UNIVERSITY OF CAMBRIDGE COMPUTER LABORATORY

RAINBOW RESEARCH GROUP

High dynamic range and tone mapping

Advanced Graphics

Rafał Mantiuk
Computer Laboratory, University of Cambridge

Cornell Box: need for tone-mapping in graphics

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

Dynamic range

Dynamic range (contrast)

- As ratio:

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

  - Usually written as C:1, for example 1000:1.

- 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)

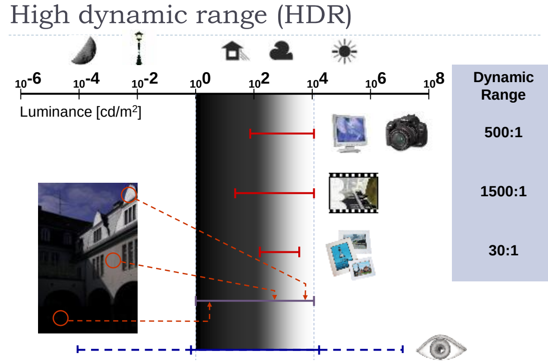

HDR in products

- UltraHD + HDR
  - Because the resolution alone gives only small improvement in image quality
- 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

7

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
  - The current standards not fully follow to this principle

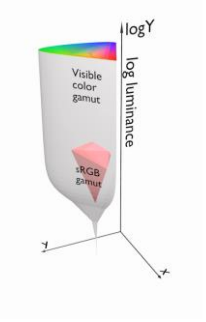

8

Tone-mapping problem

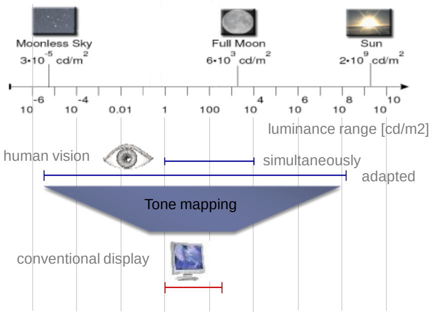

9

Tone-mapping in rendering

- Any physically-based rendering requires tone-mapping
- "HDR rendering" in games is pseudo-physically-based rendering
- Goal: to simulate a camera or the eye
- Greatly enhances realism

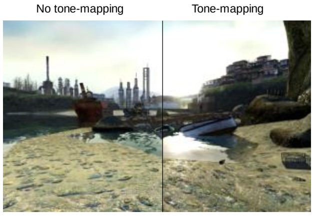

Half-Life 2: Lost coast

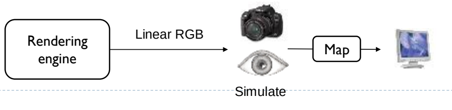

10

Three intents of tone-mapping

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

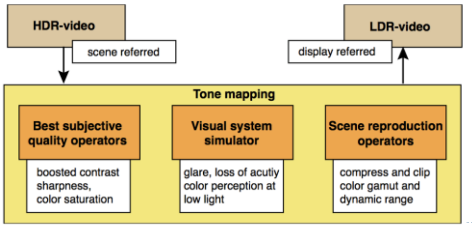

11

Intent #1: Best subjective quality

- Often interactive
- Software
  - Photoshop
  - Lightroom
  - Photomatix
- Techniques
  - Color-grading
- Artistic intent

12

Intent #2: Visual system simulator

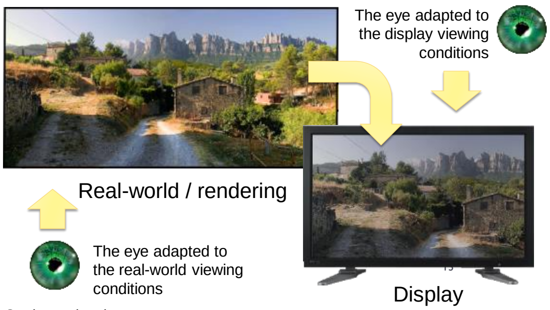

Goal: match color appearance

Visual system simulator - example

- Simulation of glare

14

Intent #3: Scene reproduction problem

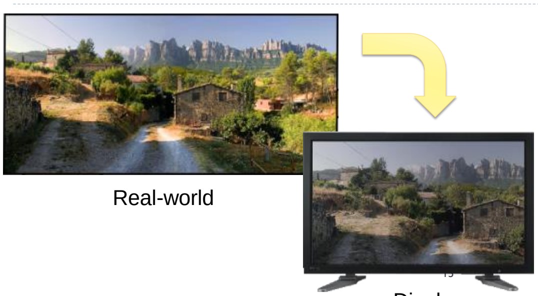

Goal: map colors to a restricted color space

Mapping problem

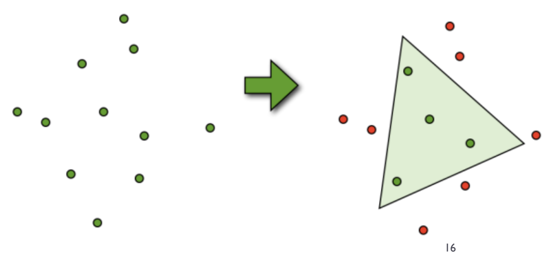

16

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

17

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)
- The same operations can be applied to linear RGB
  - Or to luminance-only and the color can be transferred

18

Multiplication – brightness change

Resulting luminance / Input luminance

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

Brightness change parameter

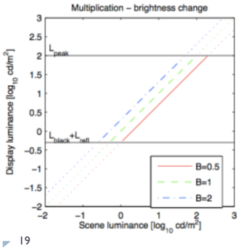

- Multiplication makes the image brighter or darker
- It does not change the dynamic range!

Power function – contrast change

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

Contrast change (gamma) / Luminance of white

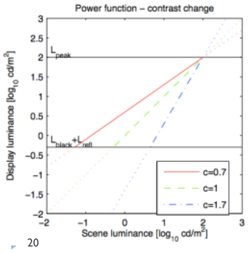

- Power function stretches or shrinks image dynamic range
- It is usually performed relative to reference white
- Apparent brightness changes is the side effect of pushing tones towards or away from the white point

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

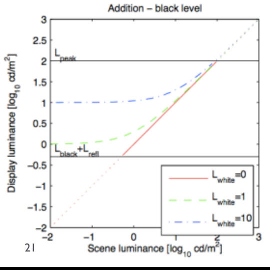

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

Two ways to do tone-mapping

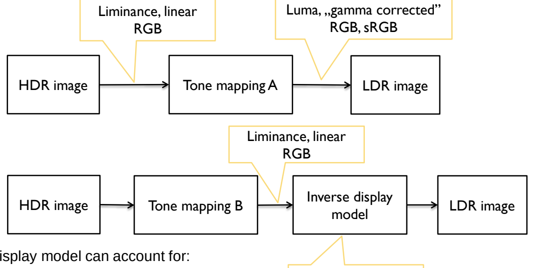

Display model can account for:

- Display peak luminance
- Display dynamic range (contrast)
- Ambient light

Sometimes known as "gamma"

Display model

- Tone-mapping needs to account for the physical model of the display
  - How a display transforms pixel values into emitted light

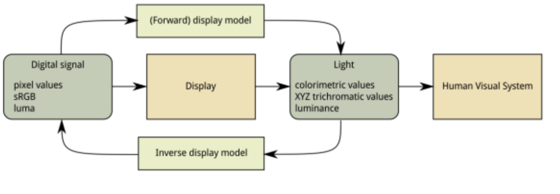

(Forward) Display model

- GOG: Gain-Gamma-Offset

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

Luminance: L; Peak luminance: L_{peak}; Display black level: L_{black}; Gamma: γ; Screen reflections: L_{refl}; Gain: $(L_{peak} - L_{black})$; Pixel value 0-1: V; Offset: $L_{black} + L_{refl}$

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

Reflectance factor (0.01): k; Ambient illumination (in lux): E_{amb}

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.

Ambient illumination compensation

Ambient illumination compensation

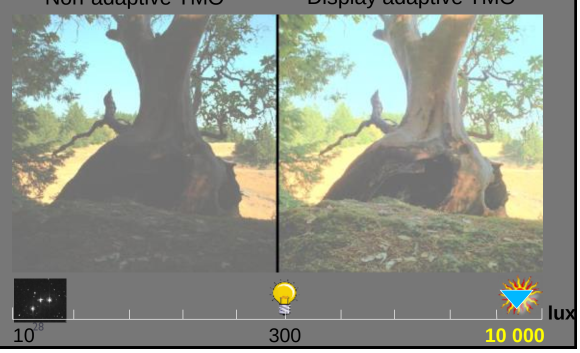

Example: Ambient light compensation

- We are looking at the screen in bright light

$L_{peak} = 100\ [cd \cdot m^{-2}]$ $\quad k = 0.005$ (Modern screens have reflectivity of around 0.5%)

$L_{black} = 0.1\ [cd \cdot m^{-2}]$

$E_{amb} = 2000\ [lux]$ $\quad L_{refl} = \frac{0.005}{2\pi} 2000 = 1.59\ [cd \cdot m^{-2}]$

- We assume that the dynamic of the input is 2.6 (≈400:1)

$$r_{in} = 2.6 \qquad r_{out} = \log_{10} \frac{L_{peak}}{L_{black} + L_{refl}} = 1.77$$

- First, we need to compress contrast to fit the available dynamic range, then compensate for ambient light

$$L_{out} = \left(\frac{L_{in}}{L_{peak}} \right)^{\frac{r_{out}}{r_{in}}} - L_{refl}$$

The resulting value is in luminance, must be mapped to display luma / gamma corrected values

Simplest, but not the best tone mapping

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

Tone-curve

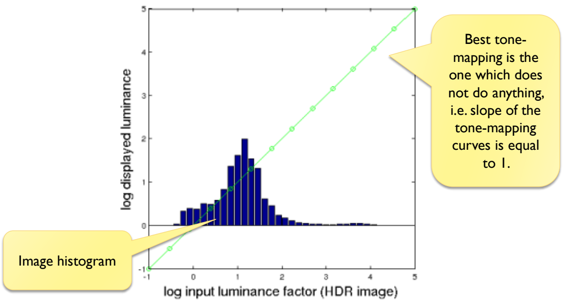

Tone-curve

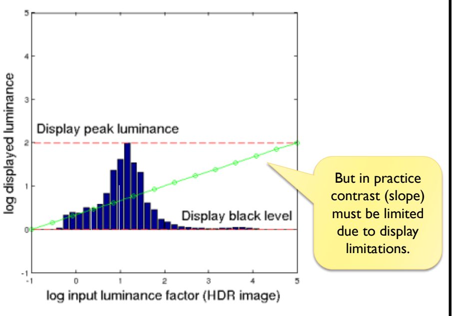

Tone-curve

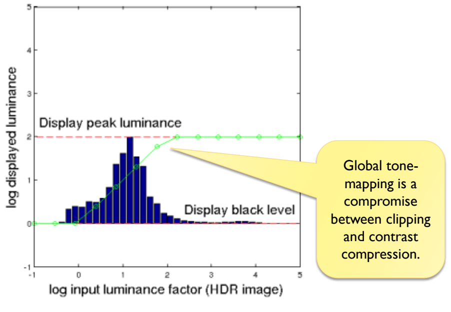

Sigmoidal tone-curves

- Very common in digital cameras
  - Mimic the response of analog film
  - Analog film has been engineered over many years to produce good tone-reproduction (given that he tone curve must not change)
- In practice - the most commonly used tone-mapping!

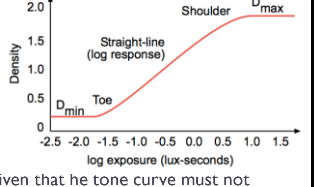

Photoreceptor response

- 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:

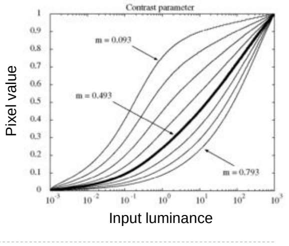

Photoreceptor TMO

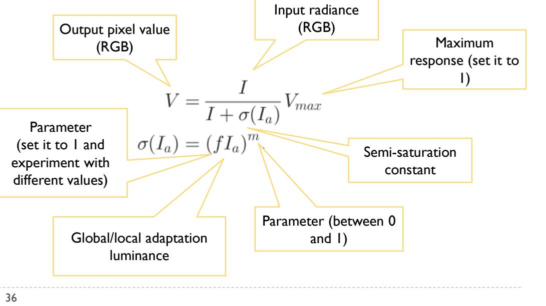

$$V = \frac{I}{I + \sigma(I_a)} V_{max}$$

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

Results: photoreceptor TMO

Histogram equalization

- 1. Compute cummulative image histogram

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

  - For HDR, operate in the log domain
- 2. Use the cummulative histogram as a tone-mapping function

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

  - For HDR, map the log-10 values to the $[dr_{out}; 0]$ range
    - where dr_{out} is the target dynamic range (of a display)

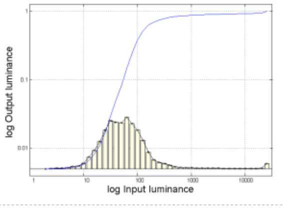

38

Histogram equalization

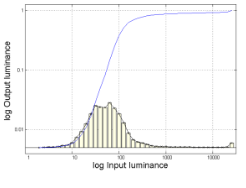

- Steepest slope for strongly represented bins
  - If many pixels have the same value - enhance contrast
  - Reduce contrast, if few pixels
- Histogram Equalization distributes contrast distortions relative to the "importance" of a brightness level

39

Histogram adjustment with a linear ceiling

- [Larson et al. 1997, IEEE TVCG]

Linear mapping | Histogram equalization | Histogram equalization with ceiling

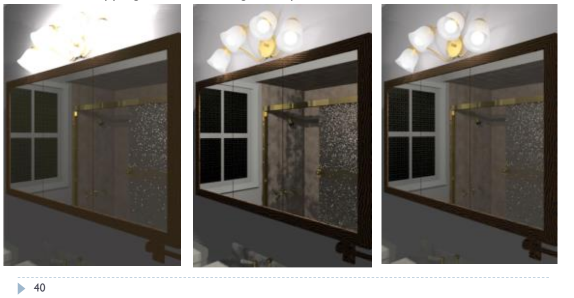

40

Histogram adjustment with a linear ceiling

- Truncate the bins that exceed the ceiling;
- Distribute the removed counts to all bins;
- Repeat until converges

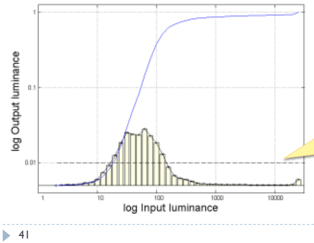

Ceiling, based on the maxiumum permissibble contrast

41

Histogram adjustment with a linear ceiling

- Truncate the bins that exceed the ceiling;
- Distribute the removed counts to all bins;
- Repeat until converges

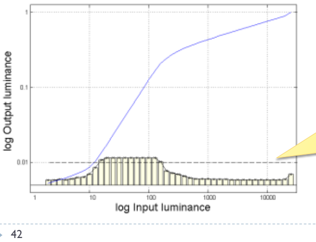

Ceiling, based on the maxiumum permissibble contrast

42

Tone-curve as an optimization problem

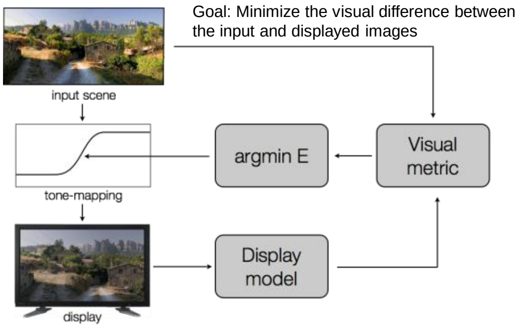

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

Colour transfer in tone-mapping

- Many tone-mapping operators work on luminance
  - For speed
  - To avoid colour artefacts
- Colours must be transferred later form the original image
- Colour transfer in the linear RGB colour space:

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

Output color channel (red) — R_{out}; Saturation parameter — s; Resulting luminance — L_{out}

- The same formula applies to green (G) and blue (B) linear colour values.

Colour transfer: out-of-gamut problem

- Colours often fall outside the colour gamut when contrast is compressed


Original image


Contrast reduced (s=1)

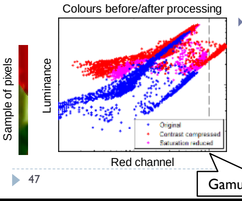

Gamut boundary

- Reduction in saturation is needed to bring the colors into gamut

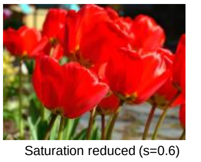
Saturation reduced (s=0.6)

Colour transfer: alternative method

- Colour transfer in linear RGB will alter resulting luminance
- Colours can be also transferred and saturation adjusted using CIE u'v' chromatic coordinates

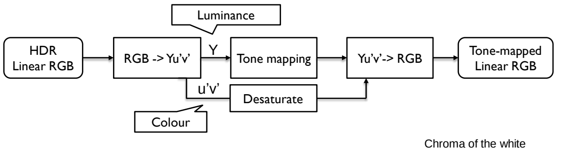

- To correct saturation:

$$u'_{out} = (u'_{in} - u'_w) \cdot s + u'_w$$
$$v'_{out} = (v'_{in} - v'_w) \cdot s + v'_w$$

Chroma of the white

$$u'_w = 0.1978$$
$$v'_w = 0.4683$$

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

Illumination & reflectance separation

Input

Illumination

Reflectance

$$Y = I \times R$$

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

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

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

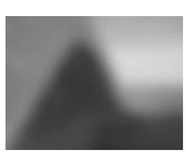

- Blurs sharp boundaries
- Causes halos

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

Tone mapping result

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

[Durand & Dorsey, SIGGRAPH 2002]

Weighted-least-squares (WLS) filter

- Stronger smoothing and still distinct edges

- Can produce stronger effects with fewer artifacts
- See image processing lecture

Tone mapping result

[Farbman et al., SIGGRAPH 2008]

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

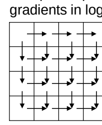

2nd step: set to 0 gradients less than the threshold

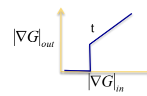

3rd step: reconstruct an image from the vector field

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

For example by solving the Poisson equation

Retinex examples

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

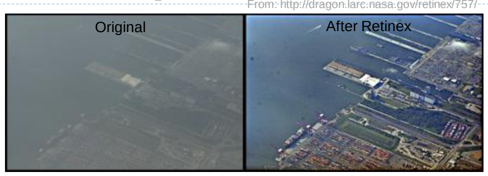

From:http://www.ipol.im/pub/algo/lmps_retinex_poisson_equation/#ref_1

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

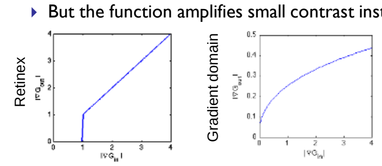

- Contrast compression achieved by global contrast reduction
  - Enhance reflectance, then compress everything

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

Glare

Glare Illusion

Photography   Painting

Computer Graphics
HDR rendering in games

Scattering of the light in the eye

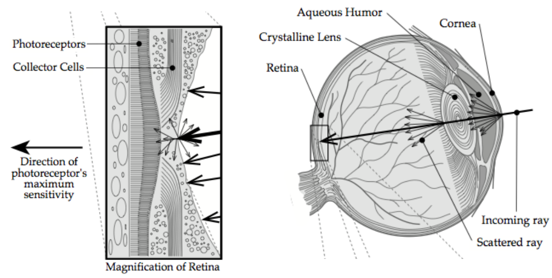

From: Sekuler, R., and Blake, R. Perception, second ed. McGraw- Hill, New York, 1990

Ciliary corona and lenticular halo

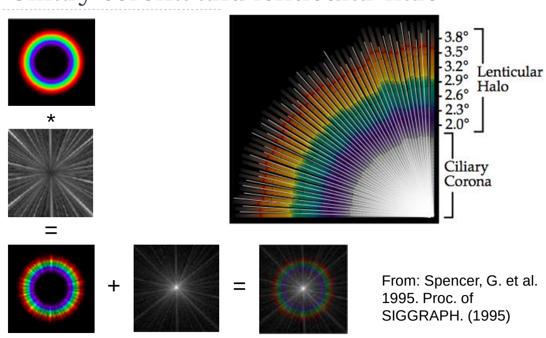

From: Spencer, G. et al. 1995. Proc. of SIGGRAPH. (1995)

Examples of simulated glare

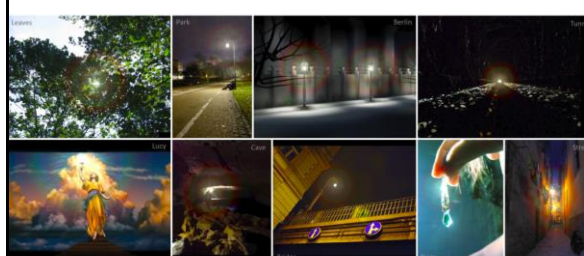

[From Ritschel et al, Eurographics 2009]

Temporal model of glare (low level)

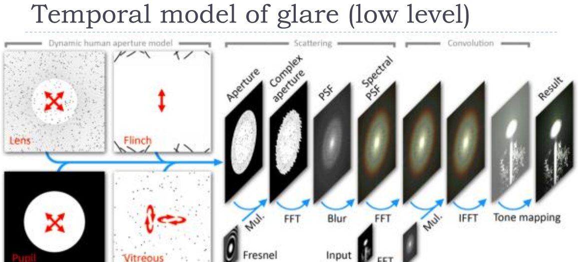

- The model assumes that glare is mostly caused by diffraction and scattering
- Can simulate temporal effects

[From Ritschel et al, Eurographics 2009]

Temporal glare

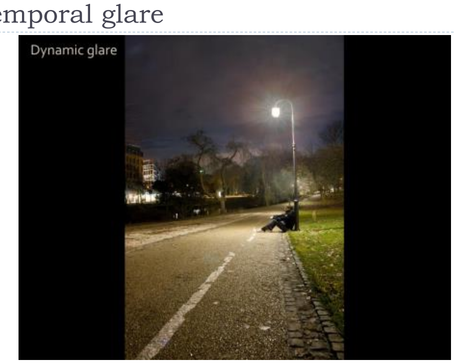

Point Spread Function of the eye

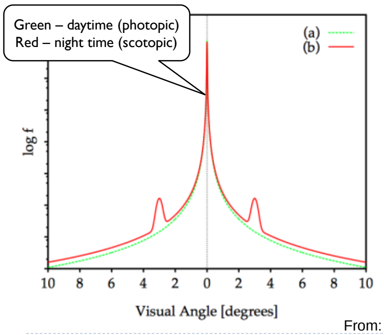

- What portion of the light is scattered towards a certain visual angle
- To simulate:
  - construct a digital filter
  - convolve the image with that filter

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

The problem of double processing

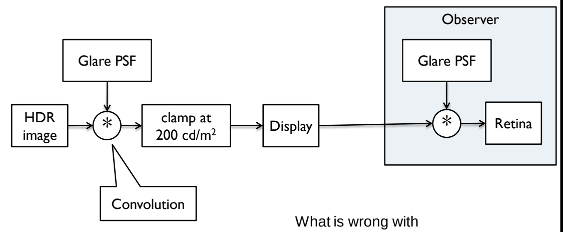

What is wrong with simulating glare this way?

The problem of double processing

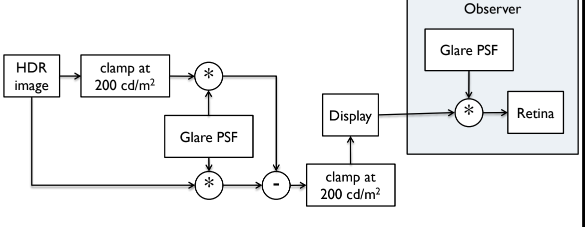

How does the diagram above avoid the problem of double processing?
Write down the operations as equations.
Can the processing be simplified?

PSF vs. OTF (Optical Transfer Function)

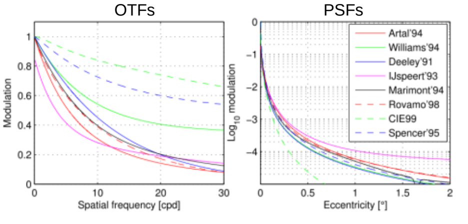

- An OTF is the Fourier transform of a PSF
- Convolution with larger kernels is faster in the Fourier domain

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

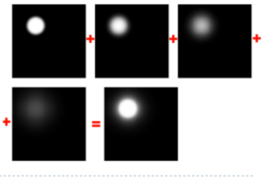

Does the exact shape of the PSF matter?

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

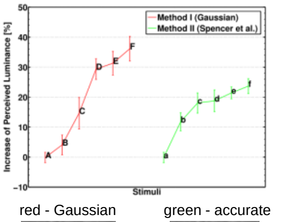

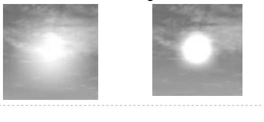

red - Gaussian　　green - accurate

[Yoshida et al., APGV 2008]

HDR rendering – motion blur

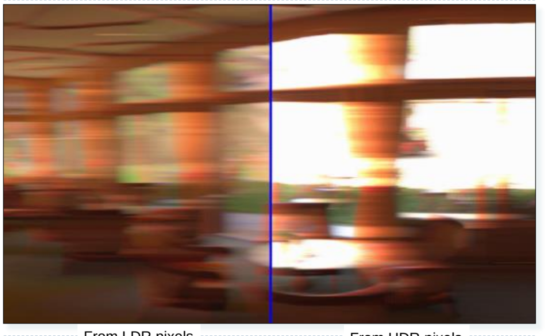

From LDR pixels — From HDR pixels

Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Color transfer
- Base-detail separation
- Glare
- Simulation of night vision

What changes at low illumination?

0.1 -> 100 cd/m²

- Global contrast
  - Relative brightness
- Local contrast
  - Visibility of small details
- Color
  - Purkinje shift
  - Saturation

Brightness reduction – tone-curve

- Perceptualy-based night-vision tone-curve
  - [Wanat et al. 2014]
  - Requires rather complex optimization

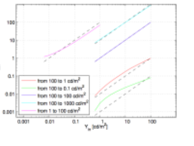

- Empirical approach (not perceptual)

Reduce brightness — Reduce contrast

$$y_{out} = b \cdot y_{in}^{\gamma} + f$$

Add „fog"

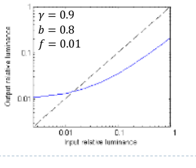

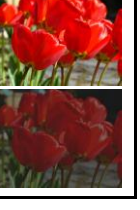

Local contrast

- Gabor patch
  - basic contrast stimulus
  - the shape matches the response pattern of the receptive fields on the retina

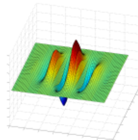

Contrast

$$G = l_{max} - l_{mean}$$

Max log luminance

Mean log luminance

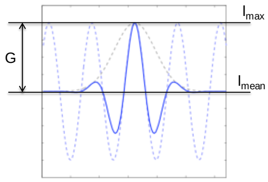

Supra-threshold contrast matching

- Kulikowski's model of matching contrast [Kulikowski 1976]
  - Contrast is perceived the same at different luminance levels when the physical contrast reduced by the corresponding detection threshold is equal at those luminance levels

Contrast at luminance A

Detection threshold at luminance A

$$G - G_T = \tilde{G} - \tilde{G}_T$$

Contrast at luminance B

Detection threshold at luminance B

The detection thresholds can be predicted by the contrast sensitivity function

Supra-threshold contrast matching

- The lines connect contrast of the same perceived magnitude

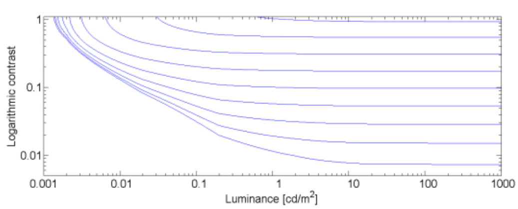

Local contrast processing

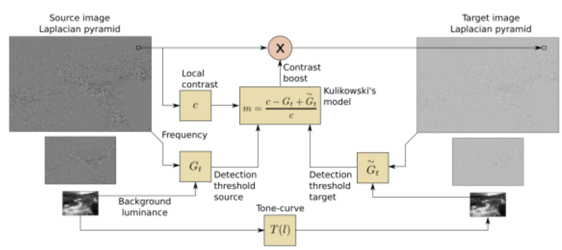

Example processing

Simulation of night vision

Target

Source

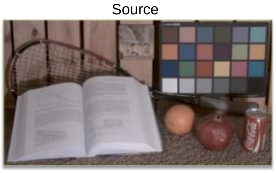

Rod contribution to colour vision

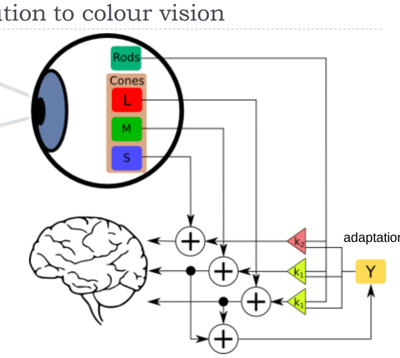

Rods and cones share the same pathway. Rods contribute to all cone responses.

[Cao et al. 2008]

Purkinje shift (effect)

- A shift in spectral sensitivity associated with the transition of cone to rod vision
  - Blue appears brighter and red appears darker in twilight
  - And the reverse is observed in daylight
- The shift to bluish hues is sometimes attributed to the Purkinje effect
  - In practice the blue-shift is very subtle
  - Much more pronouced in movies

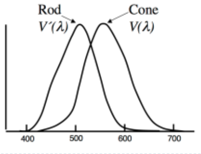

Loss of colour saturation with luminance

- Cones become less sensitive at low light
- Colours become less saturated
- Empirical formula [Wanat 2014]

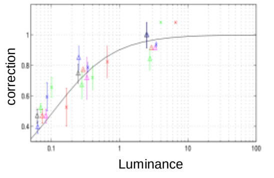

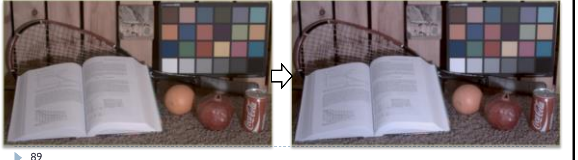

Age-adaptive night vision

Video 4

Rivoli
Simulation of age-adaptive night vision

References

- Comprehensive book on HDR Imaging
  - E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski, High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting, 2nd editio. Morgan Kaufmann, 2010.
- Overview of HDR imaging & tone-mapping
  - http://www.cl.cam.ac.uk/~rkm38/hdri_book.html
- Review of recent video tone-mapping
  - A comparative review of tone-mapping algorithms for high dynamic range video
    Gabriel Eilertsen, Rafal K. Mantiuk, Jonas Unger, Eurographics State-of-The-Art Report 2017.
- Selected papers on tone-mapping:
  - G.W. Larson, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Trans. Vis. Comput. Graph.*, vol. 3, no. 4, pp. 291–306, 1997.
  - R. Wanat and R. K. Mantiuk, "Simulating and compensating changes in appearance between day and night vision," *ACM Trans. Graph. (Proc. SIGGRAPH)*, vol. 33, no. 4, p. 147, 2014.
  - 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
  - ...