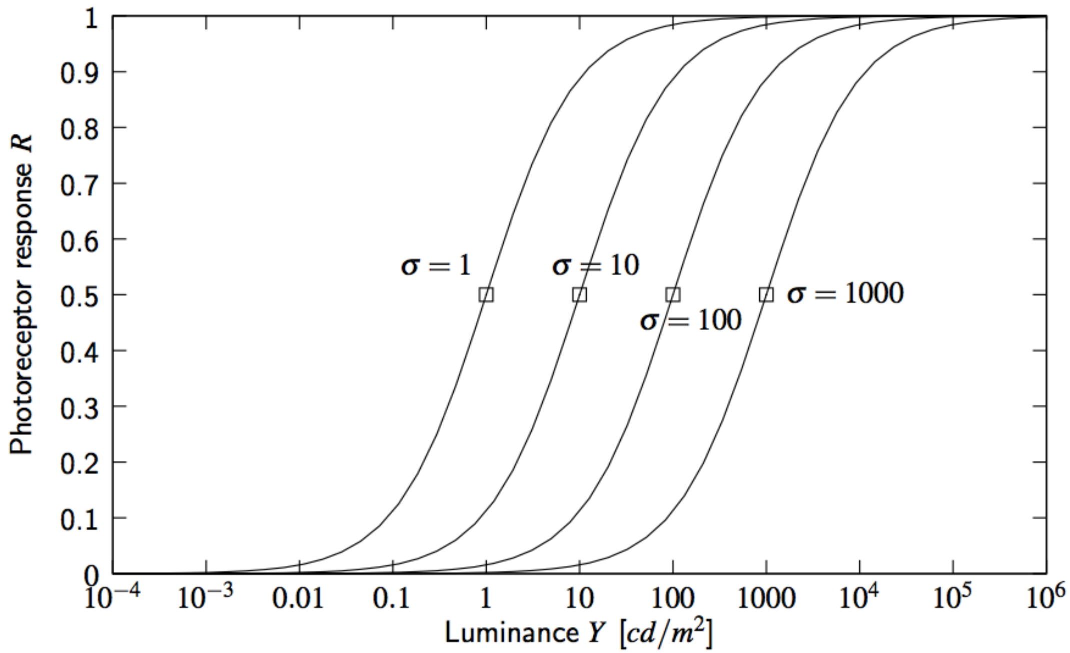

$$\frac{R}{R_{max}} = \frac{Y^n}{Y^n + \sigma^n}$$

Depends on the luminance of adaptation

How to find local adaptation?

- In physiology and psychophysics
  - Uniform adapting field
    - Controlled conditions

- In computer graphics
  - Ad hoc assumptions
    - global average luminance
    - local per-pixel luminance
    - local average computed in 1° Gaussian window
    - …

Experiment: High contrast adaptation patterns on HDR display

- Extent of pooling
- Long-range effects
- Pooling non-linearity
- Radial symmetry & contrast masking
- Natural images

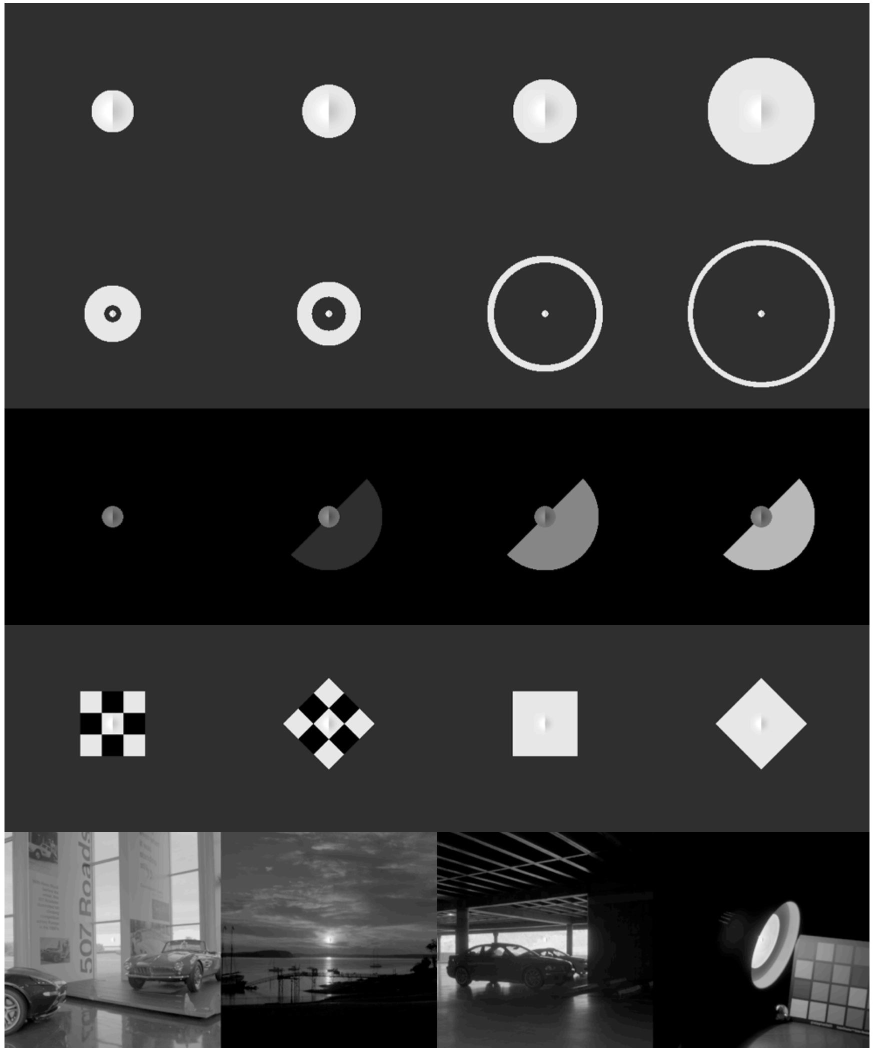

[Vangorp et al., SIGGRAPH Asia 2015]

Optimized stimuli

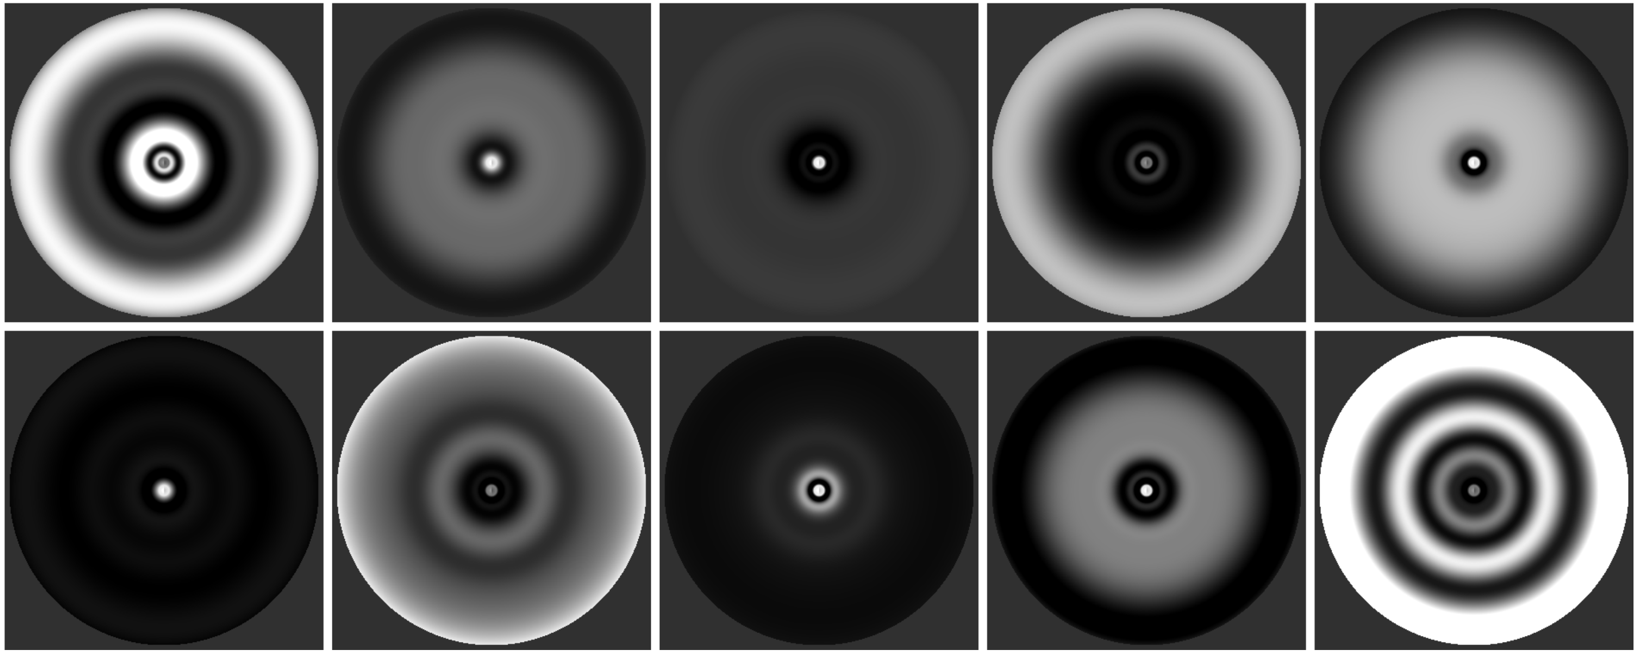

Local Adaptation Model

- Wider support at lower luminance
  - due to non-linearities
  - adaptation site shifts to postreceptoral mechanisms [Dunn et al. 2007]

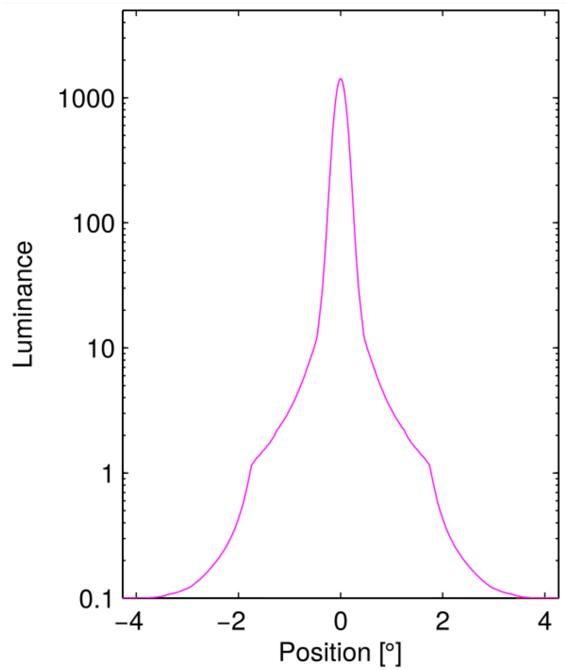

- Complex pooling mechanism
  - cross-validated to avoid overfitting
  - more complex than known retinal pooling
  - receptive fields in LGN or visual cortex?

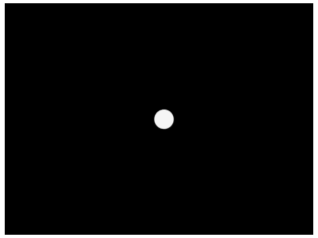

Local Adaptation Model

- Wider support at lower luminance
  - due to non-linearities
  - adaptation site shifts to postreceptoral mechanisms [Dunn et al. 2007]
- Complex pooling mechanism
  - cross-validated to avoid overfitting
  - more complex than known retinal pooling
  - receptive fields in LGN or visual cortex?

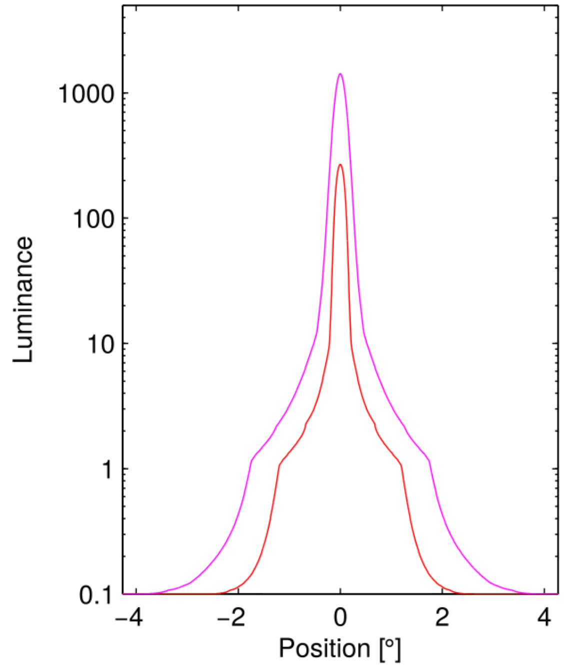

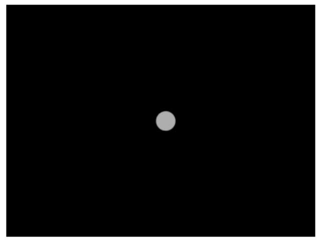

Local Adaptation Model

- Wider support at lower luminance
  - due to non-linearities
  - adaptation site shifts to postreceptoral mechanisms [Dunn et al. 2007]
- Complex pooling mechanism
  - cross-validated to avoid overfitting
  - more complex than known retinal pooling
  - receptive fields in LGN or visual cortex?

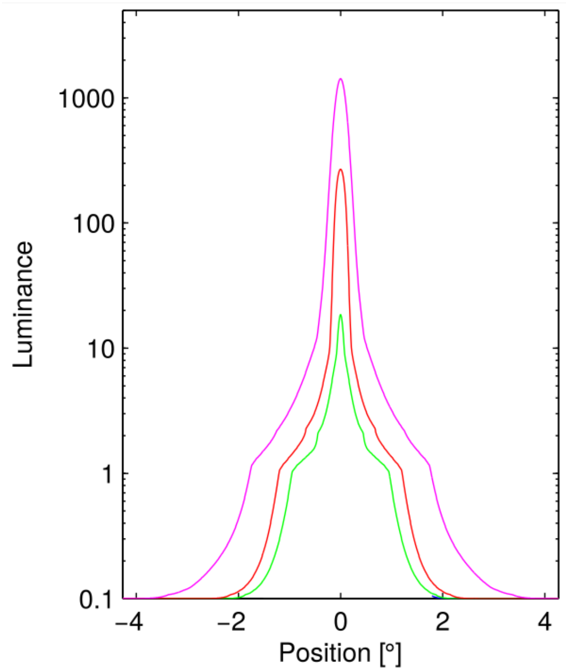

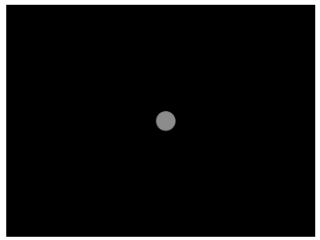

Local Adaptation Model

- Wider support at lower luminance
  - due to non-linearities
  - adaptation site shifts to postreceptoral mechanisms [Dunn et al. 2007]

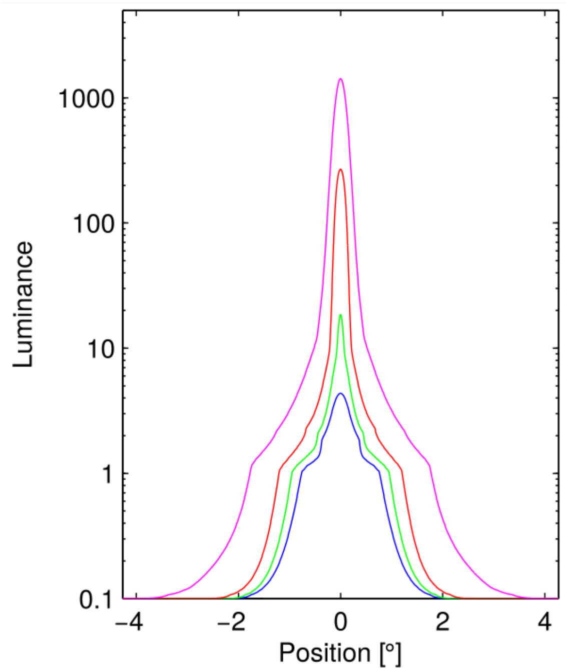

- Complex pooling mechanism
  - cross-validated to avoid overfitting
  - more complex than known retinal pooling
  - receptive fields in LGN or visual cortex?

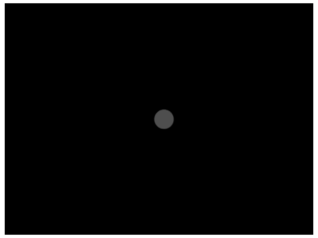

Local Adaptation Model

- Wider support at lower luminance
  - due to non-linearities
  - adaptation site shifts to postreceptoral mechanisms [Dunn et al. 2007]

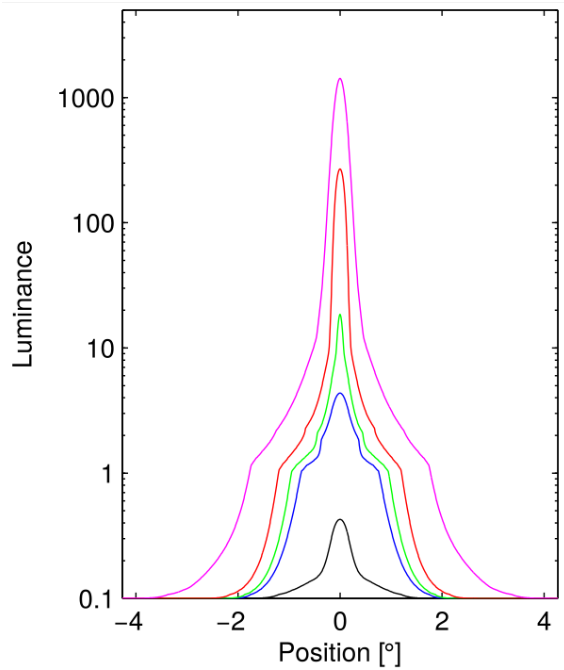

- Complex pooling mechanism
  - cross-validated to avoid overfitting
  - more complex than known retinal pooling
  - receptive fields in LGN or visual cortex?

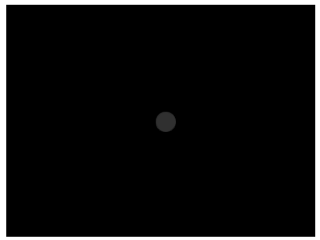

[Vangorp et al., SIGGRAPH Asia 2015]

Application: Gaze-dependent TM

Summary

- Fundamentals
  - Intents of tone-mapping
  - Display models
  - Four approaches to tone-mapping
- Perceptual issues
  - Color
  - Glare
  - Visual metrics-driven TMO
  - Light and dark adaptation

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

Age-adaptive night vision

Video 4

Rivoli

Simulation of age-adaptive night vision