Multidimensional retargeting: Tone Mapping

Rafal Mantiuk

PRIFYSGOL CYMRU UNIVERSITY OF WALES BANGOR 1884

Bangor University, UK

Research Institute of Visual Computing

http://www.bangor.ac.uk/mantiuk/

Learning outcomes

- What is tone-mapping ?
- What problem(s) does it solve ?
- Why is the problem so difficult ?
- How do we perceive high dynamic range images ?
- What are the major approaches to tone-mapping ?
- How to choose a tone-mapping for a particular application?

Tone-mapping problem

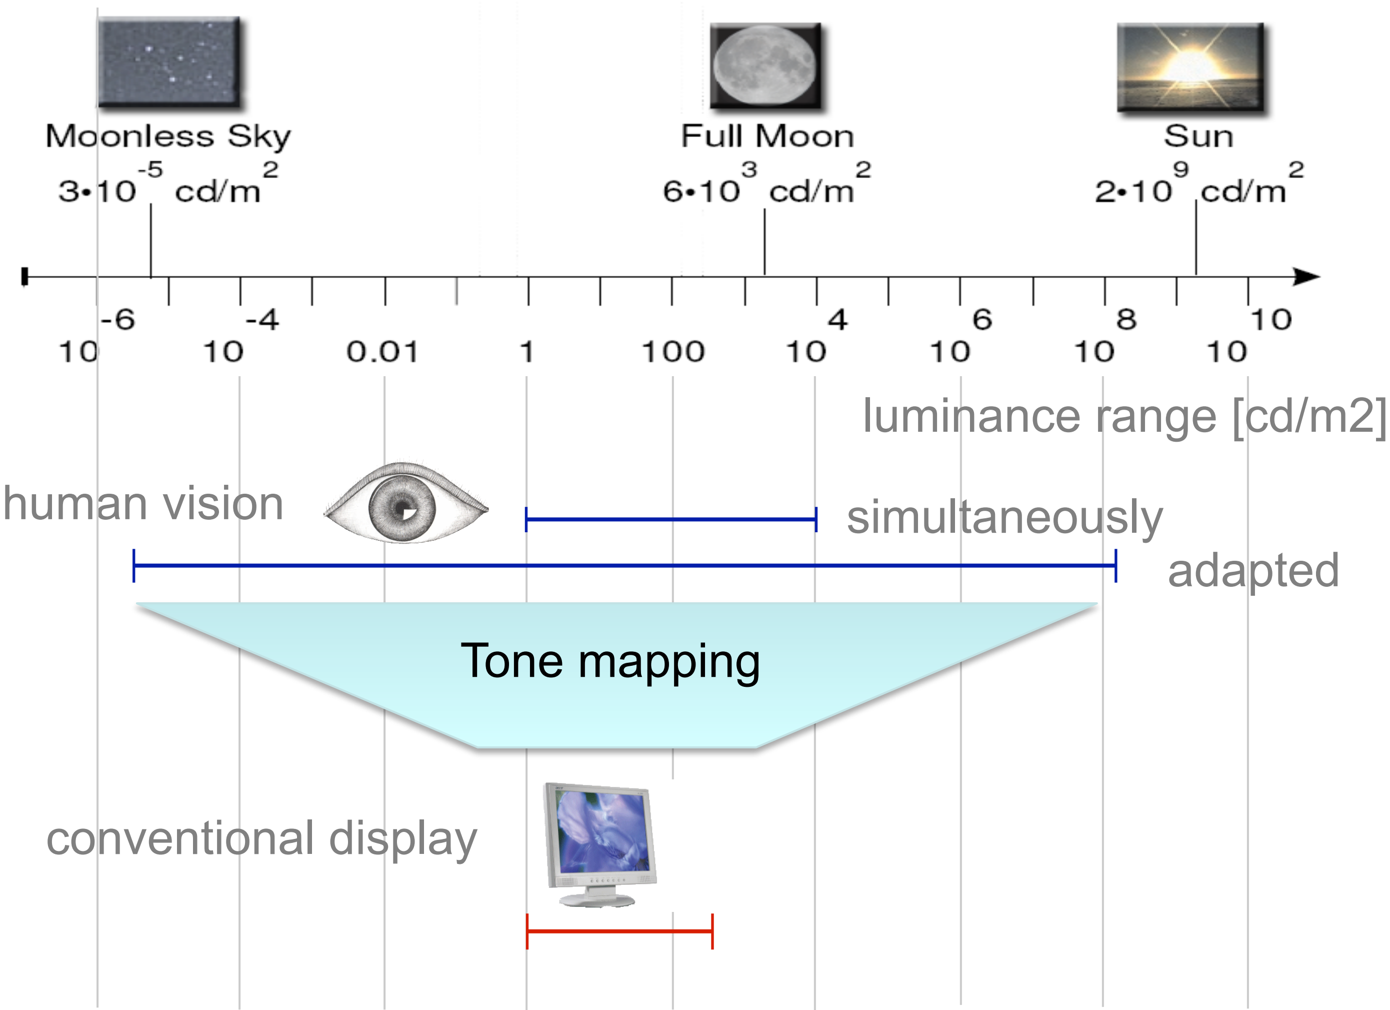

Question to the audience

- Who has never used a tone-mapping operator?

Each camera needs to tone-map a real-world captured light before it can be stored as a JPEG. This is essentially the same process as tone-mapping, although knows as ‘color reproduction’ or ‘color processing’.

Color space retargeting problem

Goal: map colors to a restricted color space

Perceptual retargeting problem

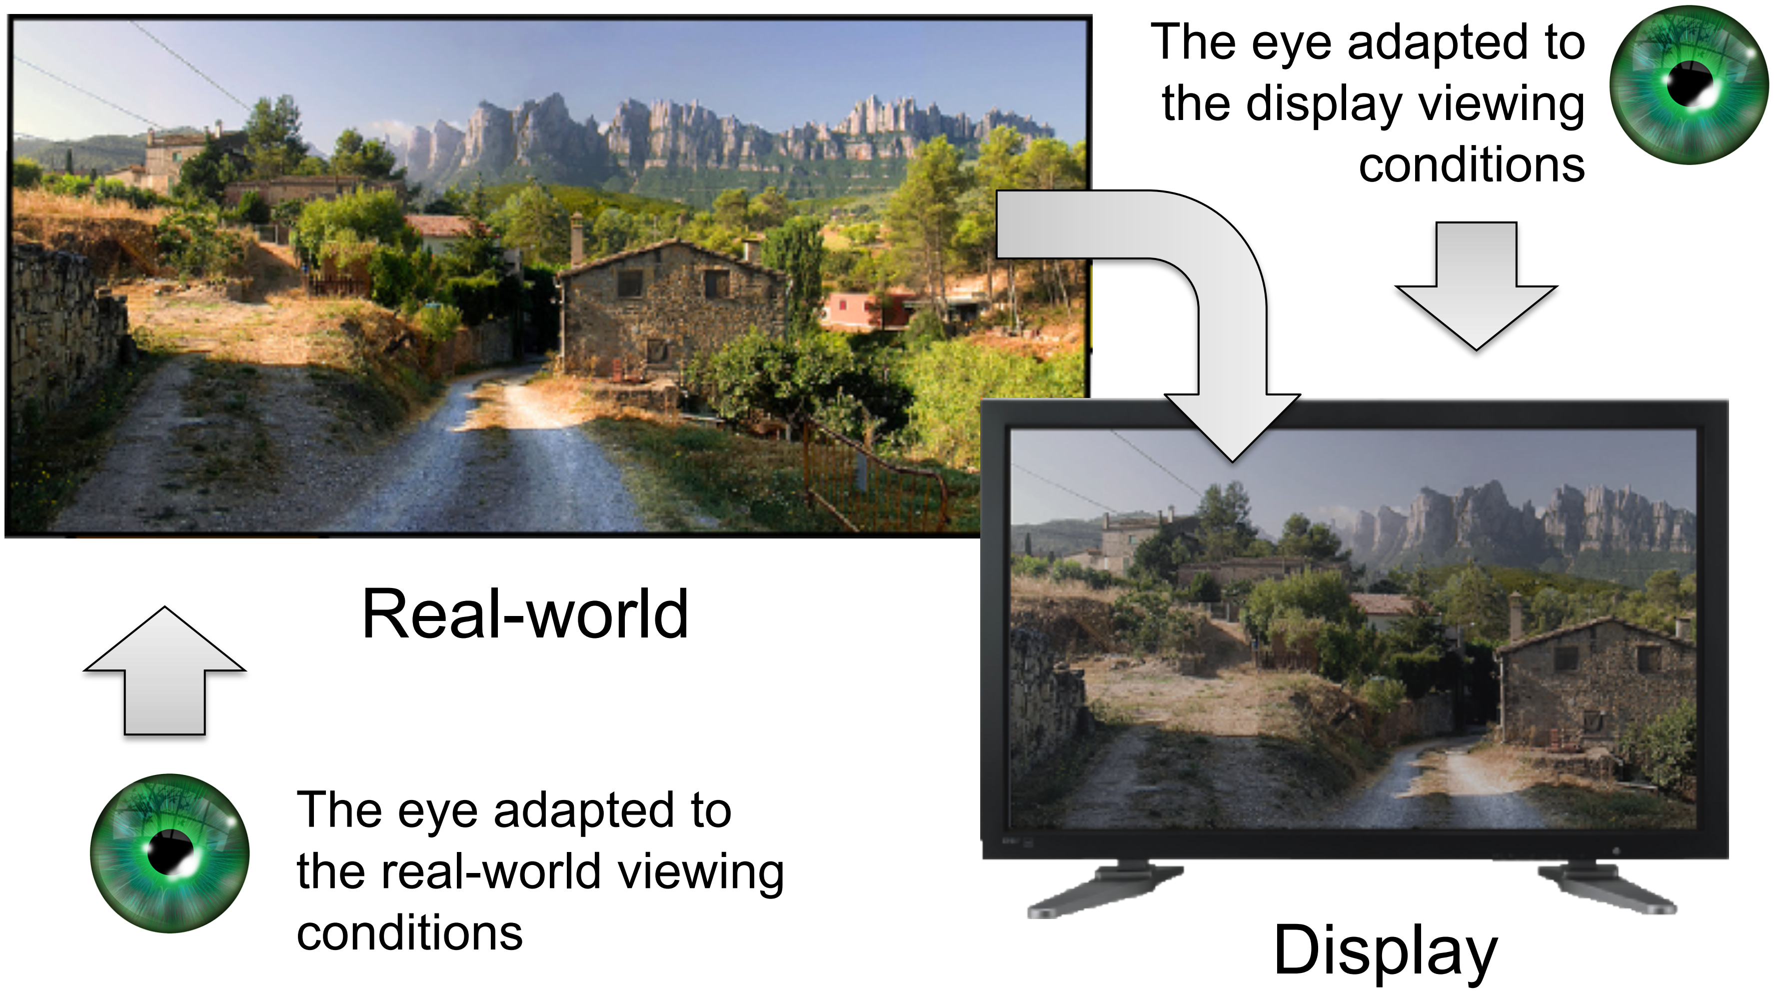

Goal: match color appearance

Tone Mapping?

- HDR ?
- Or something else ?

What is tone-mapping?

Although tone-mapping may have different meanings, this course is about:

A) Transformation of an image from an unrestricted color gamut of real world or an abstract scene to the restricted color gamut of a device

B) Retargeting the perceptual appearance from one viewing conditions to another

Input and output

- HDR
- (approximate) physical units
- luminance
- linear RGB
- scene-referred

- LDR (SDR)
- pixel values
- luma
- gamma corrected R’G’B’
- display referred

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)

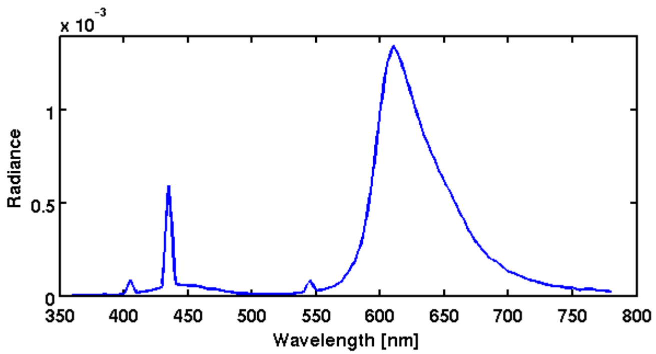

Luminous efficiency function (weighting)

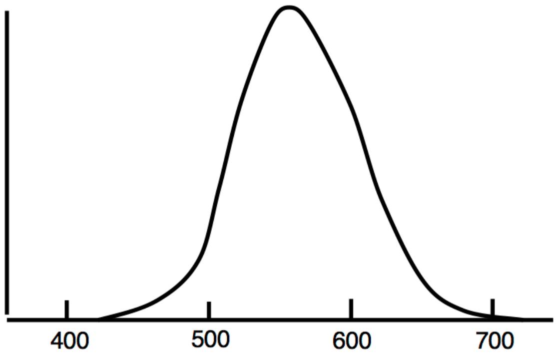

Do HDR images contain luminance values?

- Not exactly, because:
  – the combination of camera red, green and blue spectral sensitivity curves will not match the luminous efficiency function
- But they contain a good-enough approximation for most applications
  – For multi-exposure camera capture the error in luminance measurements is 10-15%

Sensitivity to luminance

- Weber-law – the just-noticeable difference is proportional to the magnitude of a stimulus

Ernst Heinrich Weber
[From wikipedia]

The smallest detectable luminance difference

Background (adapting) luminance

$$\frac{\Delta L}{L} = k$$

Constant

Typical stimuli:

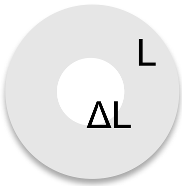

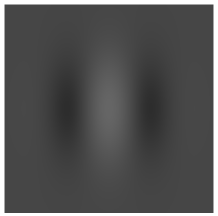

Consequence of the Weber-law

- Smallest detectable difference in luminance

$$\frac{\Delta L}{L} = k,$$

L	ΔL
100 cd/m^2	1 cd/m^2
1 cd/m^2	0.01 cd/m^2

- Adding or subtracting luminance will have different visual impact depending on the background luminance
- Unlike LDR luma values, HDR luminance values are not perceptually uniform!

How to make luminance (more) perceptually uniform?

- Using Fechnerian integration

$$dR(L) = \frac{1}{\Delta L(L)}$$

Derivative of response

Detection threshold

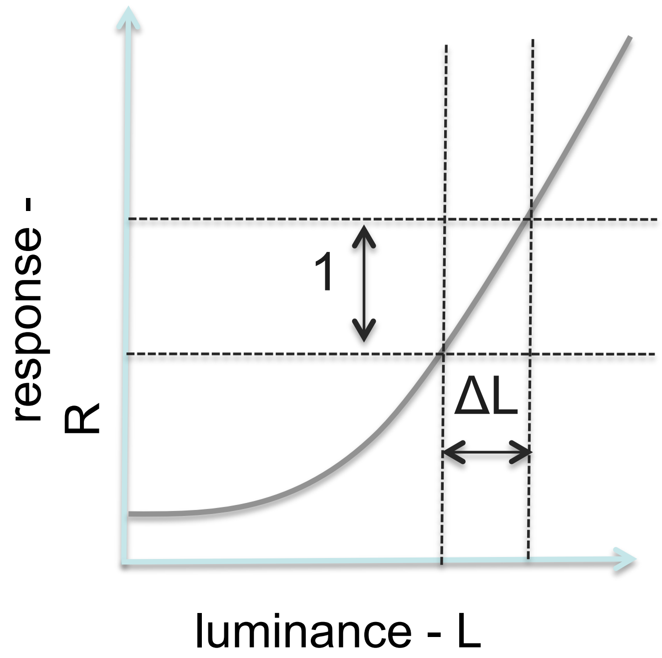

Luminance transducer:

$$R(L) = \int_0^L \frac{1}{\Delta L(l)} dl$$

Assuming the Weber law

$$\frac{\Delta L}{L} = k,$$

- and given the luminance transducer

$$R(L) = \int_0^L \frac{1}{\Delta L(l)} dl$$

- the response of the visual system to light is:

$$R(L) = \int \frac{1}{kL} dL = \frac{1}{k} \ln(L) + k_1$$

Fechner law

$$R(L) = a\ln(L)$$

- Practical insight from the Fechner law:
  - The easiest way to adopt image processing algorithms to HDR images is to convert luminance (radiance) values to the logarithmic domain

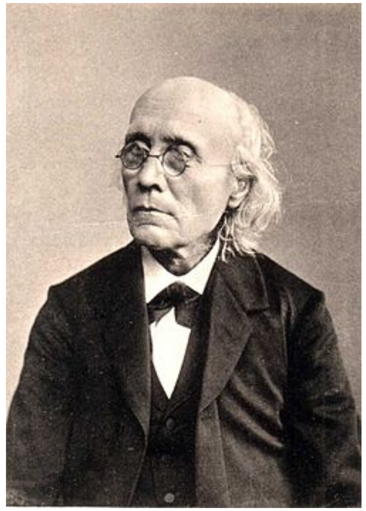

Gustav Fechner
[From Wikipedia]

But…the Fechner law does not hold for the full luminance range

- Because the Weber law does not hold either
- Threshold vs. intensity function:

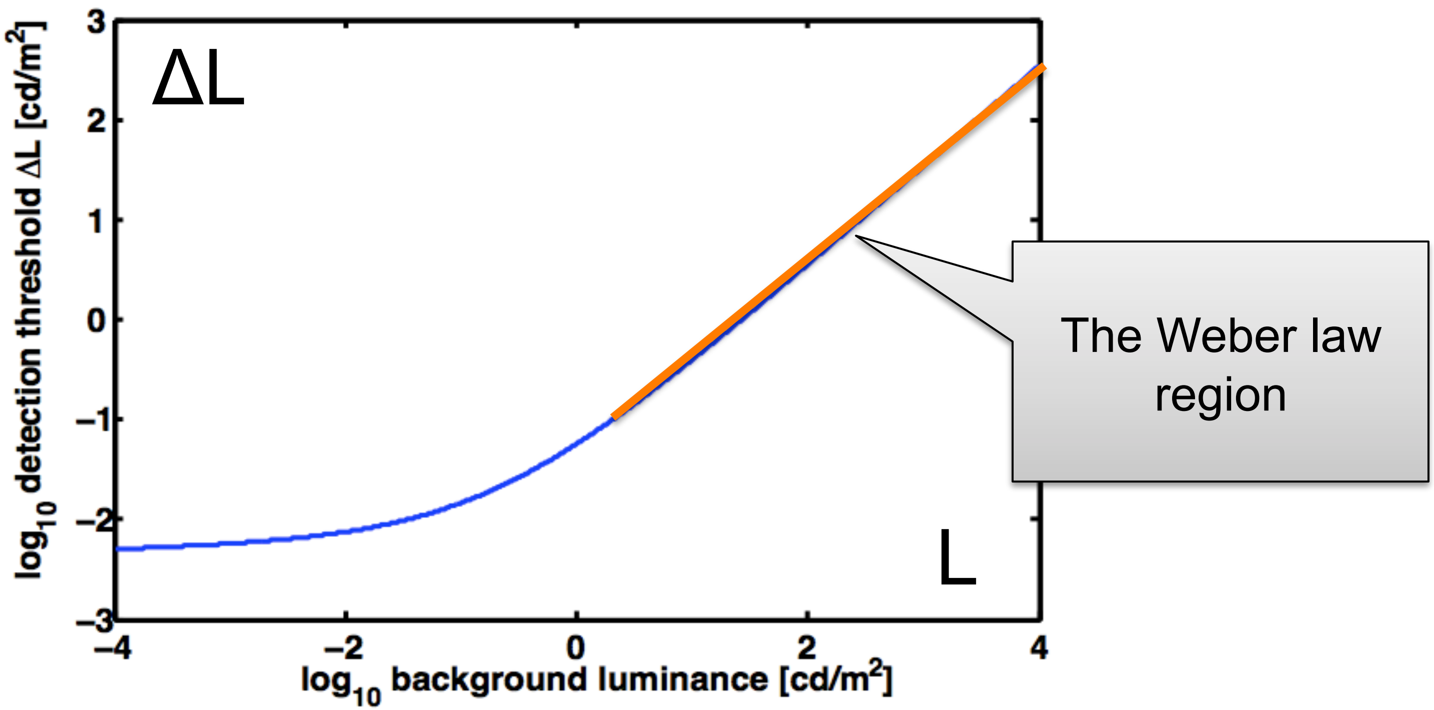

Weber-law revisited

- If we allow detection threshold to vary with luminance according to the t.v.i. function:

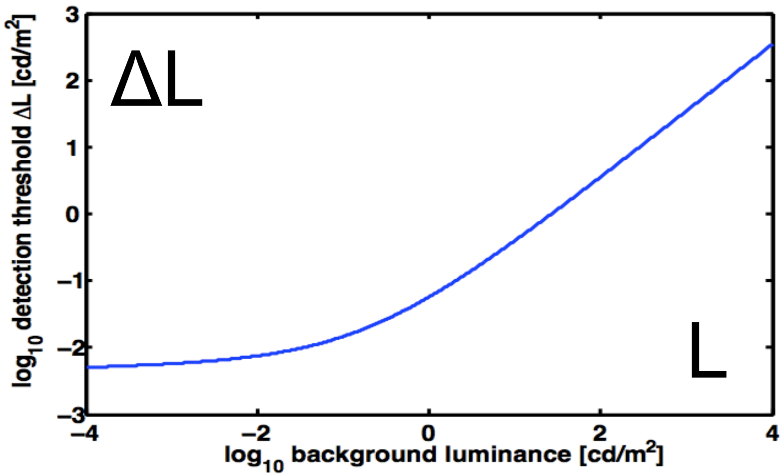

- we can get more accurate estimate of the "response":

$$R(L) = \int_0^L \frac{1}{\Delta L(l)} dl$$

Fechnerian integration and Steven's law

R(L) - function derived from the t.v.i. function

$$R(L) = \int_0^L \frac{1}{\Delta L(l)} dl$$

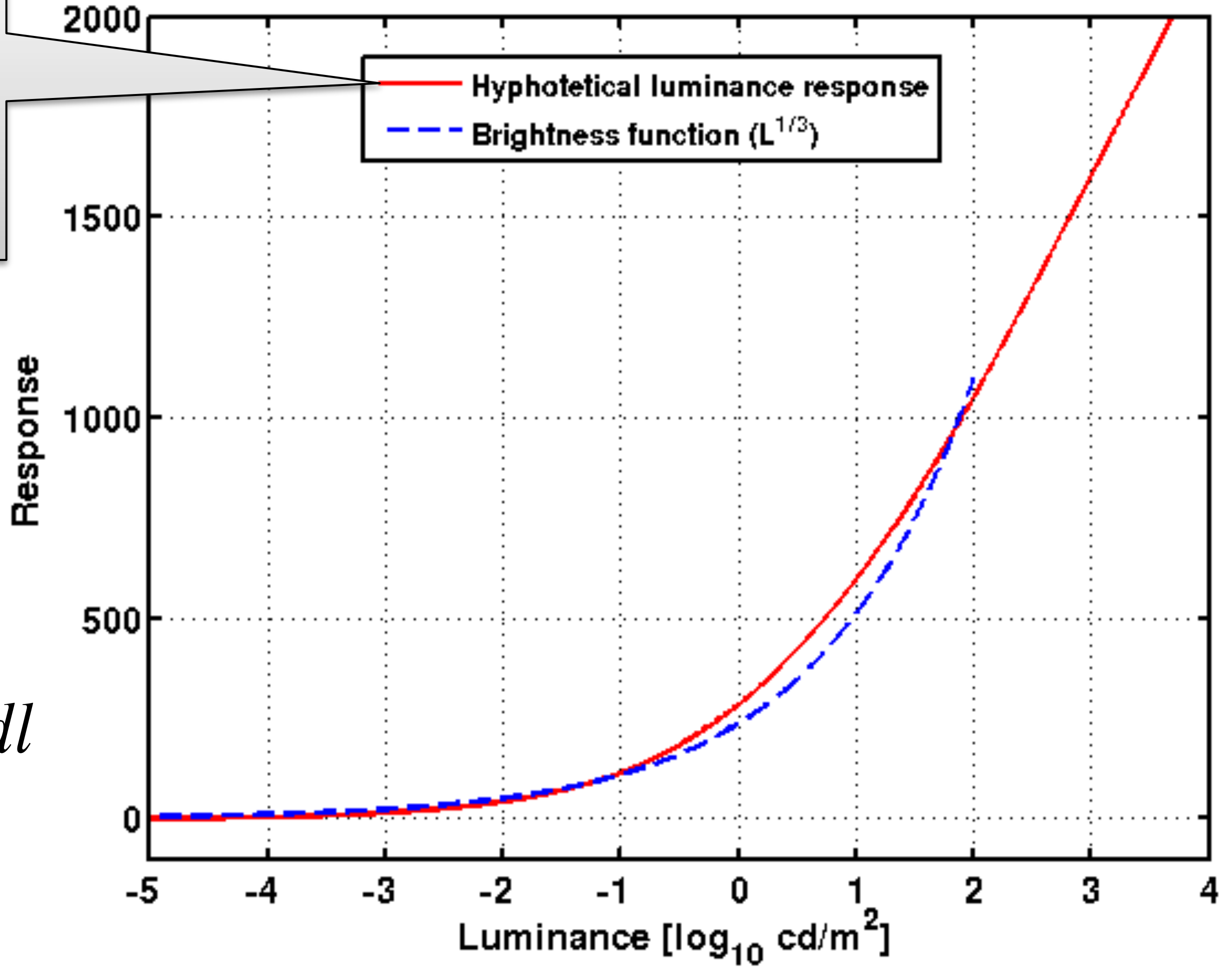

Major approaches to tone-mapping

- Illumination & reflectance separation
- Forward visual model
- Forward & inverse visual models
- Constraint mapping problem

- This is not a crisp categorization
  – Some operators combine several approaches

Major approaches to tone-mapping

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

Illumination & reflectance separation

Input

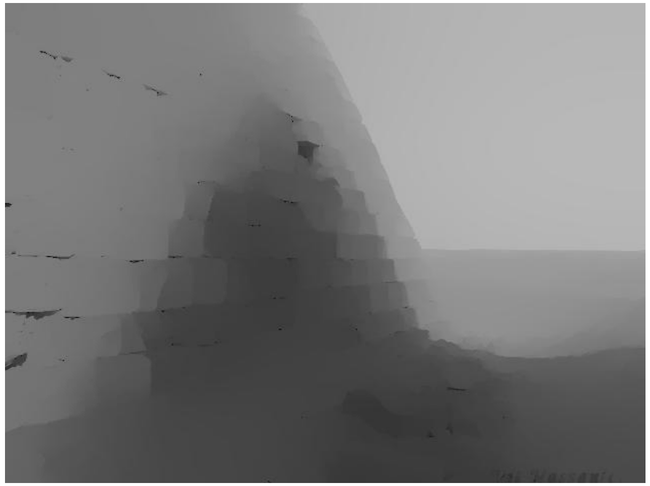

Illumination

Reflectance

Illumination and reflectance

Illumination

- 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

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

Reflectance & Illumination TMO

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

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

Tone-mapped image

Reflectance

Tone-mapping

Illumination

- for example: $I_d = R \cdot L^{1/\gamma}$

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

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

Spatially varying smoothing – less smoothing near the edges

- [Farbman et al., SIGGRAPH 2008]

WLS filter

- Stronger smoothing and still distinct edges

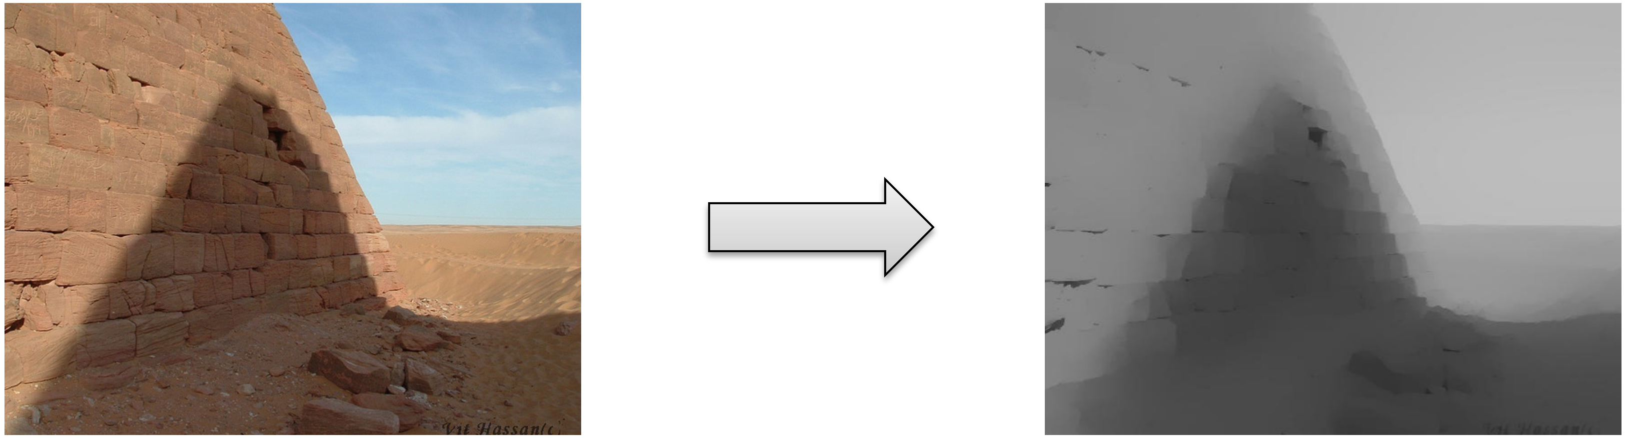

Tone mapping result

- 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

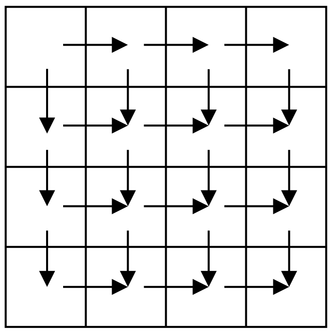

2nd step: set to 0 gradients less than the threshold

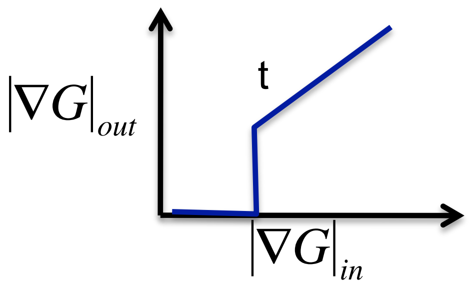

3rd step: reconstruct an image from the vector field

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

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

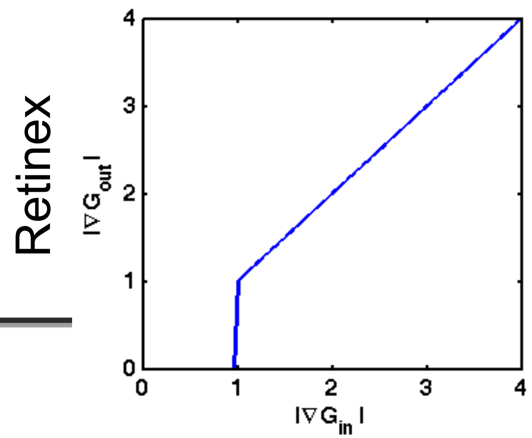

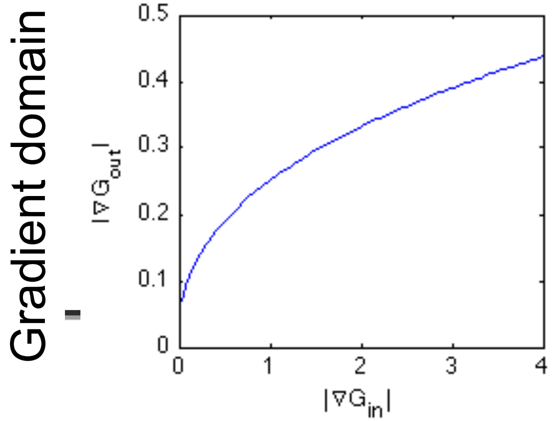

- Contrast compression achieved by global contrast reduction
  - Enhance reflectance, then compress everything

Contrast domain image processing

[Mantiuk et al., ACM Trans. Applied Perception, 2006]

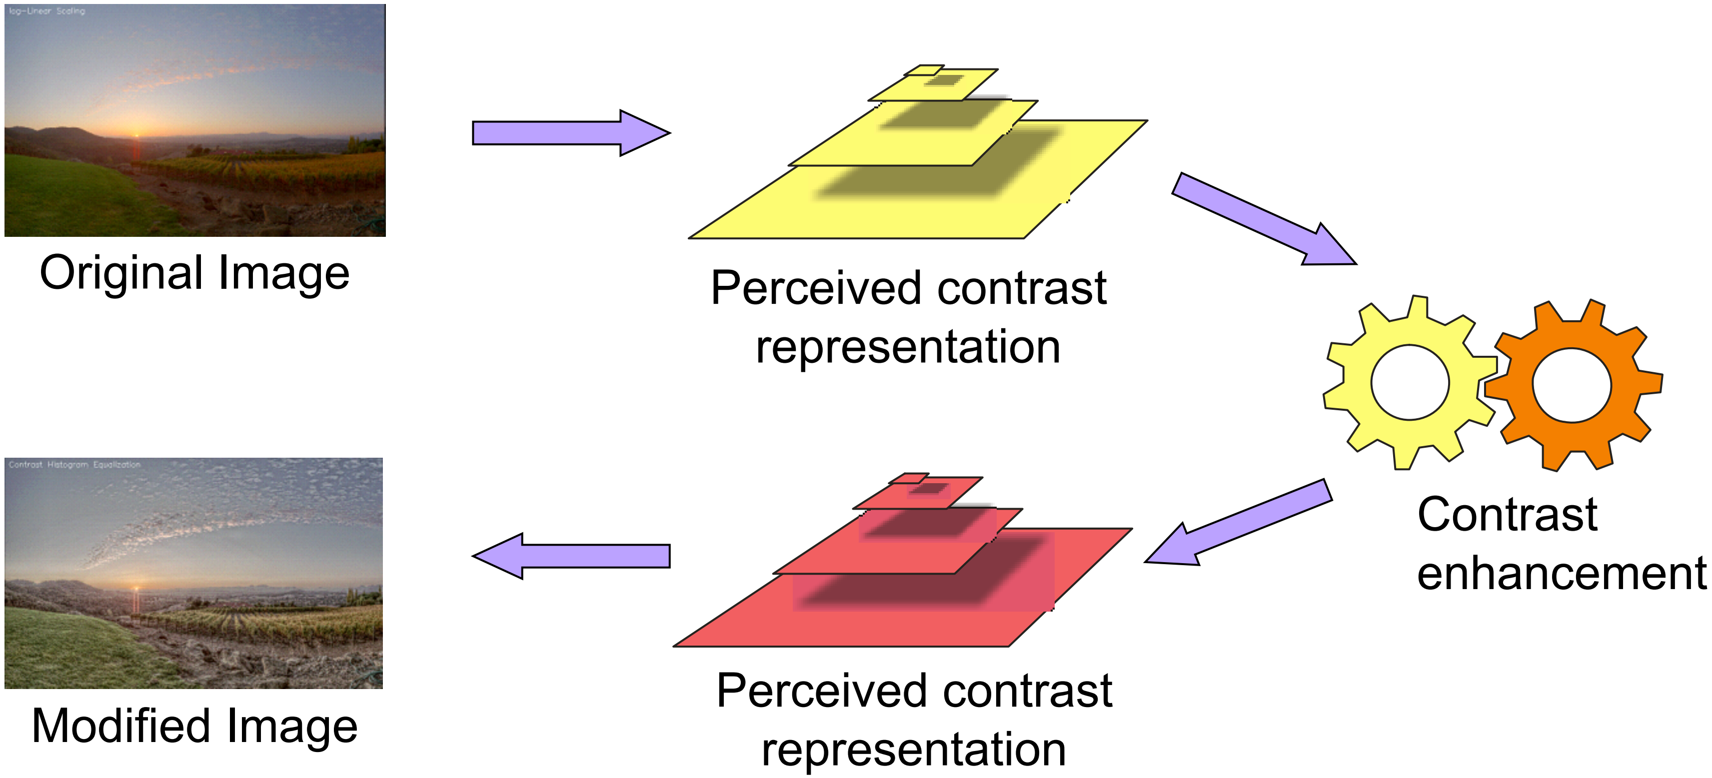

Rationale: Human eye is more sensitive to contrast than luminance

Contrast domain image processing

Wavelets

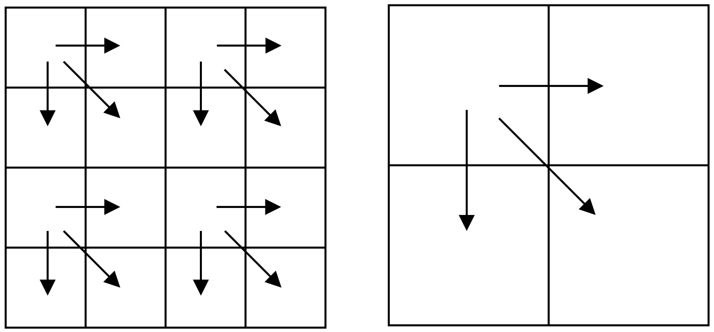

1st level 2nd level

Gradients

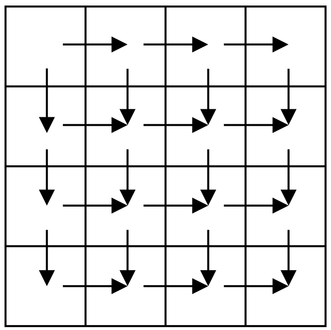

Image transform: Multi-scale contrast pyramid

Contrast pyramid

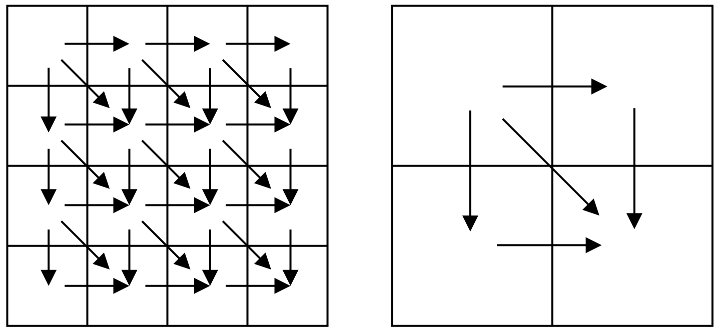

1st level 2nd level

Contrast transducer function

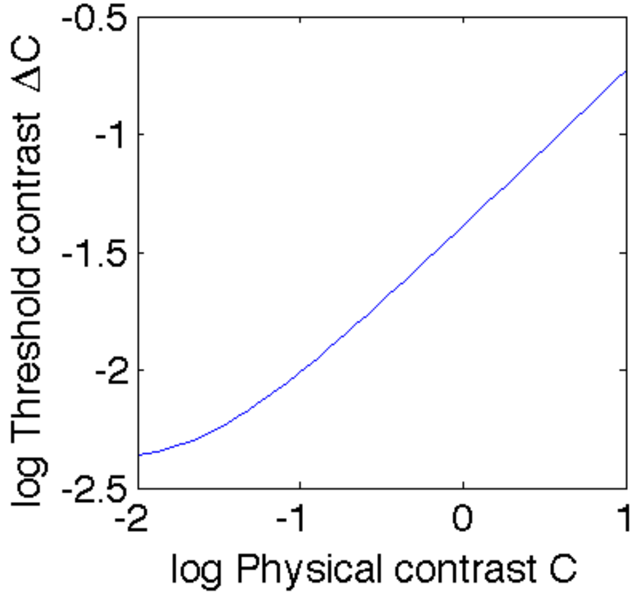

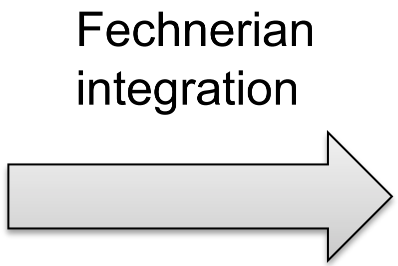

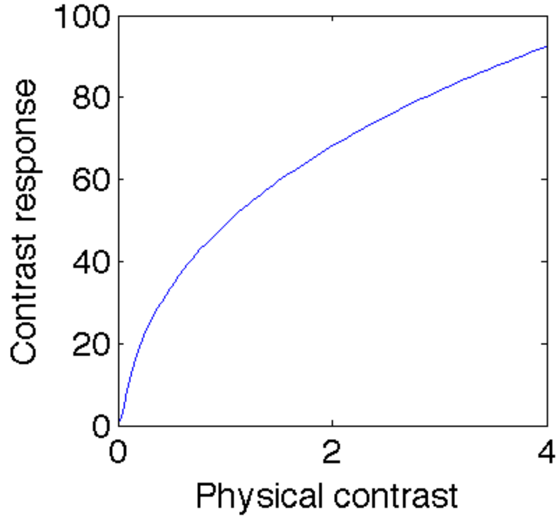

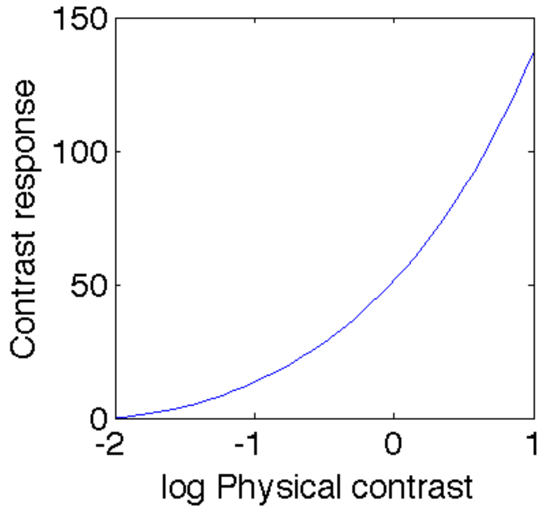

Goal: Transform contrast to the representation that is possibly perceptually uniform.

Contrast Equalization: Examples

Log-Linea

t mapping

Contrast equalization

Contrast Equalization: Examples

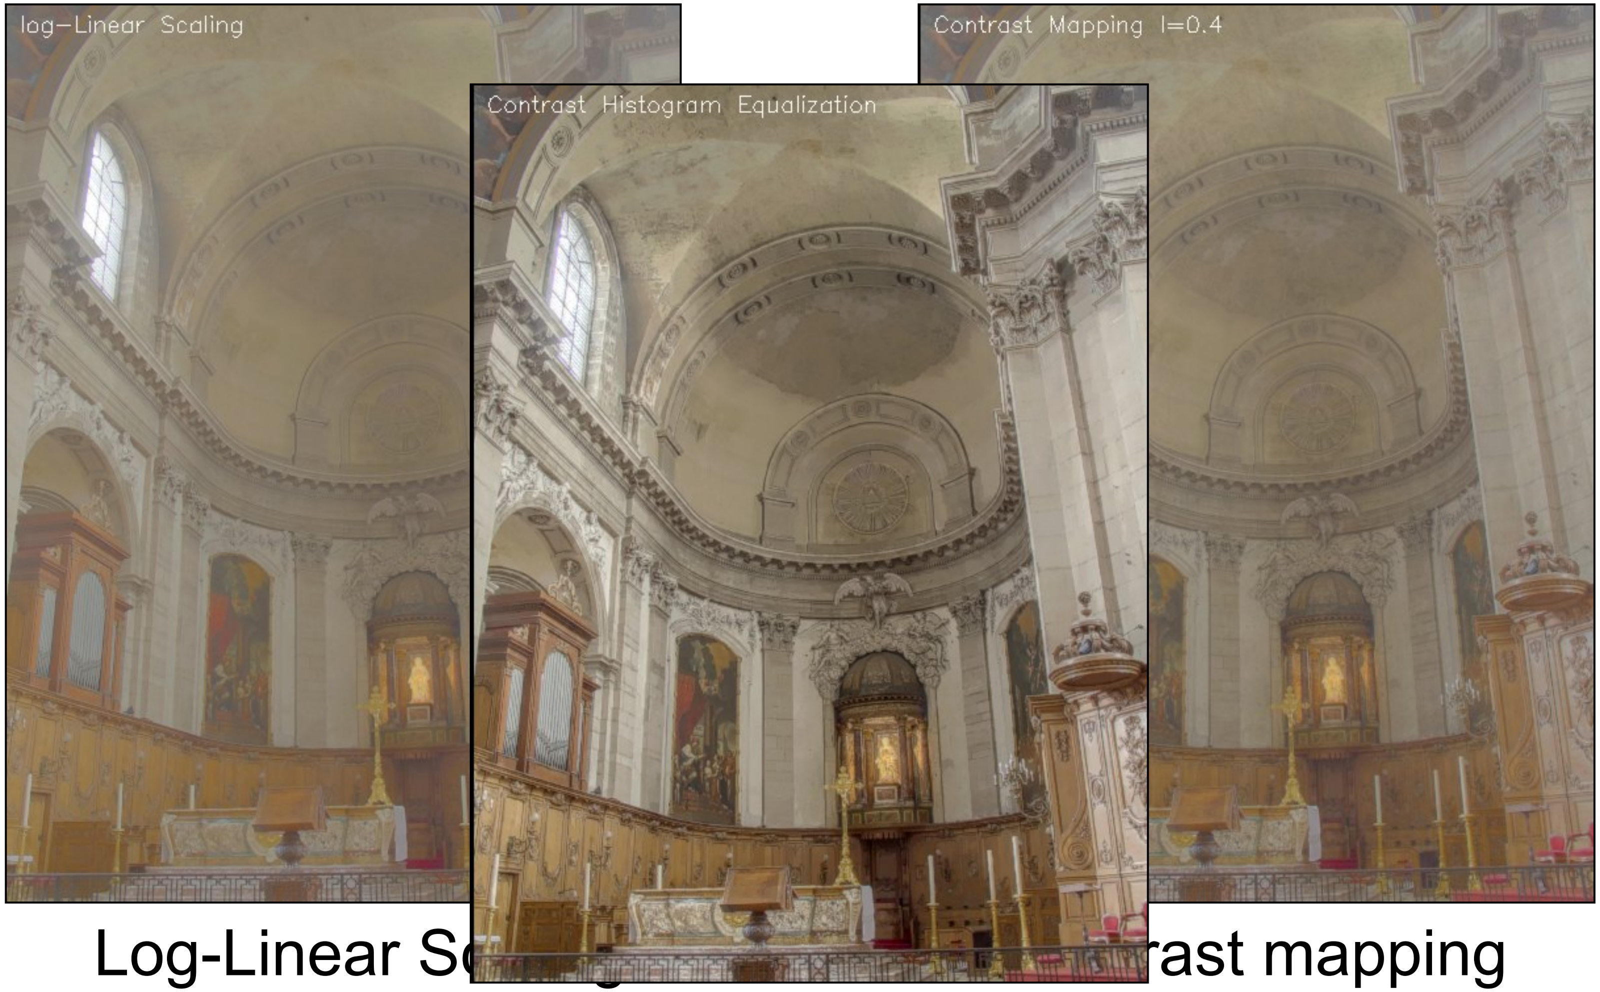

Log-Linear Sc … rast mapping

Contrast equalization

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 associated to 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

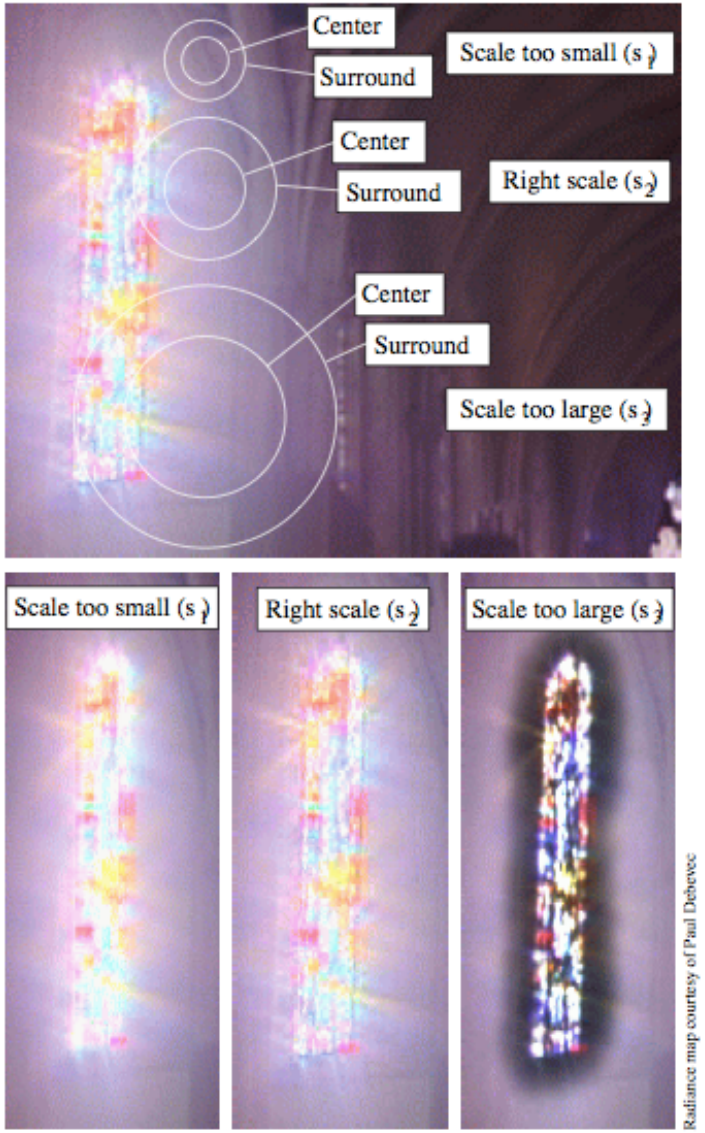

Major approaches to tone-mapping

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

Forward visual model

- Mimic the processing in the human visual system

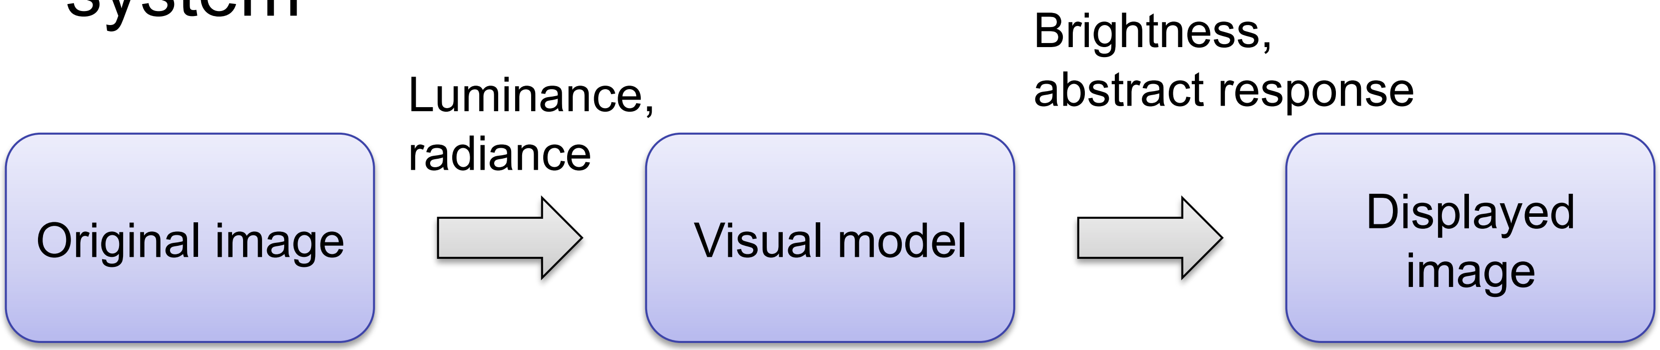

- Assumption: what is displayed is brightness or abstract response of the visual system

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:

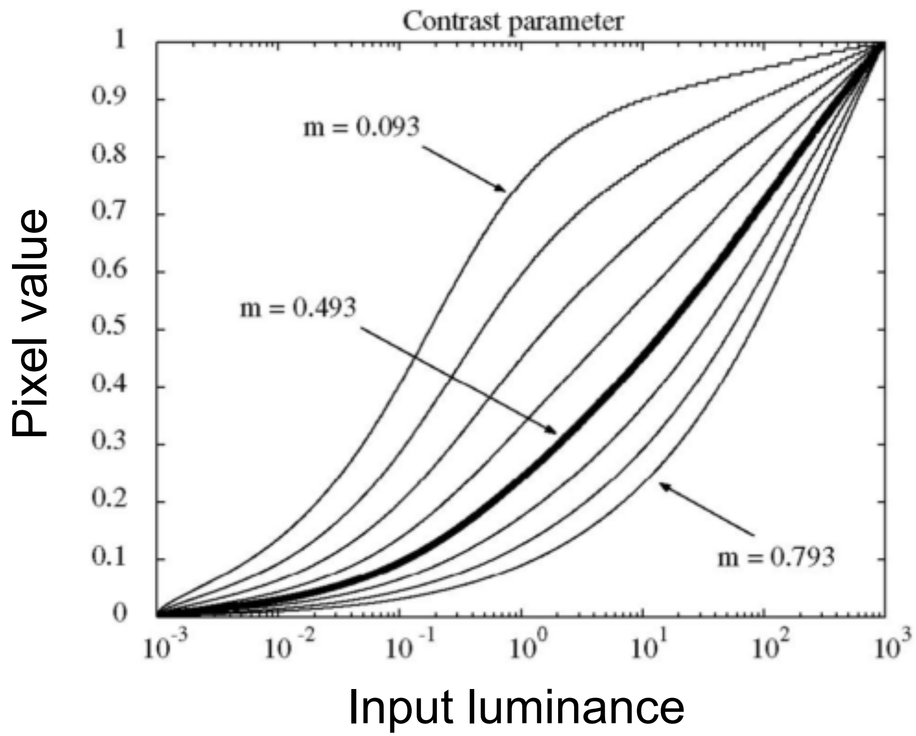

Results: photoreceptor TMO

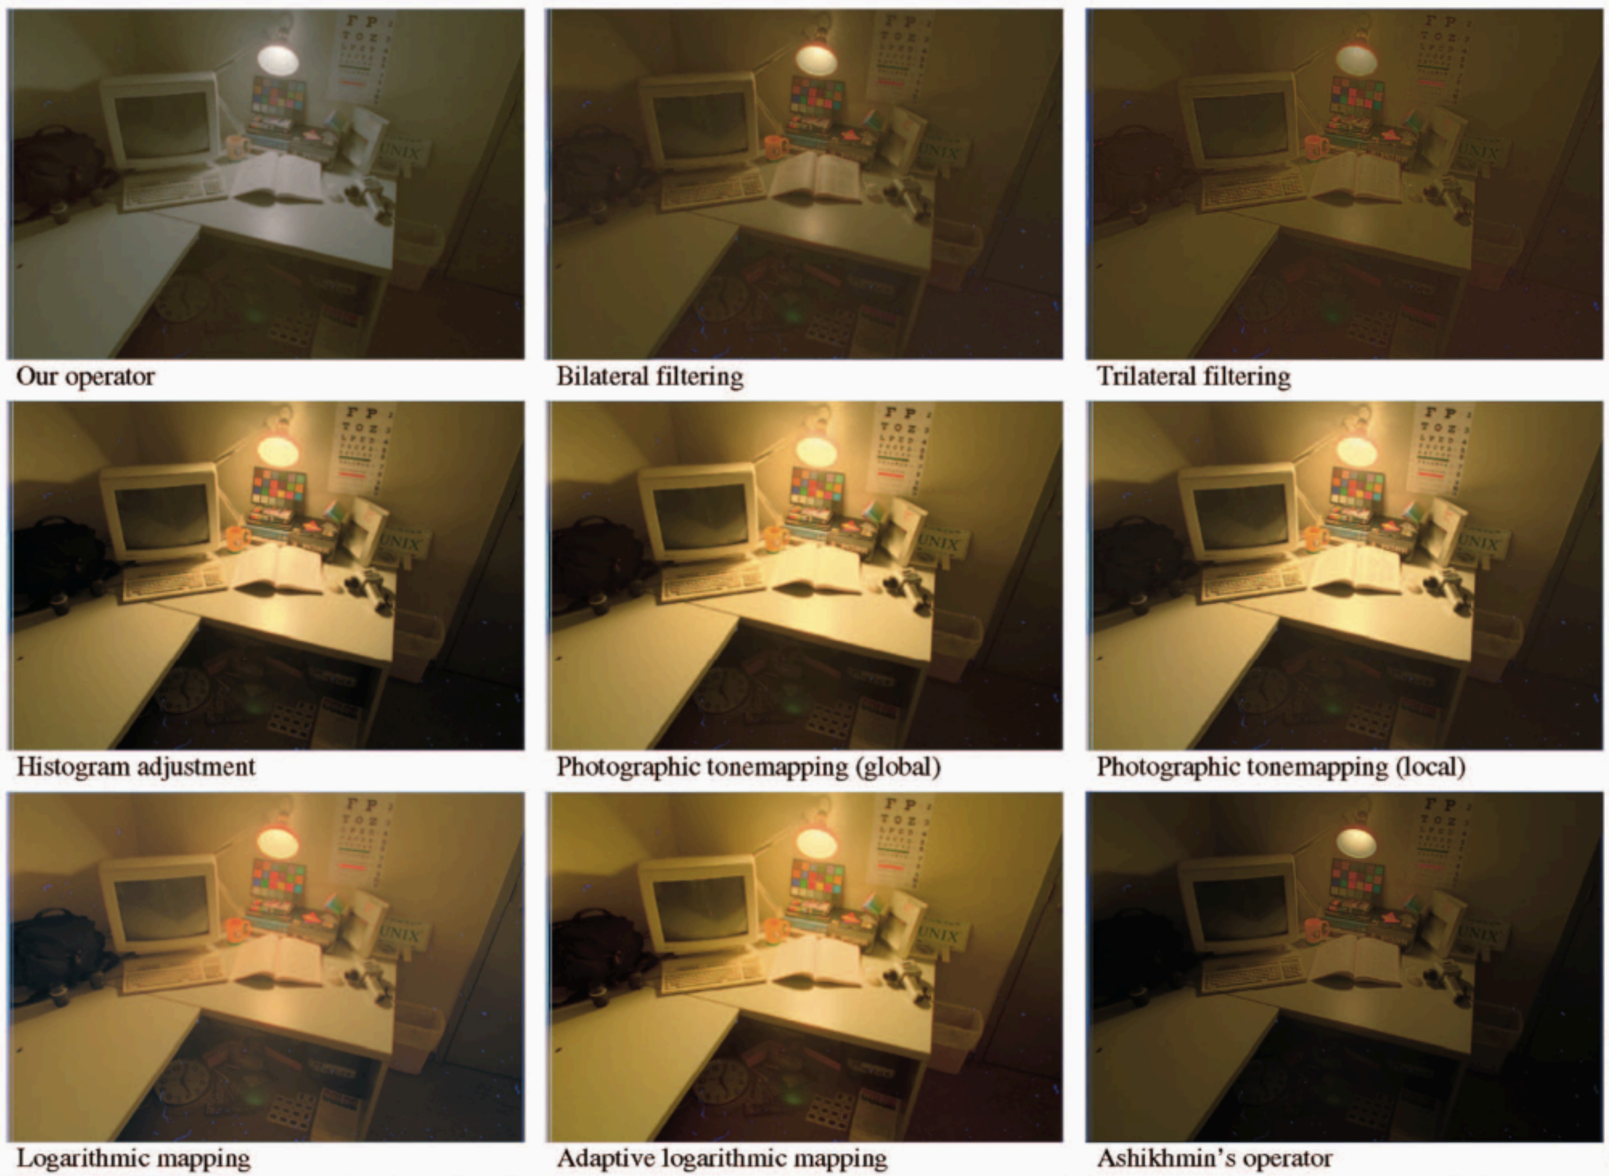

Photoreceptor models

- Naka-Rushton equation:

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

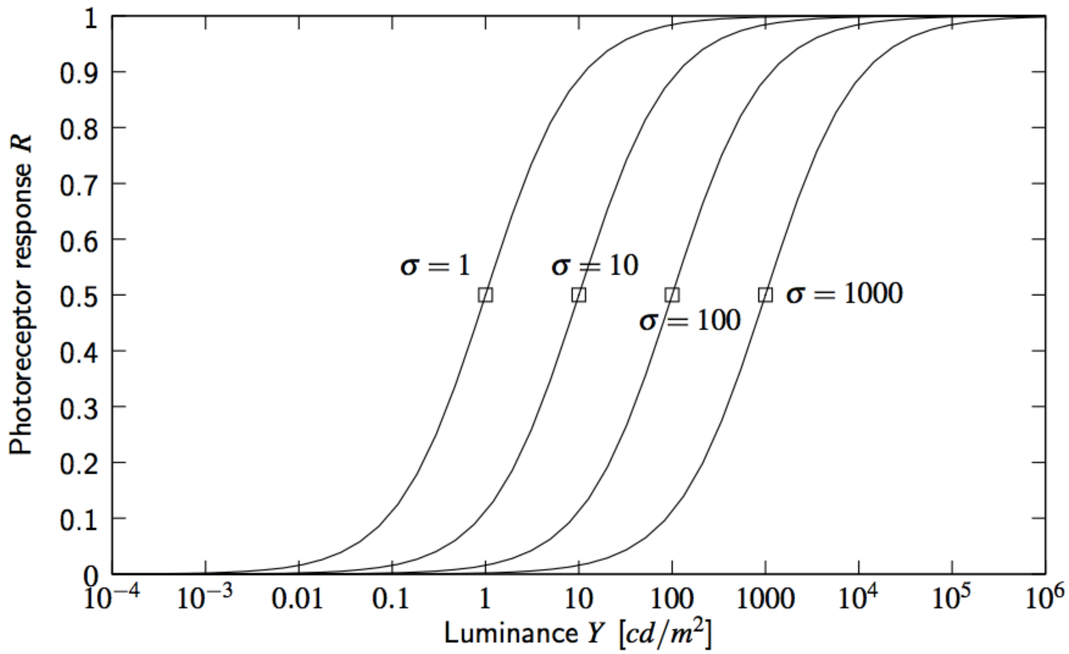

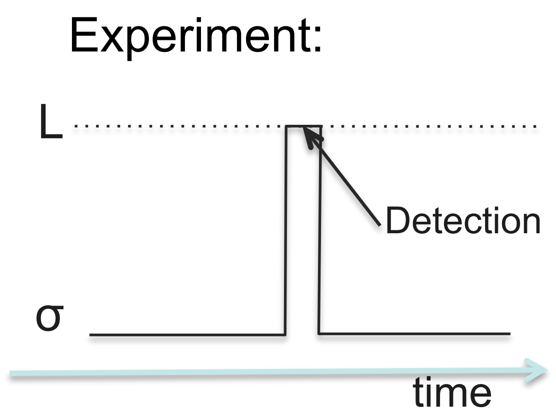

- Response of the photoreceptor to a short flicker of light - less applicable to viewing static images

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)

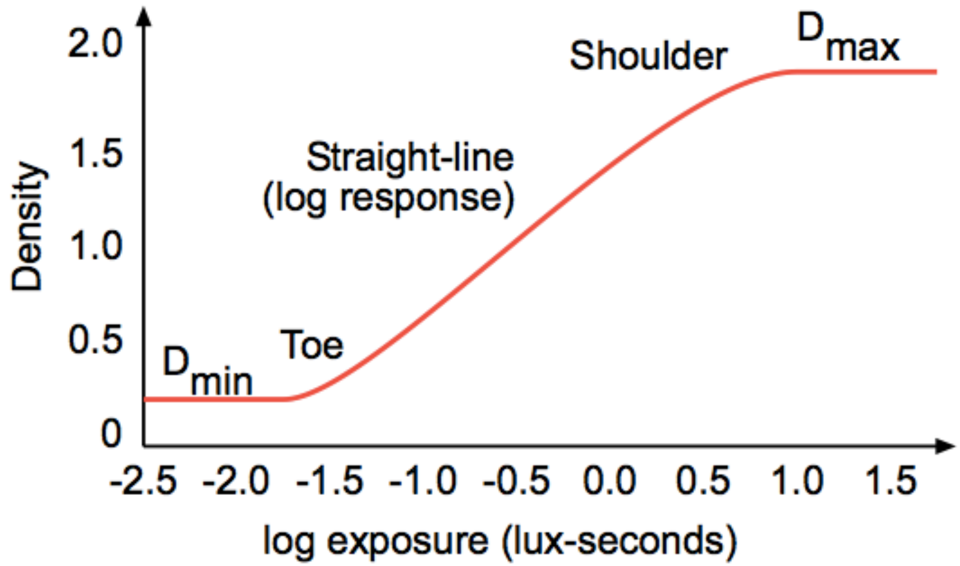

- Effectively the most commonly used tone-mapping!

Why 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

Lightness perception

- Lightness perception in tone-reproduction for high dynamic range images [Krawczyk et al. '05]
- Based on Gilchrist lightness perception theory

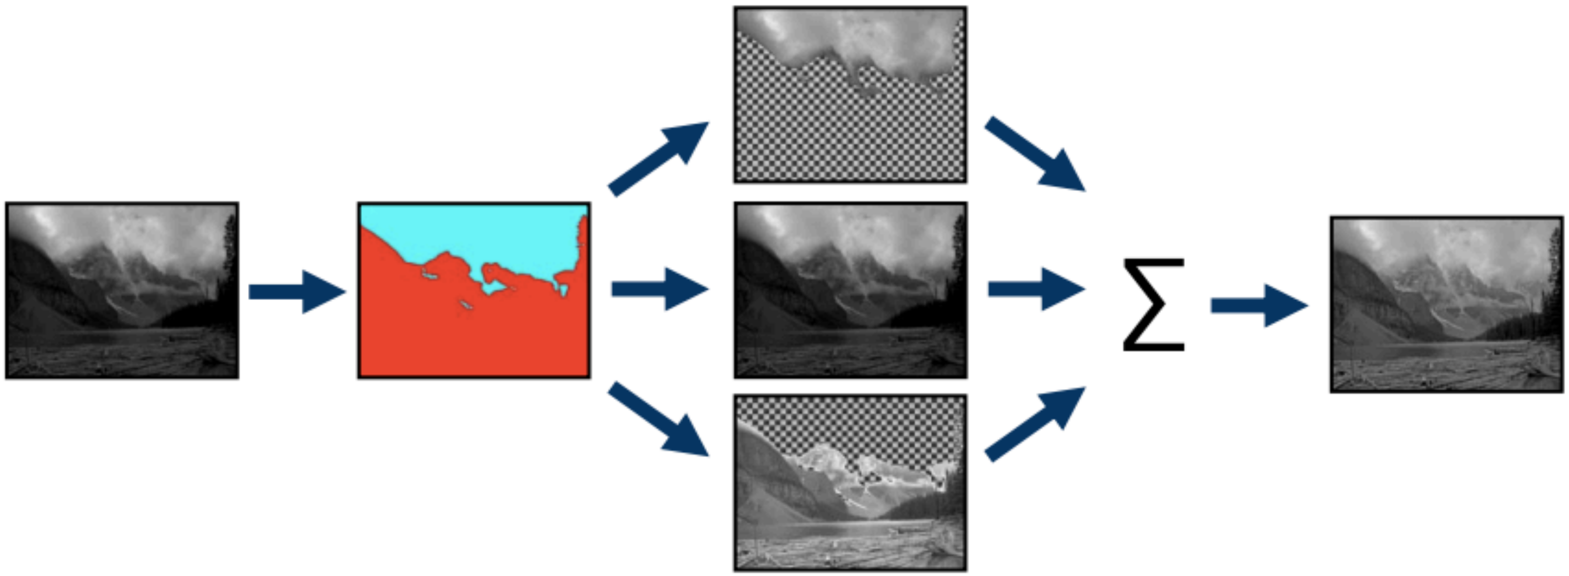

- Perceived lightness is anchored to several frameworks

Gilchrist lightness perception theory

- Frameworks – areas of common illumination
- Anchoring – the tendency of
  – highest luminance
  – largest area

  to appear white
- Tone-mapping
  – Rescale luminance in each framework to its anchor

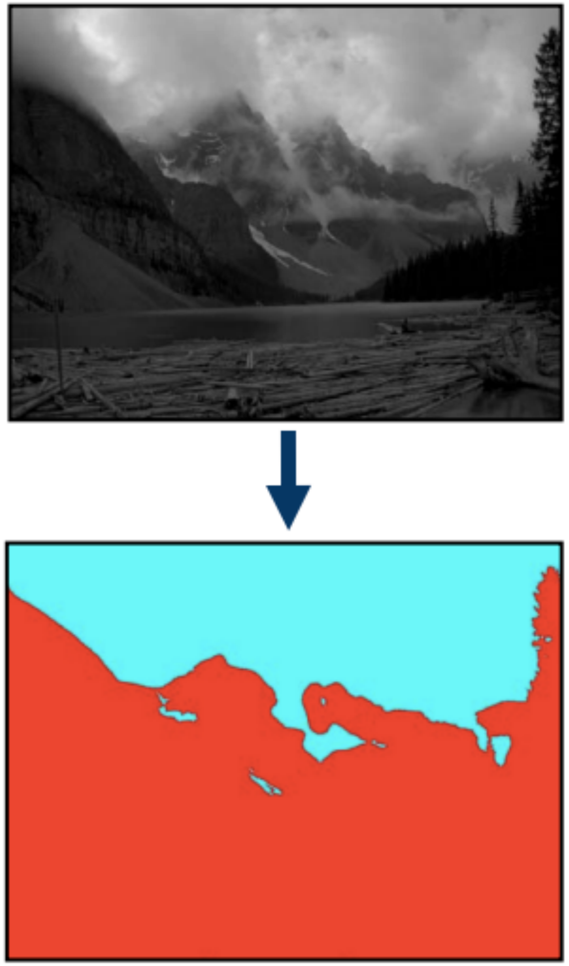

Results – lightness perception TMO

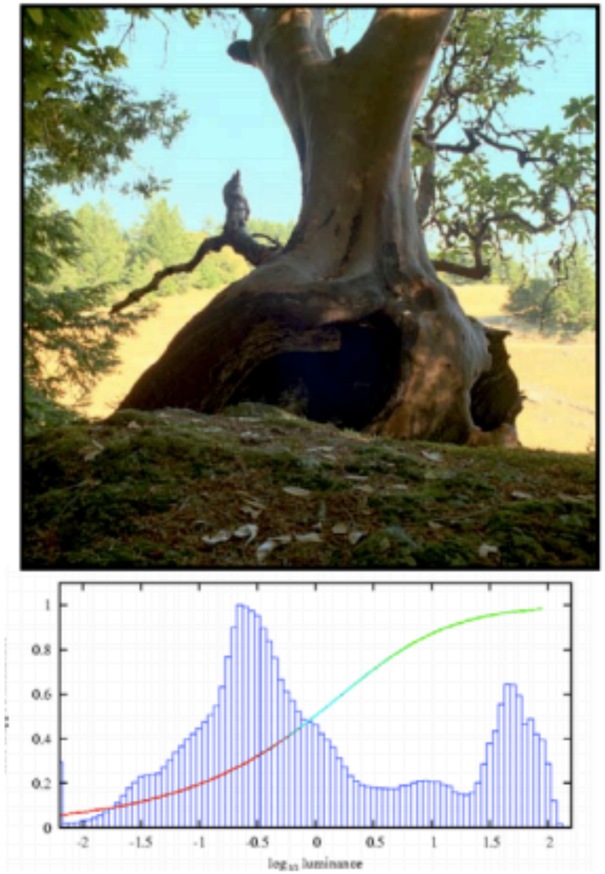

Photographic Tone Reproduction

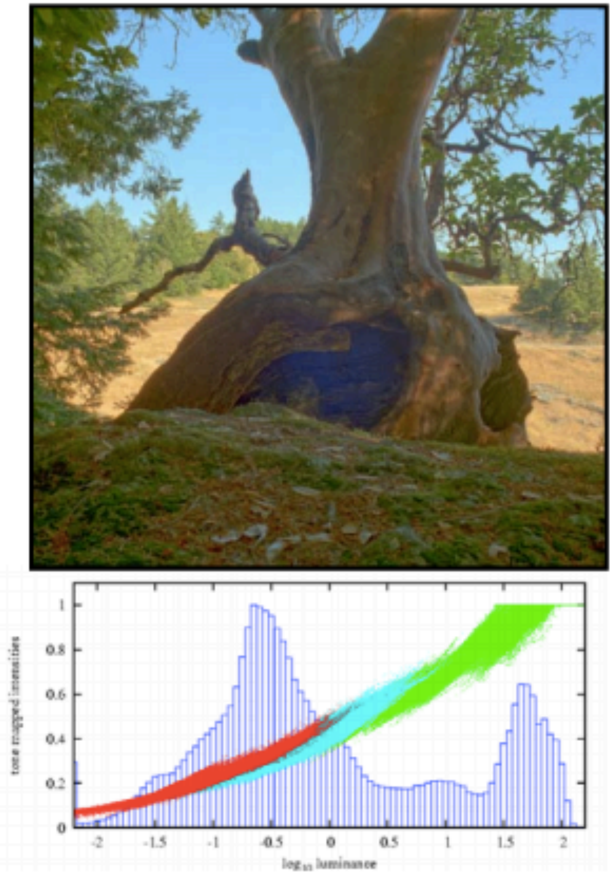

Bilateral Filtering

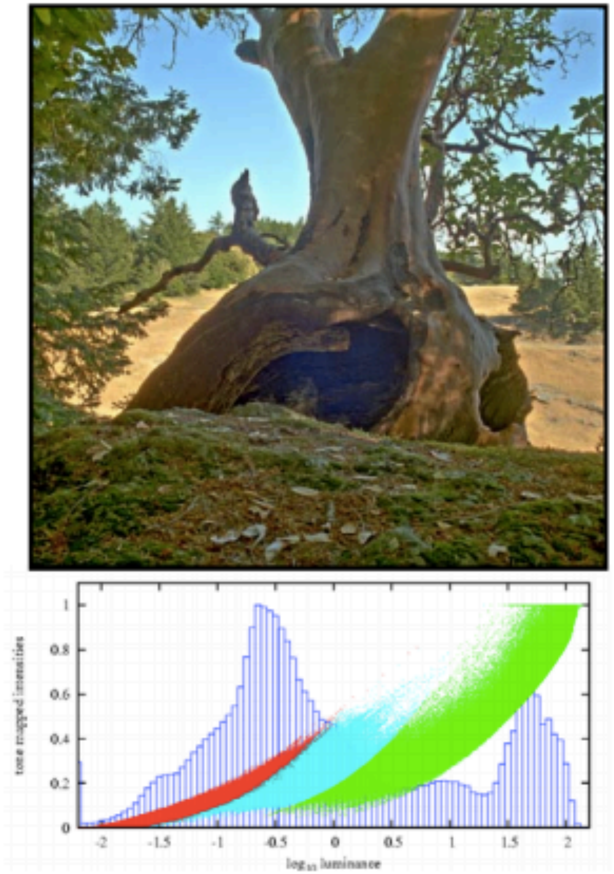

Presented Computational Model

Major approaches to tone-mapping

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

Forward and inverse visual model

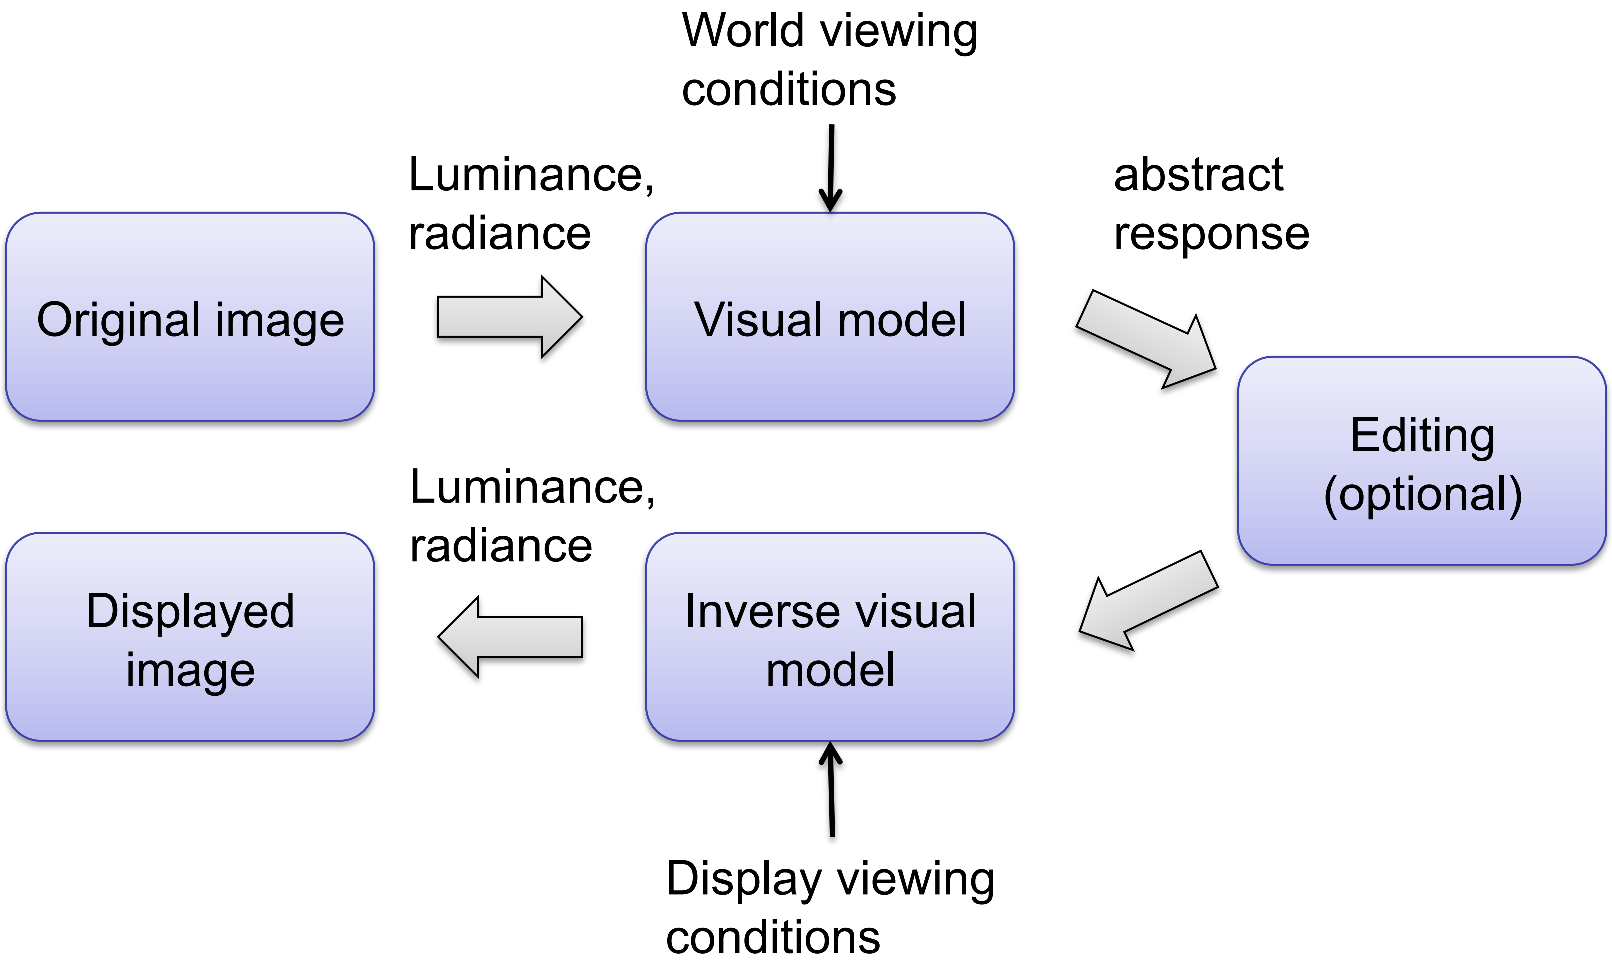

Contrast domain image processing

[Mantiuk et al., ACM Trans. Applied Perception, 2006]

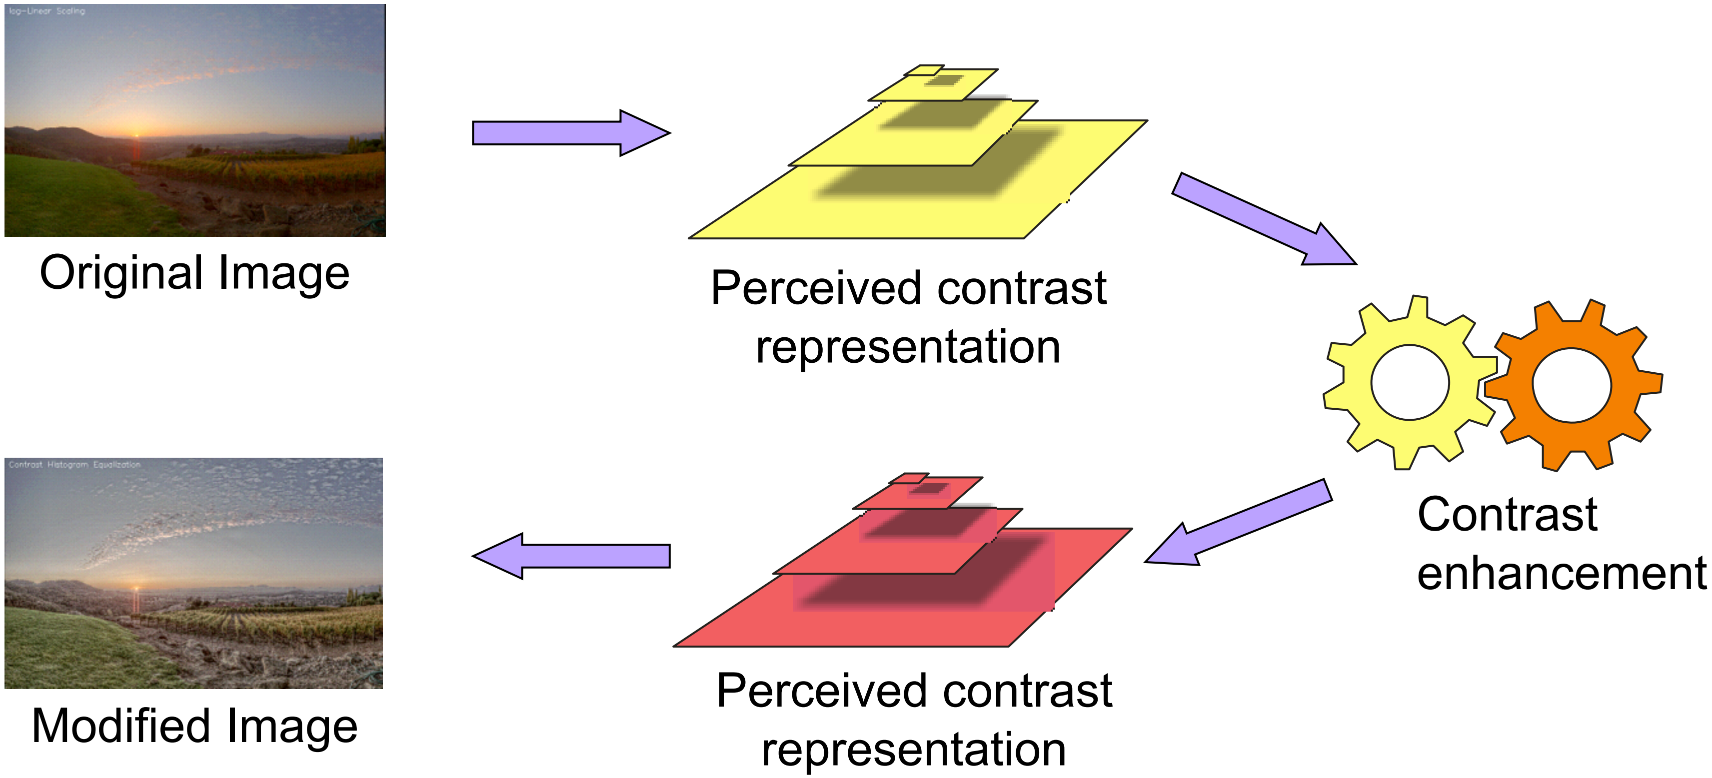

Rationale: Human eye is more sensitive to contrast than luminance

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

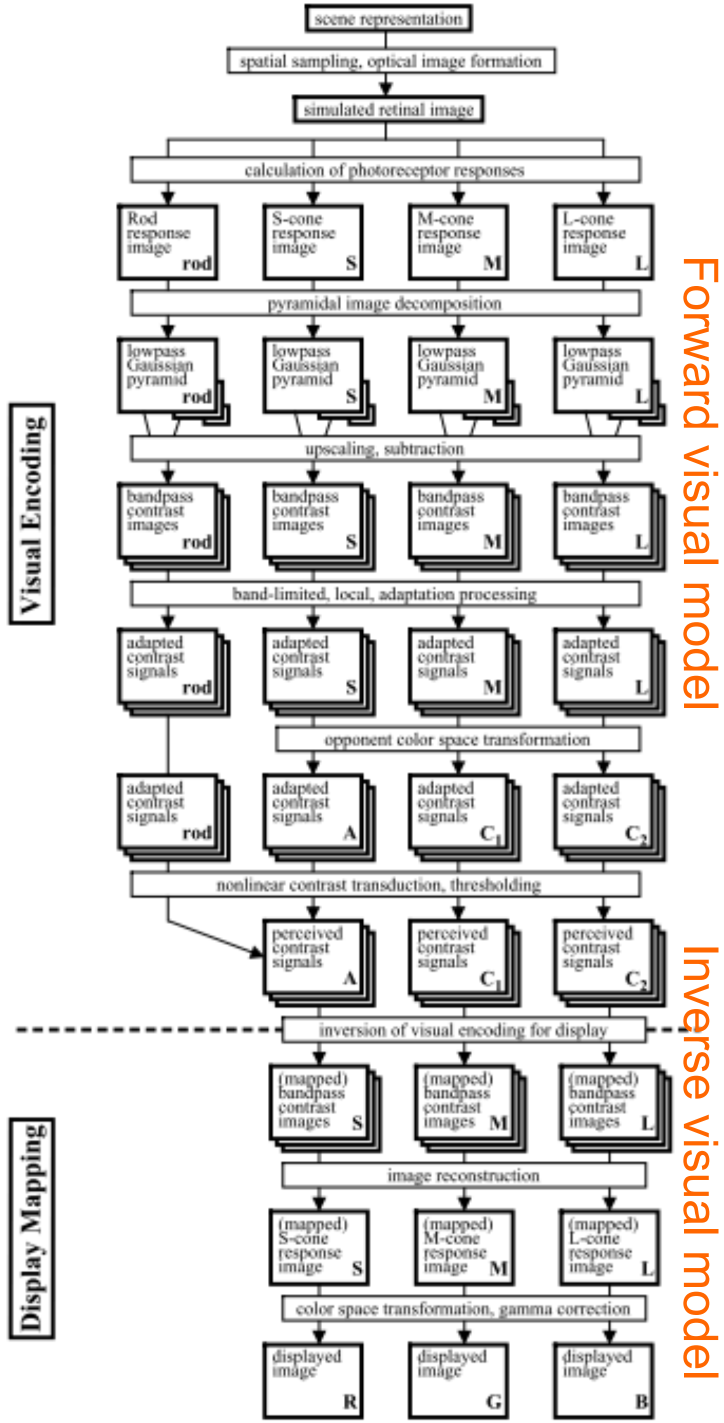

Results – multiscale model …

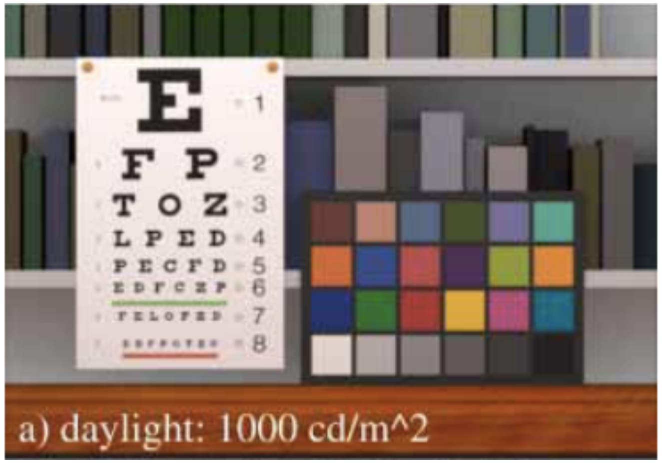

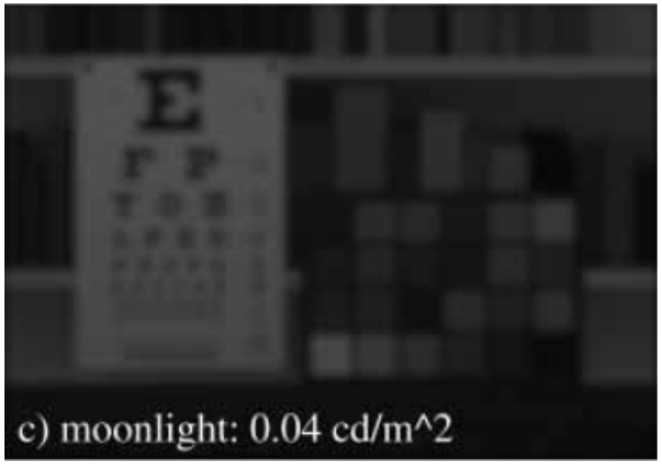

Forward and inverse visual model

- 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

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

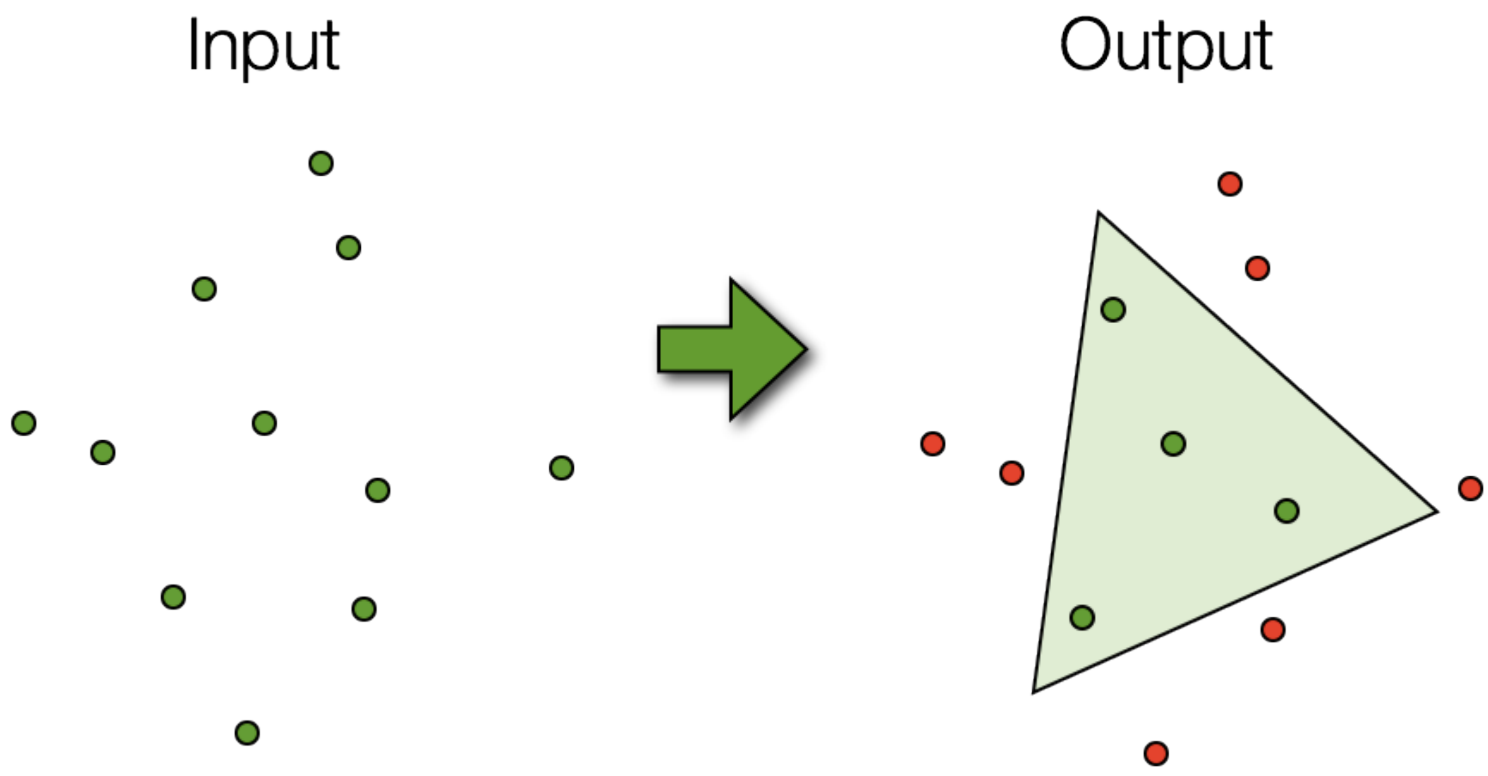

Global tone mapping operator

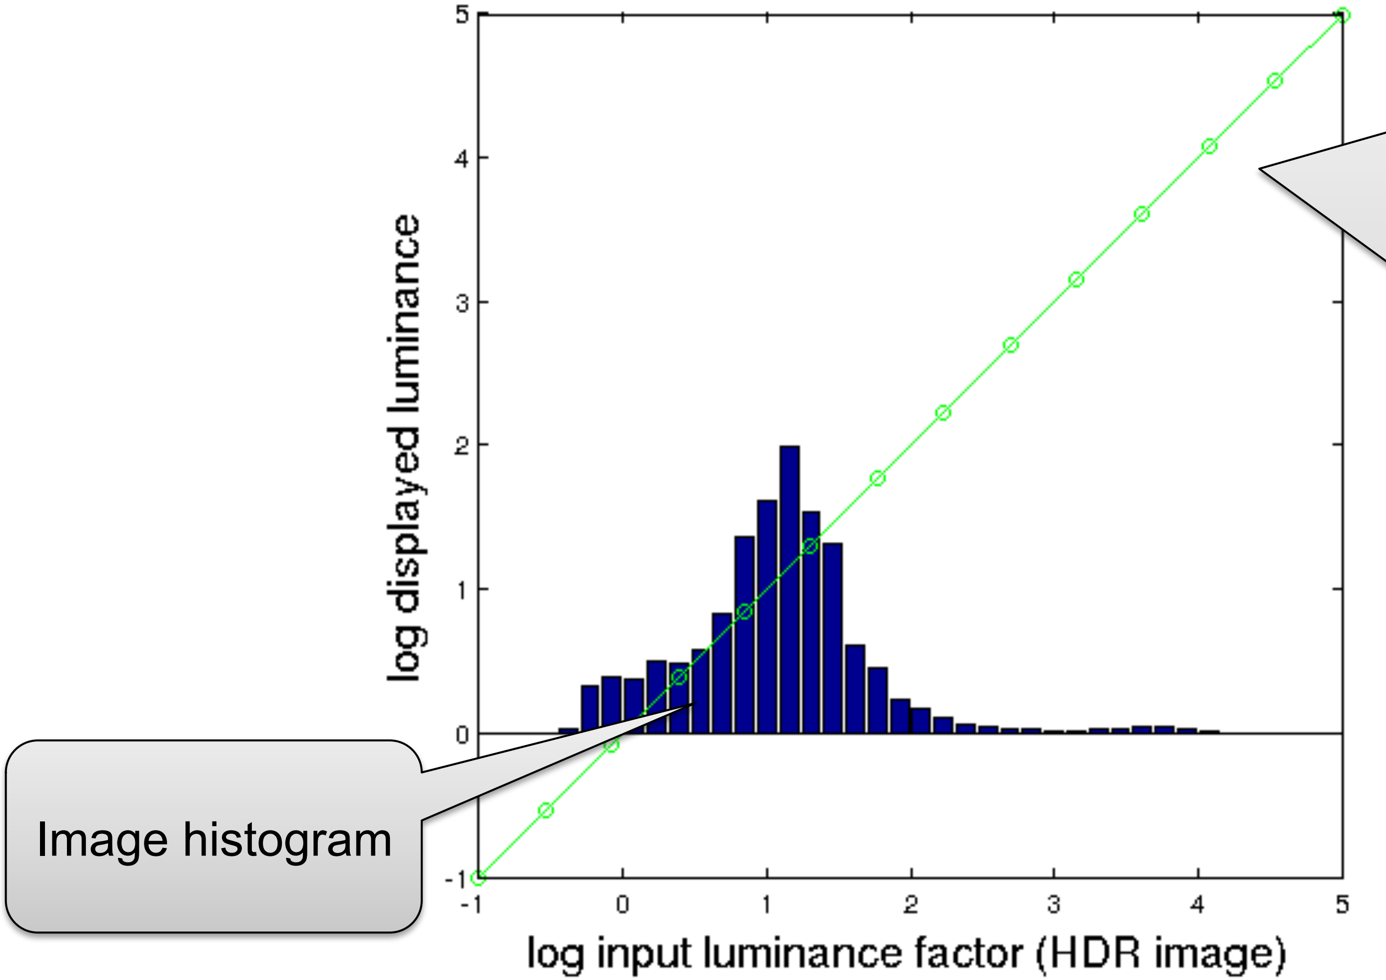

Best tone-mapping is the one which does not do anything, i.e. slope of the tone-mapping curves is equal to 1.

Display limitations

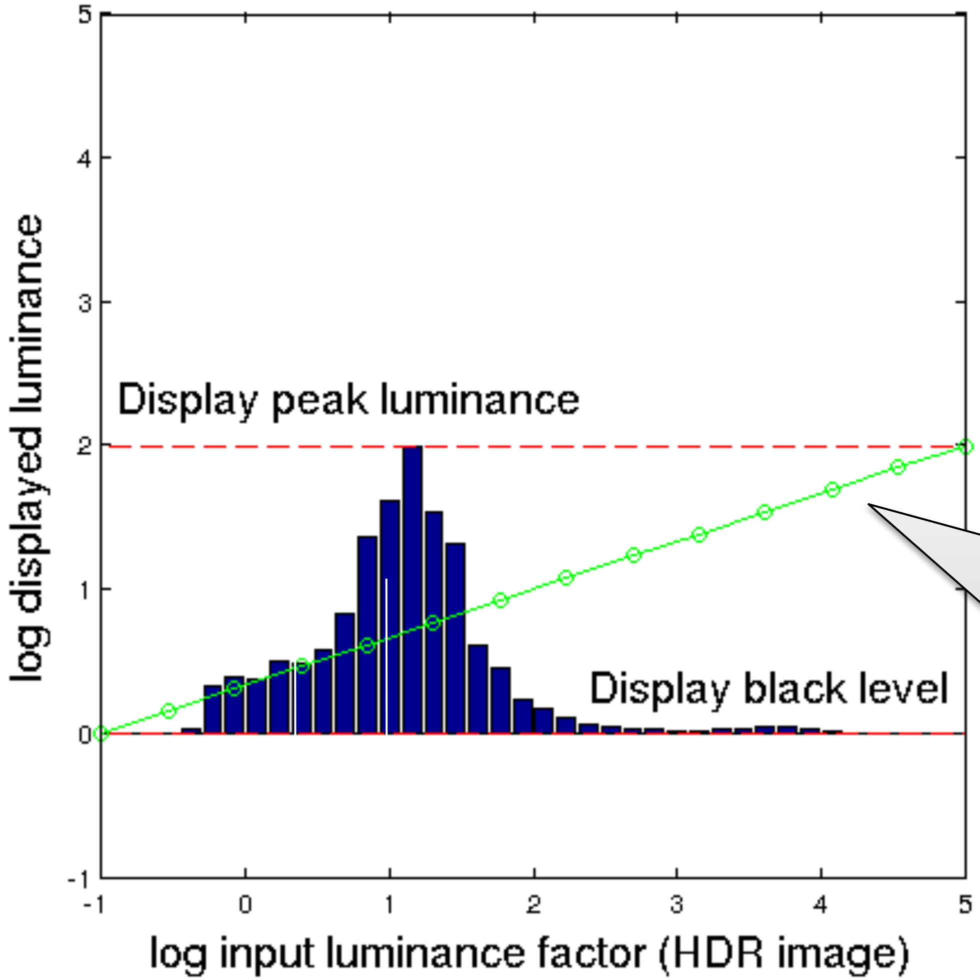

But in practice contrast (slope) must be limited due to display limitations.

Tone mapping

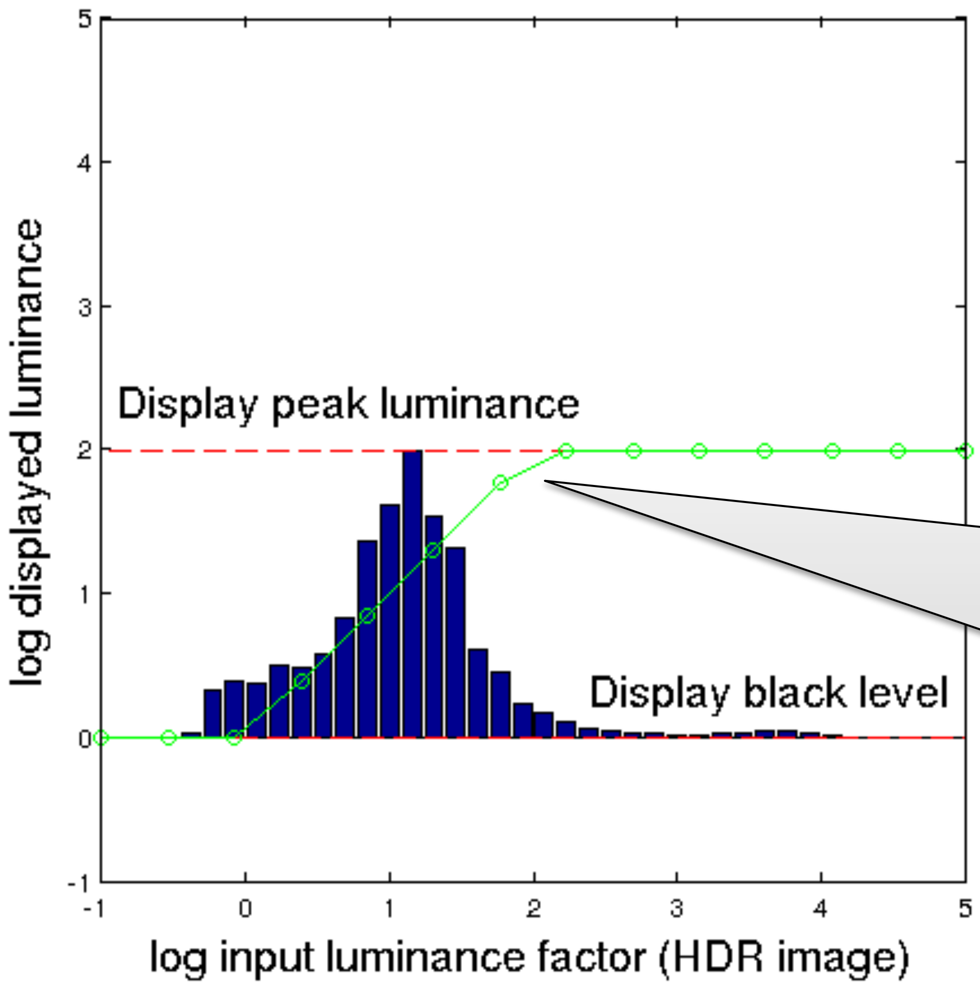

Global tone-mapping is a compromise between clipping and contrast compression.

Histogram equalization

- 1. Compute cumulative distribution function:

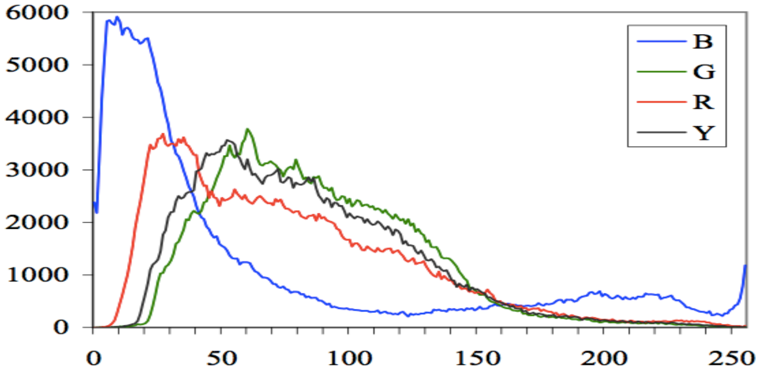

$$c(I) = \frac{1}{N}\sum_{i=0}^{I} h(i) = c(I-1) + \frac{1}{N}h(I)$$

- 2. Use that function to assign new pixel values

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

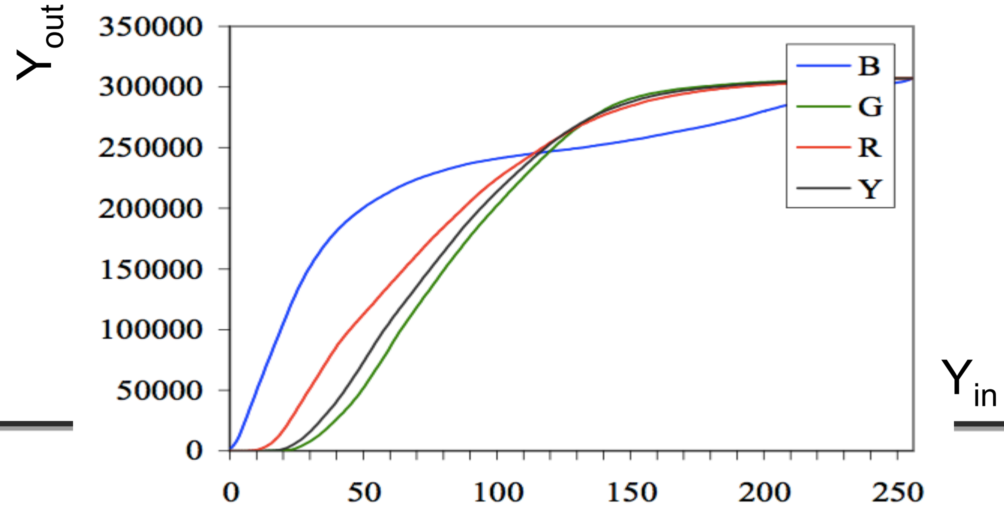

Histogram equalization

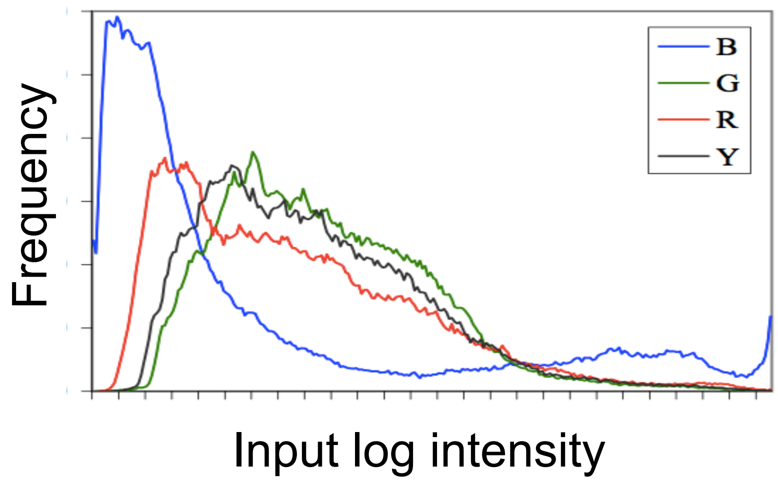

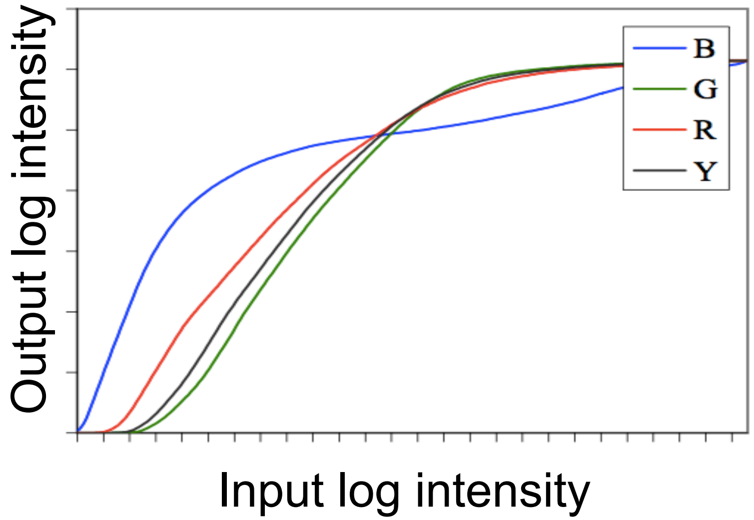

- Steepest slope for strongly represented bins
  - Enhance contrast, if many pixels
  - Reduce contrast, if few pixels
- HE distributes contrast distortions relative to the “importance” of a brightness level

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

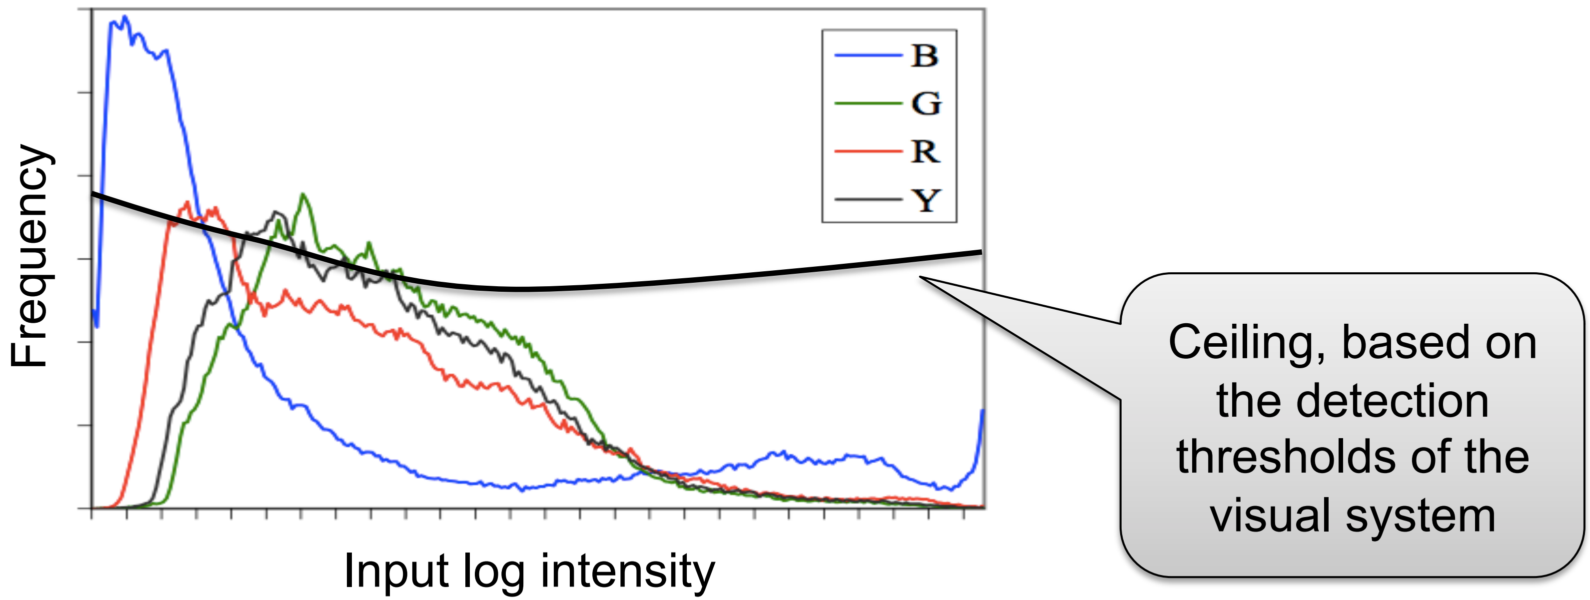

Display adaptive tone-mapping

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

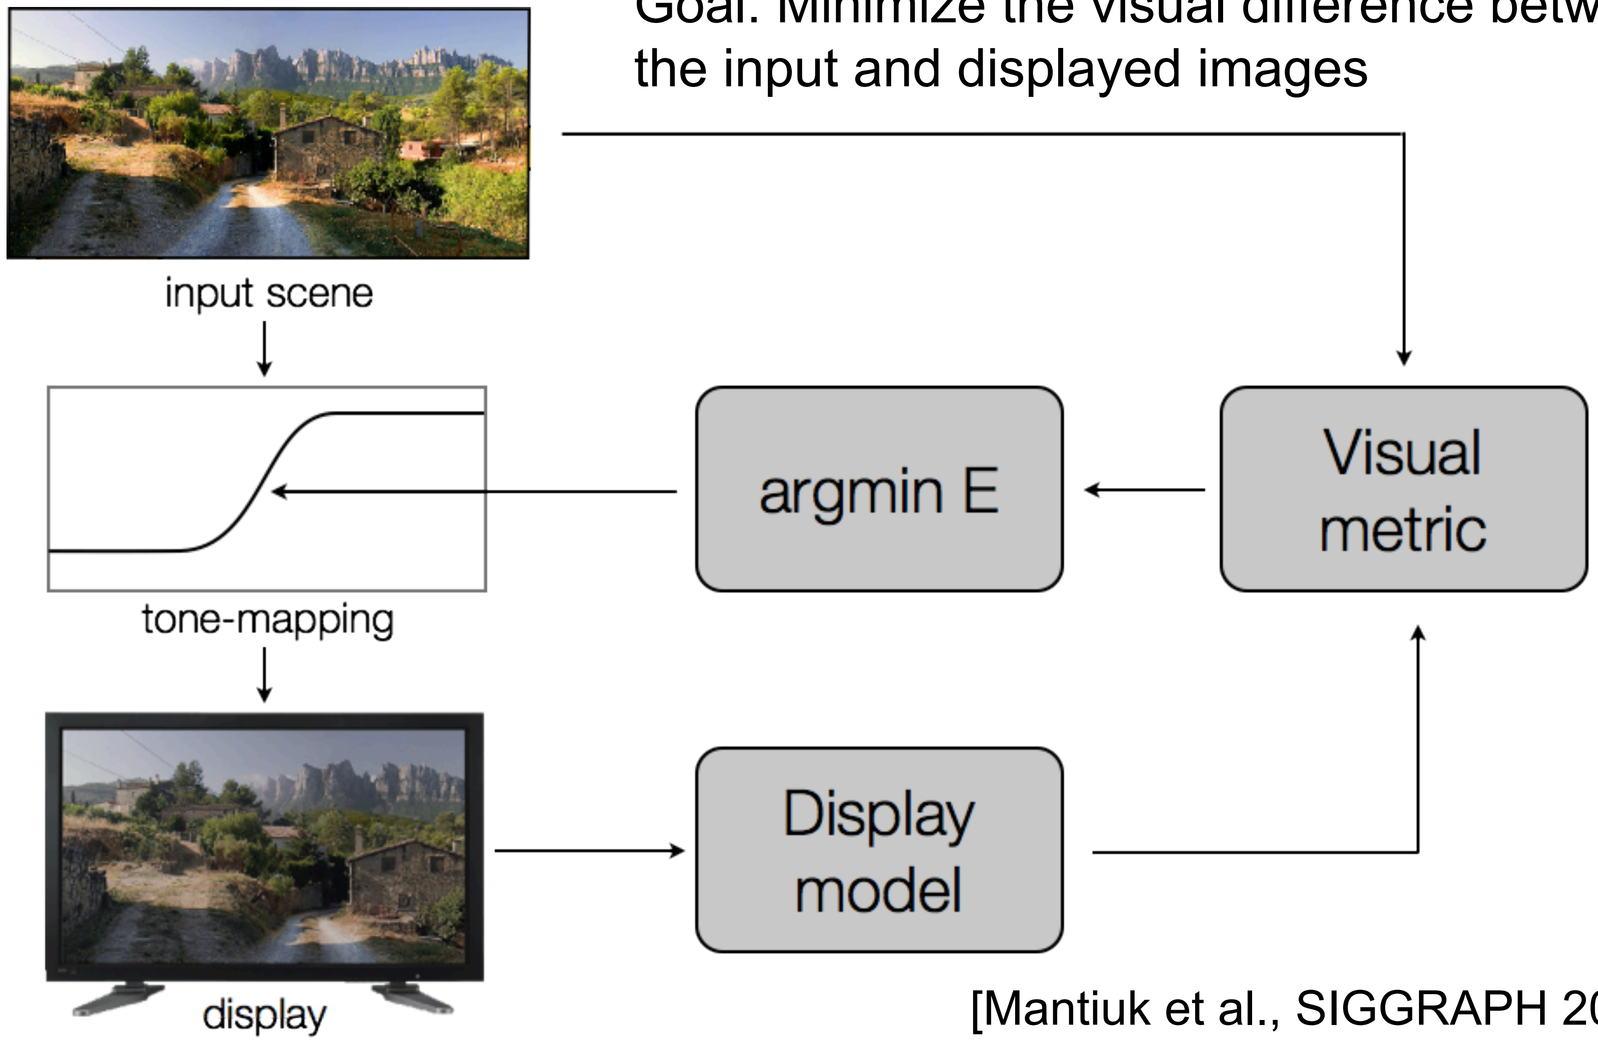

[Mantiuk et al., SIGGRAPH 2008]

Display adaptive tone-mapping

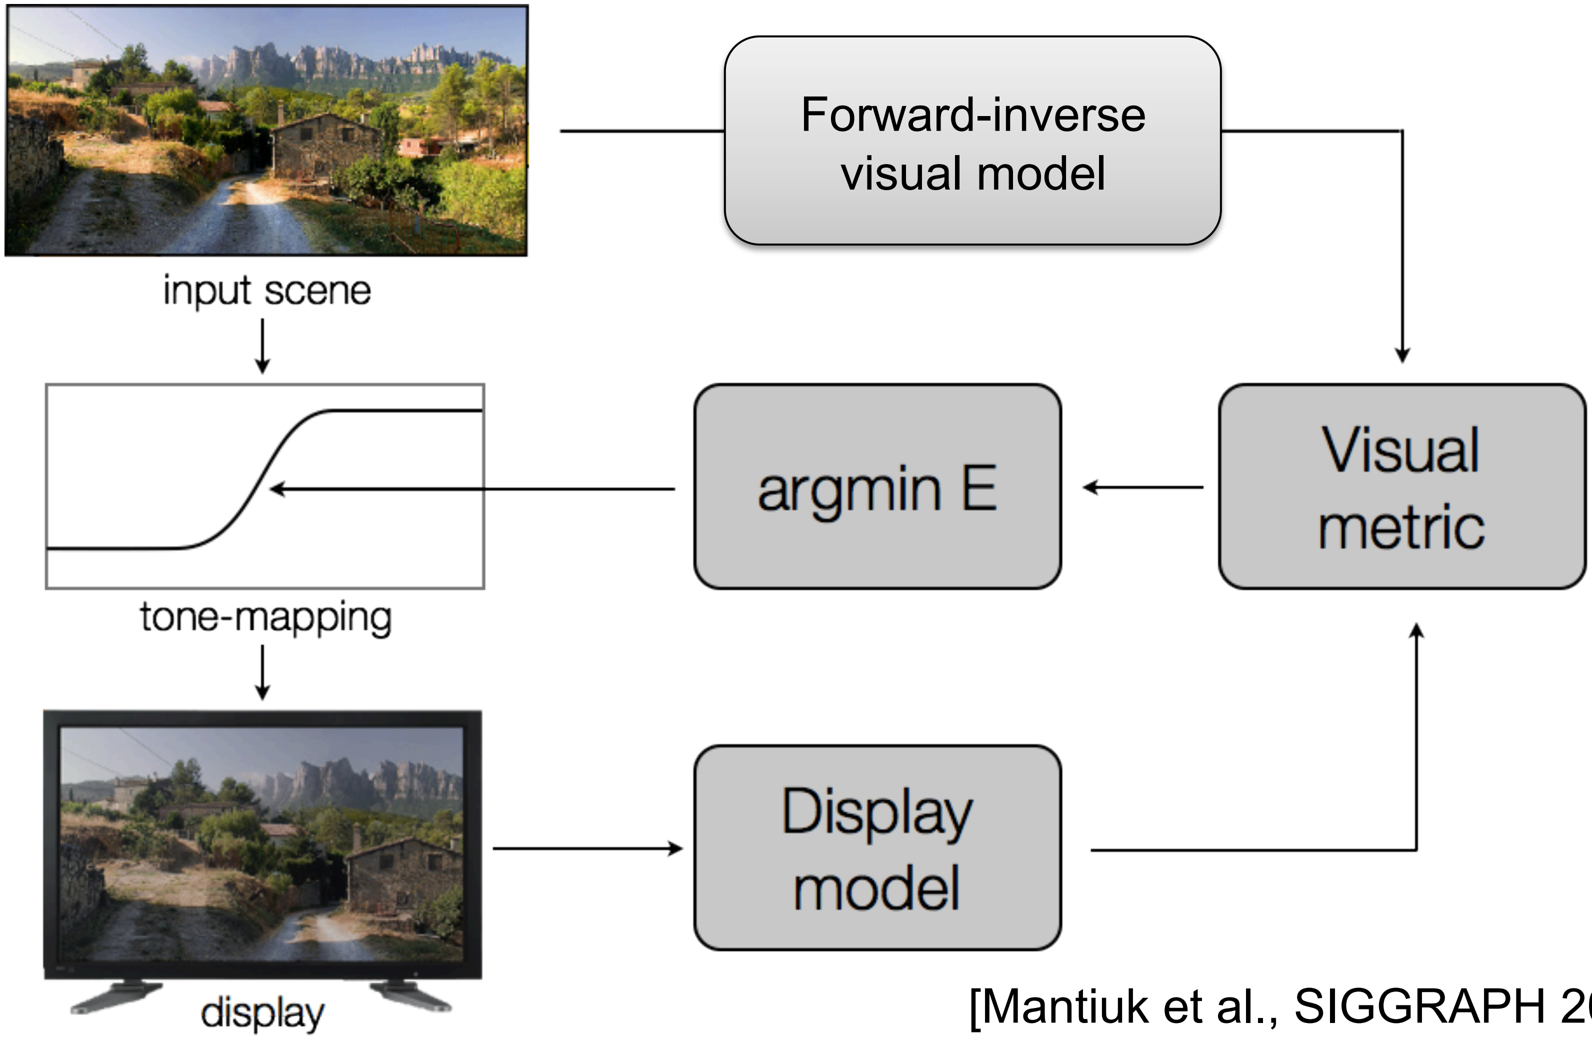

[Mantiuk et al., SIGGRAPH 2008]

Results: ambient illumination compensation

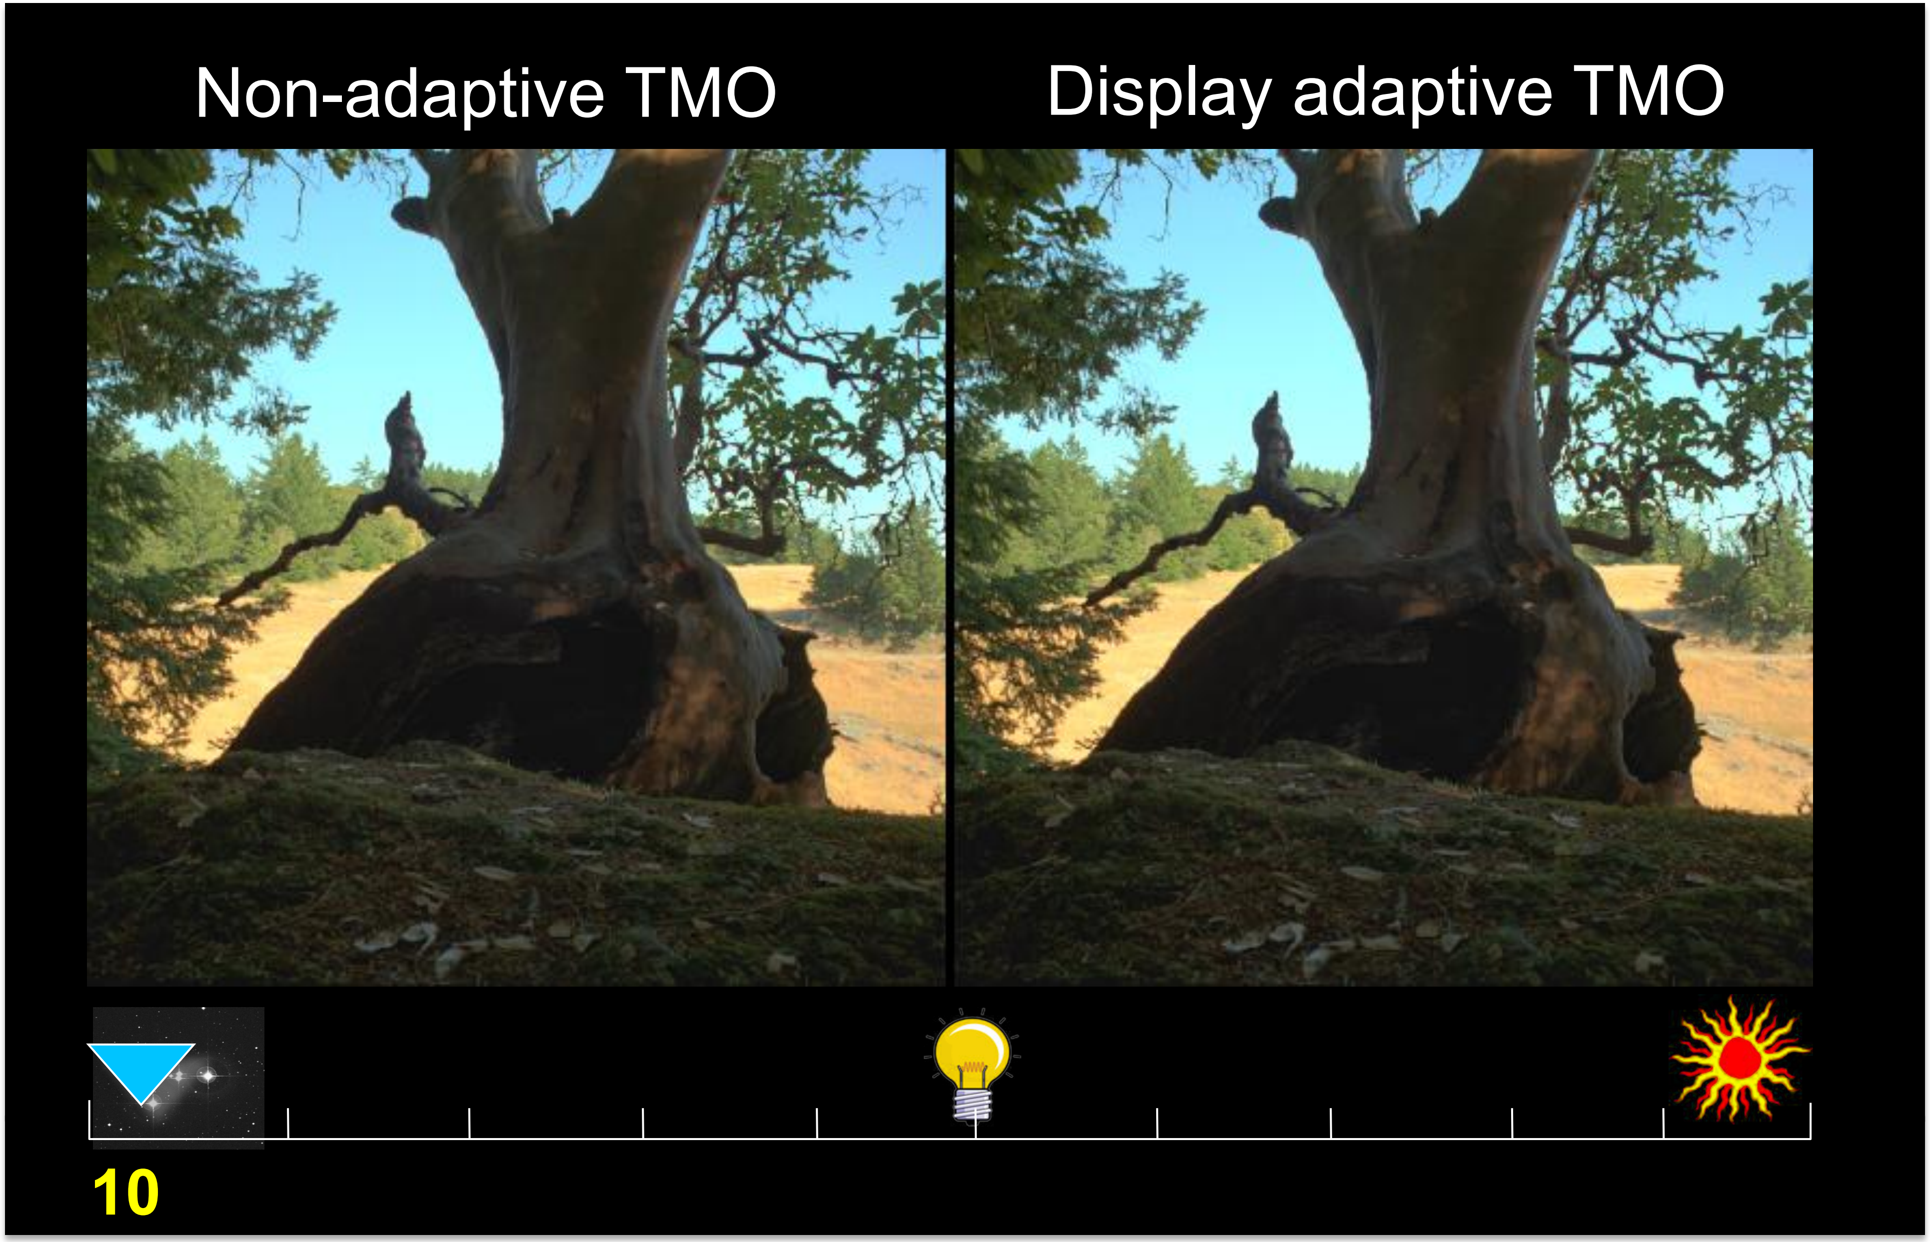

Results: ambient illumination compensation

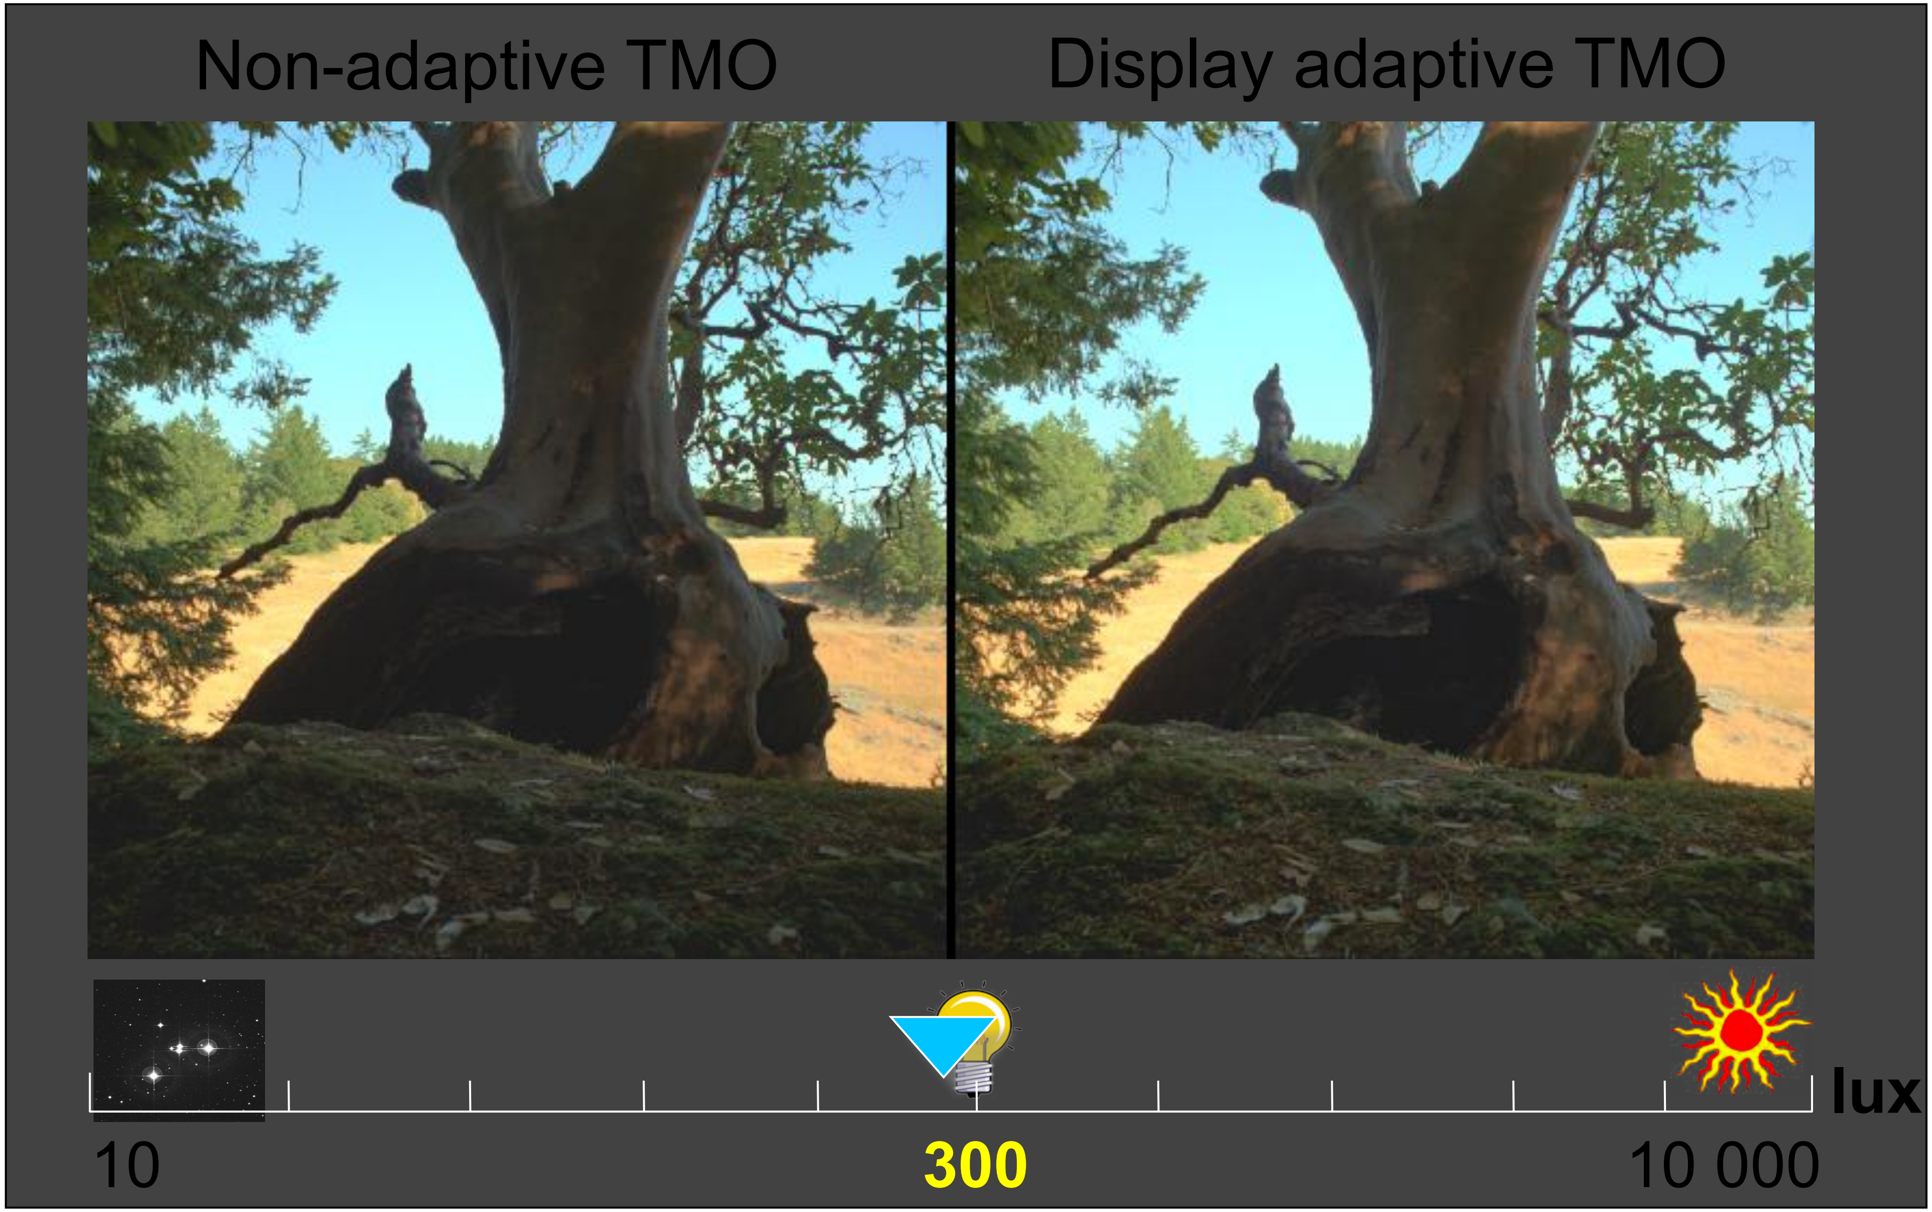

Results: ambient illumination compensation

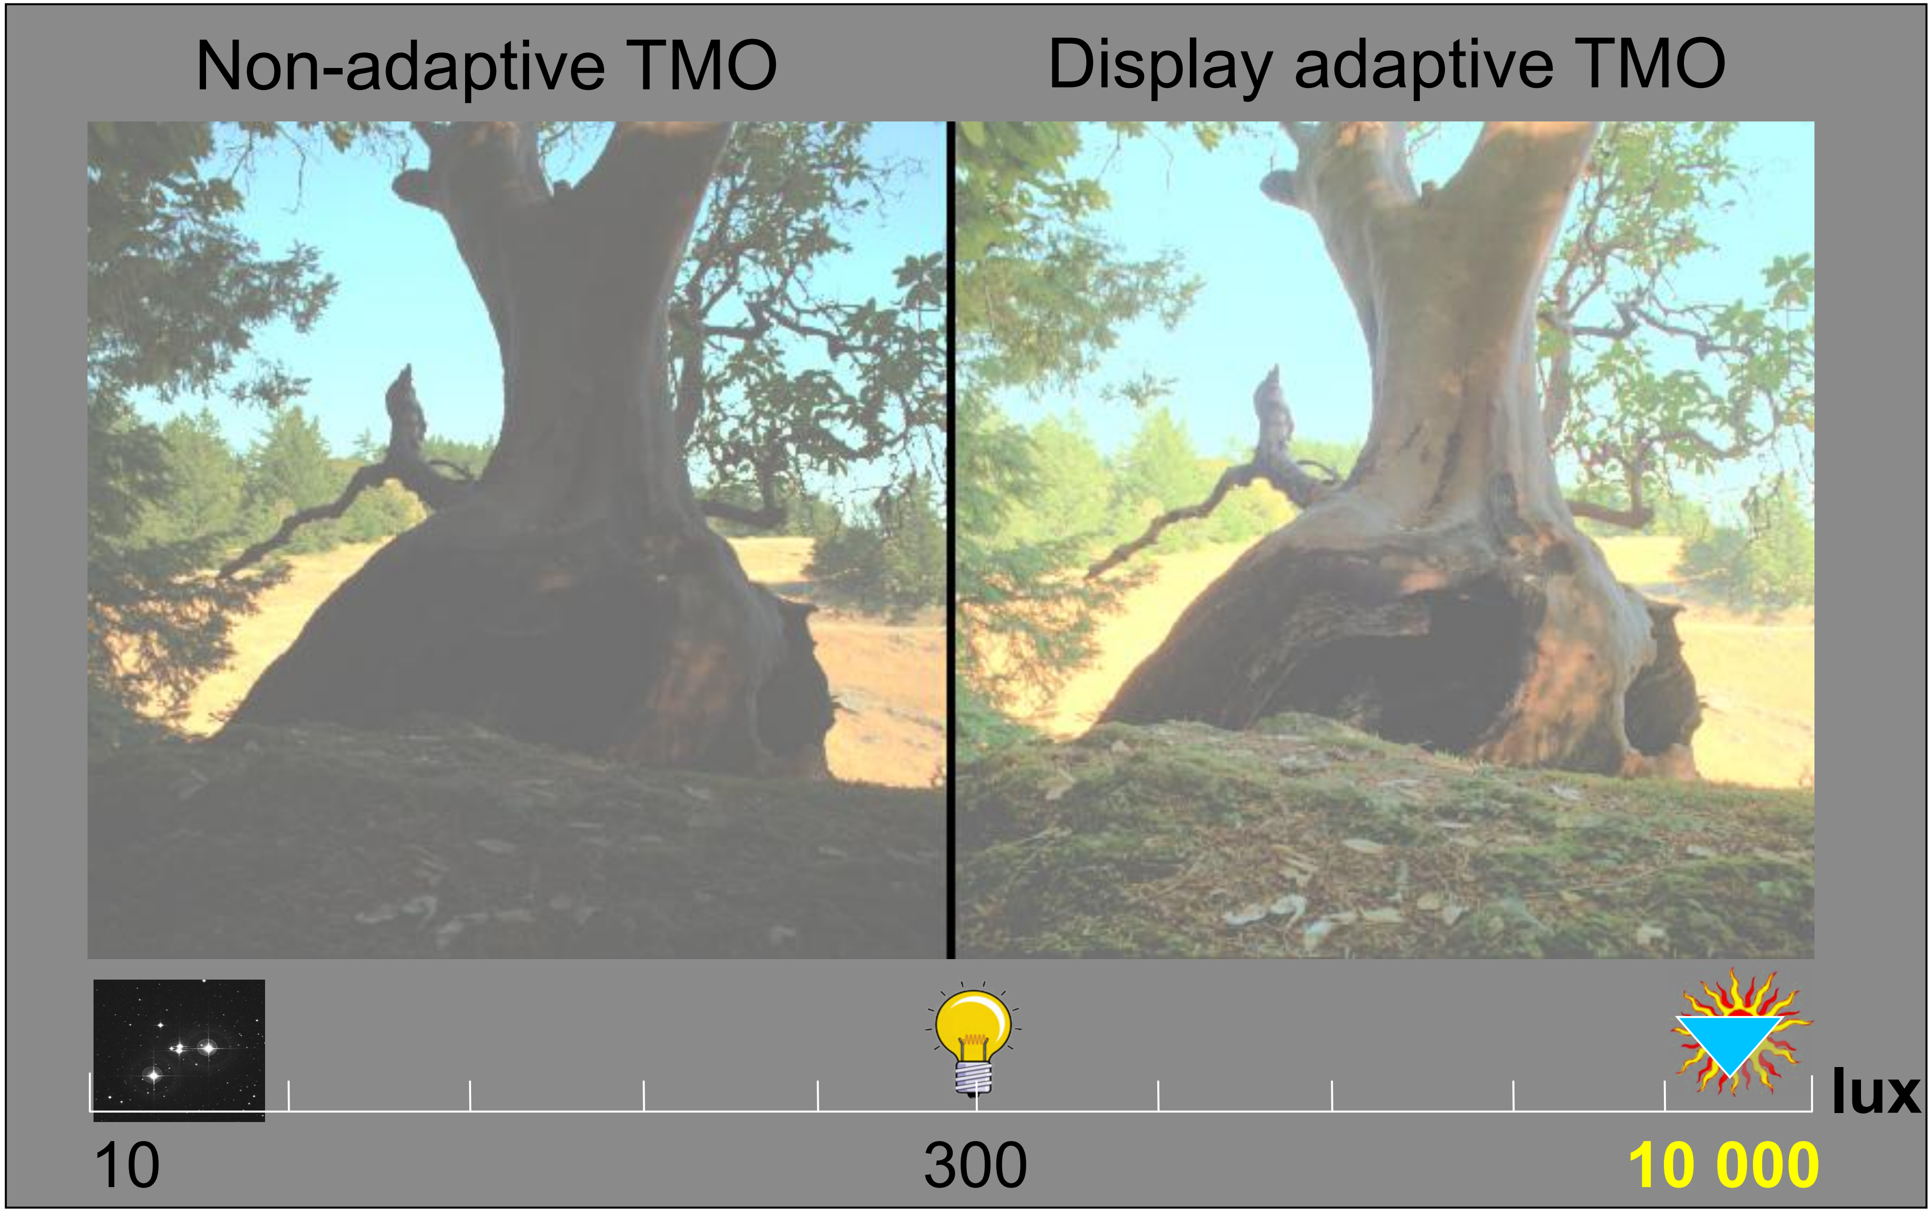

Results: display contrast

ePaper

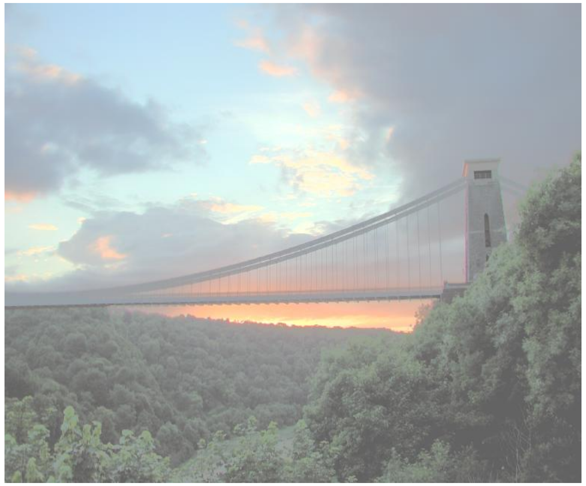

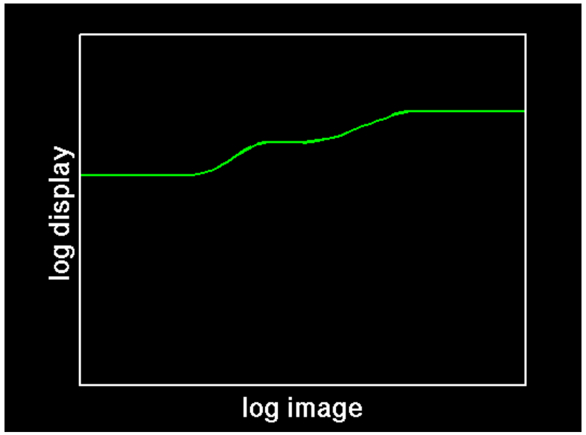

standard LCD

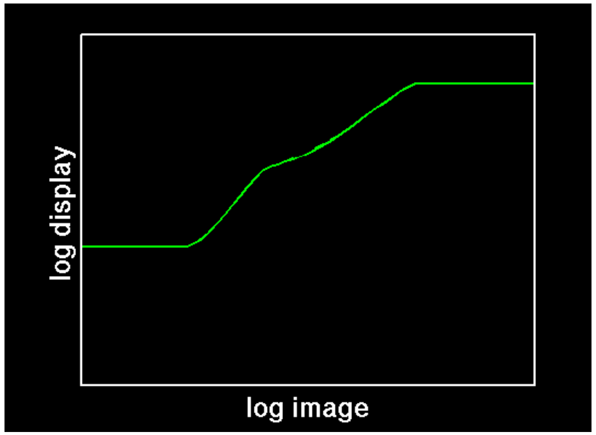

HDR display

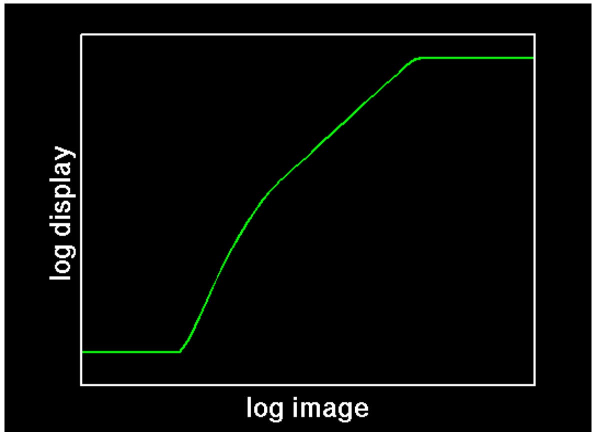

Tone-mapping for video compression

- Find the tone-curve that minimizes distortion in a backward-compatible HDR video encoding

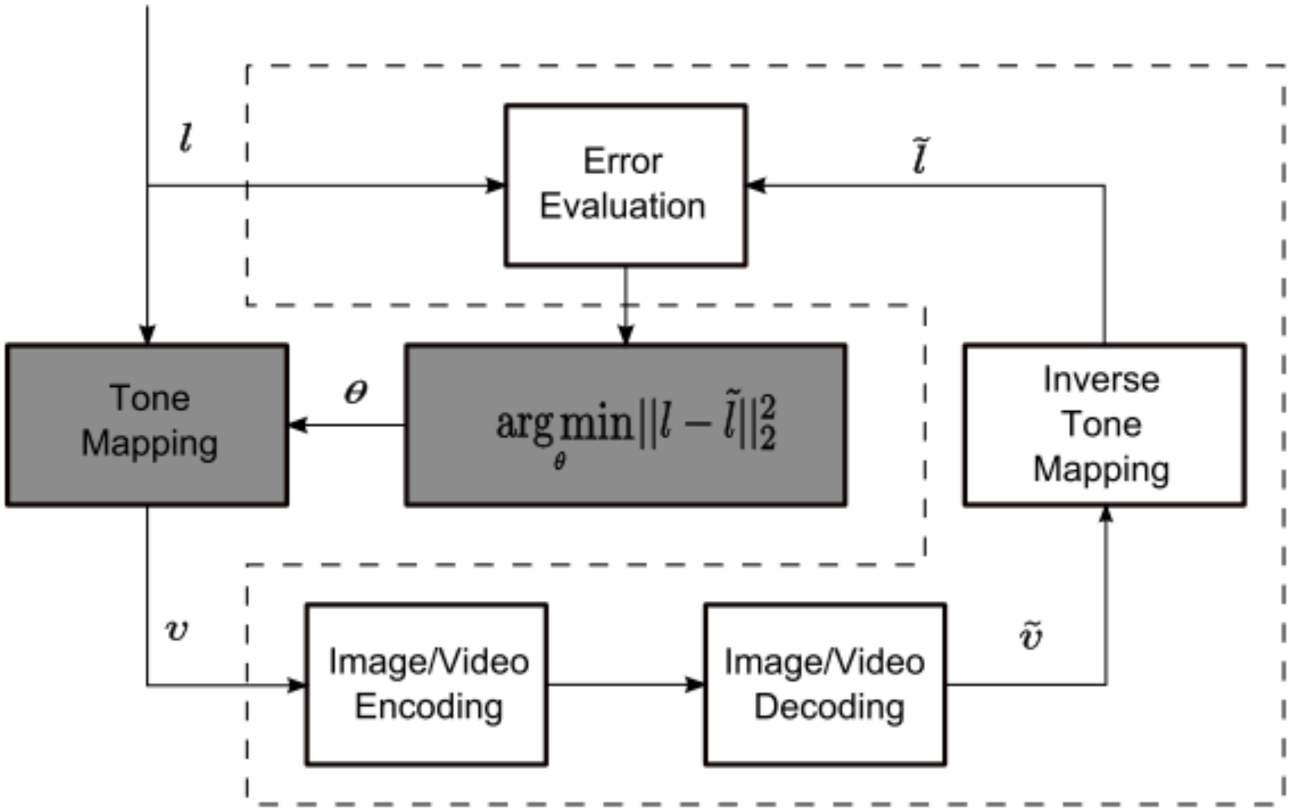

[Mai et al., IEEE TIP 2010]

Closed-form solution:

$$s_k = \frac{v_{max} \cdot p_k^{1/3}}{\delta \cdot \sum_{k=1}^{N} p_k^{1/3}}$$

Which tone-mapping to choose?

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

1. Think what is the target application
   - and thus the goal of your tone-mapping
2. Consider which tone-mapping approach(es) will deliver that goal

Future of tone-mapping

Tone-mapping of today

– Built into cameras
– Assumes that all displays are the same

Tone-mapping of tomorrow

– Display tone-maps content on demand
– Depending on viewing conditions, viewer, its capabilities
– Content recorded, stored and transmitted in an HDR format

Thank you

SIGGRAPHASIA2011 HONG KONG