

Computational Photography

Prof. Feng Liu

Spring 2022

<http://www.cs.pdx.edu/~fliu/courses/cs510/>

04/14/2022

Last Time


- Color
- Color to Gray

Today

- Re-lighting
 - Tone mapping

Dynamic Range

- ❑ Real scenes contain both very bright and very dark regions
- ❑ **Dynamic range is the ratio of the brightest to darkest points in the scene**
- ❑ Standard measurement is candelas per m^2
- ❑ For example, in the interior of the church the dynamic range is 100,000 to 1



Tone Reproduction

- The human eye can *globally* adapt to about $10^{12}:1$
 - Adjusts for the average brightness we perceive in one scene
 - Global adaptation lets us see in very low light or very bright conditions
- The human eye can *locally* adapt to about 10,000:1
 - In a single scene (global adaptation level), we can perceive contrast across this range
- Most display devices have a very limited dynamic range
 - On the order of 100:1 for a very good monitor or film
- *Tone reproduction* is the problem of making the 10,000:1 scene look right on a 100:1 display device

Different Goals

- Option 1: Reveal as much detail as possible in the image given the limited dynamic range
 - For conveying information, but not necessarily realism
- Option 2: Reveal what would be available if the viewer would really be there
 - Perceptual limits place limits on what we can see under a given set of viewing conditions





Color and Luminance

- The CIE XYZ color space is intended to encode luminance in the Y channel
- To get from RGB to XYZ, apply the following linear transform:
 - Taken from Sillion and Puech, other sources differ

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.67 & 0.21 & 0.14 \\ 0.33 & 0.71 & 0.08 \\ 0.00 & 0.08 & 0.78 \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix} \quad \begin{bmatrix} r \\ g \\ b \end{bmatrix} = \begin{bmatrix} 1.730 & -0.482 & -0.261 \\ -0.814 & 1.652 & -0.023 \\ 0.083 & -0.169 & 1.284 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

L*a*b* Color Space

- L* : luminance
- a* and b* : color

Automatic Tone Reproduction

- There are three main classes of solutions:
 - **Global operators** find a mapping from image to display luminance that is the same for every pixel
 - **Local operators** change the mapping from one pixel to the next
 - **Perceptually guided operators** use elements of human perception to guide the tone reproduction process

Linear Mappings

- The simplest thing to do is to linearly map the highest intensity in the image to the highest display intensity, or the lowest to the lowest
 - This gives very bad results, shown for mapping the maps lowest to lowest




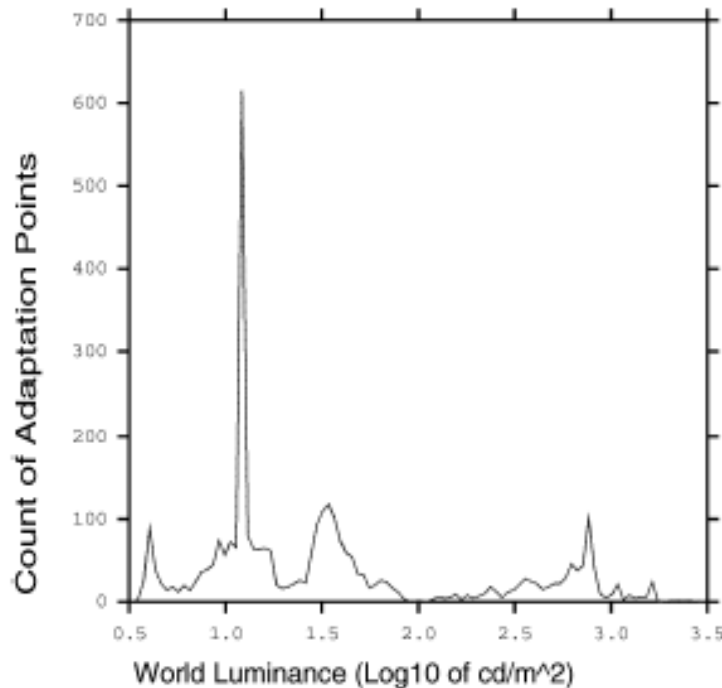
Fig. 1. A false color image showing the world luminance values for a window office in candelas per meter squared (cd/m^2 or Nits).

Non-Linear Mappings

- Instead of linear, define some other mapping:
$$L_d = M(L_w)$$
 - Display luminance is some function of the world luminance
- It is important to retain *relative* brightness, but not absolute brightness
 - If one point is brighter than another in the source, it should be brighter in the output
 - The mapping M should be strictly increasing

Histogram Methods

- In any one scene, the dynamic range is not filled uniformly
- The aim of histogram methods is to generate a mapping that “fills in the gaps” in the range



Ward, Rushmeier, and Piatko, "[A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes](#)," *IEEE Trans. on Visualization and Computer Graphics*, 1997

Slides Credit: S. Cheney

Building Histograms

- Work in brightness: $B = \log_{10}(L)$
 - Humans are more sensitive to “brightness,” but the formula is a hack
- The histogram is a count of how many pixels have each brightness
 - Choose a set of bins
 - Break the image into chunks that subtend about 1° of arc
 - Assumes you know the camera
 - Average brightness in each chunk
 - Count chunks that fall in each bin, $f(b_j)$
 - Result is graph on previous image

Cumulative Distribution

- We can consider the histogram counts as the probability of seeing a pixel in each range
- The cumulative distribution function is defined:

$$P(b) = \frac{\sum_{b_i < b} f(b_i)}{T} \quad T = \sum_{b_i} f(b_i)$$

Histogram Equalization

- The aim is an *output* histogram in which all the bins are roughly equally filled
- The naïve way to do this is to set:

$$B_{de} = \log(L_{dmin}) + [\log(L_{dmax}) - \log(L_{dmin})]P(B_w)$$

- Then, go through and convert brightness back into luminance for display

Naïve Histogram



Fig. 8. Rendering of a bathroom model mapped with a linear operator.



Fig. 9. Naive histogram equalization allows us to see the area around the light sources, but contrast is exaggerated in other areas, such as the shower tiles.

Linear left.
Histogram
right.
What went
wrong?

Avoiding Super-Sensitivity

- We should make sure that we do not increase contrast beyond the linear mapping contrast
- This imposes a constraint on the frequency in each bin:

$$\frac{f(b)}{\Delta b} \leq \frac{T}{\log(L_{dmax}) - \log(L_{dmin})}$$

- Reduce the count of bins that exceed the maximum

Better Histogram Adjustment



←
Old
New
→



Local Methods

- Histogram methods do poorly if there is too much dynamic range in the input
- Local methods exploit the local nature of contrast
 - You don't compare the brightness of two things across the room, only neighboring points

Basic Idea

- Filter to separate high and low frequencies
 - Filter to get a high-pass and low-pass version of the image
 - Recall, a smoothing filter keeps only low frequencies
- Compress the low frequencies
 - Reduce the contrast in the low pass portion
 - Using diffusion filters, typically
- Add back in the high frequencies
 - Add the reduced low-pass to the high-pass to get a new image
- *LCIS* is one algorithm

Tumblin, J., Turk, G., "Low Curvature Image Simplifiers (LCIS): A Boundary Hierarchy for Detail-Preserving Contrast Reduction." In *SIGGRAPH '99*

Interactive Local Adjustment of Tonal Values



Workflow

1. Load an image.
2. Indicate regions in the image that require adjusting.
3. Experiment with the available adjustment parameters until a satisfactory result is obtained in the desired regions.
4. Iterate steps 2 and 3 until satisfied with the entire resulting image.


An example

(a)

(b)

(c)

(d)



Constraint Propagation

□ Input

- L: input image
- g(x): user constraint

□ Output:

- f(x): constraint function over the whole image

□ Data Term

- Changes in constrained pixels should respect user specifications

$$\sum_x w(x)(f(x) - g(x))^2$$

□ Smoothness Term

- Changes between neighboring pixels should be similar

$$\sum_x h(\nabla f, \nabla L), \text{ where } h(\nabla f, \nabla L) = \frac{|f_x|^2}{|L_x|^\alpha + \varepsilon} + \frac{|f_y|^2}{|L_y|^\alpha + \varepsilon}$$

The Final System

- A quadratic minimization problem
 - Equivalent to a linear system

$$E(f) = \sum_x w(x) (f(x) - g(x))^2 + \sum_x h(\nabla f, \nabla L)$$

$$f = \arg \min_f E(f)$$

$$\frac{\partial E}{\partial f} = 0$$

⇒

$$Af = b$$

Resources Solving Large Linear System

- Python

- Numpy

- C/C++

- <http://www.cise.ufl.edu/research/sparse/CSparse/>

- Free

- Intel MKL

- Commercial tool

- One month trial available

- Matlab code:

```
A=rand(3,3);
```

```
b=rand(3,1);
```

```
x= A\b;
```

More Results



Figure 1: Three different interpretations generated from the same digital negative using our tool. Left: warm sky, high exposure in the foreground. Middle: cooler sky, medium exposure in the foreground. Right: an even cooler sky, very little exposure in the foreground leaving almost no detail but the silhouette. RAW image courtesy of Norman Koren, www.normankoren.com.

Student paper presentation

Image denoising: Can plain Neural Networks
compete with BM3D?

H. C. Burger, C. J. Schuler, and S. Harmeling
CVPR 2012

Presenter: Harish, Achyutha Nelamangala

Next Time

- Re-lighting

- HDR

- Student paper presentation

- 04/19: Chen, Raymond

- Robust color-to-gray via nonlinear global mapping
Yongjin Kim, Cheolhun Jang, Julien Demouth, and Seungyong Lee
SIGGRAPH ASIA 2009