Noise and ISO

CS 178, Spring 2013

Marc Levoy
Computer Science Department
Stanford University

Outline

- examples of camera sensor noise
  - don't confuse it with JPEG compression artifacts
- probability, mean, variance, signal-to-noise ratio (SNR)
- laundry list of noise sources
  - photon shot noise, dark current,
    hot pixels, fixed pattern noise, read noise
- SNR (again), dynamic range (DR), bits per pixel
- ISO
- denoising
  - by aligning and averaging multiple shots
  - by image processing will be covered in a later lecture

Nokia N95 cell phone at dusk

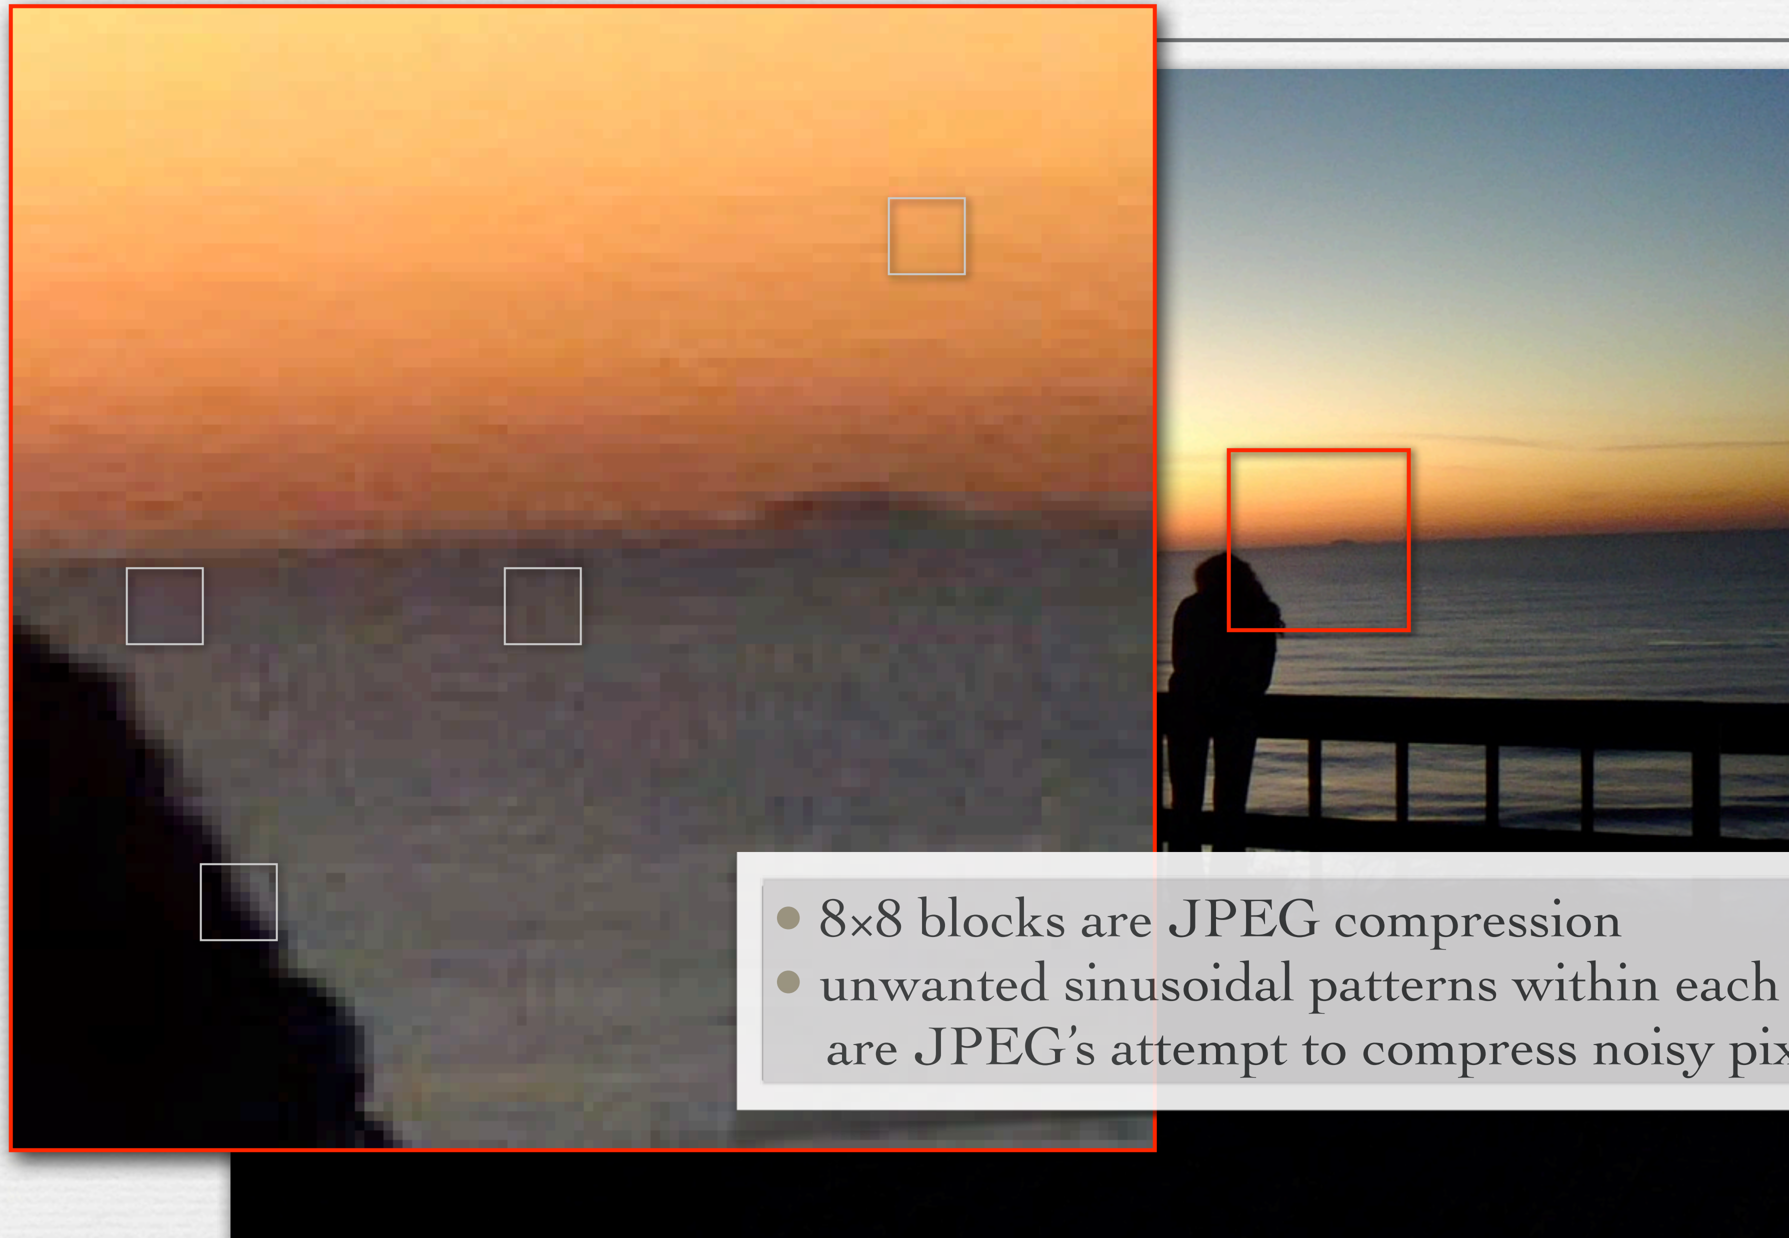

- 8×8 blocks are JPEG compression
- unwanted sinusoidal patterns within each block are JPEG's attempt to compress noisy pixels

Canon 5D II at noon

chromatic aberration!

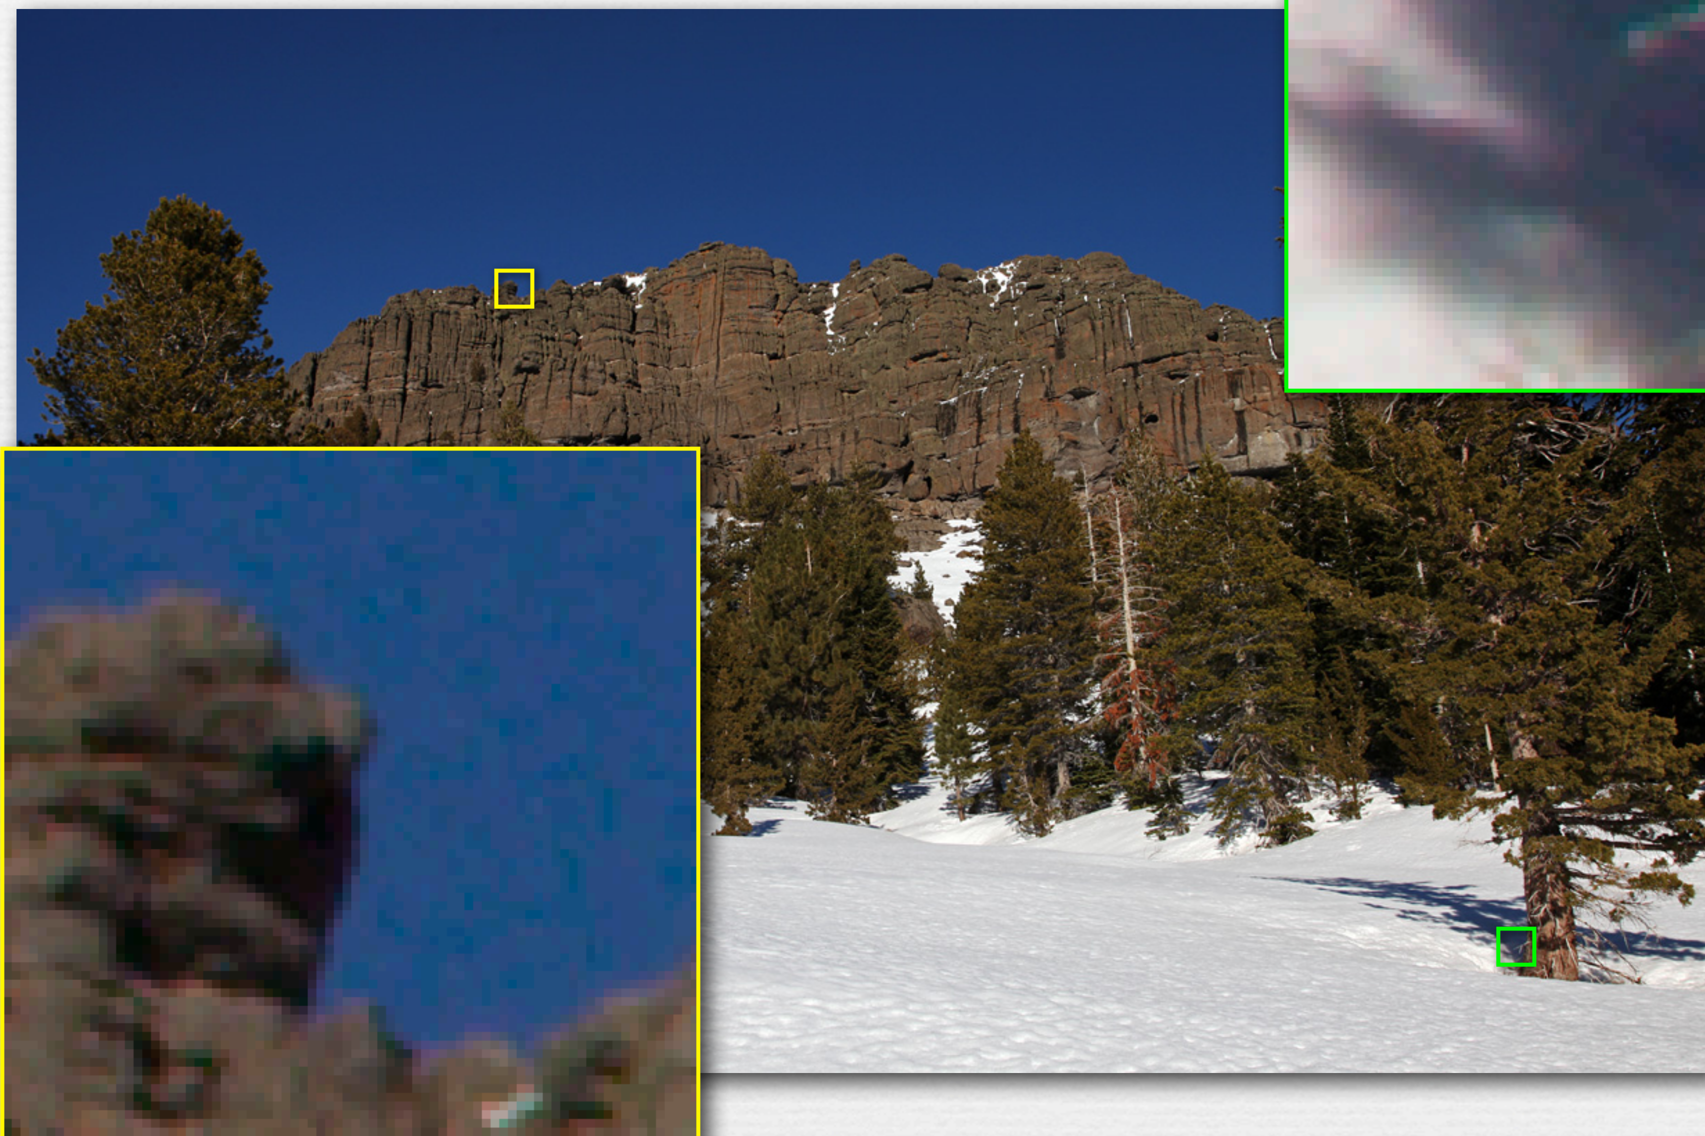

- ISO 200
- f/13.0
- 1/320 sec
- RAW w/o denoising

Canon 5D II at noon

post-processed using Canon software

- ISO 200
- f/13.0
- 1/320 sec
- RAW w/o denoising

Canon 5D II at dusk

- ISO 6400
- f/4.0
- 1/13 sec
- RAW w/o denoising

Canon 5D II at dusk

Canon 5D II at dusk

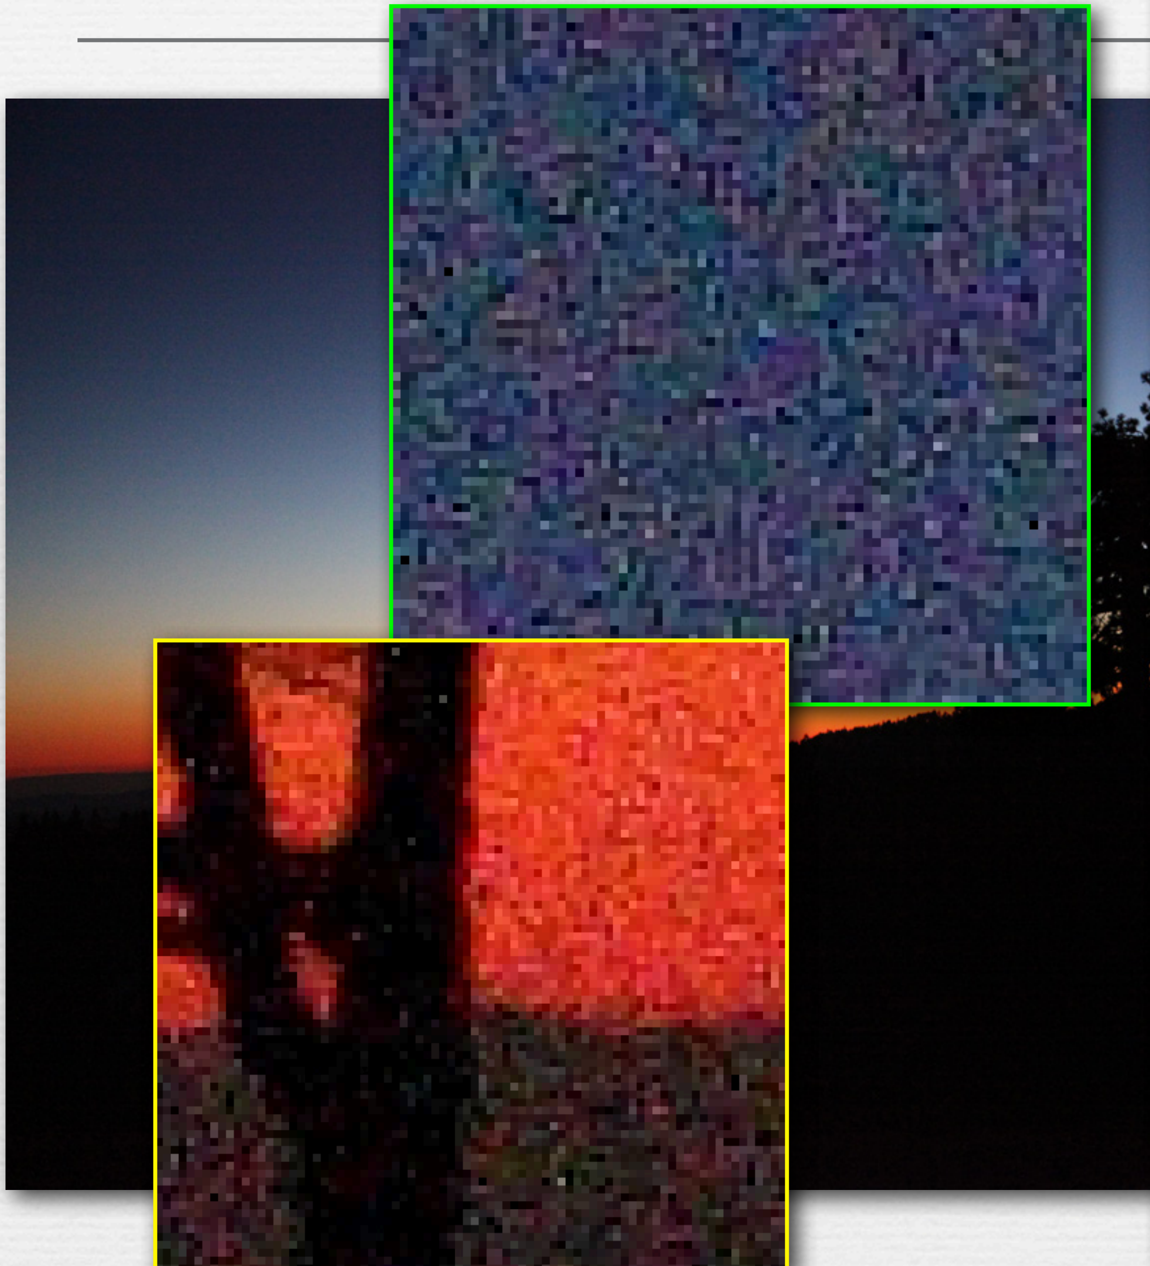

Photon shot noise

- the number of photons arriving during an exposure varies from exposure to exposure and from pixel to pixel, even if the scene is completely uniform
- this number is governed by the Poisson distribution

Poisson distribution

- expresses the probability that a certain number of events will occur during an interval of time
- applicable to events that occur
  - with a known average rate, and
  - independently of the time since the last event
- if on average λ events occur in an interval of time, the probability *p* that *k* events occur instead is

$$p(k;\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$

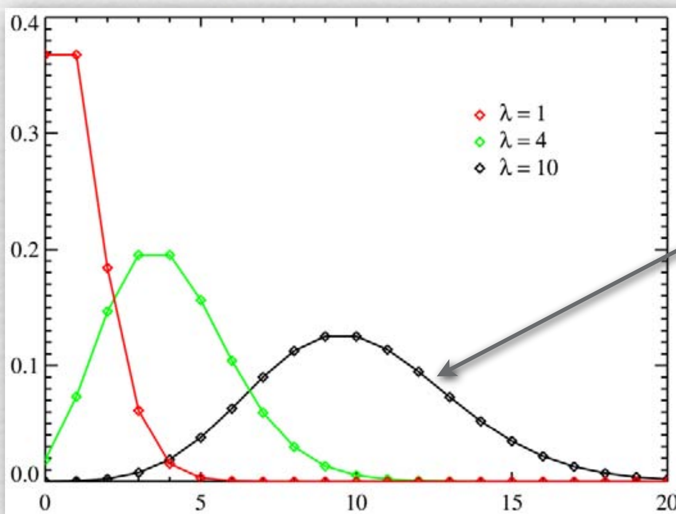

probability density function

Mean and variance

- the mean of a probability density function *p(x)* is

$$\mu = \int x\, p(x)\, dx$$

- the variance of a probability density function *p(x)* is

$$\sigma^2 = \int (x-\mu)^2\, p(x)\, dx$$

- the mean and variance of the Poisson distribution are

$$\mu = \lambda$$

$$\sigma^2 = \lambda$$

- the standard deviation is

$$\sigma = \sqrt{\lambda}$$

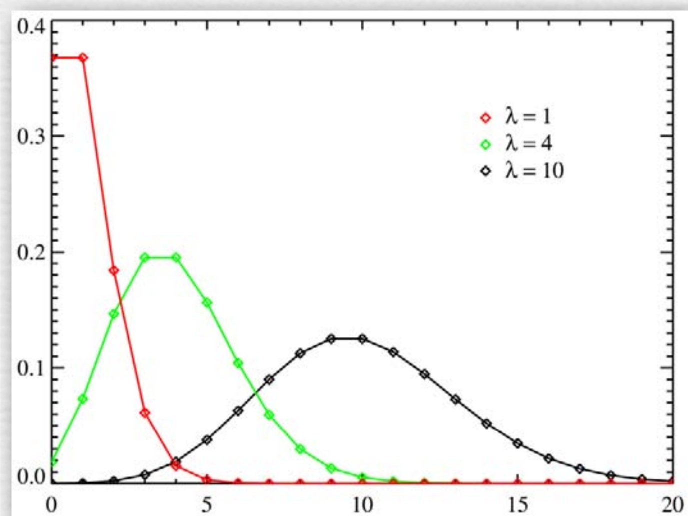

Deviation grows slower than the average.

Signal-to-noise ratio (SNR)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$SNR\ (\text{dB}) = 20 \log_{10}\left(\frac{\mu}{\sigma}\right)$$

- example
  - if SNR improves from 100:1 to 200:1,
    then it improves by $20 \log_{10}(200) - 20 \log_{10}(100) = +6$ dB

Photon shot noise (again)

- photons arrive in a Poisson distribution

$$\mu = \lambda$$
$$\sigma = \sqrt{\lambda}$$

- so $$\boxed{SNR = \frac{\mu}{\sigma} = \sqrt{\lambda}}$$

It must seem surprising that SNR could rise as a scene gets brighter (a good thing) even though noise is rising at the same time (a bad thing).

Here's a simple example. If on average 9 photons arrive at a pixel during an exposure, the standard deviation of this (according to the Poisson distribution) is sqrt(9) = 3 photons. This means that SNR = mean/stddev = 9/3 = 3:1. Now suppose instead that 100 photons arrive at the pixel, either because the scene got brighter or we increased the exposure time or we switched to a camera with bigger pixels. Now the stddev is sqrt(100) = 10, and SNR = 100/10 = 10:1. The noise got worse (stddev of 10 photons versus 3 photons), but the SNR got better (10:1 versus 3:1). The apparent image quality will be better in the second case.

- shot noise scales as square root of number of photons
- examples
  - doubling the width and height of a pixel increases its area by 4×, hence # of photons by 4×, hence SNR by 2× or +6 dB
  - opening the aperture by 1 f/stop increases the # of photons by 2×, hence SNR by √2 or +3 dB

Empirical example

- Kodak Q14 test chart

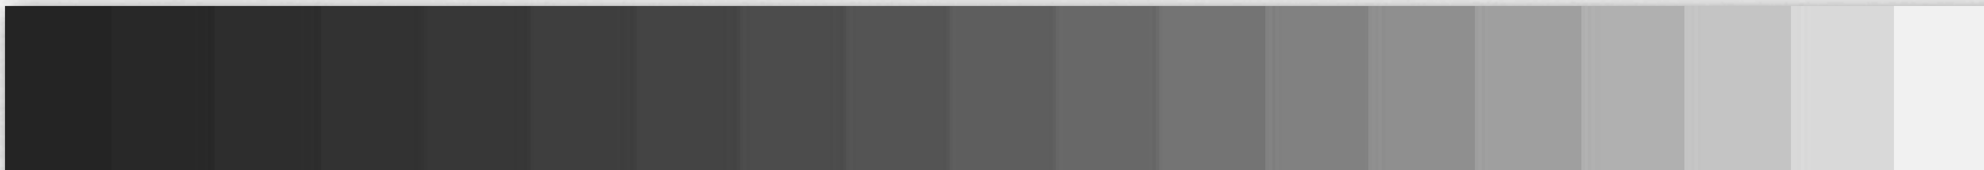

- Canon 10D, ISO 1600, crop from JPEG image

matest.com/docs/noise.html)

cropped

cropped

- noise grows as sqrt(signal), so better SNR in light tile than dark
- also, JPEG's gamma correction compresses differences among bright pixels
- these effects compound, so you observe noise only in the darkest regions

Dark current

- electrons dislodged by random thermal activity
- increases linearly with exposure time
- increases exponentially with temperature
- varies across sensor, and includes its own shot noise

don't confuse with photon shot noise

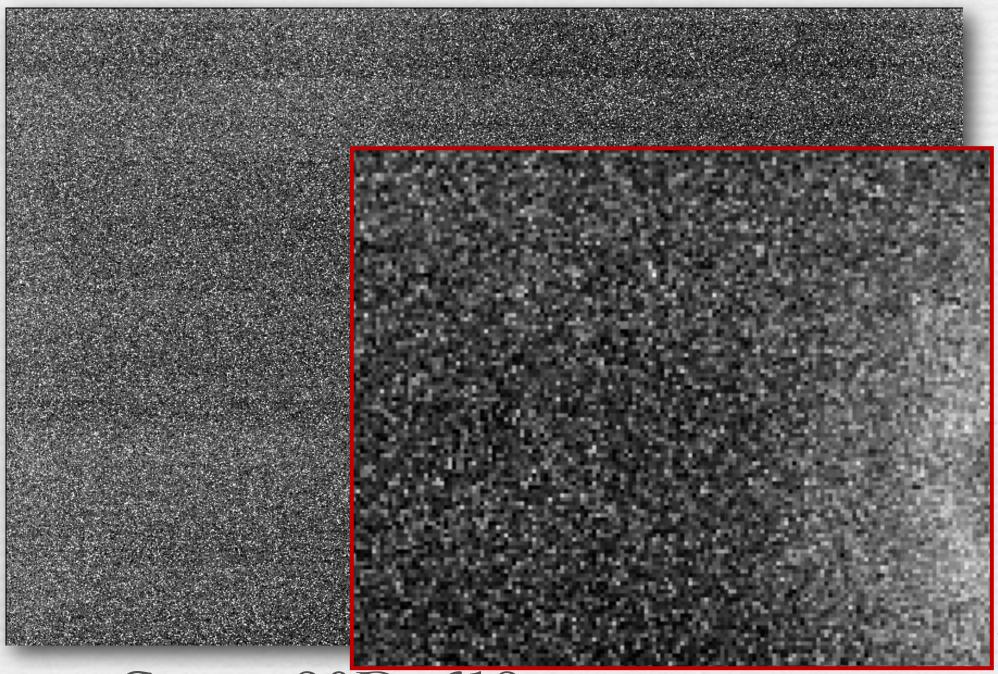

(http://theory.uchicago.edu/~ejm/pix/20d/tests/noise/)

Canon 20D, 612 sec exposure

Hot pixels

- electrons leaking into well due to manufacturing defects
- increases linearly with exposure time
- increases with temperature, but hard to model
- changes over time, and every camera has them

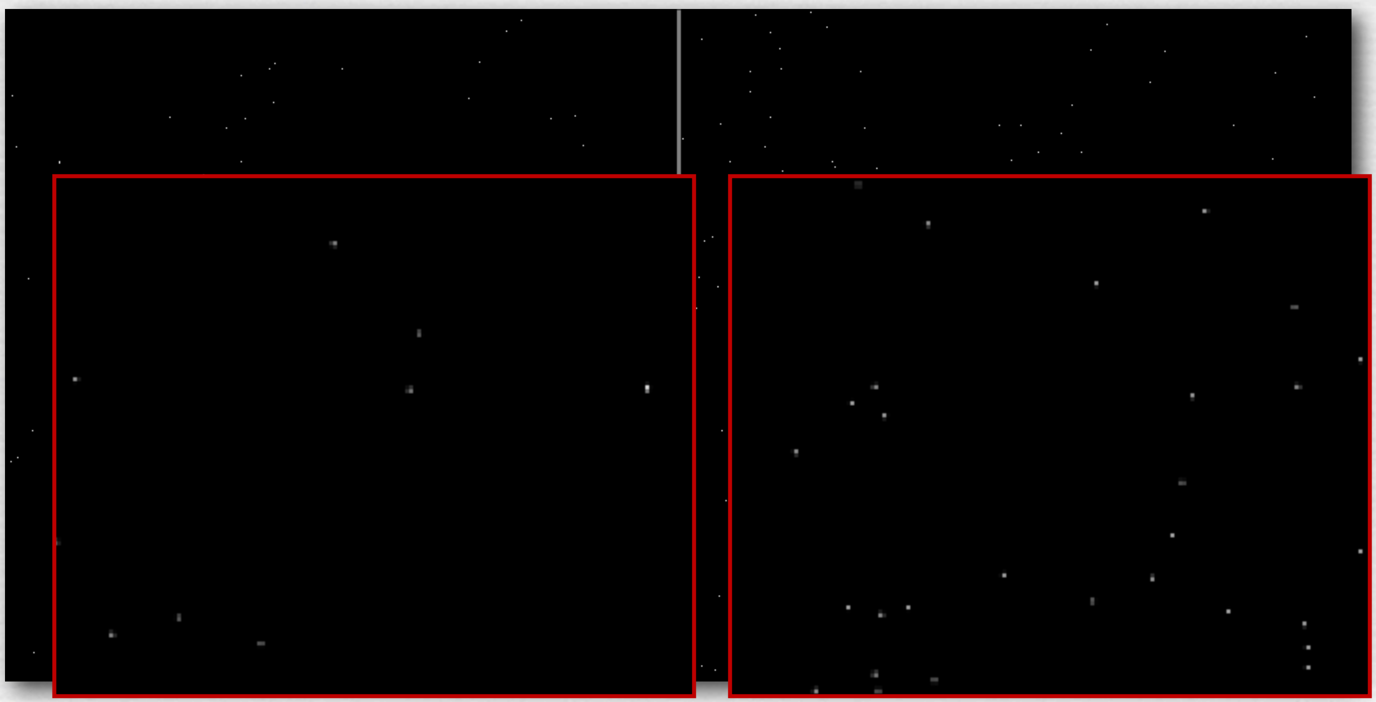

Canon 20D, 15 sec and 30 sec exposures

Fixing dark current and hot pixels

- example
  - Aptina MT9P031 (in Nokia N95 cell phone)
  - full well capacity = ~8500 electrons/pix
  - dark current = 25 electrons/pix/sec at 55°C

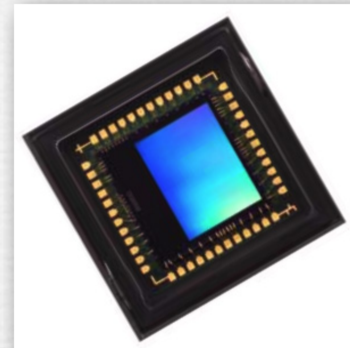

- solution #1: chill the sensor
  - Retiga 4000R bioimaging camera
  - Peltier cooled 25°C below ambient
  - full well capacity = 40,000 electrons/pix
  - dark current = 1.64 electrons/pix/sec

- solution #2: dark frame subtraction
  - available on high-end SLRs
  - compensates for average dark current
  - also compensates for hot pixels and FPN

Fixed pattern noise (FPN)

- manufacturing variations across pixels, columns, blocks
- mainly in CMOS sensors
- doesn't change over time, so read once and subtract

Canon 20D, ISO 800, cropped

Read noise

- thermal noise in readout circuitry
- again, mainly in CMOS sensors
- not fixed pattern, so only solution is cooling

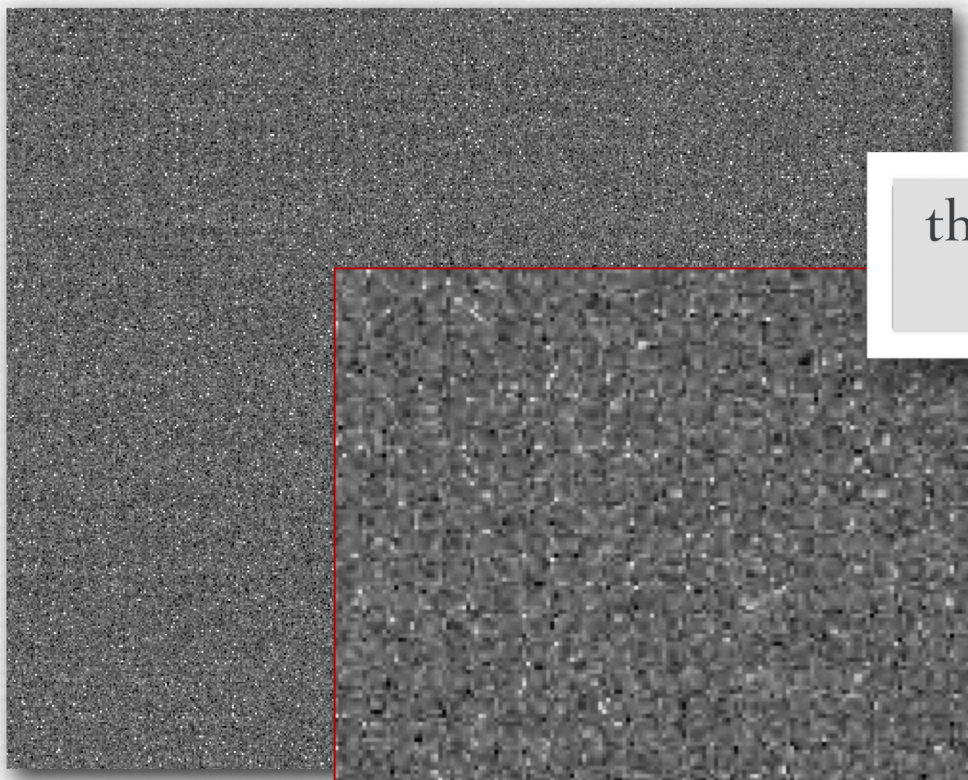

this image tainted by JPEG artifacts?

Canon 1Ds Mark III, cropped

Recap

- photon shot noise
  - unavoidable randomness in number of photons arriving
  - grows as the square root of the number of photons, so brighter lighting and longer exposures will be less noisy
- dark current noise
  - grows with exposure time and sensor temperature
  - minimal for most exposure times used in photography
  - correct by subtraction, but only corrects for average dark current
- hot pixels, fixed pattern noise
  - caused by manufacturing defects, correct by subtraction
- read noise
  - electronic noise when reading pixels, unavoidable

Questions?

Signal-to-noise ratio

(with more detailed noise model)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$= \frac{P\,Q_e\,t}{\sqrt{P\,Q_e\,t + D\,t + N_r^2}}$$

SNR changes with scene brightness, aperture, and exposure time

- where

P = incident photon flux (photons/pixel/sec)

Q_e = quantum efficiency

t = exposure time (sec)

D = dark current (electrons/pixel/sec), including hot pixels

N_r = read noise (rms electrons/pixel), including fixed pattern noise

(formula from http://learn.hamamatsu.com/articles/ccdsnr.html)

Signal-to-noise ratio
(with more detailed noise model)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$= \frac{P\,Q_e\,t}{\sqrt{P\,Q_e\,t + D\,t + N_r^2}}$$

- examples
  - Retiga 4000R = (1000 × 55%) / √(1000 × 55% + 1.64 + 12^2) = 20.8:1 assuming 1000 photons/pixel/sec for 1 second
  - Aptina MT9P031 = (1000÷11×69%) / √(1000÷11×69% + 25 + 2.6^2) = 6.5:1 assuming pixels are 1/11 as large as Retiga's
- for 10 photons/pixel/sec for 100 seconds
  - Retiga = 18.7:1
  - Aptina = 1.2:1

Don't use your cell phone for astrophotography!

Dynamic range

To reiterate the difference between SNR and DR, signal-to-noise ratio (SNR) tells you how noisy an image will be at a particular light level, and a sensor will have a different SNR for each possible light level, while dynamic range (DR) is a single number giving the maximum possible range between saturation (for bright scenes) and the noise floor (for dark scenes). DR tells you nothing about how noisy a low-light image will be; it just says that it will be (barely) distinguishable from pure noise. So a cell phone might have as large a dynamic range as an SLR, but if its low-light images are very noisy (as they typically are), you wouldn't want to use it for low-light photography.

$$DR = \frac{\text{max output swing}}{\text{noise in the dark}} = \frac{\text{saturation level - } D\,t}{\sqrt{D\,t + N_r^2}}$$

full well capacity

- examples
  - Retiga 4000R = (40,000 - 1.64) / √(1.64 + 12^2) = 3,313:1 (11.7 bits) for a 1 second exposure
  - Aptina MT9P031 = (8500 - 25) / √(25 + 2.6^2) = 1500:1 (10.5 bits) for a 1 second exposure
- determines precision required in ADC, and useful # of bits in RAW image
- any less than ~10 bits would be < 8 bits after gamma correction for JPEG encoding, and you would see quantization artifacts

Low-light cameras

- compare to 10.5 bits for Aptina
- don't use your cell phone for fluorescence microscopy!

$$DR = \frac{\text{max output swing}}{\text{noise in the dark}} = \frac{\text{saturation level} - D\,t}{\sqrt{D\,t + N_r^2}}$$

- Andor iXon+888 back-illuminated CCD
  - $40,000
- performance
  - DR = (80,000 - 0.001) / √(0.001 + 6^2)
    = 13,333:1 (13.7 bits) for a 1 second exposure

if cooled to -75° C

- "electron multiplication" mode
  - DR = (80,000 - 0.001) / √(0.001 + <1^2)
    ≈ 80,000:1 (16.2 bits)
  - "can see a black cat in a coal mine"

can reliably detect a single photon

ISO - signal gain

- doubling ISO doubles the signal
  - linear with light, so same as 2× exposure time, or –1 f/stop
  - implemented as *analog amplification* on Canon 5D II up to ISO 6400; higher ISOs are implemented using *digital multiplication* after ADC?
- you want to amplify as early as possible during readout
  - if you amplify before read noise is added, and RN is independent of signal amplitude, then the amplified signal will have better SNR
- you especially want to amplify before quantization by ADC
  - if you quantize a low signal, then brighten it in Photoshop, you will see quantization artifacts (contouring)
  - if you quantize a very low signal, you may get zero (black)
- raising exposure typically improves SNR faster than raising ISO
  - thus, you should maximize exposure time until stopped by object motion blur, camera shake blur, or saturation; if stopped by blur, then raise ISO until stopped by saturation (i.e. don't clip whites)

The signal amplification pipeline

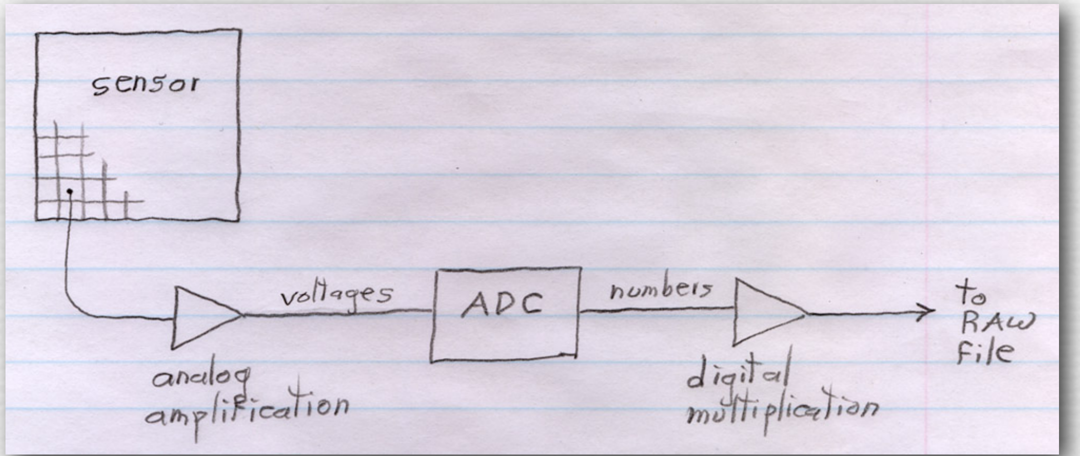

- raising the ISO is usually implemented as analog amplification (of voltages) before analog-to-digital conversion (ADC), but for high ISOs, some cameras may also perform digital multiplication (of numbers) after ADC
- analog amplification is better than digital multiplication, for the reasons given on the previous slide

To reiterate the "recipe" I gave in class, here's how to take a picture that minimizes noise:
1. Make your aperture as wide as you want it for depth of field.
2. Make your exposure as long as you dare make it, given handshake or object motion blur.
3. Raise the ISO to ensure an image that fills the range of numbers representable in the RAW or JPEG file, i.e. until the brightest object in the scene that you don't want to appear saturated just reaches white on the histogram.

All of these are done in the camera during shooting. Don't use Photoshop to brighten an image (except minor adjustments), because it will enhance noise more than raising the ISO will, and it may introduce quantization artifacts (contouring).

SNR and ISO over the years

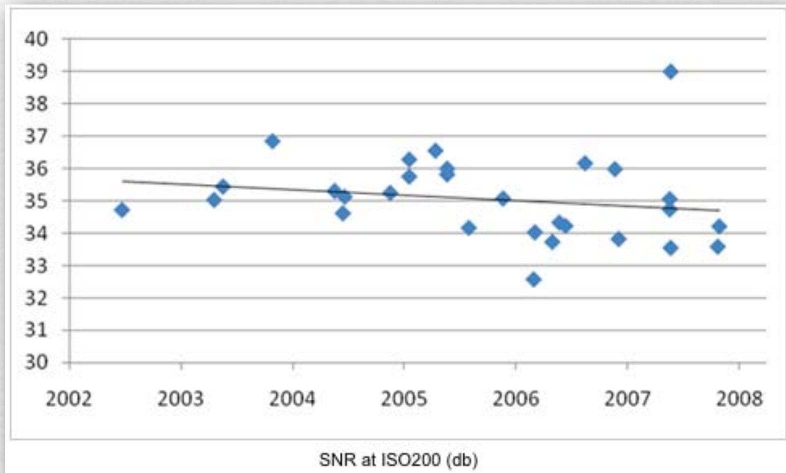

(http://www.dxomark.com/index.php/eng/Insights/SNR-evolution-over-time)

- SNR has been improving with better sensor designs
- but total # of megapixels has risen to offset these improvements, making pixels smaller, so SNR in a pixel has remained static

---

- display resolutions have not risen as fast as megapixels, so we're increasingly downsizing our images for display
- if you average 4 camera pixels to produce 1 for display, SNR doubles, so for the same display area, SNR has been improving

Effect of downsizing on image noise

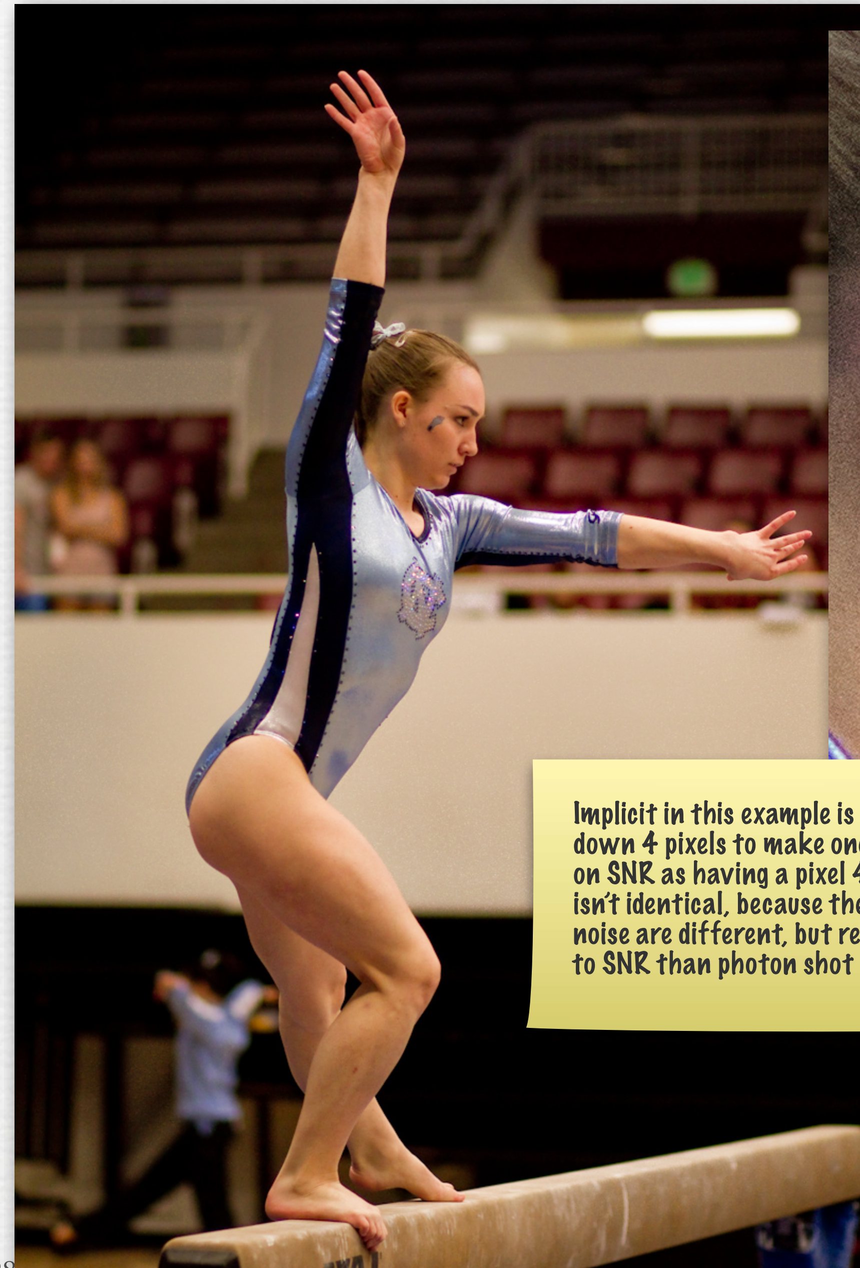

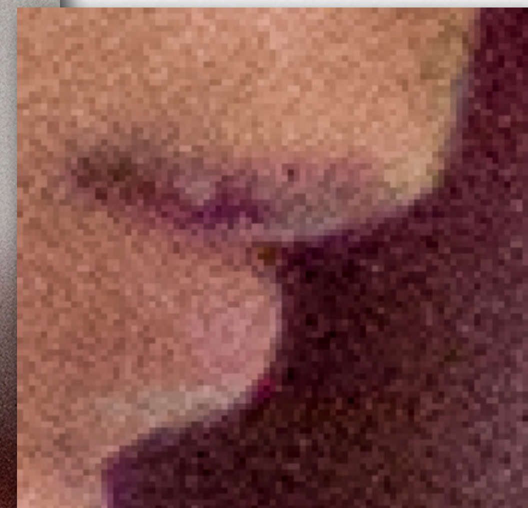

Implicit in this example is the notion that averaging down 4 pixels to make one pixel has a similar effect on SNR as having a pixel 4x as large. The effect isn't identical, because the contributions by read noise are different, but read noise is less important to SNR than photon shot noise.

averaged down

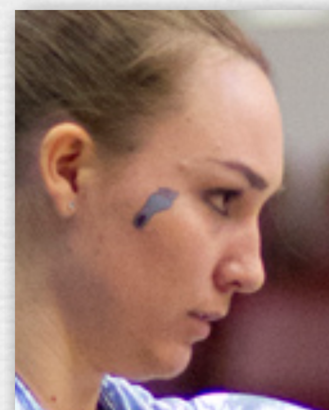

point sampled

SNR and ISO over the years

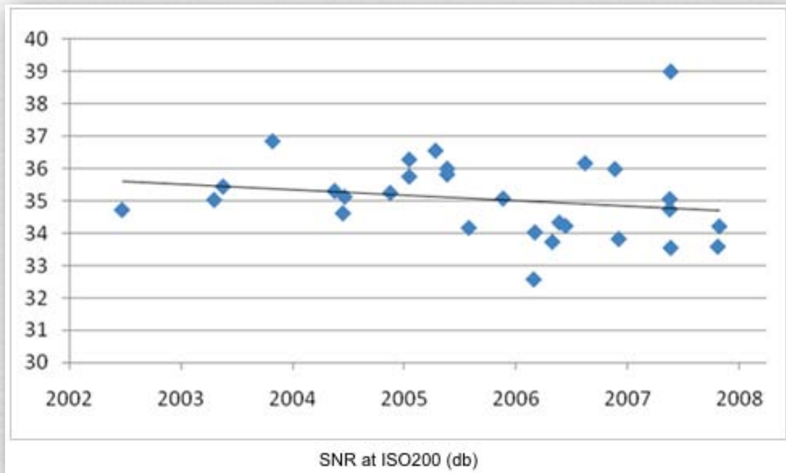

(http://www.dxomark.com/index.php/eng/Insights/SNR-evolution-over-time)

- SNR has been improving with better sensor designs
- but total # of megapixels has risen to offset these improvements, making pixels smaller, so SNR in a pixel has remained static

---

- display resolutions have not risen as fast as megapixels, so we're increasingly downsizing our images for display
- if you average 4 camera pixels to produce 1 for display, SNR doubles, so for the same display area, SNR has been improving
- this allows higher ISOs to be used in everyday photography

Nikon D3S, ISO 3200, photograph by Michael Kass

Nikon D3S, ISO 6400, photograph by Michael Kass

Nikon D3S, ISO 25,600, denoised in Lightroom 3, photograph by Fredo Durand

Nikon D3S, ISO 25,600, denoised in Lightroom 3, photograph by Fredo Durand

RAW image from camera, before denoising in Lightroom

Fredo says it was nearly too dark to read the menu, so it really looked like this (darkened)

or maybe it looked like this? (tone mapped to approximate human dark adaptation)

single frame
in dark room
using iPhone 4

average of
~30 frames
using SynthCam

SNR increases as
sqrt(# of frames)

Recap

- *signal-to-noise ratio* (SNR) is mean/stddev of pixel value
  - rises with sqrt(brightness and/or exposure time)
  - depends also on dark current and read noise
  - poor for short exposures and very long exposures
- *dynamic range* (DR) is max swing / noise in the dark
  - fixed for a particular sensor and exposure time
  - determines # of useful bits in RAW image
- *ISO* is amplification of signal before conversion to digital
  - maximize exposure time until camera or object blurs, then maximize ISO, making sure not to saturate
  - can combine multiple short-exposure high-ISO pictures

Questions?

Slide credits

- Eddy Talvala

- Filippov, A., *How many bits are really needed in the image pixels?* (sic), http://www.linuxdevices.com/articles/AT9913651997.html