Burst Photography

Computational Photography

CSE 291

Lecture 5

Announcements

- Assignment 2 is due Apr 15, 11:59 PM
- Assignment 3 will be released Apr 15
  - Due Apr 22, 11:59 PM

Burst photography

- Set of standard dynamic range (SDR) images to a high dynamic range (HDR) image

SDR images

HDR image
(mapped to 8 bits per sample)

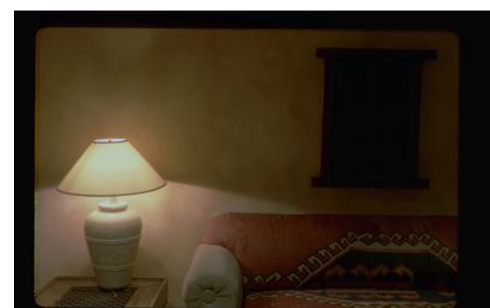

sRGB version

High dynamic range

- Definition: a dynamic range higher than what is considered standard dynamic range
- Used in several areas
  - Visual
  - Audio
    - Greater than signed 16 bits per sample
  - Radio
    - Mitigates interfering signals
  - Measurement
    - For example, high dynamic range accelerometers, which are used in seismology
  - Vision
    - Computer mediated reality
    - For example, high dynamic range welding helmet
      - Augments the image in dark areas and diminishes it in bright areas

Visual high dynamic range

- High dynamic range rendering
  - Rendering of computer graphics scenes using lighting calculations performed in high dynamic range
- High dynamic range imaging
  - Compositing and tone mapping of images to extend the dynamic range beyond the native capability of the capturing device
- High dynamic range video
  - Greater than standard dynamic range video (i.e., that found on CRT televisions produced in 1934)
- High dynamic range display
  - Ultra HD Alliance certification requirements
    - Must have either a peak brightness of over 1000 cd/m^2 and a black level less than 0.05 cd/m^2 (a contrast ratio of at least 20,000:1) **or** must have a peak brightness of over 540 cd/m^2 and a black level less than 0.0005 cd/m^2 (a contrast ratio of at least 1,080,000:1)
    - Allows for liquid crystal display (LCD) **or** organic light emitting diode (OLED) display

High dynamic range images

- Standard dynamic range images (and displays)
  - 8 bits per pixel per channel (i.e., 8 bits per sample)
- High dynamic range images
  - Greater than 8 bits per sample, usually between 10 and 16 bits per sample
    - Note: the typical human cannot differentiate beyond 13.87 bits per sample
  - “Bright things can be really bright, dark things can be really dark, and details can be seen in both”

High dynamic range images

- Computational photography
  - Primary technique: acquiring high dynamic range images
  - Secondary technique: displaying a high dynamic range image on a standard dynamic range display

Sources of high dynamic range images

- Synthetic images
  - High dynamic range rendering
- Real images
  - Imaging sensor systems producing greater than 8 bits per sample images
    - Raw image files from digital cameras, film scanners, or other image scanners
      - Note: raw image pixel values are linear, not nonlinear (e.g., gamma encoded)

Linear

Nonlinear

    - **Warning**: adding bits to the analog to digital converter does not improve the sensitivity, noise, etc. characteristics of the sensor
    - Alternatively, increase signal to noise ratio by averaging multiple images, where camera and scene are unchanged
  - High dynamic range imaging

High dynamic range imaging

- Approach
  - Acquire a set of *K* “low” dynamic range images with associated exposures (e.g., exposure times or shutter speed) *e*
  - If images are nonlinear, then
    - Estimate inverse camera response function $f(x)$
    - Nonlinear pixel values $y = f(x)$
    - Calculate high dynamic range pixel values

Estimating the inverse camera response function is more difficult than fusing the images

$$z = \frac{\sum_{k=1}^{K} \frac{f(x)w(x)}{e_k}}{\sum_{k=1}^{K} w(x)}, \text{ where weight } w(x) = \frac{f(x)}{f'(x)}$$

Camera manufacturers consider their camera response functions proprietary information (e.g., “that image looks like it was taken with a Canon, Sony, etc. camera”). There is a database online containing estimated camera response functions for hundreds of cameras.

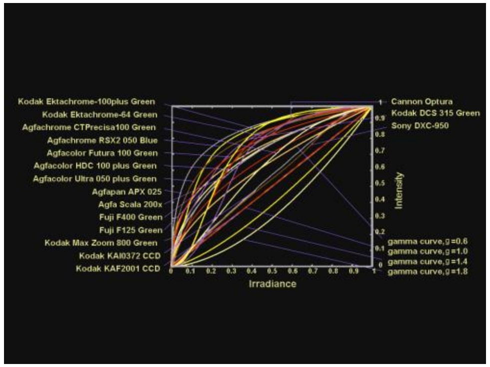

Estimation of inverse camera response function

- Model as (low degree) polynomial with *N* coefficients

$$y = f(x)$$

$$y = c_N x^N + c_{N-1} x^{N-1} + \cdots + c_1 x + c_0 = \sum_{n=0}^{N} c_n x^n$$

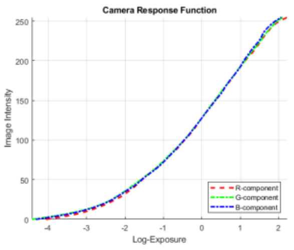

- Include the following two constraints, such that [0, 1] maps to [0, 1]

$$0 = f(0)$$
$$0 = c_0 x^0$$
$$c_0 = 0$$

$$1 = f(1)$$

$$1 = c_N + c_{N-1} + \cdots + c_1 + c_0 = \sum_{n=0}^{N} c_n = c_N + \sum_{n=0}^{N-1} c_n$$

$$c_N = 1 - \sum_{n=0}^{N-1} c_n$$

Estimation of inverse camera response function

- Given two exposures *e* and *e'*, and associated camera response values *x* and *x'*

$$e = f(x) = \sum_{n=0}^{N} c_n x^n \text{ and } e' = f(x') = \sum_{n=0}^{N} c_n x'^n$$ Solve for $c_{N-1}, \ldots, c_1$

$$1 = 1/e \sum_{n=0}^{N} c_n x^n \text{ and } 1 = 1/e' \sum_{n=0}^{N} c_n x'^n$$

$$0 = 1/e \sum_{n=0}^{N} c_n x^n - 1/e' \sum_{n=0}^{N} c_n x'^n$$ (square of this is error; used later)

$$0 = 1/e(c_N x^N + \sum_{n=0}^{N-1} c_n x^n) - 1/e'(c_N x'^N + \sum_{n=0}^{N-1} c_n x'^n)$$

$$0 = 1/e((1 - \sum_{n=0}^{N-1} c_n)x^N + \sum_{n=0}^{N-1} c_n x^n) - 1/e'((1 - \sum_{n=0}^{N-1} c_n)x'^N + \sum_{n=0}^{N-1} c_n x'^n)$$

$$\vdots$$

$$\sum_{n=0}^{N-1} c_n(1/e(x^n - x^N) - 1/e'(x'^n - x'^N)) = x'^N/e' - x^N/e$$

Estimation of inverse camera response function

- In vector form

$$(1/e(x^{N-1}-x^{N})-1/e'(x'^{N-1}-x'^{N}),\ldots,1/e(x^{1}-x^{N})-1/e'(x'^{1}-x'^{N}))\begin{pmatrix}c_{N-1}\\ \vdots \\ c_1\end{pmatrix}=x'^{N}/e'-x^{N}/e$$

- Include all combinations of pairs of exposures over all pixels in both images
  - Alternatively, samples regions within the image
- Try up to a maximum number of polynomial coefficients and use one that results is the minimum error
  - Minimum number of polynomial coefficients without overfitting to the data

Estimation of inverse camera response function

Another set of linear equations

- If color, then
  - Estimate inverse camera response function for each color channel
  - Solve for scales (one for each channel) that preserves chromaticity of nonlinear pixel values

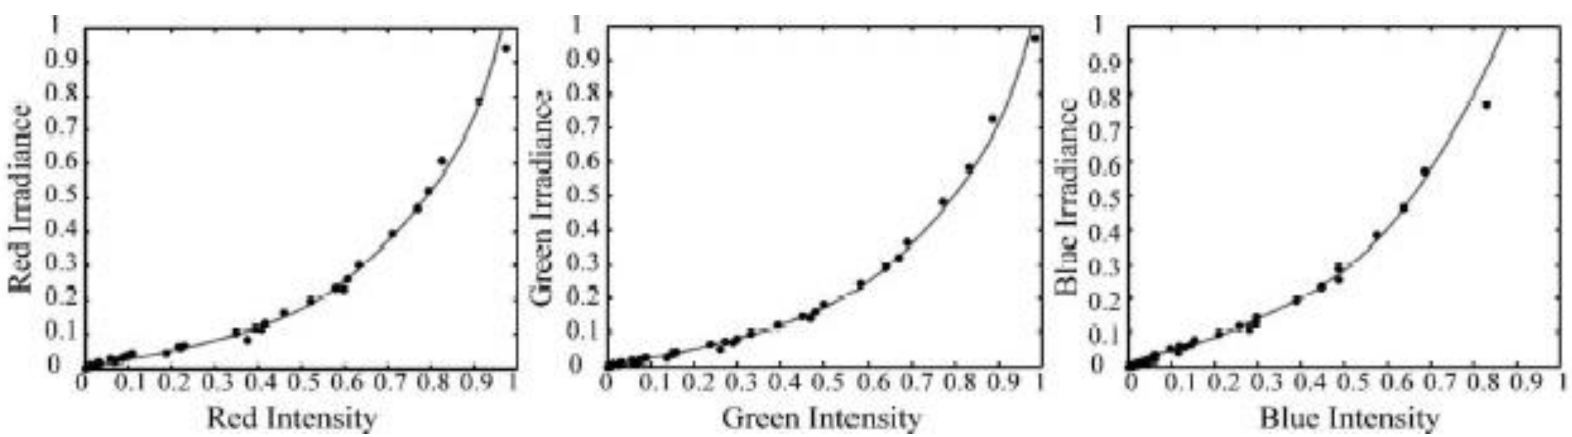

High dynamic range imaging example

- Data courtesy of Columbia University (https://www.cs.columbia.edu/CAVE/software/rascal/rrslrr.php)
- Exposures acquired using a film scanner
- All exposures are (nearly) aligned

     

High dynamic range imaging example

- Exposure 1

High dynamic range imaging example

- Exposure 2

High dynamic range imaging example

- Exposure 3

High dynamic range imaging example

- Exposure 4

High dynamic range imaging example

- Exposure 5

High dynamic range imaging example

- HDR image mapped to 8 bits per sample
  - Camera response

High dynamic range imaging example

Exposure 3

HDR image mapped to 8 bits per sample

High dynamic range imaging example

- HDR image mapped to 8 bits per sample

sRGB

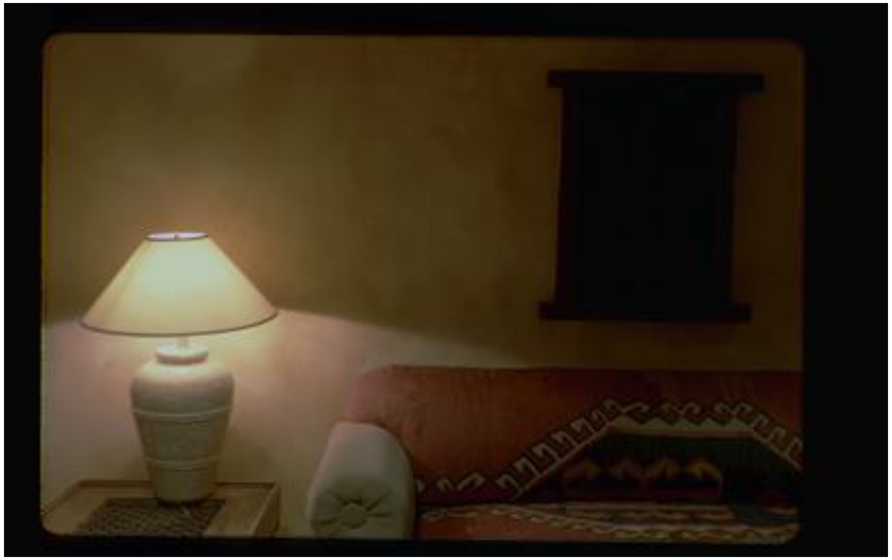

Camera response function

Representing high dynamic range images

- In the entertainment industry (target is human eye; maximum 13.87 bits per sample), high dynamic range images are processed using 32-bit (single precision) floating point and stored using 16-bit (half precision) floating point
  - Unlike integer formats, floating point has a dynamic range
    - Relatively high precision for values near zero
    - Relative precision decreasing as the value moves away from zero
- Conveniently, floating point pixel data is normalized such that it is in the range [0, 1]
- The IEEE Standard for Floating-Point arithmetic (IEEE 754) added the 16-bit binary format (commonly referred to as half) in the 2008 version of the standard

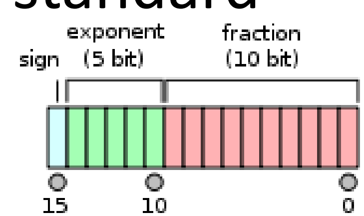

Representing high dynamic range images

- Standard file formats that support half precision floating point are JPEG XR and OpenEXR, which was developed by Industrial Light & Magic and used extensively in the visual effects industry
- In modern graphics cards, half precision is usually used, as it requires half the memory bandwidth (getting data to/from GPU RAM in half the time)
- A C/C++ data type representing IEEE 754 half precision floating-point can be found in OpenEXR, OpenGL, OpenCL, and DirectX

Displaying high dynamic range images

- Ideally, simply use a high dynamic range display device
  - Ultra HD Alliance certification requirements
    - LCD: must have either a peak brightness of over 1000 cd/m2 and a black level less than 0.05 cd/m2 (14.29 bits per sample)
    - OLED: must have a peak brightness of over 540 cd/m2 and a black level less than 0.0005 cd/m2 (20.04 bits per sample)
  - High dynamic range display devices are rarely available

Displaying high dynamic range images

- Display a high dynamic range image on a standard dynamic range display
  - Convert the high dynamic range image to a standard dynamic range image (then display the resulting image)
    - If raw image files, then use the 8 most significant bits (i.e., discard the least significant bits)
    - Map high dynamic range image to 8 bits per sample (as done in the examples shown earlier)
    - Vary the exposure manually using a high dynamic range image viewer (e.g., OpenHDR Viewer https://viewer.openhdr.org/)
    - Tone mapping

Tone mapping

- Map a high dynamic range image to a standard dynamic range image that approximates the appearance of the high dynamic range one
  - Reduce the dynamic range of an entire image while retaining localized contrast, preserving the image details and color appearance important to appreciate the original scene content
- Various tone mapping operators have been developed by the image processing and computer graphics communities

Summary

- High dynamic range image acquisition is now common
  - Higher sensitivity, lower noise with greater than 8 bits per sample imaging sensor systems
  - High dynamic range imaging from multiple exposures
    - We assumed the camera and scene are unchanged between exposures, but there are techniques for image alignment, and ghost removal (caused by moving objects in the scene), and lens flare removal
- High dynamic range images are typically stored and transmitted/received using 16-bit (half precision) floating point
  - Half the storage and bandwidth requirements compared to 32-bit (single precision) floating point

Summary

- High dynamic range image and video formats are mature, having been use for nearly 20 years
- Currently, high dynamic range displays are rare but are expected to be common in the near future
  - Until then, map high dynamic range images to standard dynamic range (e.g, using tone mapping)
- In the future, the entire image and video pipeline (acquisition → storage → display) may be high dynamic range