

Tone Reproduction

Tone corrections

HDR and other applications of convolutions

Dynamic range of Eye vs Dynamic Range of Monitor


- Brightness (informally Luminosity) how much energy (in watts)
- This is different than colors
- Crout approximation R+G+B

The World is a High Dynamic Range (HDR)

1:1 (ratio of intensities in brightest point vs. darkest point)



1:1,500



1:25,000



1:400,000



1:2,000,000,000




Most monitors: 1:255
Problem: Show images that has a very high dynamic range on our monitors

With HDR + Tone Mapping




Examples


Sun overexposed
Foreground too dark



Inside is too dark
Outside is too bright



Somehow needs to map images with HDR into monitors (LDR)



First Attempt: Linear transformation:


Need to map something to something - but does it look anywhere realistic?
Could easily change to color of the pixel - need to be careful!

1. Multiple problems:
2. Format of files of images with HDR
3. Changing intensities of images without changing colors: RGB vs HSV
4. Correct each pixel individually: (local operations:) α, β, γ corrections (gain, bias)
5. Big topic: Global operations: Convolution, high-pass low-pass filters

Briefly: How to store HDR images

- Options: PPM could use 2 bytes 2^{16} bytes for R,G,B for each pixel
- Other options: Use a long float (8 bytes) for each value of R,G,B -overkill
- Common formats - to be discussed later

Radiance RGBE Format (.hdr)



Ward, Greg. "Real Pixels," in Graphics Gems IV, edited by James Arvo, Academic Press, 1994

How not to change the color:

- We will modify the intensity $I = (R + G + B)/3$.
- For each pixel, we will map its original intensity I to a target intensity I' (devil in details)
- . Then will apply: $NewRGB = OldRGB \cdot \left(\frac{I'}{I}\right)$
- The perception of the color should stay the same.
- Terms that we will use interchangeably (though they are not identical: Intensity I, lightness L, Brightness B.

First Example: Linear Rescaling

- Brightness (informally Luminosity) how much energy (in watts)
- Rescaling** is a point processing technique that alters the **contrast** and/or **brightness** of an image.
- In photography, **exposure** is a measure of how much light is projected onto the imaging sensor.
 - Overexposure**: more light than what the sensor can measure.
 - Underexposure**: sensor is unable to detect the light.
- Images which are underexposed or overexposed can frequently be improved by brightening or darkening them.
- The contrast of an image can be altered to bring out the internal structure of the image.

Rescaling Math

- Given a sample C_{in} of the source image, rescaling computes the output sample, C_{out} , using the scaling function

$$C_{out} = \alpha C_{in} + \beta$$

- α is a real-valued scaling factor known as **gain**
- β is a real-valued scaling factor known as **bias**

Why Use Both α , β ?

- The relative change in contrast can be simplified as

$$\begin{aligned} \frac{\Delta S'}{\Delta S} &= \frac{|(\alpha S_1 + \beta) - (\alpha S_2 + \beta)|}{|S_1 - S_2|} \\ &= \frac{|\alpha| \cdot |S_1 - S_2|}{|S_1 - S_2|} \\ &= |\alpha|. \end{aligned}$$

- Thus, gain (α) controls the change in contrast.
- Whereas bias (β) does not affect the contrast
- Bias, however, controls the final **brightness** of the rescaled image. Negative bias darkens and positive bias brightens the image

Why Use Both α , β ?

- Consider two rescaled source samples of S rescaled to S'.
- Calculate the **contrast** (the absolute difference) between the source and destination, called ΔS and $\Delta S'$.
- Now consider the relative change in contrast between the source and destination.

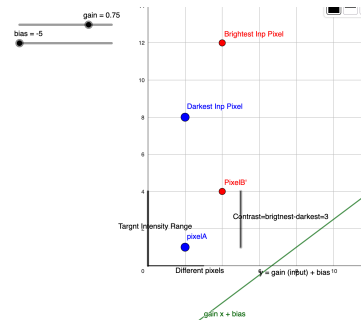
$$\begin{aligned} S'_1 &= \alpha S_1 + \beta, \\ S'_2 &= \alpha S_2 + \beta. \end{aligned}$$

$$\begin{aligned} \Delta S' &= |S'_1 - S'_2|, \\ \Delta S &= |S_1 - S_2|. \end{aligned}$$

$$\frac{\Delta S'}{\Delta S} = \frac{|S'_1 - S'_2|}{|S_1 - S_2|}$$

Why use gain and bias

<https://www.geogebra.org/m/fskd5rpb>



Want to maximize contrast =
brightness point -darkest non-zero one

Another idea: γ -correction

- Instead of
 - $New_Intensity(p) = \alpha \cdot Old_Intensity + \beta$
 - Do
 - $New_Intensity(p) = (\alpha \cdot Old_Intensity + \beta)^\gamma$


Here $0 \leq \gamma \leq 1$

This is equivalent to
Calculate $\log_2(Intensity)$
Multiply by γ
Calculate $2^{this\ value}$


Sidebar: Relating Contrast Sensivities to Signal Processing

Spatial Frequency - how many pixels do we need to move (geographically) until the intensity changes significantly.

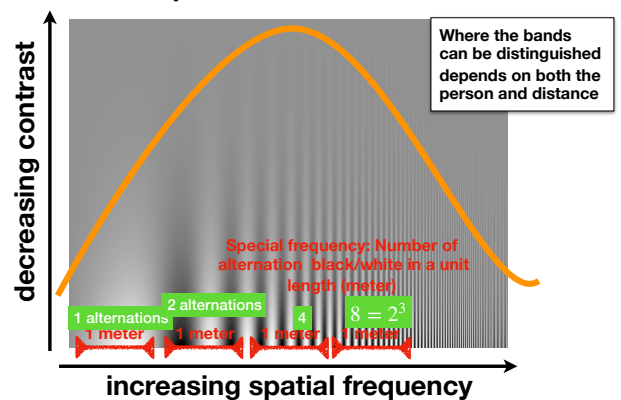
If the spatial frequency is high, but the display cannot reflect it, the image looks blurry.



Contrast Sensitivity Function Campbell-Robson Chart



Contrast Sensitivity Function Campbell-Robson Chart



Contrast Sensitivities Vary by Channel

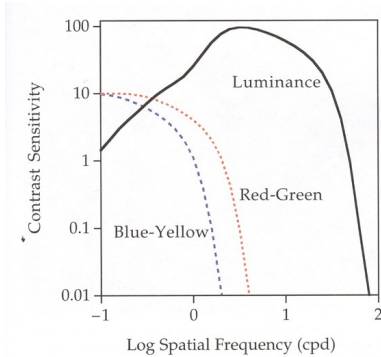



Figure 1-18. Spatial contrast sensitivity functions for luminance and chromatic contrast.

Rescaling Examples



gain = 1, bias = 55



gain = 1, bias = -55



gain = 2, bias = 0



gain = .5, bias = 0

Human Perception


- Eye distinguishes color intensities as a function of the ratio between intensities.
- Consider $I_1 < I_2 < I_3$, for the step between I_1 and I_2 to look like the step from I_2 to I_3 , it must be that:

$$I_2 / I_1 = I_3 / I_2$$

- As opposed to the differences in contrast! $I_2 - I_1 \neq I_3 - I_2$

Perceived (I_p) vs. Actual (I_a) Intensity

- Perceived light actually behaves like $I_p = (I_a)^\gamma$



Linear

Gamma2.2


<http://www.anywhere.com/gward/hdrenc/>

Example: Gamma Correction

FIGURE 3.9
(a) Aerial image.
(b)-(d) Results of applying the transformation in Eq. (3.2-3) with $c = 1$ and $\gamma = 3.0, 4.0,$ and 5.0 , respectively. (Original image for this example courtesy of NASA.)



$$s = r^\gamma$$



Putting it all together: Gain, Bias, and Gamma

Not these operations still work on each pixel individually

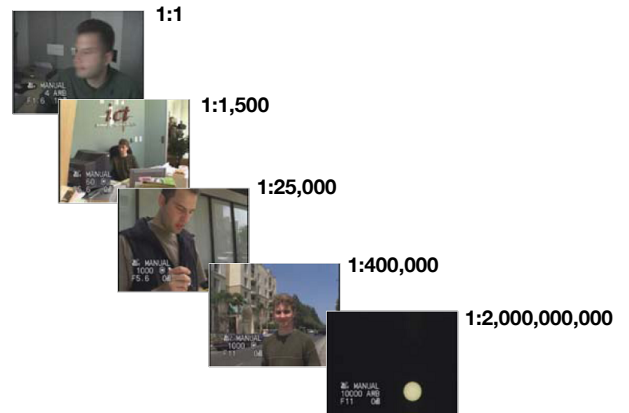
- $C_{out} = (\alpha C_{in} + \beta)^\gamma$
- α is known as **gain** (exposure)
- β is known as **bias** (offset)
- γ maps to a non-linear curve (**gamma** correction)



Image of Photoshop from
Christian Bloch - The HDR Handbook 2.0

Dynamic Range

The World is a High Dynamic Range (HDR)



Eyes and Dynamic Range


- We're sensitive to change (multiplicative)
 - A ratio of 1:2 is perceived as the same contrast as a ratio of 100 to 200
 - Use the log domain as much as possible
- But, eyes are **not** photometers
 - Dynamic adaptation (very local in retina)
 - Different sensitivity to spatial frequencies



Headlights are ON in both photos →



Approach: Visual Matching





- We do not need to reproduce the true radiance as long as it gives us a visual match.

What if rescaling and gamma-correction are not sufficient

- A useful collection of algorithms use the idea that in multiple scenarios, the highlights-illuminated regions
- Applicable for regions where changes from high brightness to low brightness occurs is slow: Big regions are "high", big regions are "low" (low spatial frequency)
- Note that this does not mean the "high" or "low" regions are short in details.
- Examples: Sky, windows during day light
- Or the way chatGPT rephrased it:
 - A collection of algorithms that is useful in many scenarios uses the idea that in regions where there is a slow change from bright to dark, large regions can be identified as either "high" or "low" brightness.
 - This is because these regions have low spatial frequency. However, this does not mean that the "high" or "low" regions are lacking in detail. Examples of this include the sky and windows during daylight hours.

Without HDR + Tone Mapping






With HDR + Tone Mapping



Approach: Visual Matching

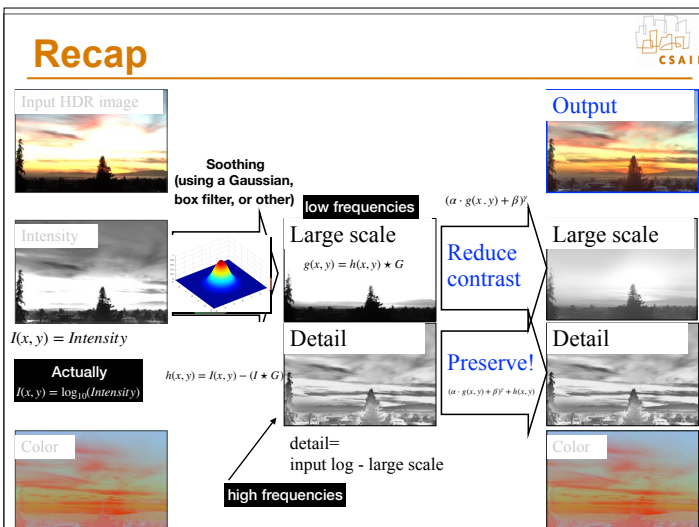


- We do not need to reproduce the true radiance as long as it gives us a visual match.
- Much more important to keep contrast between each pixel and its neighborhood pixels
- Means: Some neighborhoods must be treated different than others

Consider only the radiance map (the intensity at each pixel)

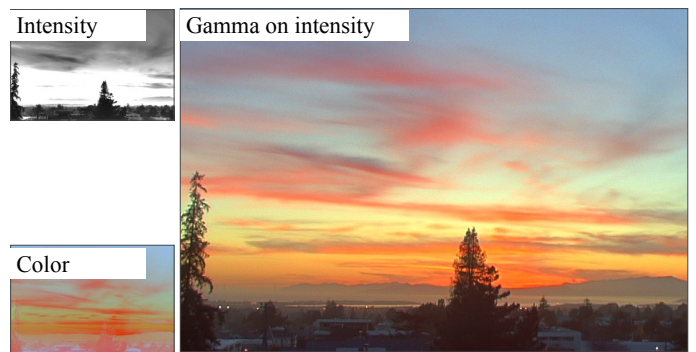
- Will produce two maps of intensities:
 - $I_{slow}(x, y)$ - changes slow, but large amplitude and large contrast
 - $I_{fast}(x, y)$ - changes rapidly, but small amplitude.
- (α, β, γ) -correction will be applied only to I_{slow}
- Put them back together

<https://www.geogebra.org/m/w6qnw9g>



Gamma compression on Intensity

- Colors ok, but details in intensity are blurry



Given - an image with a large dynamic range (e.g. from image of outdoor scene)
Need to compress it to much smaller dynamic range (monitor)
but avoid decreasing local contrast

In the GG demo, the input image represents the input image
Option 1 - compress $C_{out}(p_i) = \alpha C_{in}(p_i) + \beta$

<https://www.geogebra.org/m/w6qnvw9g>

g - (gain) = 0.34
β (bias) = 2.61
local contrast = 1.2

The discretized image (only 6 levels)

Original Image scaled
Local contrast $\# = 1.034$
highest intensity the monitor could display. Contrast = 9-4=5
Monitor's Dynamic Range
Lowest intensity the monitor could display

Given - an image with a large dynamic range (e.g. from image of outdoor scene)
Need to compress it to much smaller dynamic range (monitor)
but avoid decreasing local contrast

In the GG demo, the input image represents the input image
Option 1 - compress $C_{out}(p_i) = \alpha C_{in}(p_i) + \beta$

g - (gain) = 0.34
β (bias) = 2.61
local contrast = 1.2

The discretized image (only 6 levels)

Original Image scaled
Local contrast $\# = 1.034$
highest intensity the monitor could display. Contrast = 9-4=5
Monitor's Dynamic Range
Lowest intensity the monitor could display

Recap

Input HDR image

Output

Soothing (using a Gaussian, box filter, or other)

low frequencies

Large scale
 $g(x, y) = h(x, y) * G$

Detail
 $h(x, y) = I(x, y) - (I * G)$

Reduce contrast
 $(\alpha * g(x, y) + \beta)$

Preserve!
 $(\alpha * g(x, y) + \beta) * h(x, y)$

Large scale

Detail

Color

detail = input log - large scale

high frequencies

Actually $I(x, y) = \log_2(Intensity)$

Oppenheim 1968, Chiu et al. 1993

- Reduce contrast of low-frequencies
- Keep mid and high frequencies

Low-freq. Reduce low frequency

High-freq.

Color

A better idea

- $f(x)$ = Brightness level at pixel x
- Compute $s(x)$ - a low pass filter:
-For example, compute $s(x) = (f(x-1) + f(x) + f(x+1)) / 3$ (mean filter)
better idea: Use convolution with a tent or a gaussian
- in $s(x)$, details that change frequently (high special frequencies) are averaged.
- compute $h(x) = f(x) - s(x)$. Here only the details that have high special frequencies appear.
- Hopefully, the contest (max intensity - min intensity) are no nearly as large as the contrast in $s(x)$
- Rescale $s(x)$ by using the the gain and bias (α, β) such that it fits the monitor dynamic range.
 $s(x) \rightarrow \alpha \cdot s(x) + \beta$. Usually $\alpha < 1$.
- Place back the missing details. The output is $f_{out}(x) = \alpha s(x) + \beta + h(x)$

The halo nightmare

- For strong edges
- Because they contain high frequency

Low-freq.

High-freq.

Color

Reduce low frequency

Start with Gaussian filtering

- Here, input is a step function + noise

$$J = f \otimes I$$

Weight (eg $\frac{1}{\text{distance}(x, \xi)}$)

Intensity at ξ

output ← input

Gaussian filter as weighted average

- Weight of ξ depends on distance to x

$$J(x) = \sum_{\xi} \text{Weight} \left(\frac{1}{\text{distance}(x, \xi)} \right) f(x, \xi) I(\xi)$$

output ← input

Gaussian filter as weighted average

- Weight of ξ depends on distance to x

Non Bilateral

$$J(x) = \sum_{\xi} \text{Weight} \left(\frac{1}{\text{distance}(x, \xi)} \right) f(x, \xi) I(\xi)$$

output ← input

The problem of edges

- Here, $I(\xi)$ "pollutes" our estimate $J(x)$
- It is too different

$$J(x) = \sum_{\xi} f(x, \xi) I(\xi)$$

output ← input

Principle of Bilateral filtering

[Tomasi and Manduchi 1998]

If $I(\xi) - I(x)$ is large, then $g(I(\xi) - I(x)) \approx 0$

- Penalty g on the intensity difference

Remember that the sum of weights $\sum f(x, \xi)$ must be 1.
 What to do if we skip some terms? (that is We will divide the total sum by $k(x)$ - see next slide)

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x, \xi) g(I(\xi) - I(x)) I(\xi)$$

$g(|a - b|) = 1$ if a is close to b , and zero otherwise


output ← input

Bilateral filtering

[Tomasi and Manduchi 1998]

- **Spatial Gaussian f** Remember that the sum of weights must be 1.
What to do if we skip some terms?
We will divide the total sum by $k(x)$

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x, \xi) g(I(\xi) - I(x)) I(\xi)$$



output


input

Bilateral filtering

[Tomasi and Manduchi 1998]

- **Spatial Gaussian f**
- **Gaussian g on the intensity difference**

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x, \xi) g(I(\xi) - I(x)) I(\xi)$$



output


input

Normalization factor

[Tomasi and Manduchi 1998]

- $k(x) = \sum_{\xi} f(x, \xi) g(I(\xi) - I(x))$

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x, \xi) g(I(\xi) - I(x)) I(\xi)$$



output


input

Bilateral filtering is non-linear

[Tomasi and Manduchi 1998]

- **The weights are different for each output pixel**

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x, \xi) g(I(\xi) - I(x)) I(\xi)$$



output

input