


Exposure and high dynamic range imaging



15-463, 15-663, 15-862
Computational Photography
Fall 2026, Lecture 5

Course announcements

- Homework assignment 2 is out.
 - Due September 25th.
 - Bonus component worth an extra half homework assignment.

Overview of today's lecture

- Exposure control.
- Light metering.
- Our devices do not match the world.
- High dynamic range imaging.
- Radiometric calibration.
- Other aspects of HDR imaging.
- Tonemapping.
- A few thoughts on HDR imaging and tonemapping.

Slide credits

Many of these slides were inspired or adapted from:

- James Hays (Georgia Tech).
- Fredo Durand (MIT).
- Gordon Wetzstein (Stanford).
- Marc Levoy (Stanford, Google).
- Sylvain Paris (Adobe).
- Sam Hasinoff (Google).

Exposure control


What is exposure?

Roughly speaking, the “brightness” of a captured image given a fixed scene.


$$\text{Exposure} = \text{Gain} \times \text{Flux} \times \text{Time}$$

- Flux is controlled by the aperture.
- Time is controlled by the shutter speed.
- Gain is controlled by the ISO.

Exposure controls brightness of image




Exposure controls brightness of image



Shutter speed


Controls the **length of time** that shutter remains open.



closed shutter

Shutter speed

Controls the **length of time** that shutter remains open.




open shutter



Shutter speed

Controls the **period of time** that shutter remains open.




open shutter

What happens to the image as we increase shutter speed?

Side-effects of shutter speed

Moving scene elements appear blurry.




How can we “simulate” decreasing the shutter speed?

Motion deblurring




Shah et al. High-quality Motion Deblurring from a Single Image, SIGGRAPH 2008

Exposure controls brightness of image




Aperture size

Controls area of lens that lets light pass through.



Aperture size

Controls area of lens that lets light pass through.



Aperture size

Most lenses have apertures of variable size.

- The size of the aperture is expressed as the “f-number”: The bigger this number, the smaller the aperture.



f / 1.4



f / 2.8



f / 4



f / 8



f / 16

You can see the aperture by removing the lens and looking inside it.

Side-effects of aperture size

Depth of field decreases as aperture size increases.

- Having a very sharp depth of field is known as “bokeh”.



How can we simulate bokeh?


How can we simulate bokeh?

Infer per-pixel depth, then blur with depth-dependent kernel.

- Example: Google camera “lens blur” feature




Exposure controls brightness of image




The (in-camera) image processing pipeline

The sequence of image processing operations applied by the camera's image signal processor (ISP) to convert a RAW image into a "conventional" image.



Analog front-end



analog amplifier (gain):

- gets voltage in range needed by A/D converter
- accommodates ISO settings.
- accounts for vignetting.

analog-to-digital converter (ADC):

- depending on sensor, output has 10-16 bits.
- most often (?) 12 bits.

look-up table (LUT):

- corrects non-linearities in sensor's response function (within proper exposure).
- corrects defective pixels.

Side-effects of increasing ISO

Image becomes very grainy because noise is amplified.



ISO 80



ISO 800



ISO 1600

Note about the name ISO

ISO is not an acronym.

- It refers to the International **O**rganization for **S**tandardization.
- ISO comes from the Greek word *ἴσος*, which means *equal*.
- It is pronounced (roughly) *eye-zo*, and should not be spelled out.

Camera modes

Aperture priority (“A”): you set aperture, camera sets everything else.

- Pros: Direct depth of field control.
- Cons: Can require impossible shutter speed (e.g. with f/1.4 for a bright scene).

Shutter speed priority (“S”): you set shutter speed, camera sets everything else.


- Pros: Direct motion blur control.
- Cons: Can require impossible aperture (e.g. when requesting a 1/1000 speed for a dark scene)

Automatic (“AUTO”): camera sets everything.

- Pros: Very fast, requires no experience.
- Cons: No control.

Manual (“M”): you set everything.

- Pros: Full control.
- Cons: Very slow, requires a lot of experience.



generic camera mode dial

Camera modes

Aperture priority (“A”): you set aperture, camera sets everything else.

- Pros: Direct depth of field control.
- Cons: Can require impossible shutter speed (e.g. with f/1.4 for a bright scene).

Shutter speed priority (“S”): you set shutter speed, camera sets everything else.


- Pros: Direct motion blur control.
- Cons: Can require impossible aperture (e.g. when requesting a 1/1000 speed for a dark scene)

Automatic (“AUTO”): camera sets everything.

- Pros: Very fast, requires no experience.
- Cons: No control.

Manual (“M”): you set everything.

- Pros: Full control.
- Cons: Very slow, requires a lot of experience.




generic camera mode dial


Light metering

Light metering in modern cameras

- SLR cameras use a separate low-resolution sensor that is placed at the focusing screen.





- Mirrorless cameras use measurements directly from the main sensor.



Light metering in modern cameras

- Measurements are averaged to produce a single intensity estimate, which is assumed to correspond to 18% reflectance (the “key”).
- Exposure is set so that this average is exposed at the middle of the sensor’s dynamic range.
- Averaging can be done in many ways:
 1. Center-weighted.
 2. Spot.
 3. Scene-specific preset (portrait, landscape, horizon).
 4. “Intelligently” using proprietary algorithm.



Our devices do not match the world

The world has a high dynamic range



1



1500



25,000




400,000




2,000,000,000


The world has a high dynamic range




(Digital) sensors also have a low dynamic range



(Digital) images have an even lower dynamic range




(Digital) images have an even lower dynamic range



(Digital) images have an even lower dynamic range

Any guesses about the dynamic range of a standard 0-255 image?




pure black

pure white

(Digital) images have an even lower dynamic range

Any guesses about the dynamic range of a standard 0-255 image?




pure black


pure white

about 50x
brighter

(Digital) images have an even lower dynamic range



(Digital) images have an even lower dynamic range



Our devices do not match the real world

- 10:1 photographic print (higher for glossy paper)
- 20:1 artist's paints
- 200:1 slide film
- 500:1 negative film
- 1000:1 LCD display
- 2000:1 digital SLR (at 12 bits)
- 100000:1 real world

Two challenges:

1. HDR imaging – which parts of the world do we measure in the 8-14 bits available to our sensor?
2. Tonemapping – which parts of the world do we show in the 4-10 bits available to our display?

Our devices do not match the real world

- 10:1 photographic print (higher for glossy paper)
- 20:1 artist's paints
- 200:1 slide film
- 500:1 negative film
- 1000:1 LCD display
- 2000:1 digital SLR (at 12 bits)
- 100000:1 real world

HDR imaging and tonemapping are distinct techniques with different goals

Two challenges:

HDR imaging compensates for *sensor* limitations

1. HDR imaging – which parts of the world do we measure in the 8-14 bits available to our sensor?
2. Tonemapping – which parts of the world do we show in the 4-10 bits available to our display?

Tonemapping compensates for *display* limitations

High dynamic range imaging














Key idea

1. Exposure bracketing: Capture multiple LDR images at different exposures




2. Merging: Combine them into a single HDR image



Key idea

1. Exposure bracketing: Capture multiple LDR images at different exposures




2. Merging: Combine them into a single HDR image




Ways to vary exposure


1. Shutter speed



2. F-stop (aperture, iris)



3. ISO



4. Neutral density (ND) filters



Pros and cons of each for HDR?

Ways to vary exposure

1. Shutter speed

- Range: about 30 sec to 1/4000 sec (6 orders of magnitude)
- Pros: repeatable, linear
- Cons: noise and motion blur for long exposure

2. F-stop (aperture, iris)

- Range: about f/0.98 to f/22 (3 orders of magnitude)
- Pros: fully optical, no noise
- Cons: changes depth of field

3. ISO

- Range: about 100 to 1600 (1.5 orders of magnitude)
- Pros: no movement at all
- Cons: noise

4. Neutral density (ND) filters

- Range: up to 6 densities (6 orders of magnitude)
- Pros: works with strobe/flash
- Cons: not perfectly neutral (color shift), extra glass (interreflections, aberrations), need to touch camera (shake)

Exposure bracketing with shutter speed

Note: shutter times usually obey a power series – each “stop” is a factor of 2

1/4, 1/8, 1/15, 1/30, 1/60, 1/125, 1/250, 1/500, 1/1000 sec

usually really is

1/4, 1/8, 1/16, 1/32, 1/64, 1/128, 1/256, 1/512, 1/1024 sec

Questions:

1. How many exposures?
2. What exposures?

Exposure bracketing with shutter speed

Note: shutter times usually obey a power series – each “stop” is a factor of 2

1/4, 1/8, 1/15, 1/30, 1/60, 1/125, 1/250, 1/500, 1/1000 sec

usually really is

1/4, 1/8, 1/16, 1/32, 1/64, 1/128, 1/256, 1/512, 1/1024 sec


Questions:

1. How many exposures?
2. What exposures?

Answer: Depends on the scene, but a good default is 5 exposures, the metered exposure and +/- 2 stops around that.

Key idea

1. Exposure bracketing: Capture multiple LDR images at different exposures




2. Merging: Combine them into a single HDR image




The image processing pipeline

The sequence of image processing operations applied by the camera's image signal processor (ISP) to convert a RAW image into a "conventional" image.




The image processing pipeline

The sequence of image processing operations applied by the camera's image signal processor (ISP) to convert a RAW image into a "conventional" image.




RAW images have a linear response curve

Colorchecker: Great tool for radiometric and color calibration.




Patches at bottom row have log-reflectance that increases linearly.

linear RAW



Over/under exposure



in highlights we are limited by clipping



in shadows we are limited by noise




RAW (linear) image formation model

Real scene flux for image pixel (x,y) : $\Phi(x, y)$

Exposure time:


t_5



t_4



t_3



t_2



t_1




What is an expression for the image $I_{\text{linear}}(x,y)$ as a function of $\Phi(x,y)$?

RAW (linear) image formation model

Real scene flux for image pixel (x,y) : $\Phi(x, y)$

Exposure time:


t_5




t_4



t_3



t_2



t_1



What is an expression for the image $I_{\text{linear}}(x,y)$ as a function of $\Phi(x,y)$?

$$I_{\text{linear}}(x,y) = \text{clip}[t_i \cdot \Phi(x,y) + \text{noise}]$$

How would you merge these images into an HDR one?

Merging RAW (linear) exposure stacks

For each pixel:

1. Find “valid” images
2. Weight valid pixel values appropriately
3. Form a new pixel value as the weighted average of valid pixel values

How would you implement steps 1-2?

t_5



t_4



t_3



t_2



t_1



Merging RAW (linear) exposure stacks

For each pixel:

1. Find “valid” images

← (noise) $0.05 < \text{pixel} < 0.95$ (clipping)

2. Weight valid pixel values appropriately

● noise

● valid


3. Form a new pixel value as the weighted average of valid pixel values

● clipped


t_5




t_4




t_3



t_2



t_1



Merging RAW (linear) exposure stacks

For each pixel:

1. Find “valid” images ← (noise) $0.05 < \text{pixel} < 0.95$ (clipping)
2. Weight valid pixel values appropriately ← (pixel value) / t_i
3. Form a new pixel value as the weighted average of valid pixel values

t_5



t_4



t_3




t_2



t_1



Merging result (after tonemapping)




What if I cannot use RAW?

Radiometric calibration


The image processing pipeline

- Can you foresee any problem when we switch from RAW to rendered images?



The image processing pipeline

- Can you foresee any problem when we switch from RAW to rendered images?
- How do we deal with the nonlinearities?



Radiometric calibration


The process of measuring the camera's response curve. Can be done in three ways:

- Take images of scenes with different flux while keeping exposure the same.
- Takes images under different exposures while keeping flux the same.
- Takes images of scenes with different flux and under different exposures.


Same camera exposure, varying scene flux

Colorchecker: Great tool for radiometric and color calibration.

e.g. JPEG



Patches at bottom row have log-reflectance that increases linearly.



Different values correspond to patches of increasing reflected flux.


Same scene flux, varying camera exposure

White balance card: Great tool for white balancing and radiometric calibration.



All points on (the white part of) the target have the same reflectance.

e.g. JPEG




Different values correspond to images taken under increasing camera exposure.

Varying both scene flux and camera exposure

You can do this using the LDR exposure stack itself.

- More information in homework assignment 2.

Different scene flux, same camera exposure




Same scene flux, different camera exposure

Non-linear image formation model

Real scene flux for image pixel (x,y) : $\Phi(x, y)$

Exposure time: t_i



$$I_{\text{linear}}(x,y) = \text{clip}[t_i \cdot \Phi(x,y) + \text{noise}]$$


$$I_{\text{non-linear}}(x,y) = f[I_{\text{linear}}(x,y)]$$

How would you merge the non-linear images into an HDR one?

Non-linear image formation model

Real scene flux for image pixel (x,y) : $\Phi(x, y)$

Exposure time: t_i



$$I_{\text{linear}}(x,y) = \text{clip}[t_i \cdot \Phi(x,y) + \text{noise}]$$


$$I_{\text{non-linear}}(x,y) = f[I_{\text{linear}}(x,y)]$$

$$I_{\text{est}}(x,y) = f^{-1}[I_{\text{non-linear}}(x,y)]$$

Use inverse transform to estimate linear image, then proceed as before

Linearization

$$I_{\text{non-linear}}(x,y) = f[I_{\text{linear}}(x,y)]$$



$$I_{\text{est}}(x,y) = f^{-1}[I_{\text{non-linear}}(x,y)]$$



Merging non-linear exposure stacks

1. Calibrate response curve

2. Linearize images

For each pixel:

3. Find “valid” images

← (noise) $0.05 < \text{pixel} < 0.95$ (clipping)

4. Weight valid pixel values appropriately

← $(\text{pixel value}) / t_i$

5. Form a new pixel value as the weighted average of valid pixel values


→ Same steps as in the RAW case.

What if I cannot measure the response curve?

You may find information in the image itself

If you cannot do calibration, take a look at the image's EXIF data (if available).

Often contains information about tone reproduction curve and color space.



Tone reproduction curves

The exact tone reproduction curve depends on the camera.


- Often well approximated as L^γ , for different values of the power γ (“gamma”).
- A good default is $\gamma = 1/2.2$.



before gamma



after gamma



If nothing else, take the square of your image to approximately remove effect of tone reproduction curve.

Other aspects of HDR imaging

Relative vs absolute flux

Final fused HDR image gives flux only up to a global scale


- If we know exact flux at one point, we can convert relative HDR image to absolute flux map



HDR image
(relative flux)



bhphoto.com



W/sr/m2
121.741
28.869
6.846
1.623
0.384
0.091
0.021
0.005

absolute
flux map

spotmeter (absolute
flux at one point)

Basic HDR approach

1. Capture multiple LDR images at different exposures
2. Merge them into a single HDR image

Any problems with this approach?

Basic HDR approach

1. Capture multiple LDR images at different exposures
2. Merge them into a single HDR image

Problem: Very sensitive to movement

- Scene must be completely static
- Camera must not move

Most modern automatic HDR solutions include an alignment step before merging exposures

How do we store HDR images?


- Most standard image formats store integer 8-bit images
- Some image formats store integer 12-bit or 16-bit images
- HDR images are floating point 32-bit or 64-bit images

How do we store HDR images?

Use specialized image formats for HDR images


portable float map (.pfm)

- very simple to implement




Radiance format (.hdr)

- supported by Matlab



OpenEXR format (.exr)

- multiple extra features



Another type of HDR images

Light probes: place a chrome sphere in the scene and capture an HDR image

- Used to measure real-world illumination environments (“environment maps”)




Application: image-based relighting
(later lecture)

Another way to create HDR images

Physics-based renderers simulate flux maps (relative or absolute)

- Their outputs are very often HDR images



Our devices do not match the real world

- 10:1 photographic print (higher for glossy paper)
- 20:1 artist's paints
- 200:1 slide film
- 500:1 negative film
- 1000:1 LCD display
- 2000:1 digital SLR (at 12 bits)
- 100000:1 real world

HDR imaging and tonemapping are distinct techniques with different goals

Two challenges:


HDR imaging compensates for *sensor* limitations

1. HDR imaging – which parts of the world do we measure in the 8-14 bits available to our sensor?
2. Tonemapping – which parts of the world do we show in the 4-10 bits available to our display?

Tonemapping compensates for *display* limitations



Tonemapping

How do we display our HDR images?



Linear scaling

Scale image so that maximum value equals 1.




HDR image *looks* underexposed because of the display's limited dynamic range, but is *not* actually underexposed.



Linear scaling

Scale image so that 10% value equals 1.



HDR image *looks* saturated because of the display's limited dynamic range, but is *not* actually saturated.




Can you think of something better?

Photographic tonemapping

Apply the same non-linear scaling to all pixels in the image so that:

- Bring everything within range $\rightarrow$ asymptote to 1
- Leave dark areas alone $\rightarrow$ slope = 1 near 0



$$I_{\text{display}} = \frac{I_{\text{HDR}}}{1 + I_{\text{HDR}}}$$


(exact formula more complicated)

- Photographic because designed to approximate film zone system.
- Perceptually motivated, as it approximates our eye's response curve.


What is the zone system?

- Technique formulated by Ansel Adams for film development.
- Still used with digital photography.

Zone		Description
	0	Pure black
	I	Near black, with slight tonality but no texture
	II	Textured black; the darkest part of the image in which slight detail is recorded
	III	Average dark materials and low values showing adequate texture
	IV	Average dark foliage, dark stone, or landscape shadows
	V	Middle gray: clear north sky; dark skin, average weathered wood
	VI	Average Caucasian skin; light stone; shadows on snow in sunlit landscapes
	VII	Very light skin; shadows in snow with acute side lighting
	VIII	Lightest tone with texture: textured snow
	IX	Slight tone without texture; glaring snow
	X	Pure white: light sources and specular reflections




Examples



photographic tonemapping




linear scaling (map 10% to 1)



linear scaling (map 100% to 1)

Compare with LDR images



photographic tonemapping



high exposure



low exposure

Dealing with color

If we tonemap all channels the same, colors are washed out



Can you think of a way to deal with this?

Intensity-only tonemapping

tonemap
intensity (e.g.,
luminance Y in
 xyY)



leave color
the same
(e.g., xy in
 xyY)



How would you implement this?

Comparison


Color now OK, but some details are washed out due to loss of contrast



Can you think of a way to deal with this?

Low-frequency intensity-only tonemapping

tonemap low-frequency
intensity component



leave high-frequency
intensity component
the same



leave color the same



How would you implement this?


Comparison

We got nice color and contrast, but now we've run into the *halo* plague



Can you think of a way to deal with this?

Edge-aware filtering and tonemapping



Comparison

We fixed the halos without losing contrast






Gradient-domain processing and tonemapping

Compute gradients, scale and merge them, then integrate (solve Poisson problem).

- More in later lecture.




Comparison (which one do you like better?)



photographic



bilateral filtering



gradient-domain

Comparison (which one do you like better?)



photographic



bilateral filtering



gradient-domain

Comparison (which one do you like better?)



There is no ground-truth: which one looks better is entirely subjective



photographic

bilateral filtering

gradient-domain

Tonemapping for a single image

Modern DSLR sensors capture about 3 stops of dynamic range.

- Tonemap single RAW file instead of using camera's default rendering.

result from image
processing pipeline
(basic tone
reproduction)




tonemapping using
bilateral filtering (I
think)

Tonemapping for a single image

Modern DSLR sensors capture about 3 stops of dynamic range.

- Tonemap single RAW file instead of using camera's default rendering.



Careful not to “tonemap” noise.

Some notes about HDR imaging and tonemapping

Our devices do not match the real world

- 10:1 photographic print (higher for glossy paper)
- 20:1 artist's paints
- 200:1 slide film
- 500:1 negative film
- 1000:1 LCD display
- 2000:1 digital SLR (at 12 bits)
- 100000:1 real world

HDR imaging and tonemapping are distinct techniques with different goals

Two challenges:

HDR imaging compensates for *sensor* limitations

1. HDR imaging – which parts of the world do we measure in the 8-14 bits available to our sensor?
2. Tonemapping – which parts of the world do we show in the 4-10 bits available to our display?

Tonemapping compensates for *display* limitations

A note about terminology

“High-dynamic-range imaging” is used to refer to a lot of different things:

1. Using single RAW images.
2. Performing radiometric calibration.
3. Merging an exposure stack.
4. Tonemapping an image (linear or non-linear, HDR or LDR).
5. Some or all of the above.

Technically, HDR imaging and tonemapping are distinct processes:

- HDR imaging is the process of creating a radiometrically linear image, free of overexposure and underexposure artifacts. This is achieved using some combination of 1-3, depending on the imaging scenario.
- Tonemapping (step 4) process of mapping the intensity values in an image (linear or non-linear, HDR or LDR) to the range of tones available in a display.


But:

- In consumer photography, “HDR photography” is often used to refer to both HDR imaging (steps 1-3) and tonemapping (step 4).

Another note about terminology

Tonemapping is just another form of tone reproduction.

- Many ISPs implement the tonemapping algorithms we discussed for tone reproduction.



A note of caution

- HDR photography can produce very visually compelling results.







A note of caution

- HDR photography can produce very visually compelling results.
- It is also a very routinely abused technique, resulting in awful results.









A note of caution

- HDR photography can produce very visually compelling results.
- It is also a very routinely abused technique, resulting in awful results.
- The problem typically is tonemapping, not HDR imaging itself.

A note about HDR today

- Most cameras (even phone cameras) have automatic HDR modes/apps.
- Popular-enough feature that phone manufacturers are actively competing about which one has the best HDR.
- The technology behind some of those apps (e.g., Google's HDR+) is published in SIGGRAPH and SIGGRAPH Asia conferences.

Burst photography for high dynamic range and low-light imaging on mobile cameras

Samuel W. Hasinoff
Jonathan T. Barron

Dillon Sharlet
Florian Kainz
Google Research

Ryan Geiss
Jiawen Chen

Andrew Adams
Marc Levoy



Figure 1: A comparison of a conventional camera pipeline (left, middle) and our burst photography pipeline (right) running on the same cell-phone camera. In this low-light setting (about 0.7 lux), the conventional camera pipeline underexposes (left). Brightening the image (middle) reveals heavy spatial denoising, which results in loss of detail and an unpleasantly blotchy appearance. Fusing a burst of images increases the signal-to-noise ratio, making aggressive spatial denoising unnecessary. We encourage the reader to zoom in. While our pipeline excels in low-light and high-dynamic-range scenes (for an example of the latter see figure 10), it is computationally efficient and reliably artifact-free, so it can be deployed on a mobile camera and used as a substitute for the conventional pipeline in almost all circumstances. For readability the figure has been made uniformly brighter than the original photographs.

Abstract

Cell phone cameras have small apertures, which limits the number of photons they can gather, leading to noisy images in low light. They also have small sensor pixels, which limits the number of electrons each pixel can store, leading to limited dynamic range. We describe a computational photography pipeline that captures, aligns, and merges a burst of frames to reduce noise and increase dynamic range. Our system has several key features that help make it robust and efficient. First, we do not use bracketed exposures. Instead, we capture frames of constant exposure, which makes alignment more robust, and we set this exposure low enough to avoid blowing out highlights. The resulting merged image has clean shadows and high bit depth, allowing us to apply standard HDR tone mapping methods. Second, we begin from Bayer raw frames rather than the demosaicked RGB (or YUV) frames produced by hardware Image Signal Processors (ISPs) common on mobile platforms. This gives us more bits per pixel and allows us to circumvent the ISP's unwanted tone mapping and spatial denoising. Third, we use a novel FFT-based alignment algorithm and a hybrid 2D/3D Wiener filter to denoise and merge the frames in a burst. Our implementation is built atop Android's Camera2 API, which provides per-frame camera control and access to raw imagery, and is written in the Halide domain-specific language (DSL). It runs in 4 seconds on device (for a 12 Mpix image), requires no user intervention, and ships on several mass-produced cell phones.

Keywords: computational photography, high dynamic range

Concepts: •Computing methodologies → Computational photography; Image processing;

1 Introduction

The main technical impediment to better photographs is lack of light. In indoor or night-time shots, the scene as a whole may provide insufficient light. The standard solution is either to apply analog or digital gain, which amplifies noise, or to lengthen exposure time, which causes motion blur due to camera shake or subject motion. Surprisingly, daytime shots with high dynamic range may also suffer from lack of light. In particular, if exposure time is reduced to avoid

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 Copyright held by the owner/authors. Publication rights licensed to ACM. SA '16 Technical Papers, December 05 - 08, 2016, Macao ISBN: 978-1-4503-4514-9/16/12 DOI: <http://dx.doi.org/10.1145/2980179.2980254>

References

Basic reading:

- Szeliski textbook, Sections 10.1, 10.2.
- Debevec and Malik, “Recovering High Dynamic Range Radiance Maps from Photographs,” SIGGRAPH 1997.
- Mitsunaga and Nayar, “Radiometric self calibration,” CVPR 1999.
The two classical papers on radiometric calibration and HDR imaging, which more or less started HDR imaging research in computer vision and graphics.
- Reinhard et al., “Photographic Tone Reproduction for Digital Images,” SIGGRAPH 2002.
The photographic tonemapping paper, including a very nice discussion of the zone system for film.

Additional reading:

- Reinhard et al., “High Dynamic Range Imaging, Second Edition: Acquisition, Display, and Image-Based Lighting,” Morgan Kaufmann 2010.
A very comprehensive book about everything relating to HDR imaging and tonemapping.
- Durand and Dorsey, “Fast bilateral filtering for the display of high-dynamic-range images,” SIGGRAPH 2002.
The paper on tonemapping using bilateral filtering.
- Fattal et al., “Gradient Domain High Dynamic Range Compression,” SIGGRAPH 2002.
The paper on tonemapping using gradient-domain processing.
- Debevec, “Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography,” SIGGRAPH 1998.
The original HDR light probe paper.
- Hasinoff et al., “Burst photography for high dynamic range and low-light imaging on mobile cameras,” SIGGRAPH Asia 2016.
The paper describing Google’s HDR+.
- Ward, “The radiance lighting simulation and rendering system,” SIGGRAPH 1994.
- Ward, “High Dynamic Range Image Encodings,” http://www.anywhere.com/gward/hdrenc/hdr_encodings.html
The paper that introduced (among other things) the .hdr image format for HDR images. The website has a very detailed discussion of HDR image formats.
- Shah et al., “High-quality Motion Deblurring from a Single Image,” SIGGRAPH 2008.
- Fergus et al., “Removing Camera Shake from a Single Image,” SIGGRAPH 2006.
Two standard papers on motion deblurring for dealing with long shutter speeds.
- Barron et al., “Fast Bilateral-Space Stereo for Synthetic Defocus,” CVPR 2015.
The paper that introduced the lens blur algorithm.
- Kuang et al., “Evaluating HDR rendering algorithms,” TAP 2007.
One of many, many papers trying to do a perceptual evaluation of different tonemapping algorithms.
- Levoy, “Extreme imaging using cell phones,” <http://graphics.stanford.edu/talks/seeinthedark-public-15sep16.key.pdf>
A set of slides by Marc Levoy on the challenges of HDR imaging and modern approaches for addressing them.