CS 89.15/189.5, Fall 2015

COMPUTATIONAL ASPECTS OF DIGITAL PHOTOGRAPHY

Noise & Denoising

Wojciech Jarosz
wojciech.k.jarosz@dartmouth.edu

Dartmouth

Today's agenda

Info on paper discussions

Noise

Denoising with multiple images

- Probability review

Noise characteristics

Sources of noise

Class portraits!

https://youtu.be/o9BqrSAHbTc

Presentation guidelines

30 minutes per paper (presentation + discussion)

You may use Powerpoint/Keynote, blackboard, etc.

Focus on getting across the main points of the paper first

- *Present the paper as if everyone skimmed it but forgot it, or didn't understand.*
- First present the problem that the paper solves and the general approach.
- You should then give a clear and concise description of the main technical parts of the paper (algorithms, equations, etc).

After a webpage by Kyros Kutulakos

Presentation guidelines

Everyone will have read the paper before class.

- Your job should not be simply reciting what is in the paper
- Go beyond that, working out exactly how the algorithm (or theory) works and deciding how to present this in class.
- The best way to present an approach may not be the order in which things are described in the paper.

After a webpage by Kyros Kutulakos

Leave no stone unturned

A paper's content may not be sufficient to fully describe how a technique works.

- may depend on prior papers/techniques

A major goal of your presentation is to fill in these gaps and present a complete picture of the paper in class.

If there is something you don't understand, you must either work it out yourself, or come to office hours so that we can resolve it together.

After a webpage by Kyros Kutulakos

Presentation guidelines

Some authors provide presentations and other material online.

- Proper attribution rules apply

Practice, practice, practice

- You should practice your presentation at home, and time yourself, before coming to class.
- Pay attention to what you did (and did not) like about your classmates' presentation style, level of preparation, etc. with an eye toward improving your own presentation skills.

After a webpage by Kyros Kutulakos

Paper discussion

Everyone else will not be a passive observer

Discussant will initiate and facilitate the discussion

Everyone is expected to participate in the discussions

Things to think about

Limitations

- Do you think everything will work as described?
- What are the corner/failure cases?
- The paper may not be forthcoming about limitations

Future work?

Relations/comparisons

- How does the paper relate to other papers we have read?
- Can you imagine applying the ideas to a different problem?

Noise

Noisy image

Usually for dark conditions

After a slide by Frédo Durand

Noise

Fluctuation when taking multiple shots

After a slide by Frédo Durand

Canon 1D mark IIN at ISO 3200

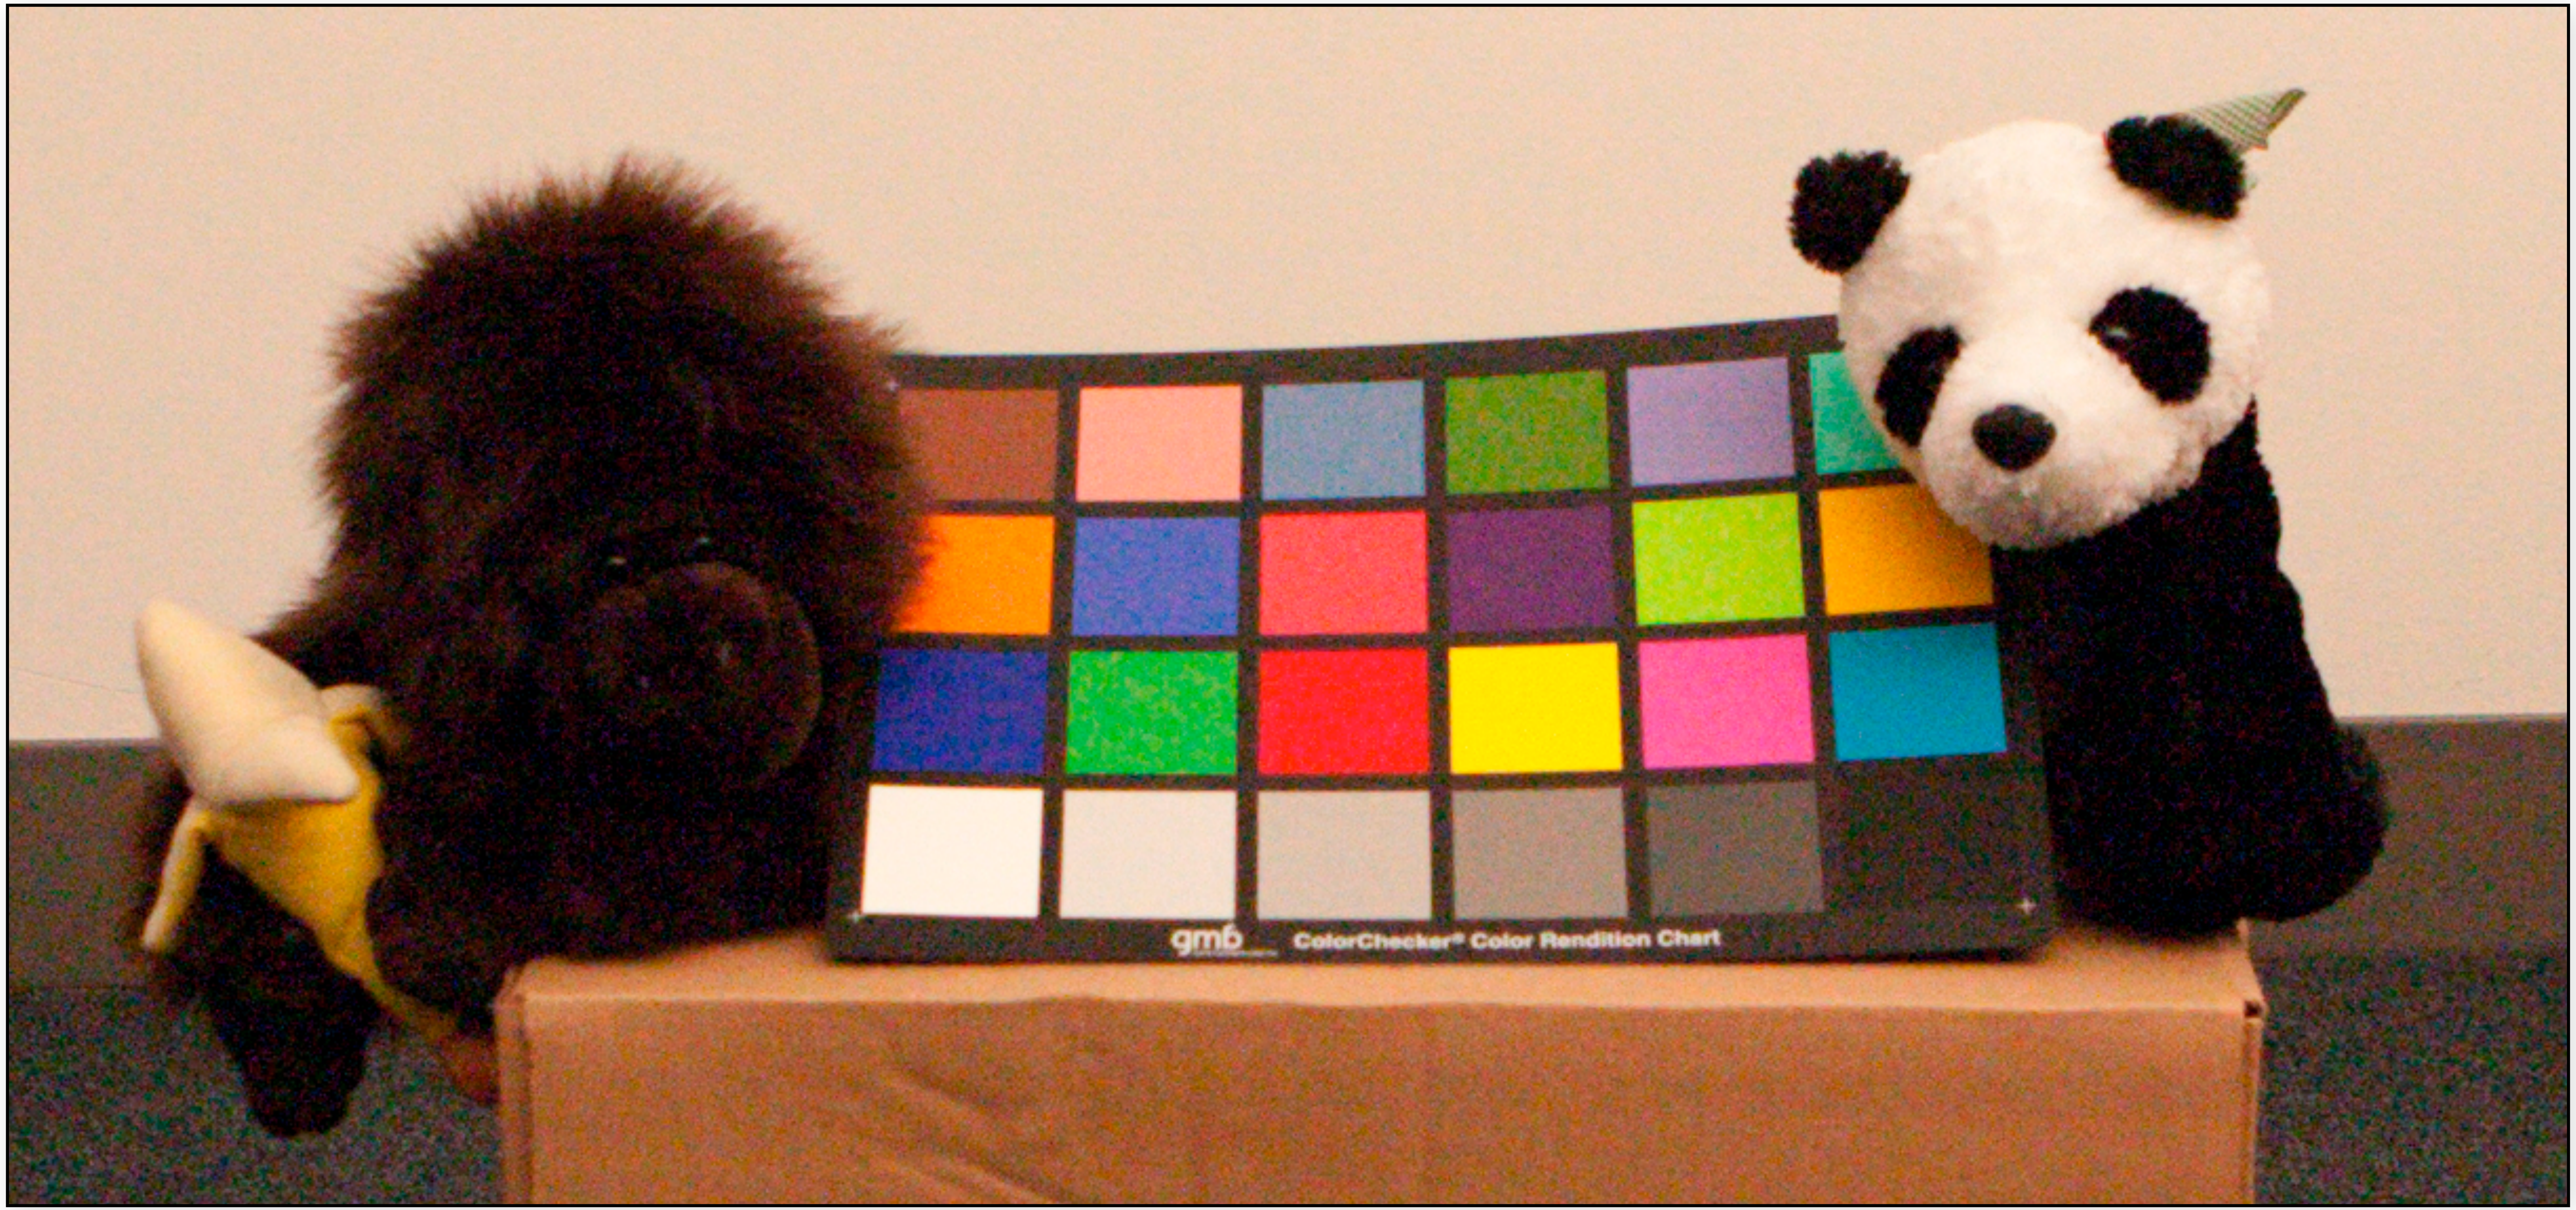

After a slide by Frédo Durand

Canon 1D mark IIN at ISO 3200

What should the histogram be within this box?

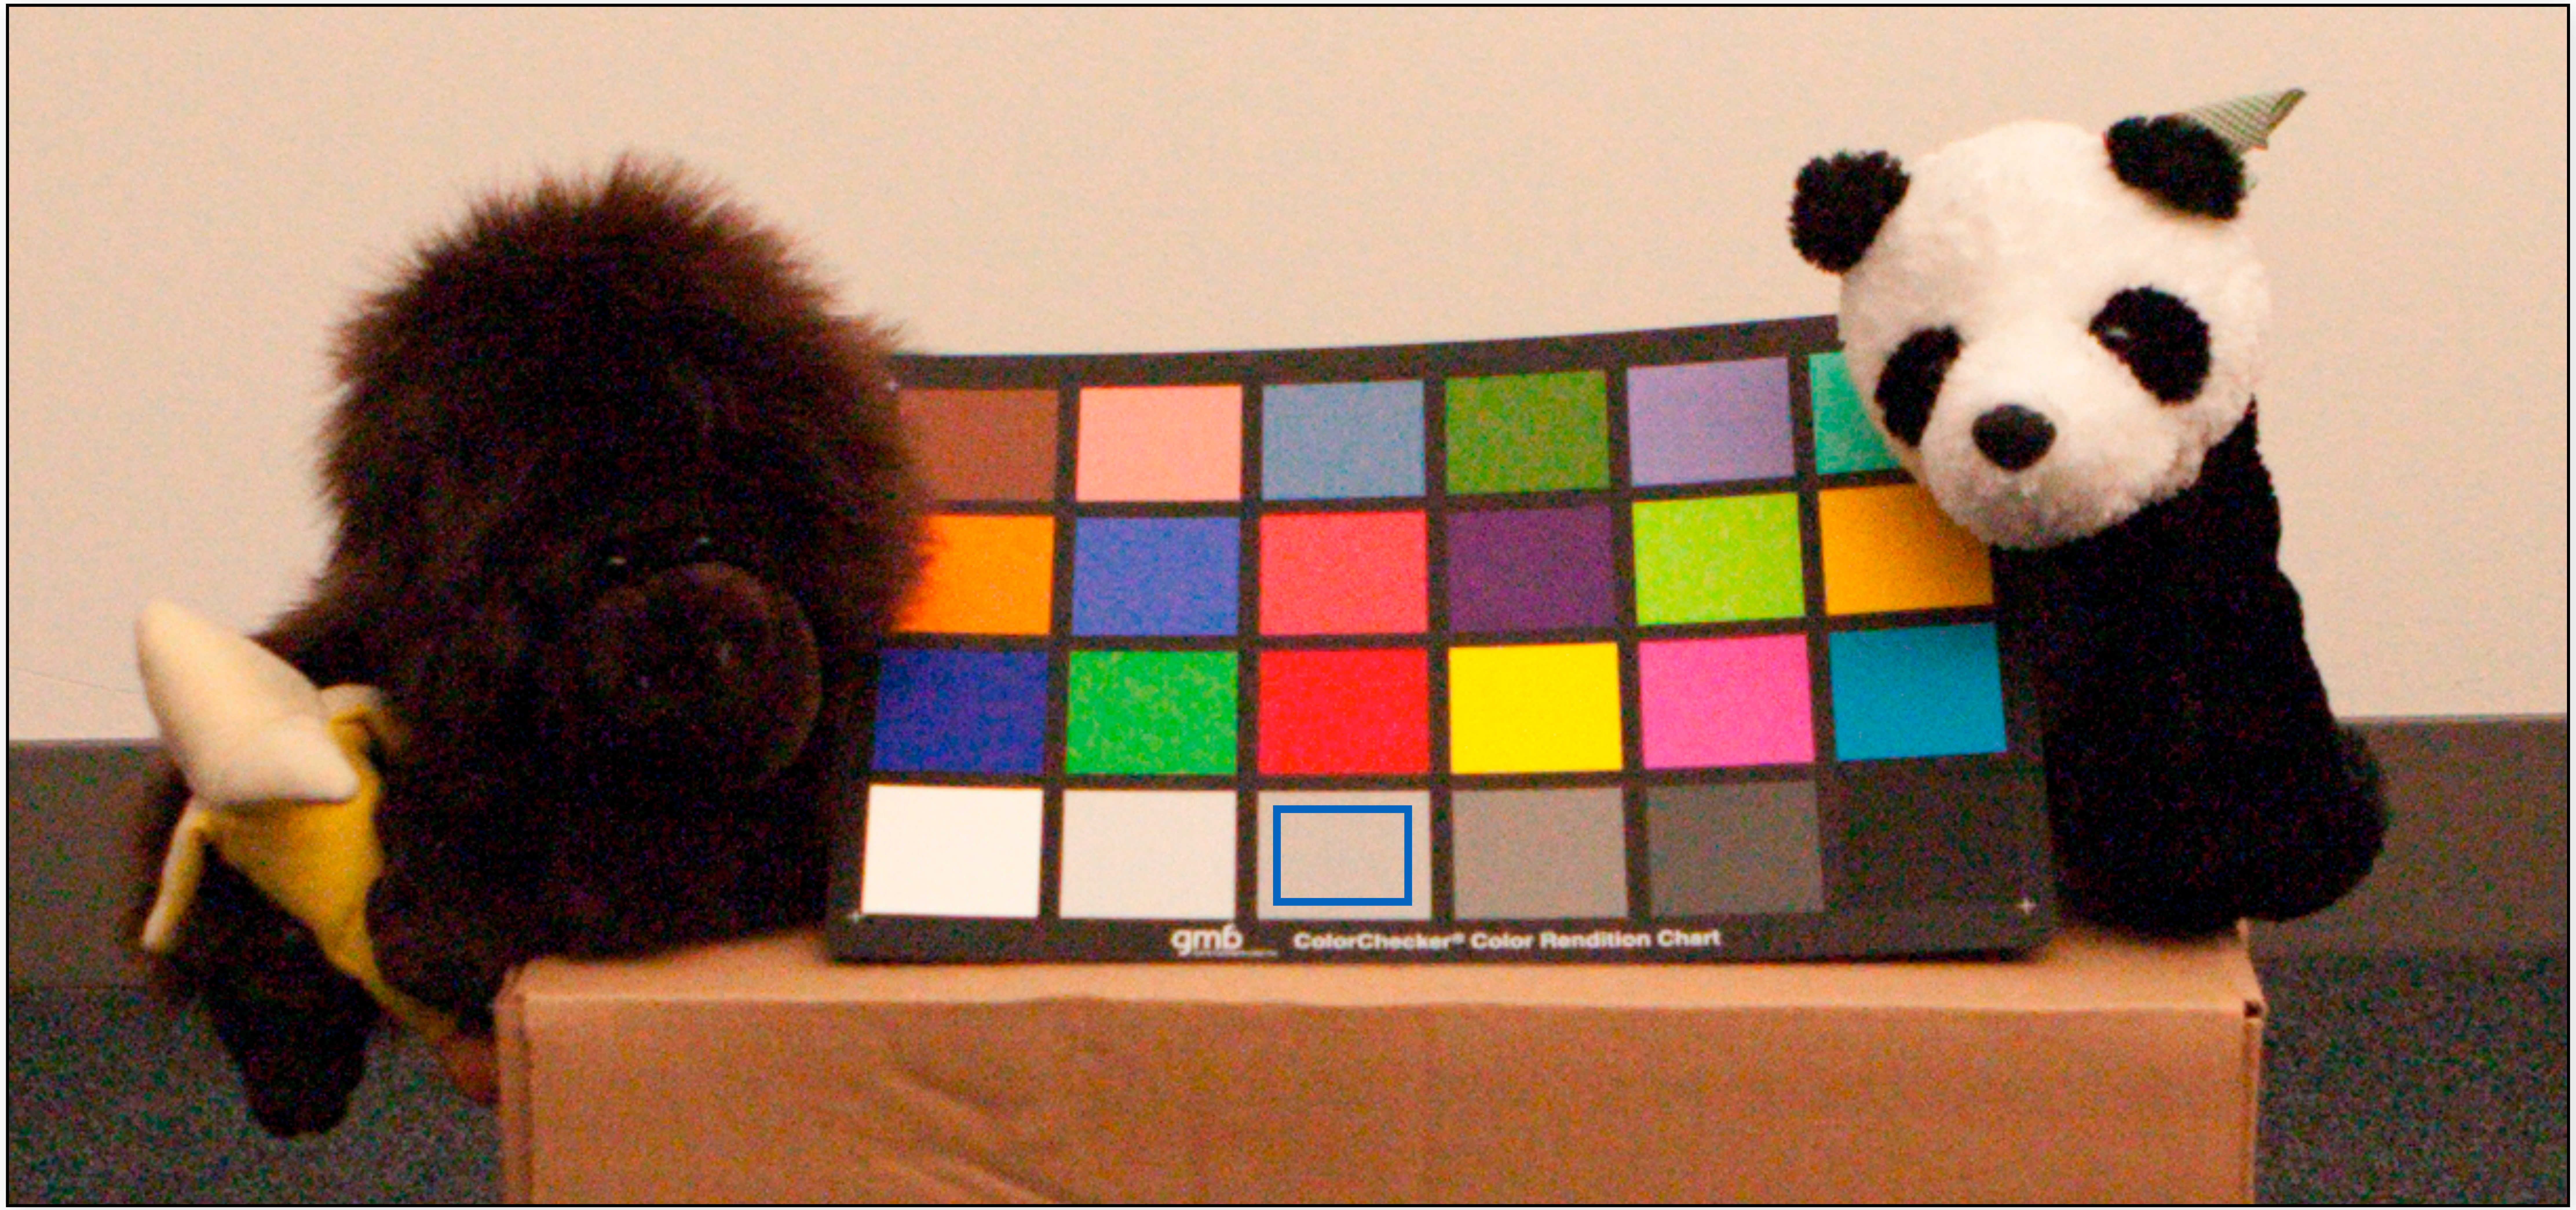

After a slide by Frédo Durand

Photoshop demo

Histogram of grey patch

Should be single values for RGB (constant color)

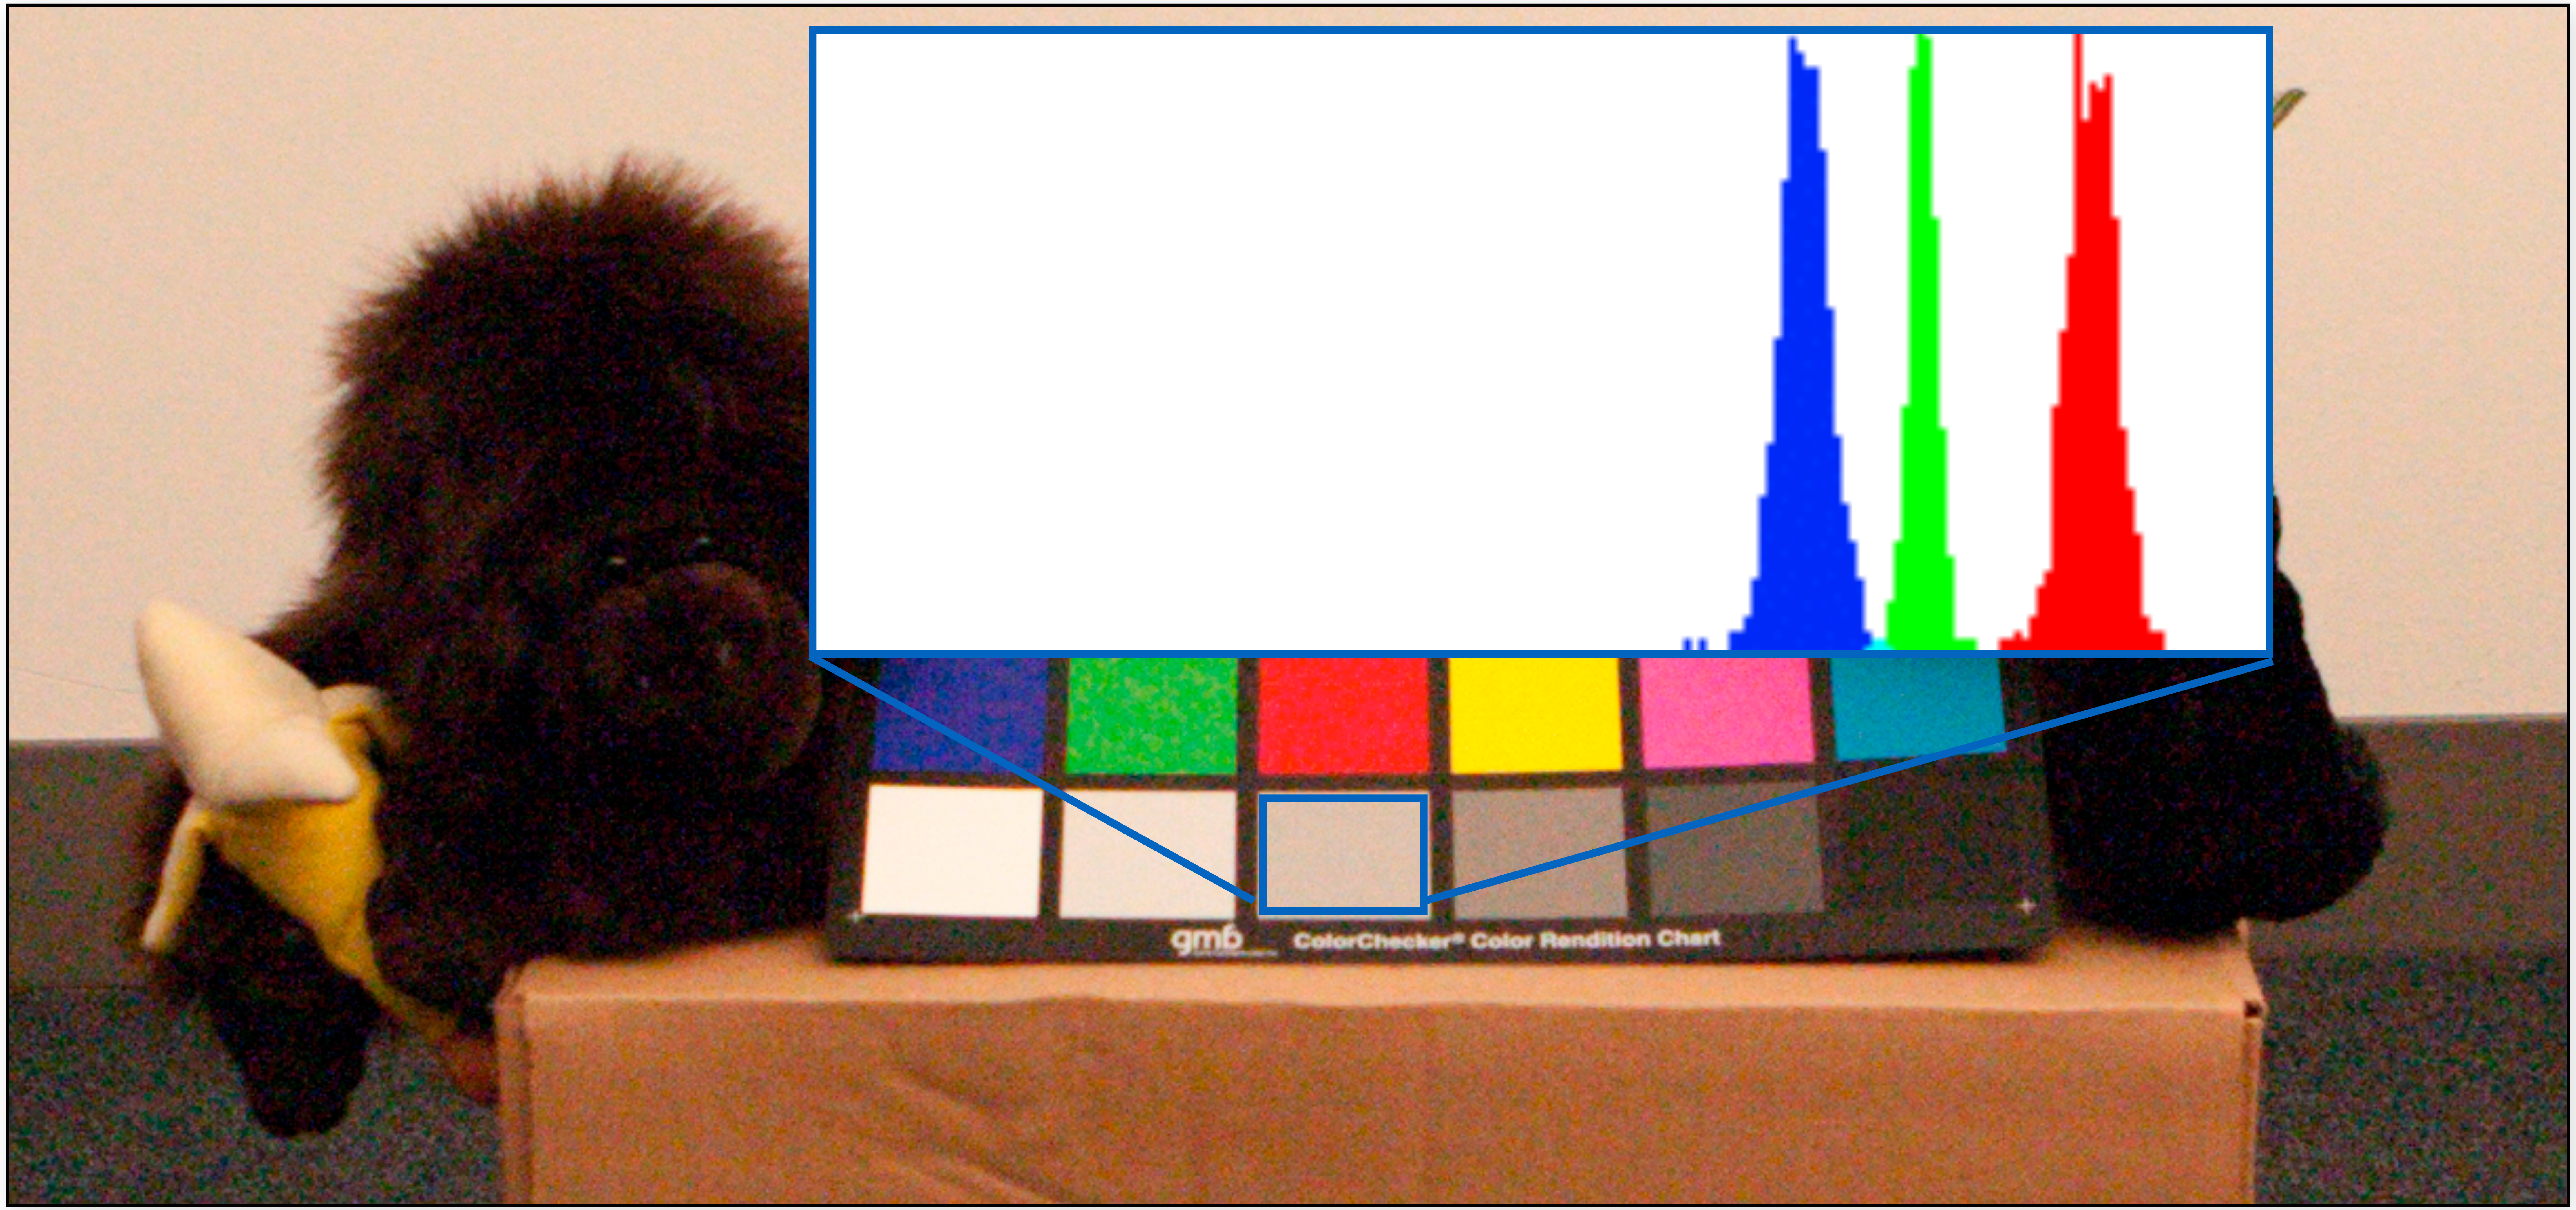

After a slide by Frédo Durand

Recap

Noise exists

Noise can be observed as:

- fluctuation over time
- fluctuation over space when should be constant

After a slide by Frédo Durand

Denoising by averaging

Averaging pseudo-code

```
mean = imSeq[1]
for i = 2 to imSeq.size():
   mean += imSeq[i]
mean /= imSeq.size()
```

After a slide by Frédo Durand

1 image

After a slide by Frédo Durand

3 images

After a slide by Frédo Durand

5 images

After a slide by Frédo Durand

Probabilistic perspective

Noise statistics / probability

Denote pixel values like random variables X

Mean: μ or E[x], the true measurement

Variance: $\sigma^2[x]$ = ~ average squared error

- more precisely: average squared difference to mean

$$\sigma^2[x] = E[(E[x] - x)^2]$$

$$\sigma^2 = E[x^2] - E[x]^2$$

Standard deviation: σ[x] = square root of variance

- In same unit as measurement

After a slide by Frédo Durand

Estimating the sample mean

Say we have N measurements x_i

How would you estimate their mean?

$$\mu_N = \frac{1}{N} \sum x_i$$

Estimating the sample variance

Say we have N measurements x_i

How would you estimate their variance?

Use original definition:

$$\sigma^2[x] = E[(E[x] - x)^2]$$

$$\mu_N = \frac{1}{N}\sum x_i \qquad \sigma_N^2 = \frac{1}{N}\sum(\mu_N - x_i)^2$$

- this underestimates variance!

Estimating the sample variance

$$\sigma_N^2 = \frac{1}{N}\sum(\mu_N - x_i)^2 \qquad \sigma_N^2 = \frac{1}{N-1}\sum(\mu_N - x_i)^2$$

Divide by N-1, not by N

- Otherwise, variance would be underestimated on average
- called Bessel correction: removes bias
- Intuition: we use the same samples for estimating the mean and variance, which introduces correlation that underestimates variance

After a slide by Frédo Durand

Example with coin flip, N=2

We do 2 coin flips

Try to estimate mean & variance

Sometimes we'll be wrong

- e.g. if we get 0 twice, we'll think variance is zero
- but we'd like to be right on average (called unbiased)

After a slide by Frédo Durand

Example with coin flip, N=2

4 scenarios: (0,0) ; (0, 1) ; (1, 0) ; (1, 1)

- mean estimates: 0 ; 0.5 ; 0.5 ; 1
- average of the mean estimations: 0.5, equal to true mean (unbiased)

true variance: 0.25

sum of squared differences to sample mean: 0; 0.5; 0.5; 0

estimator $\sigma_N^2 = \frac{1}{N} \sum (\mu_N - x_i)^2$ 0 ; 0.25 ; 0.25 ; 0

- 0.125 on average, biased

estimator $\sigma_N^2 = \frac{1}{N-1} \sum (\mu_N - x_i)^2$ 0 ; 0.5 ; 0.5 ; 0

- 0.25 on average, unbiased

After a slide by Frédo Durand

Signal-to-noise ratio (SNR)

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

$$\log SNR(dB) = 10 \log_{10}\left(\frac{\mu^2}{\sigma^2}\right) = 20 \log_{10}\left(\frac{\mu}{\sigma}\right)$$

After a slide by Marc Levoy

SNR in practice

Be careful. Sometimes variance is zero (for no good reason) and will break things

- practical hack: take the max of σ^2 and a small number, e.g. 1e-6

After a slide by Frédo Durand

Basic probability tools

Goal

Analyze how the mean & variance evolve when we denoise by averaging multiple frames

Formula for average: $\frac{1}{N}\sum x_i$

- addition
- multiply by scalar

After a slide by Frédo Durand

Expected value

$E[kx] =$

$E[x+y] =$

$E[xy] =$

After a slide by Frédo Durand

Expected value

$E[kx] = kE[x]$

$E[x+y] =$

$E[xy] =$

After a slide by Frédo Durand

Expected value

$E[kx] = kE[x]$

$E[x+y] = E[x]+E[y]$

$E[xy] =$

After a slide by Frédo Durand

Expected value

$E[kx] = kE[x]$

$E[x+y] = E[x]+E[y]$

$E[xy] = E[x]E[y]$?

After a slide by Frédo Durand

Expected value

E[kx] = kE[x]

E[x+y] = E[x]+E[y]

E[xy] = E[x]E[y]

- only when they are uncorrelated!

$$
\begin{aligned}
E[xy] &= \int_y \int_x xy\, p(x,y) dx dy \\
&= \int_y \int_x xy\, p(x)p(y) dx dy \\
&= \int_y p(y) y \int_x x\, p(x) dx dy \\
&= \int_y p(y) y E[x] dy \\
&= E[y]E[x]
\end{aligned}
$$

After a slide by Frédo Durand

Variance identity

$$\sigma^2[x] = E[(E[x] - x)^2]$$
$$= E[E[x]^2 - 2xE[x] + x^2]$$
$$= E[x]^2 - 2E[x]E[x] + E[x^2]$$
$$= -E[x]^2 + E[x^2]$$

$$\sigma^2 = E[x^2] - E[x]^2$$

After a slide by Frédo Durand

Variance properties

Multiplication by k:

$$\sigma^2[kx] = E[(kx)^2] - E[kx]^2 = k^2\sigma^2[x]$$

not linear, quadratic!

Addition of two random variables

$$\sigma^2[x+y] = E[(x+y)^2] - E[x+y]^2$$

$$= E[x^2 + 2xy + y^2] - (E[x] + E[y])^2$$

$$= E[x^2] + E[y^2] + \cancel{E[2xy]} - E[x]^2 - E[y]^2 - \cancel{2E[x]E[y]}$$

uncorrelated:
E[xy]=E[x]E[y]

$$= E[x^2] - E[x]^2 + E[y^2] - E[y]^2$$

$$= \sigma^2[x] + \sigma^2[y]$$

variance is additive!

After a slide by Frédo Durand

Take home message

Noise/measurement as random variable

Mean, variance, standard deviation

Variance:

- multiplication by k => k^2
- addition => addition

SNR, log of SNR

After a slide by Frédo Durand

Convergence

Convergence

Assume images are IID random measurements

Variance for one image: $\sigma^2[x_i]$

Average: $\frac{1}{N}\sum_{i=1}^{N} x_i$

What is the variance of the average?

After a slide by Frédo Durand

Convergence

Assume images are IID random measurements

Variance for one image: $\sigma^2[x_i]$

Average: $\frac{1}{N}\sum_{i=1}^{N} x_i$

$$\sigma^2\left[\frac{1}{N}\sum_{i=1}^{N} x_i\right] = \left(\frac{1}{N}\right)^2 \sum_{i=1}^{N} \sigma^2[x_i]$$

$$= \left(\frac{1}{N}\right)^2 N\sigma^2[x_i]$$

$$= \frac{1}{N}\sigma^2[x_i]$$

After a slide by Frédo Durand

IMPORTANT RESULT

Denoising by averaging:

- variance is reduced as 1/N
- standard deviation (error) is reduced by sqrt(N)

After a slide by Frédo Durand

Alignment

Brute force

Assignment 3!

Try all possible shifts within +/- maxOffset

Keep the one with minimum sum of square differences

After a slide by Frédo Durand

Casio EXF1, Google glass

Can do denoising by aligning and averaging N images

After a slide by Frédo Durand

Noise characteristics

Analyzing noise

Camera on tripod, many pictures

Compute mean, variance, stddev, SNR

After a slide by Frédo Durand

Exposure

Get the right amount of light to sensor/film

Two main parameters:

- Shutter speed
- Aperture (area of lens)

\+ sensor sensitivity (ISO)

In what follows, I kept the exposure the same and explored the tradeoff between shutter speed and ISO

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

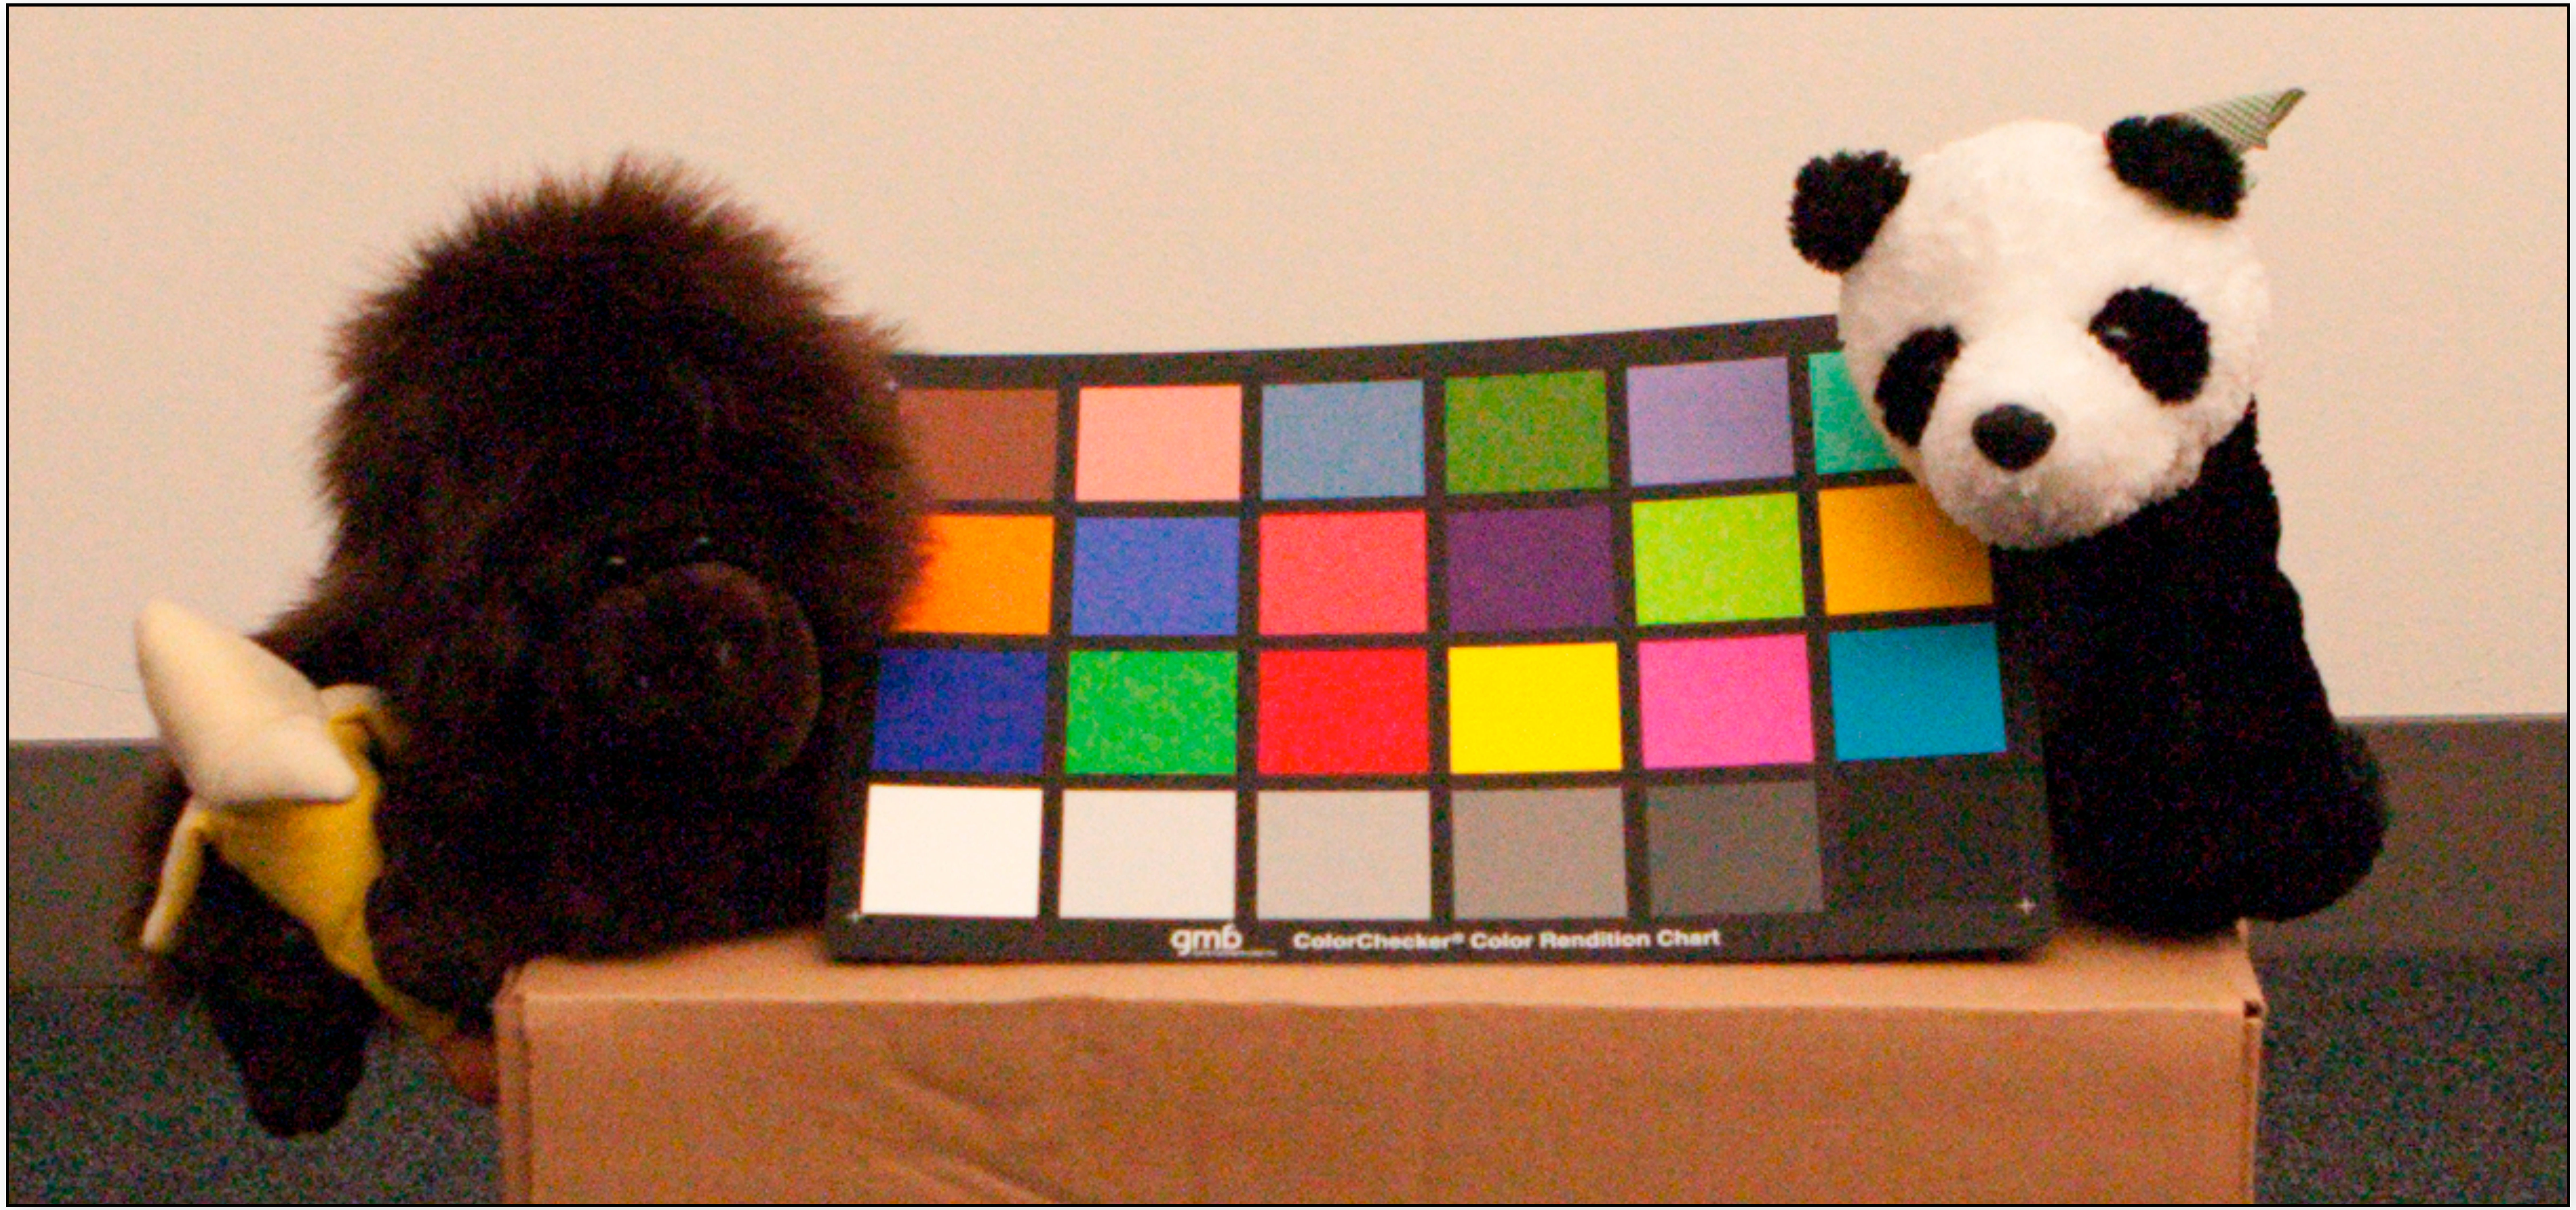

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

Looks noisy, especially in dark areas

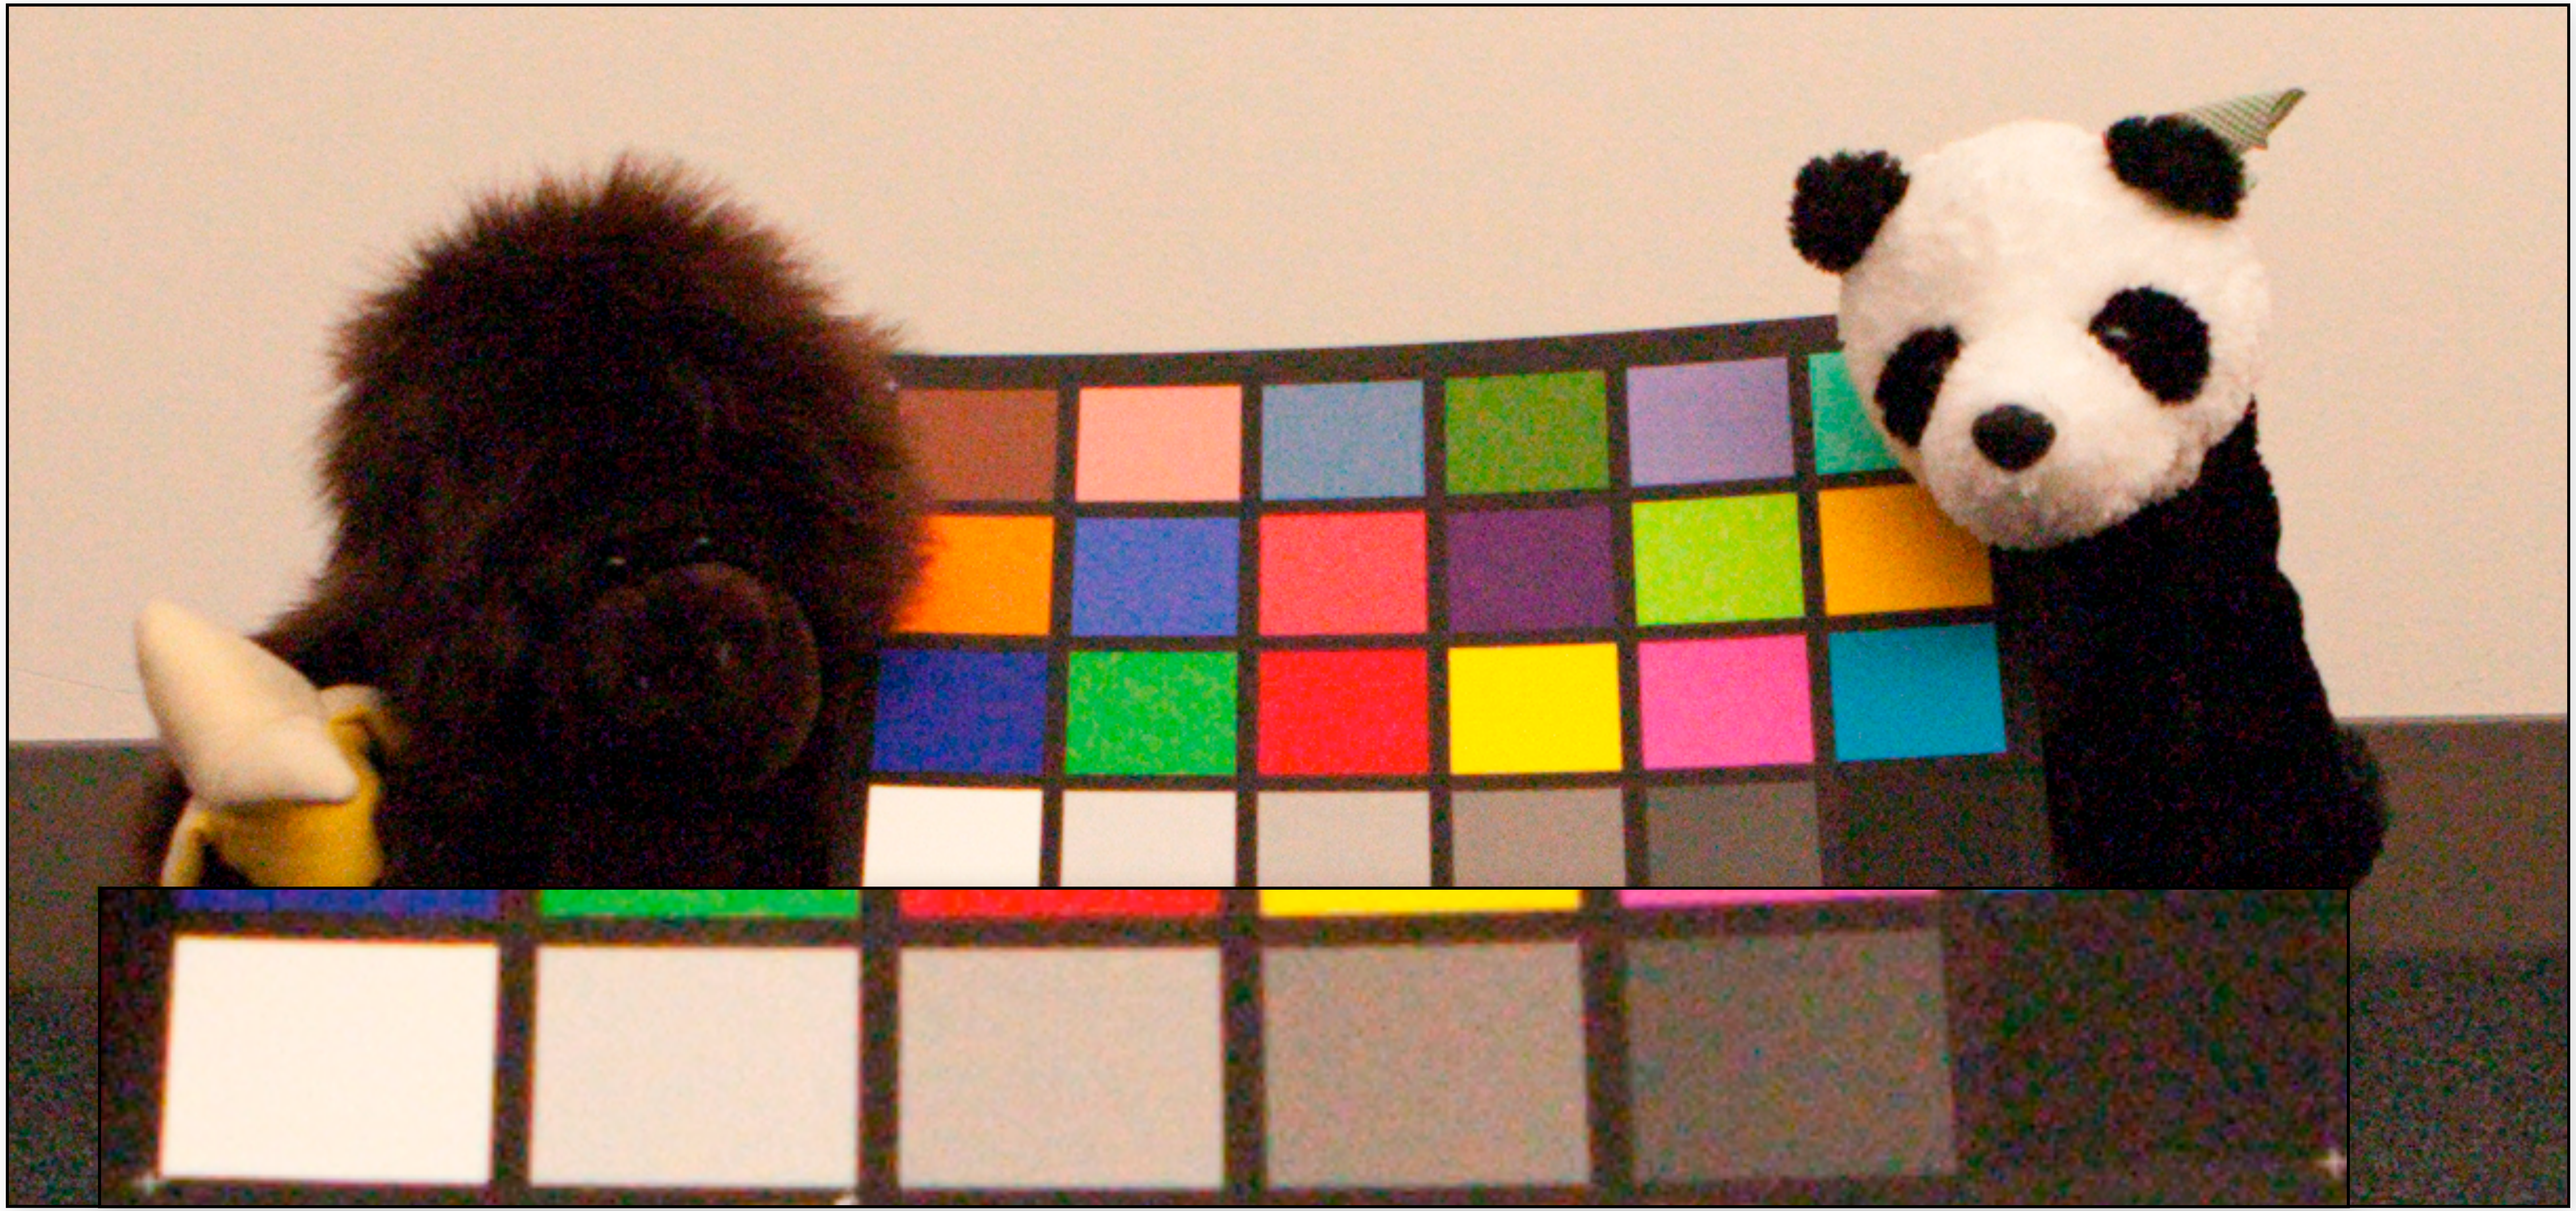

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

Denoised with 45 images (estimator of mean)

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

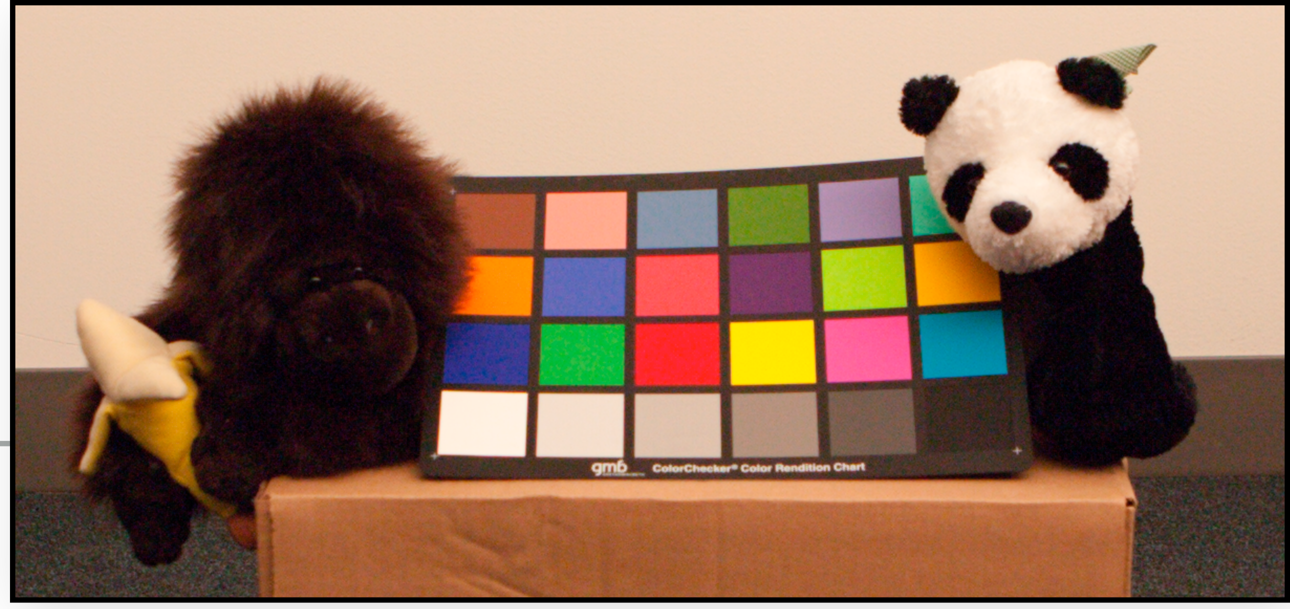

Standard deviation (some alignment issues...)

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

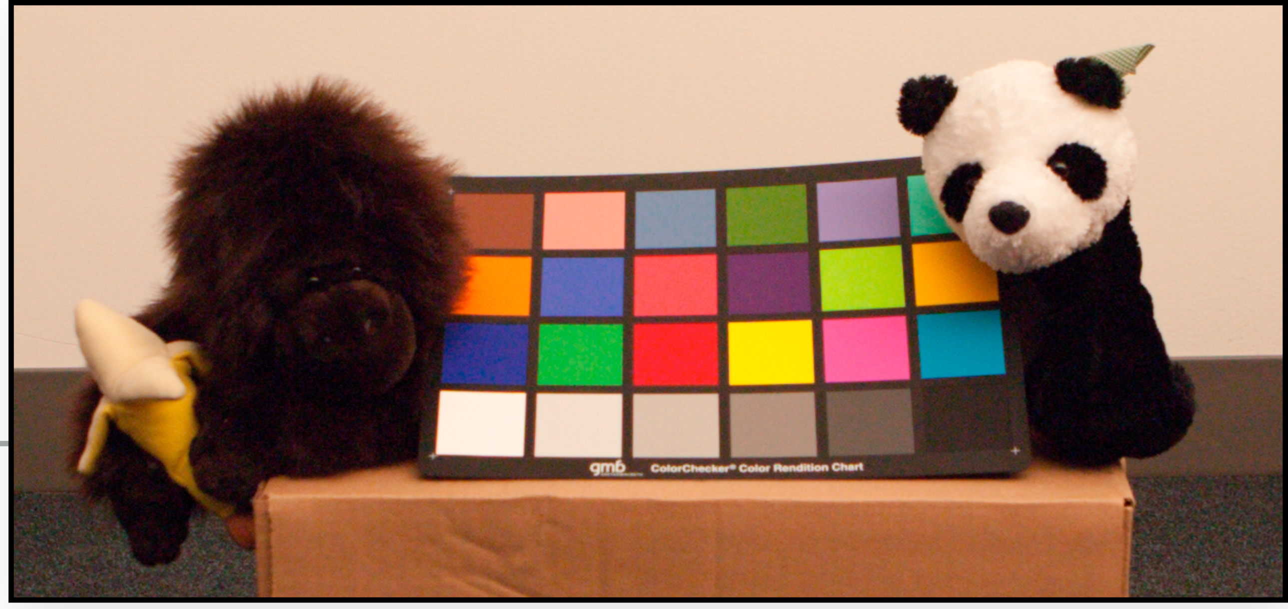

Standard deviation (some alignment issues...)

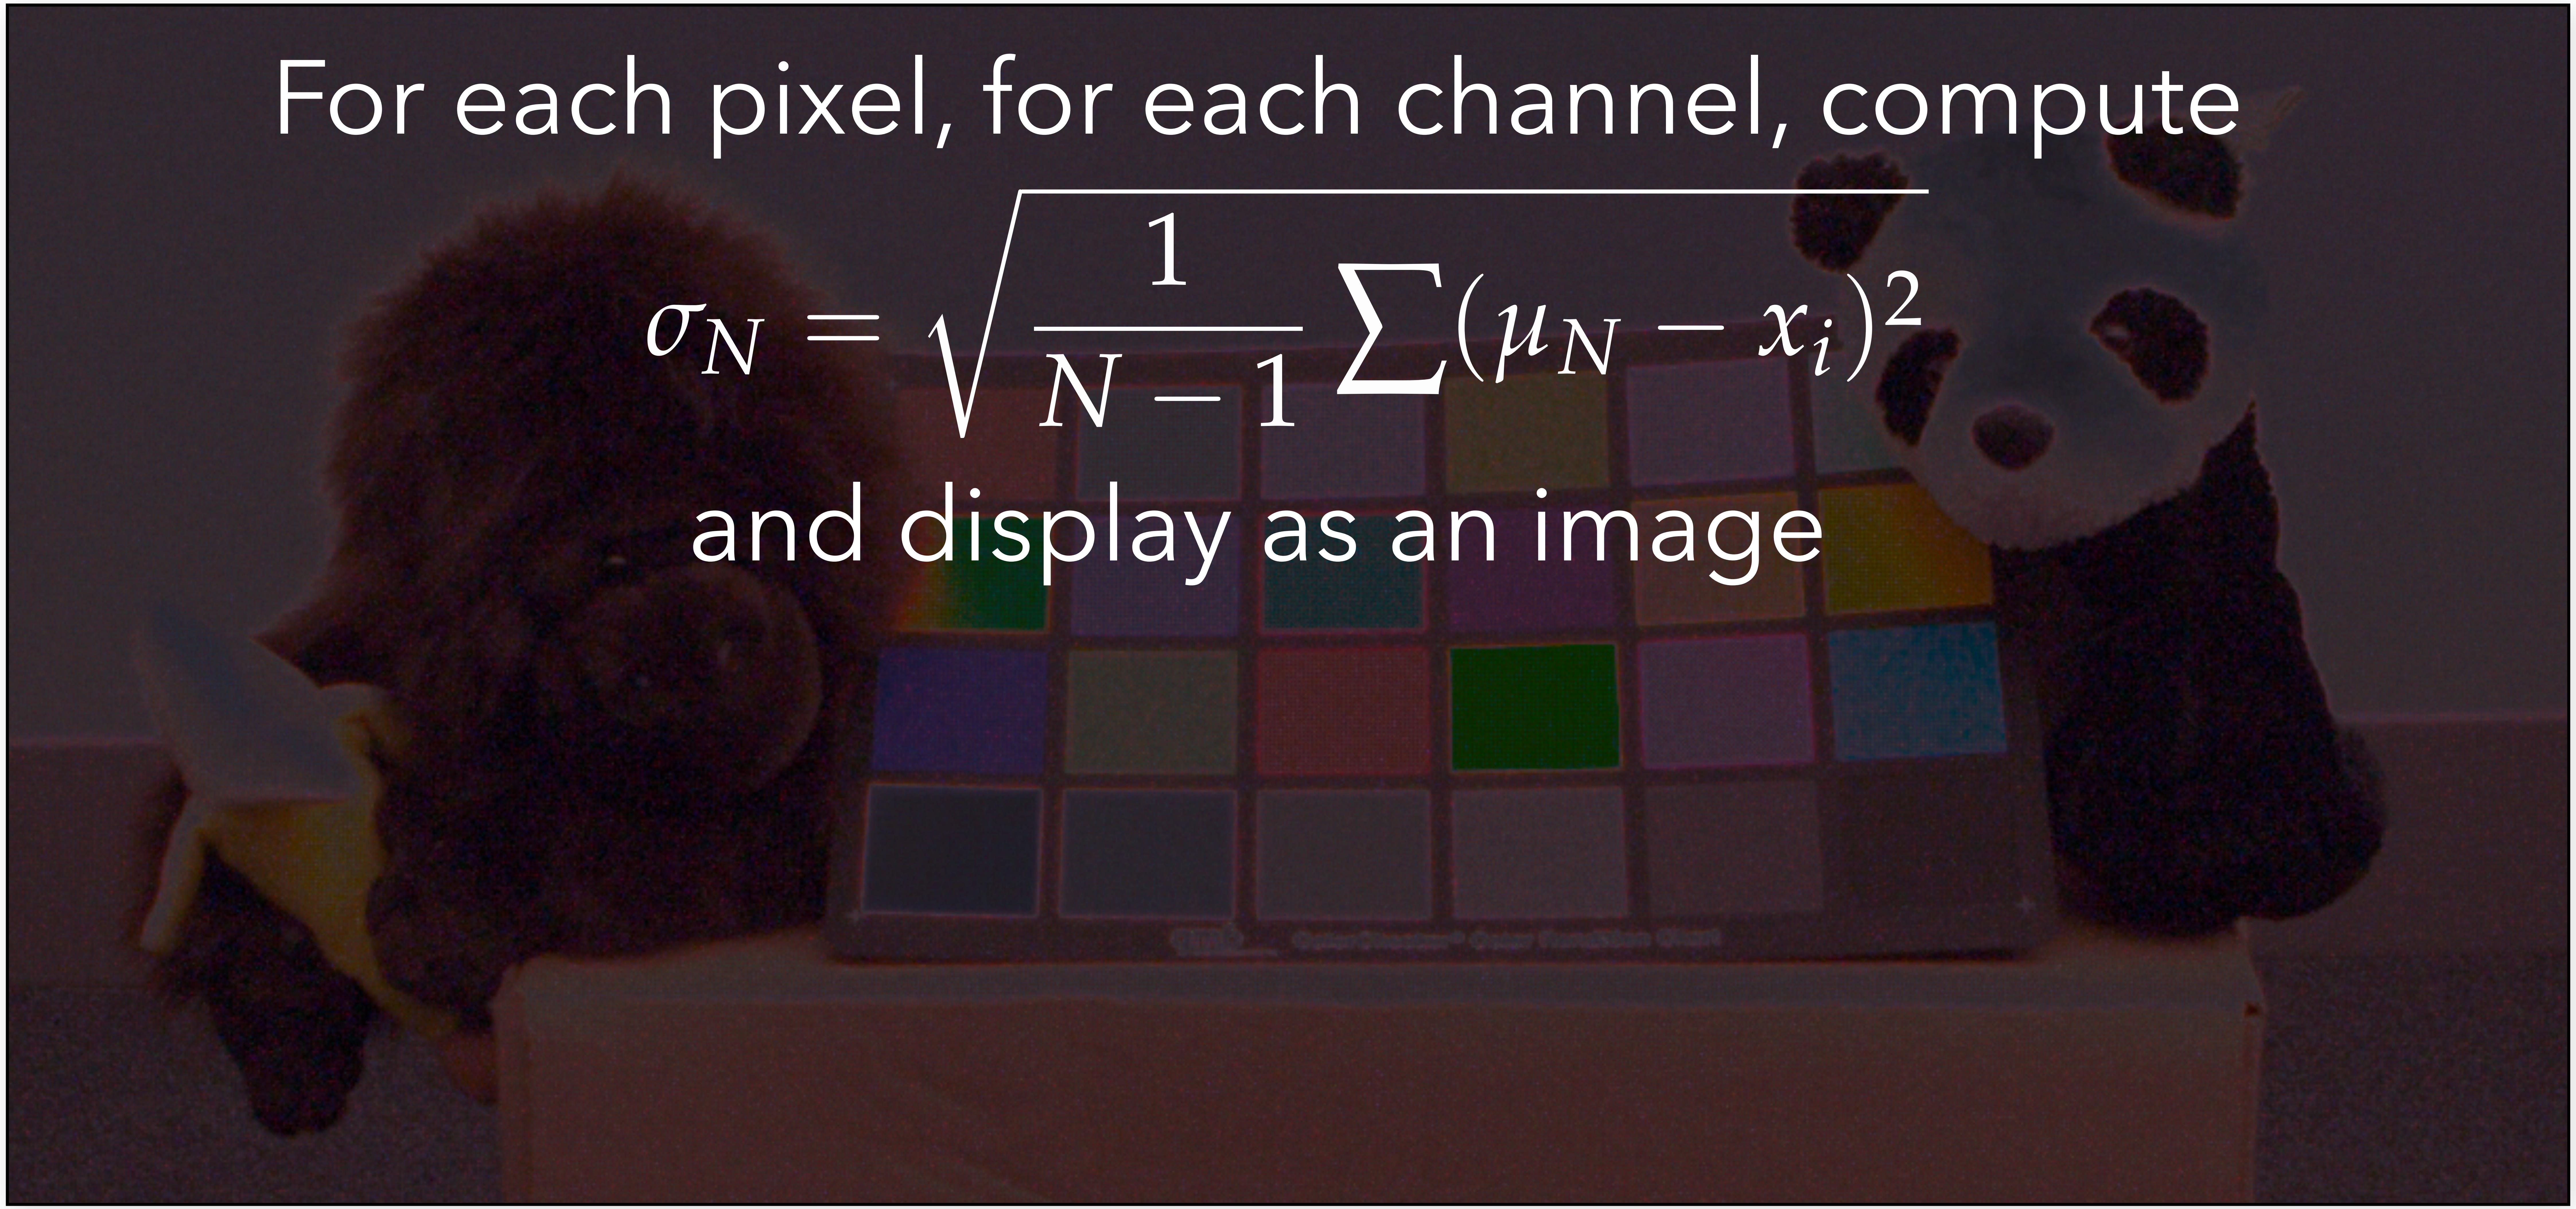

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

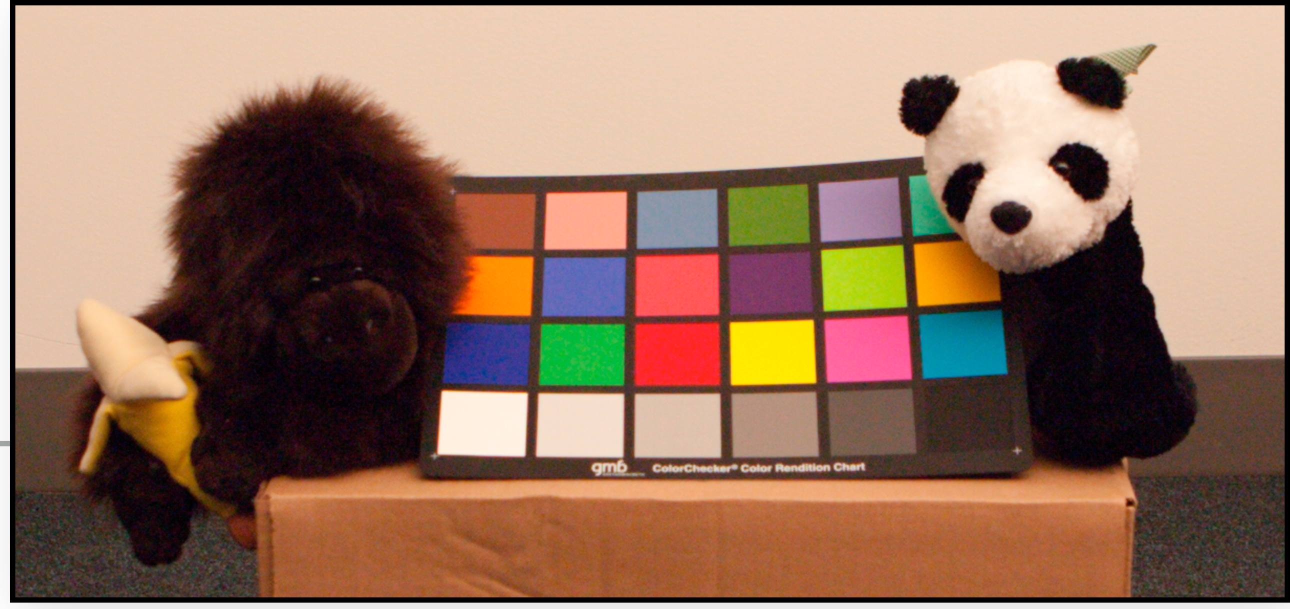

Standard deviation (some alignment issues...)

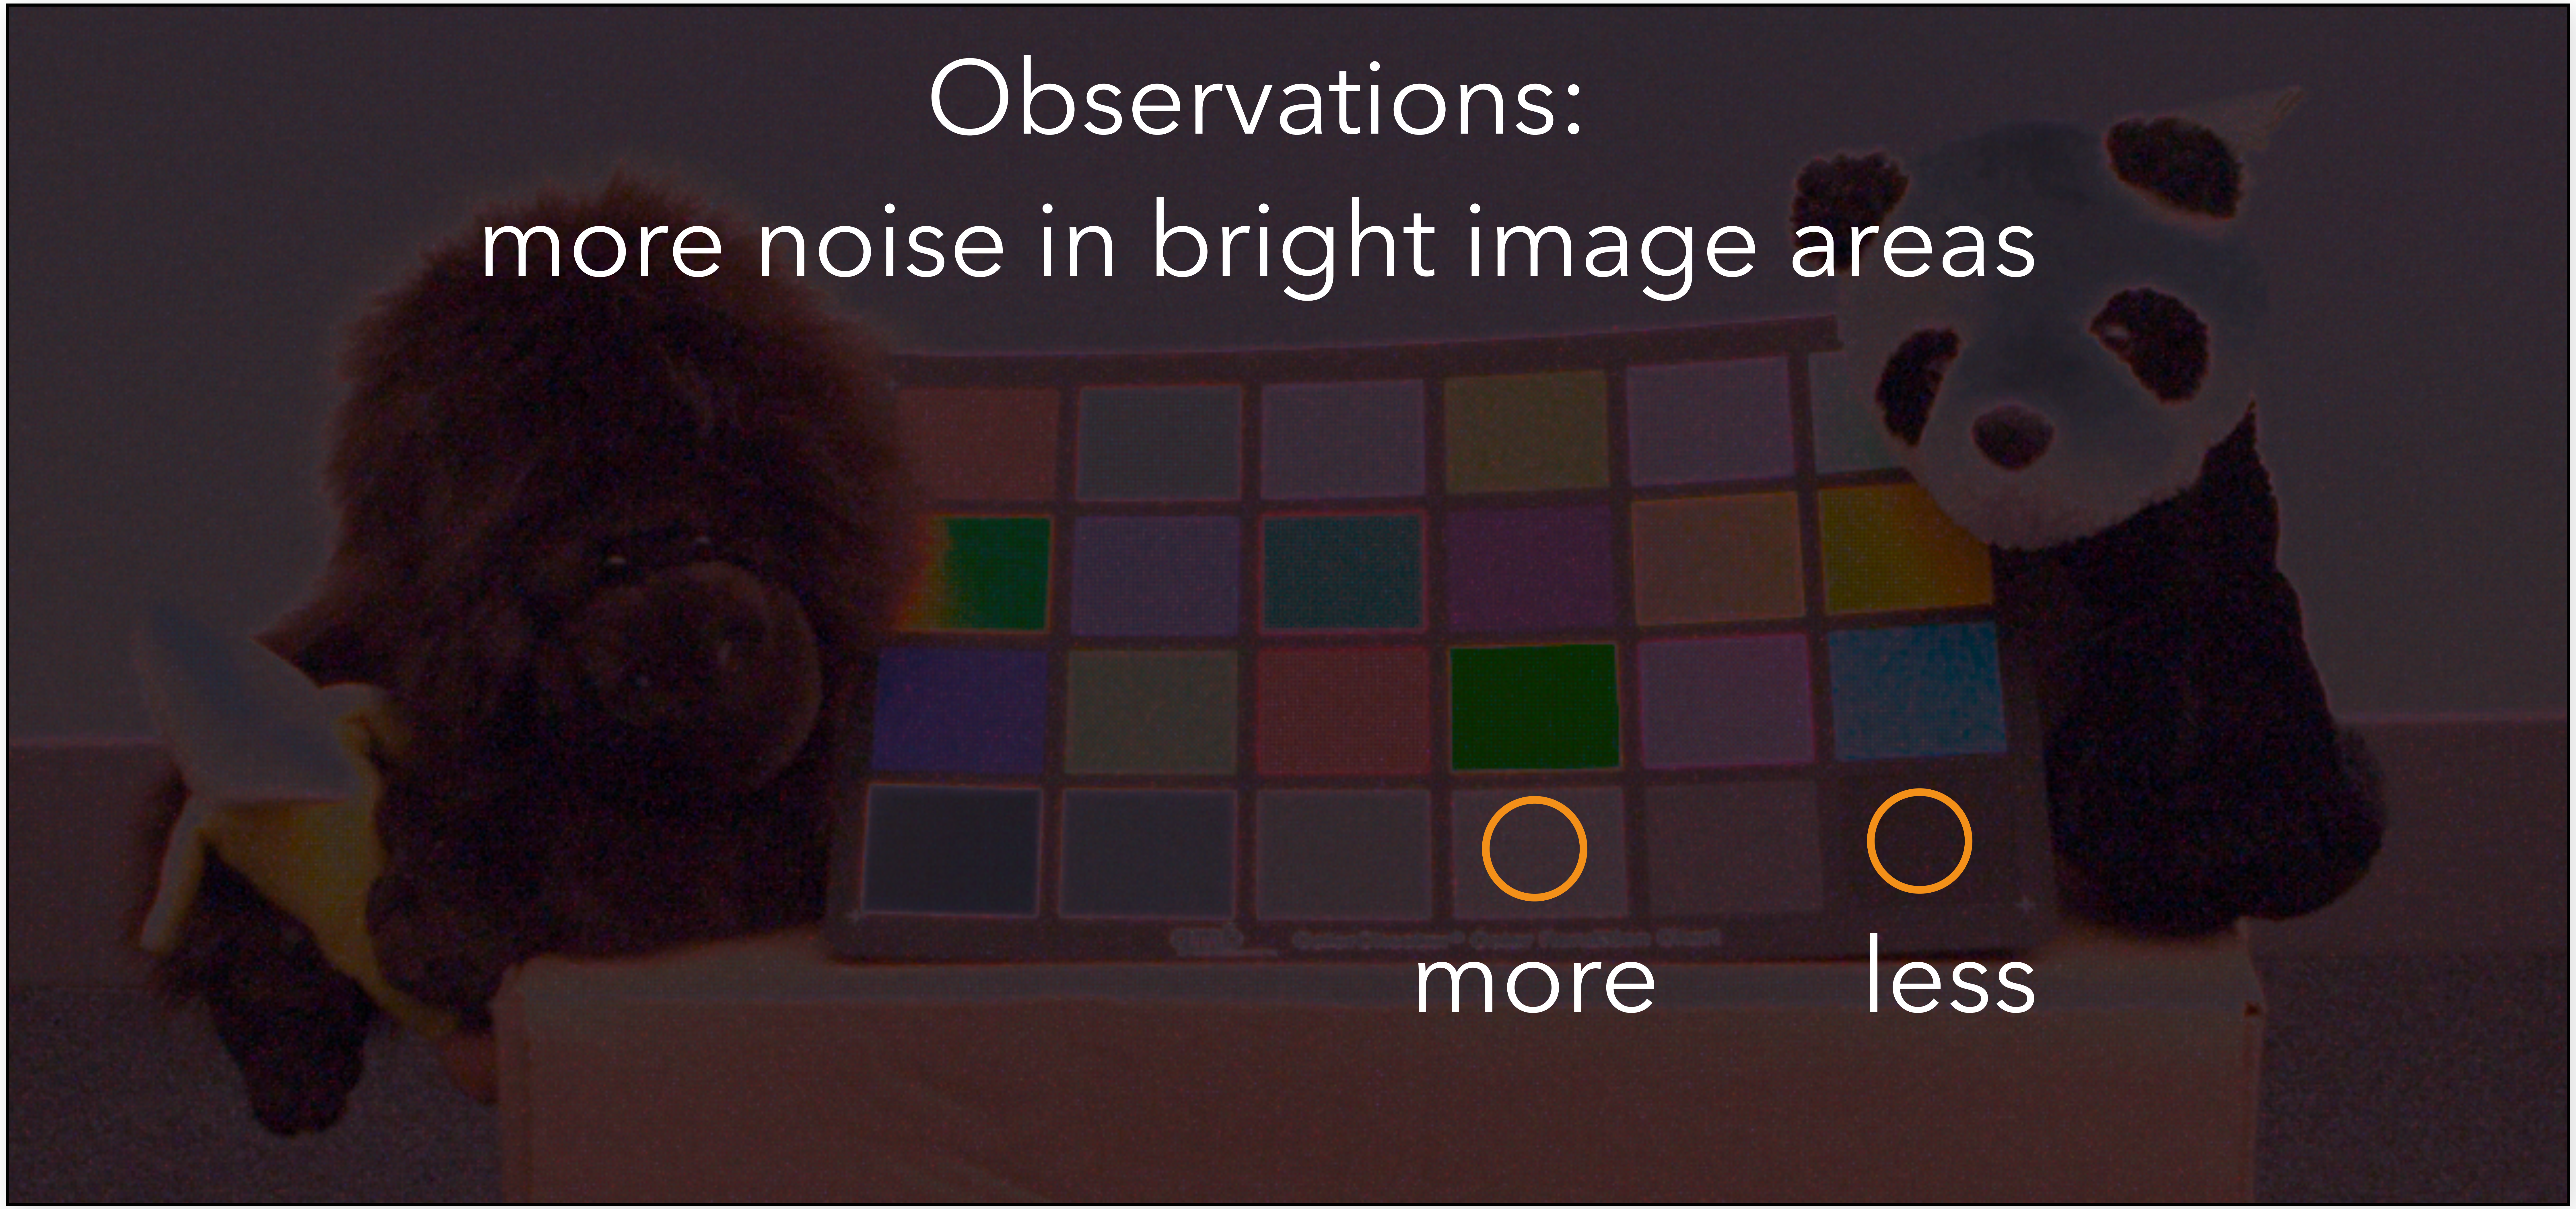

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

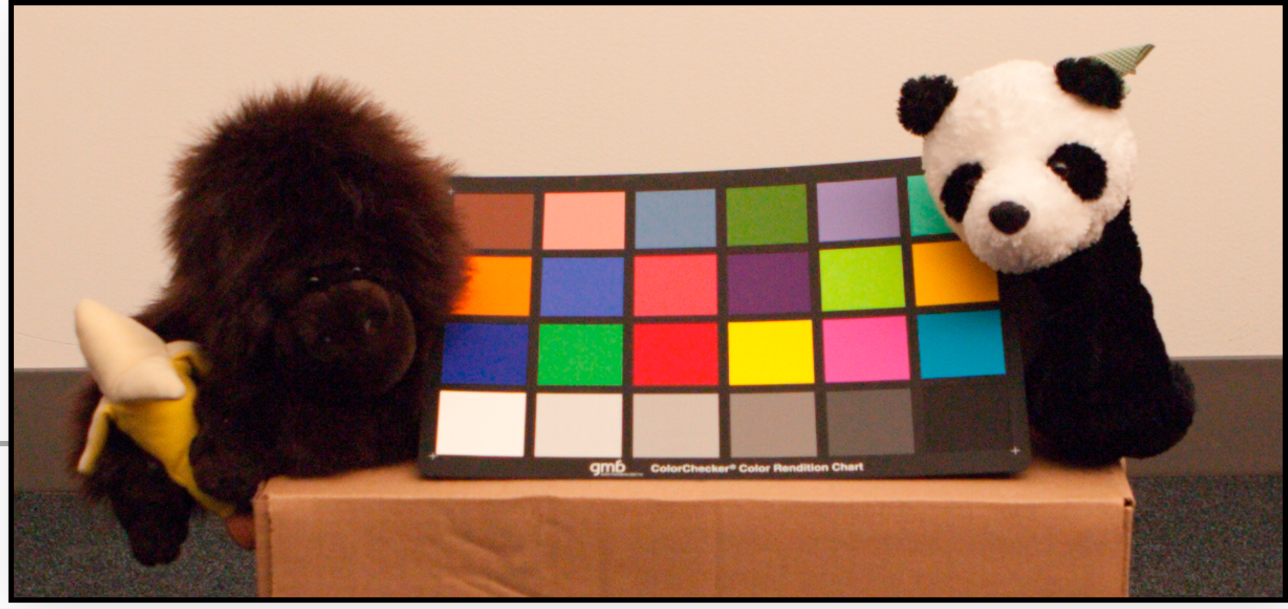

log SNR

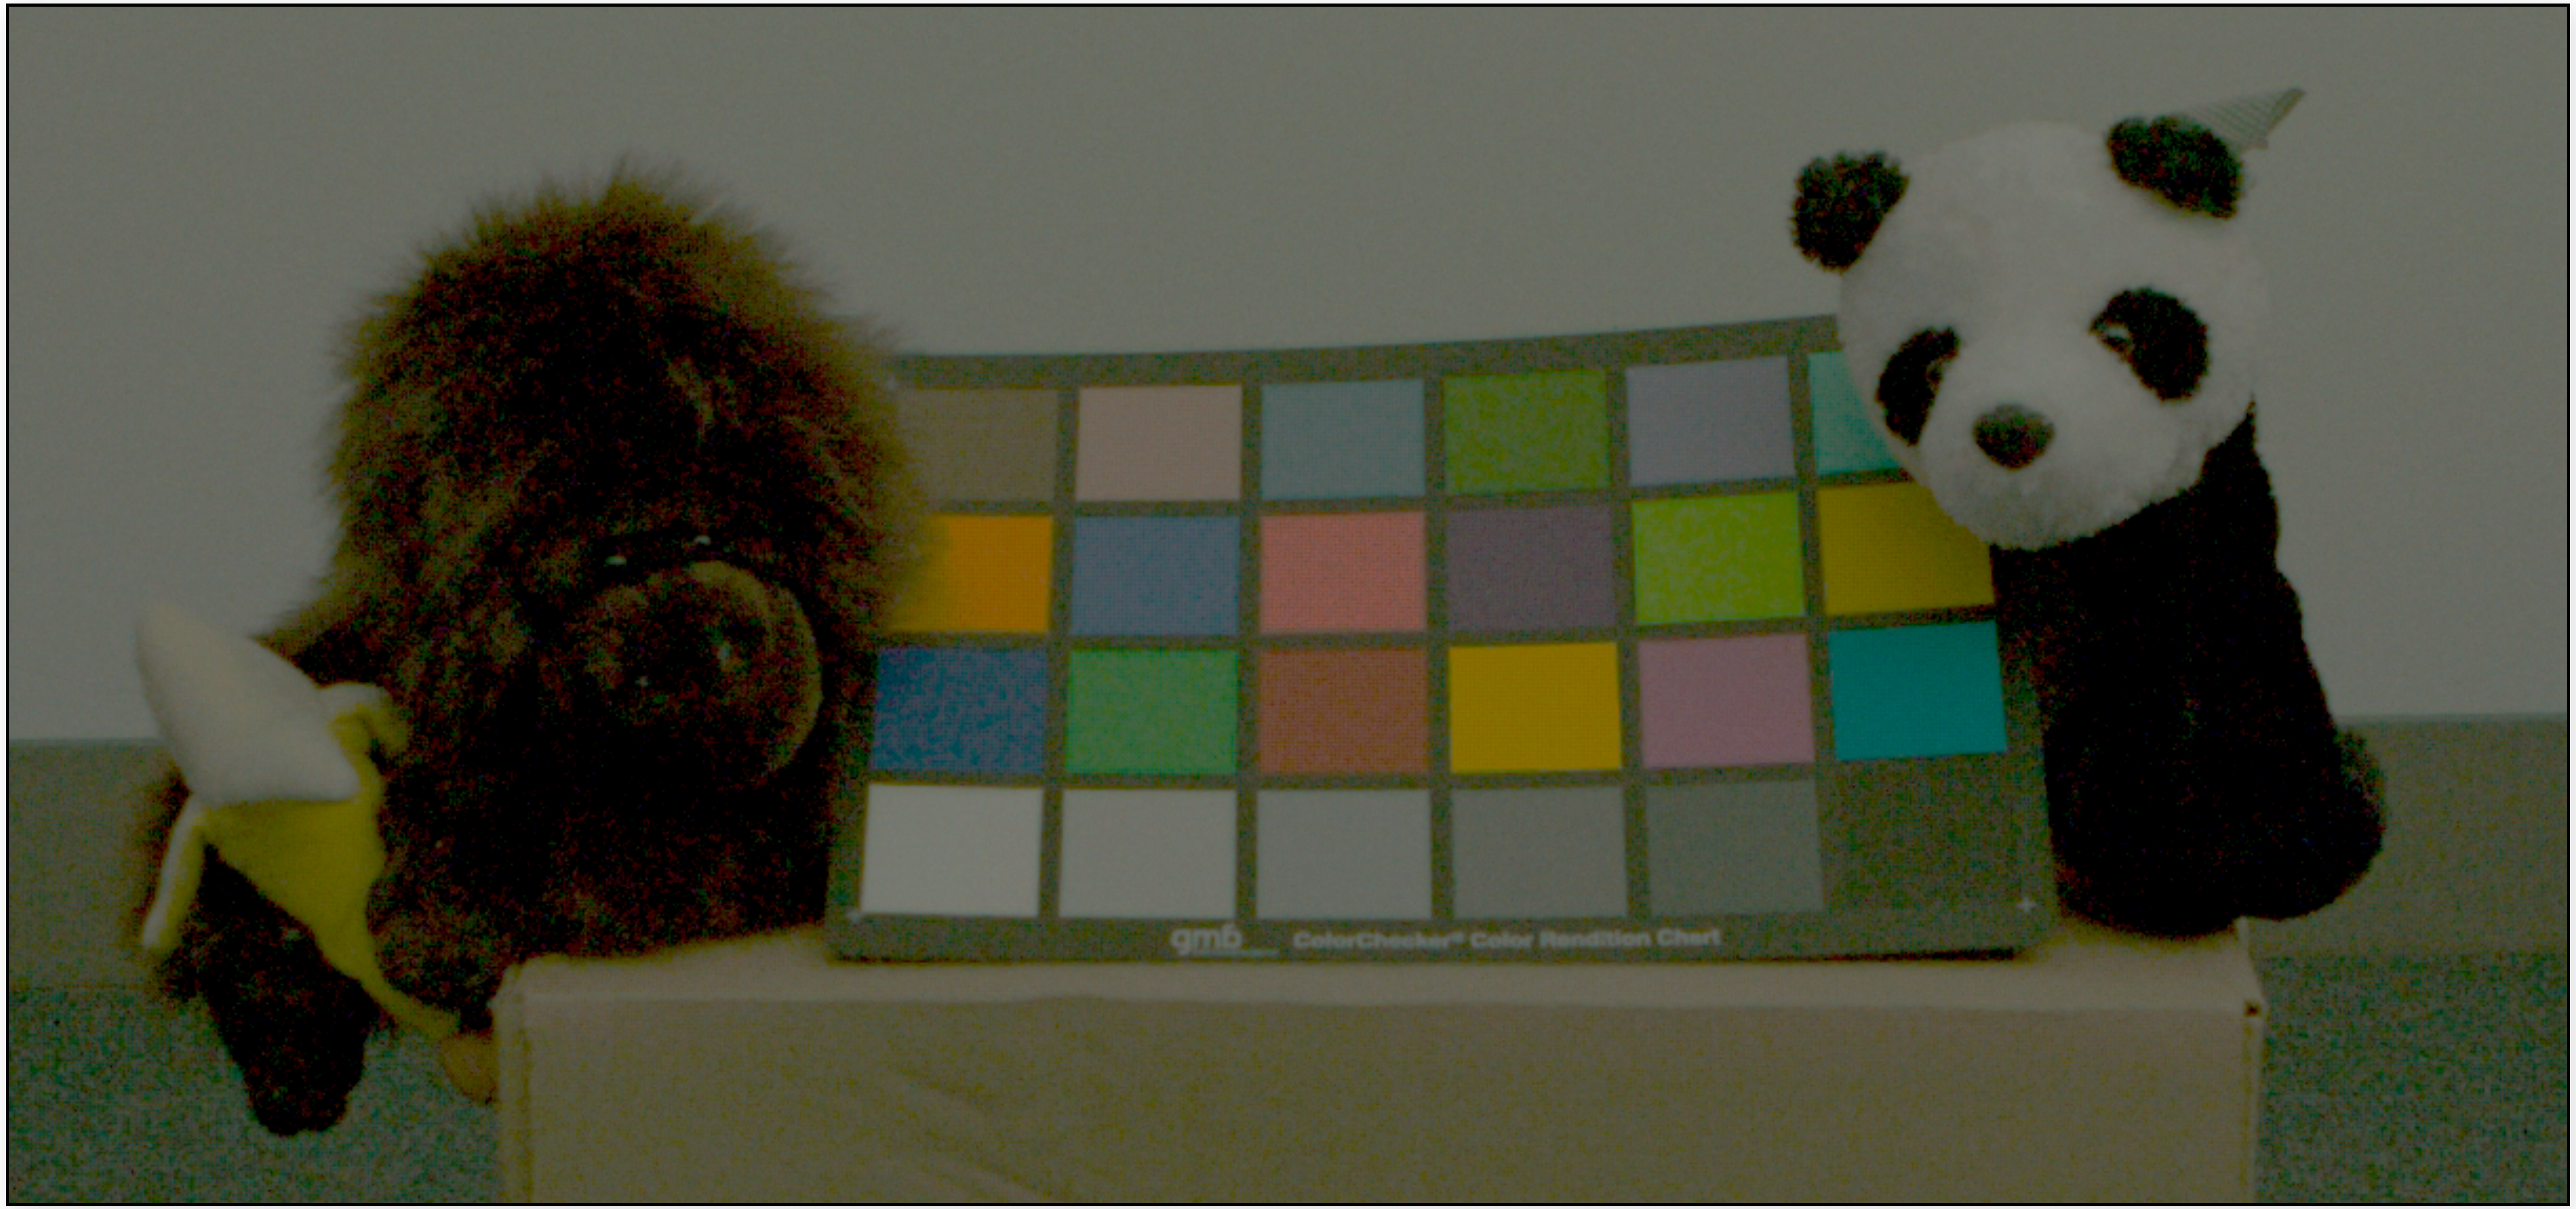

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

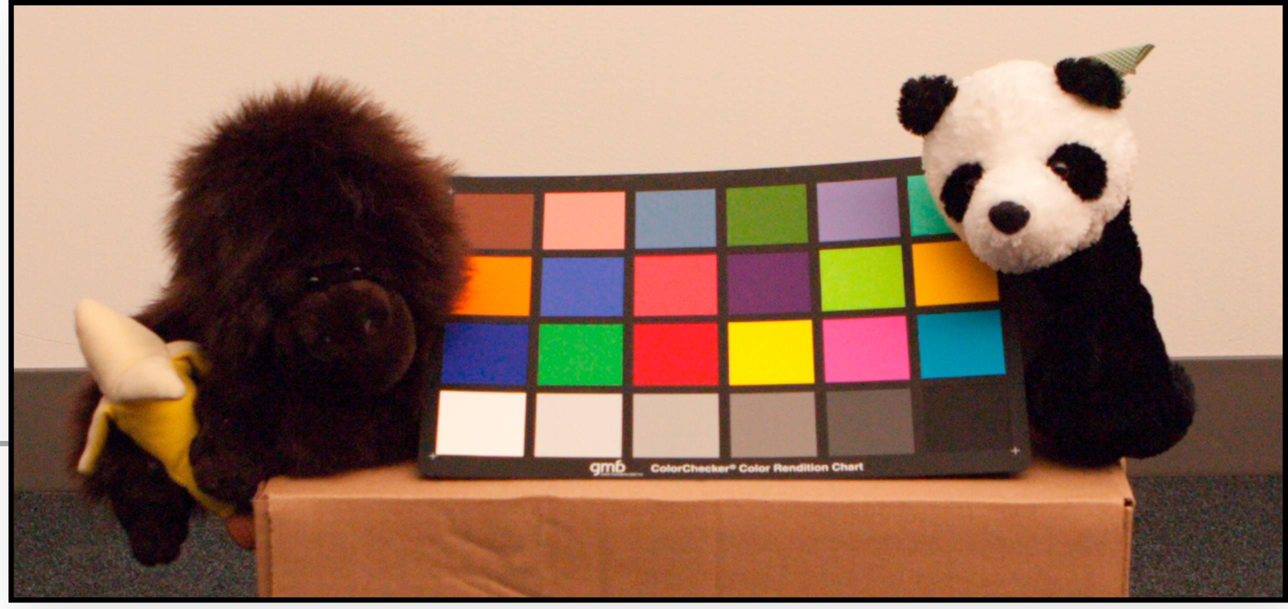

log SNR – looks a lot like the image!

After a slide by Frédo Durand

Observations

Noise is more visible in dark areas

Noise is numerically higher in bright areas

SNR is better in bright areas

Canon 1D IIN at ISO 3200

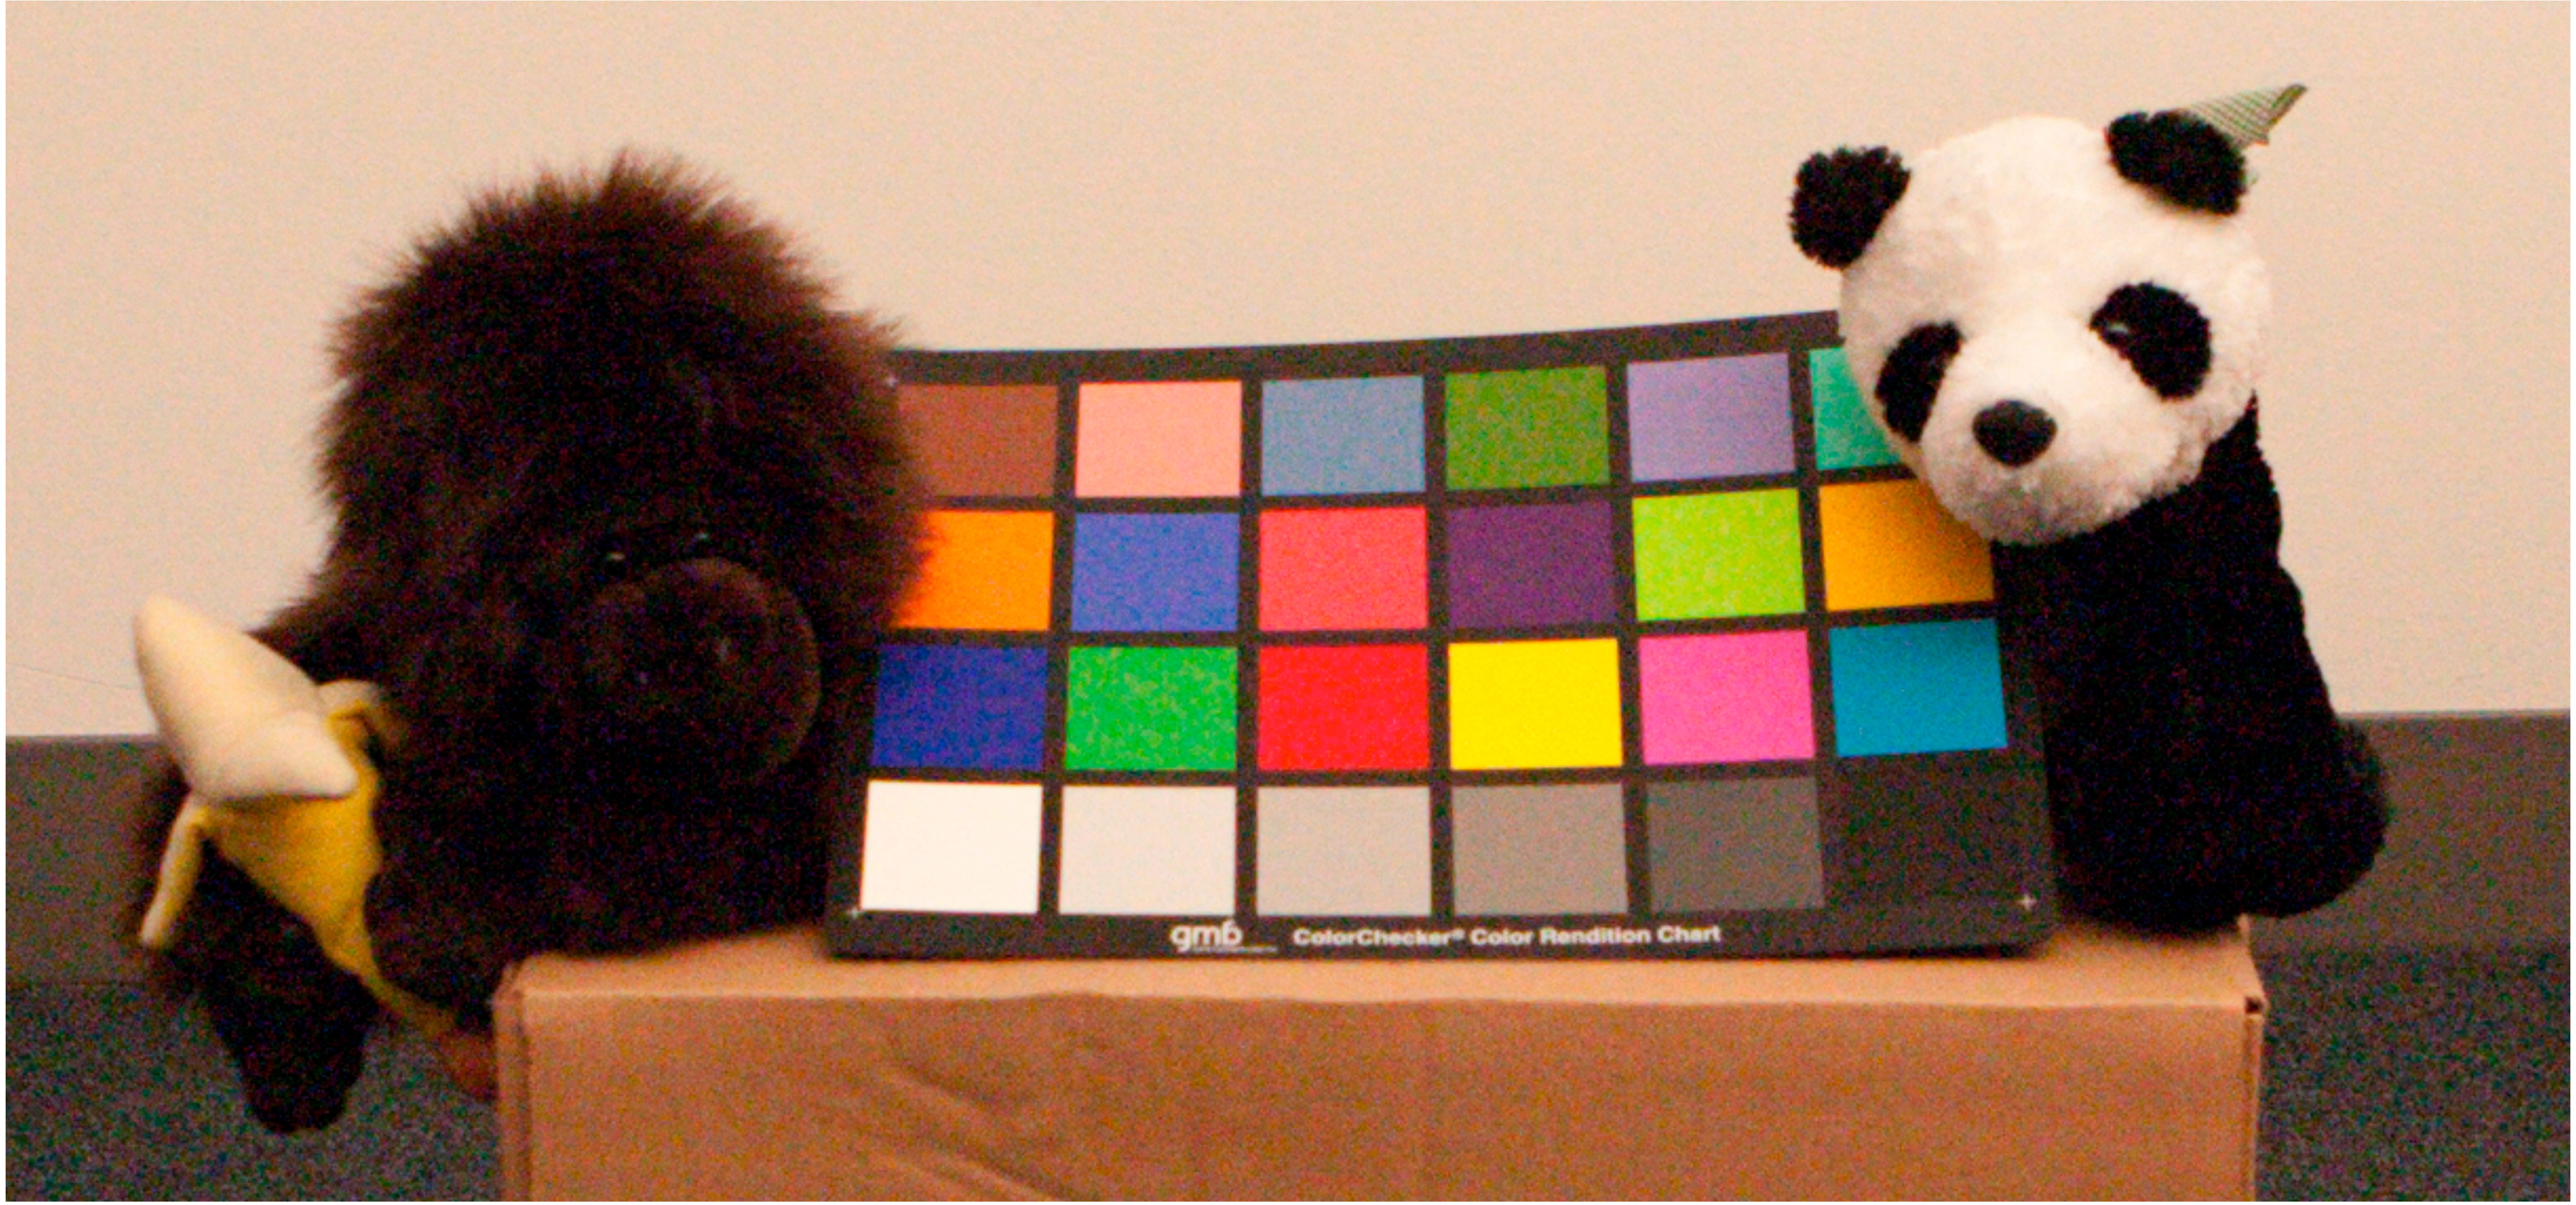

After a slide by Frédo Durand

Canon 1D II, ISO 100

A lot less noisy!

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

Standard deviation (some alignment issues…)

After a slide by Frédo Durand

Canon 1D II at ISO 100

Standard deviation (some alignment issues…)

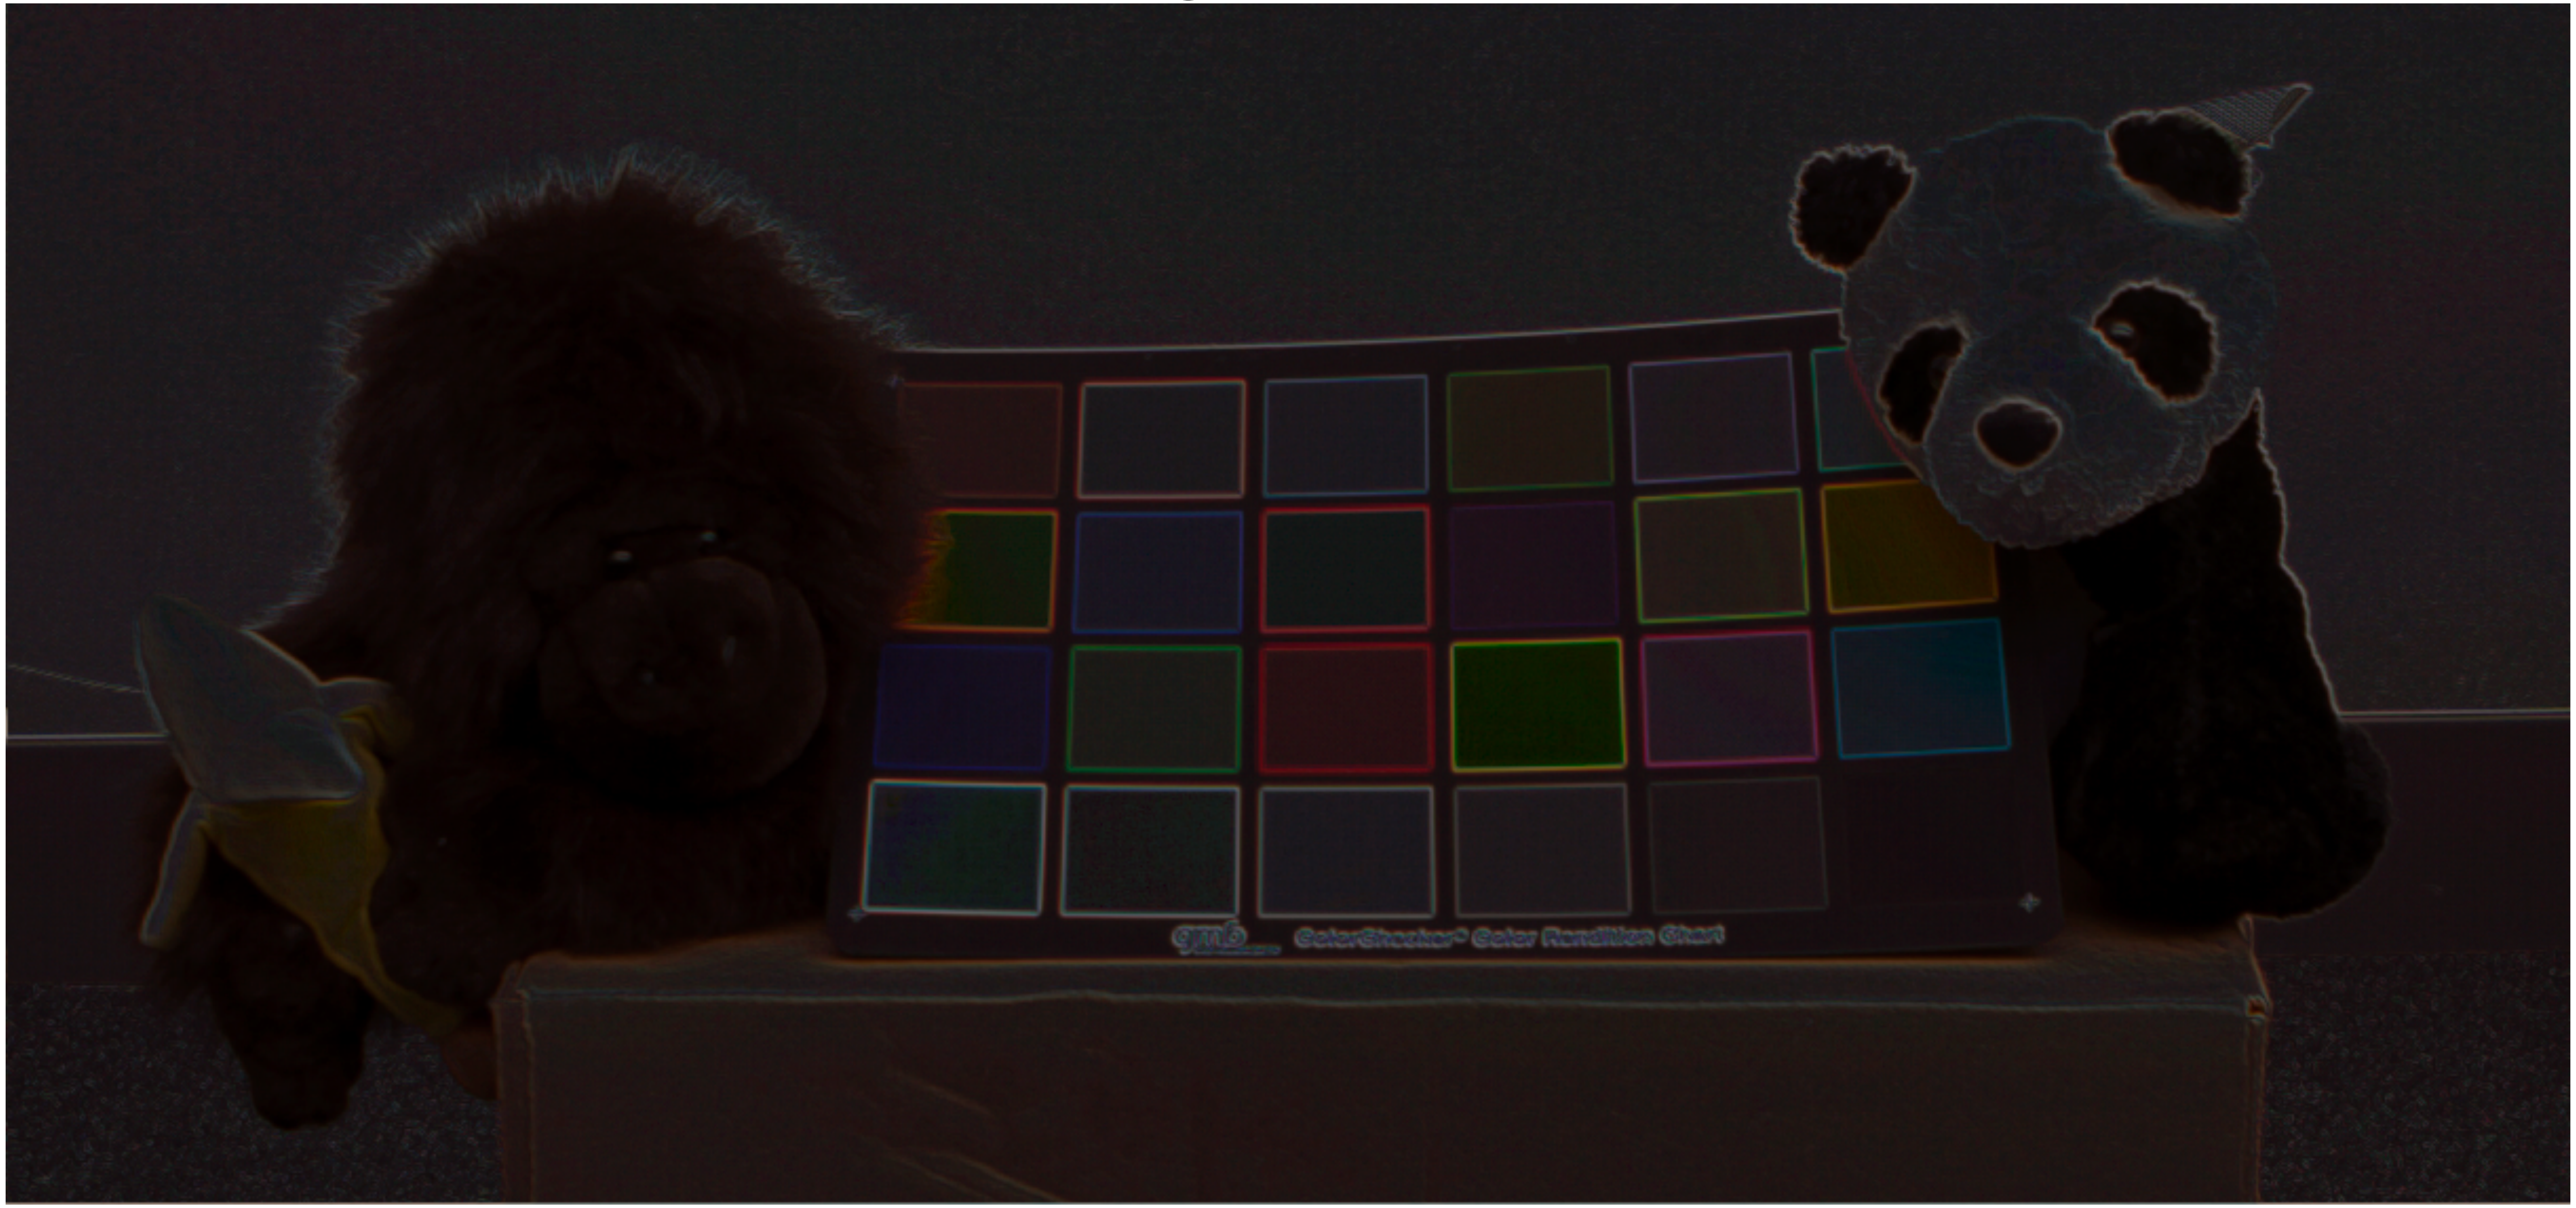

After a slide by Frédo Durand

Canon 1D IIN at ISO 3200

log SNR

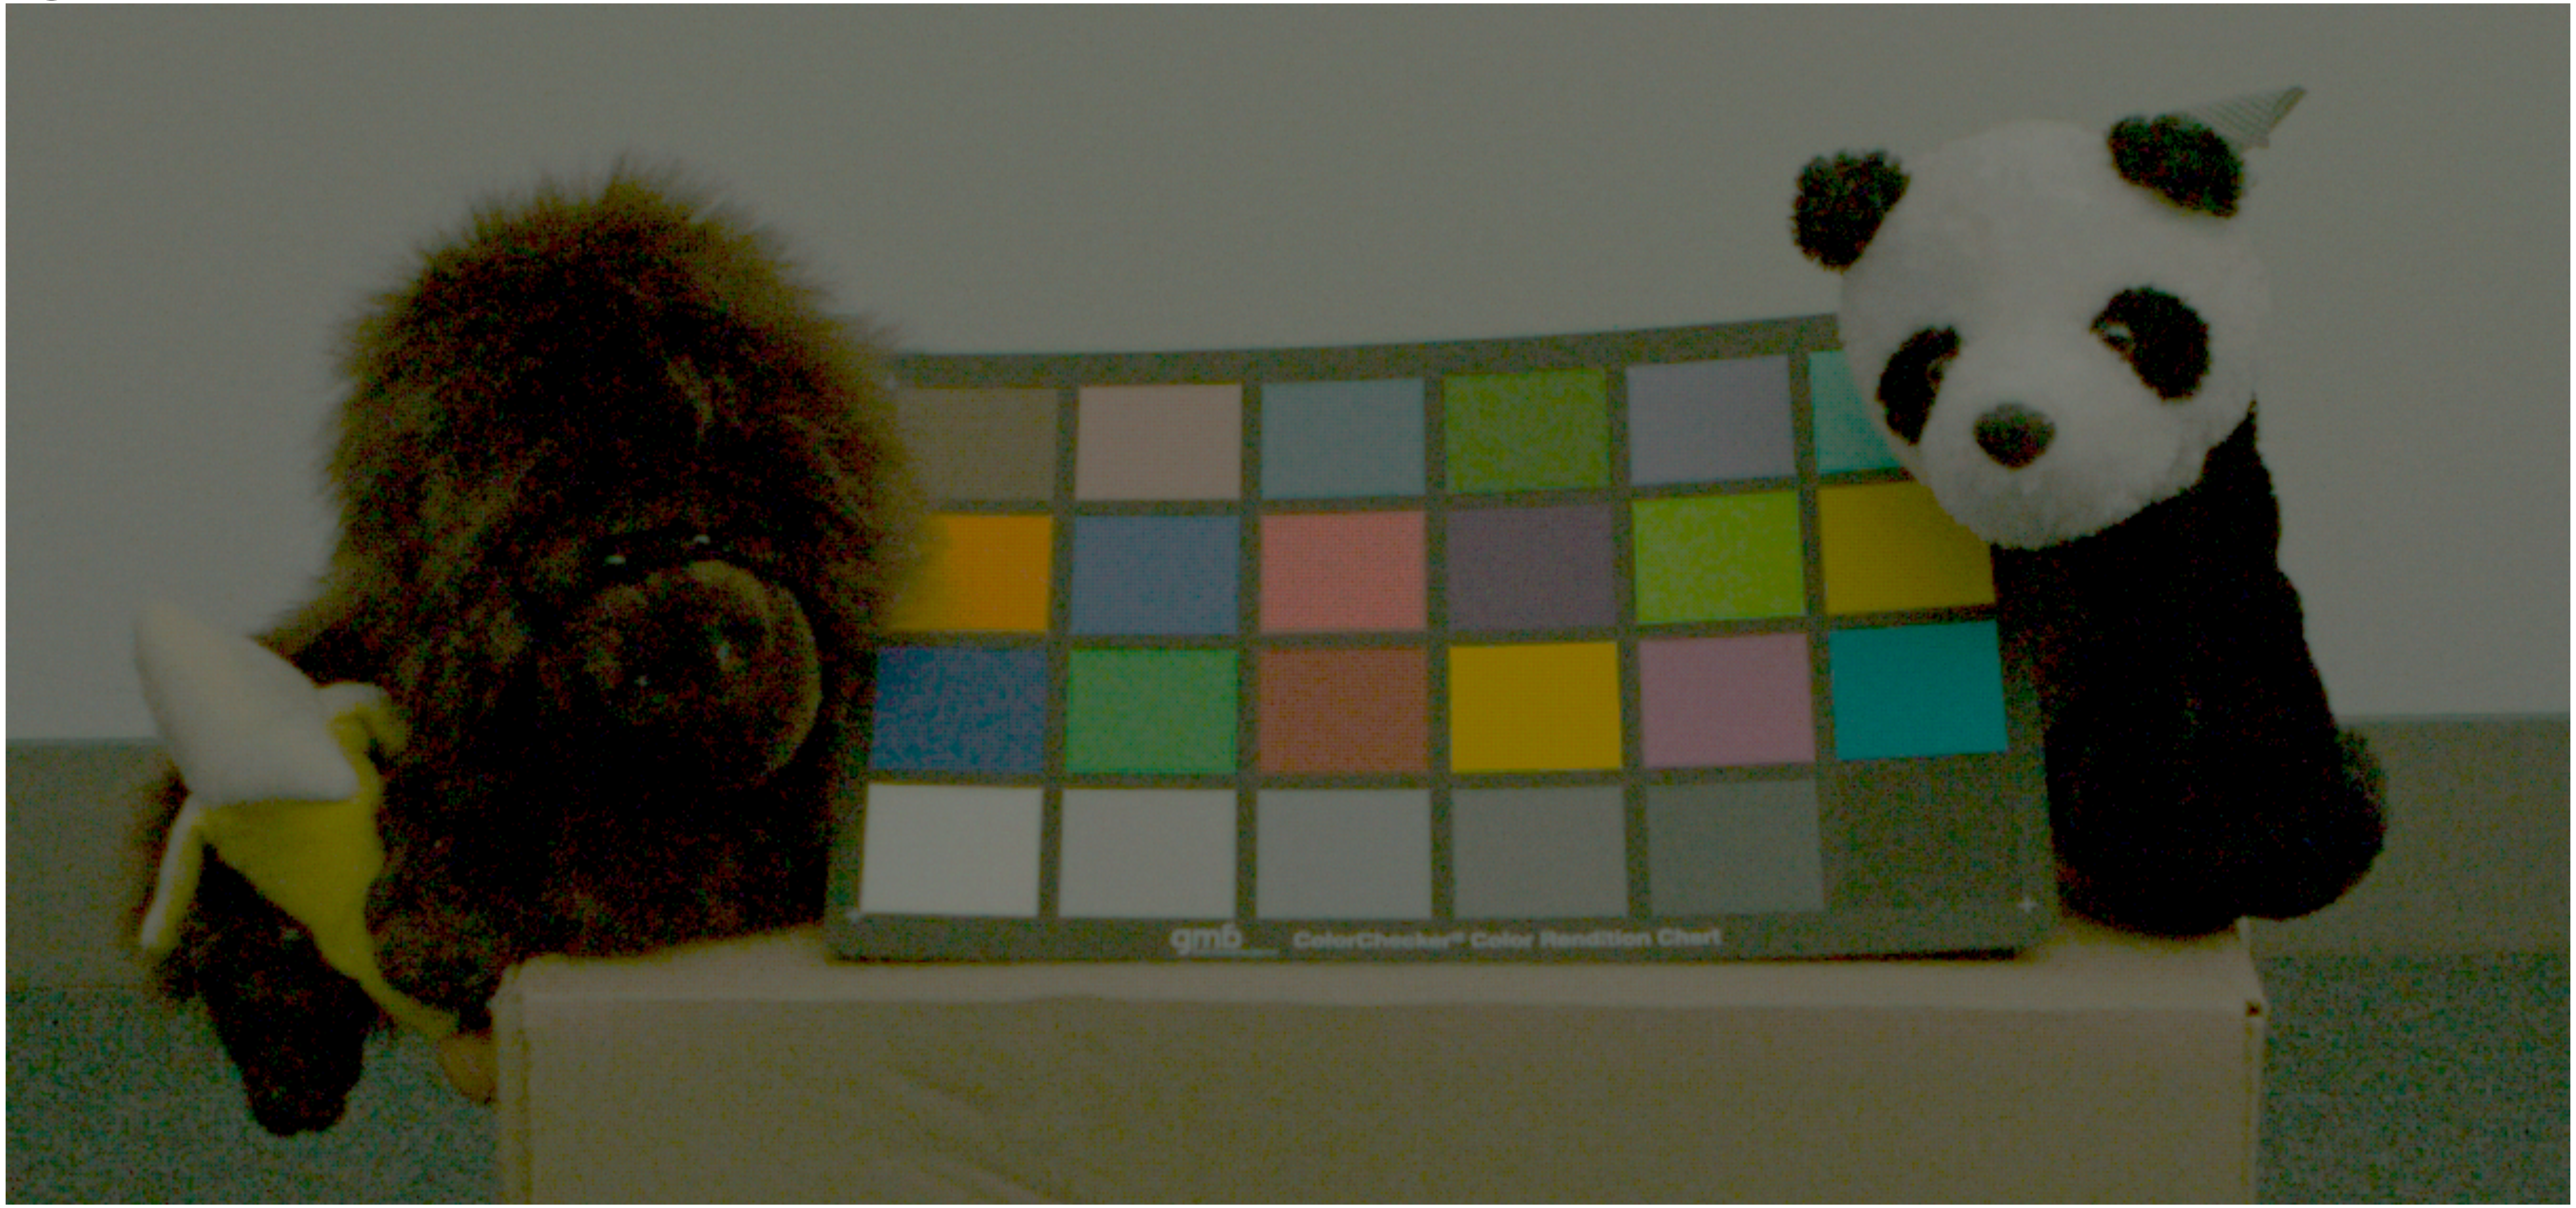

After a slide by Frédo Durand

Canon 1D IIN at ISO 1600

log SNR

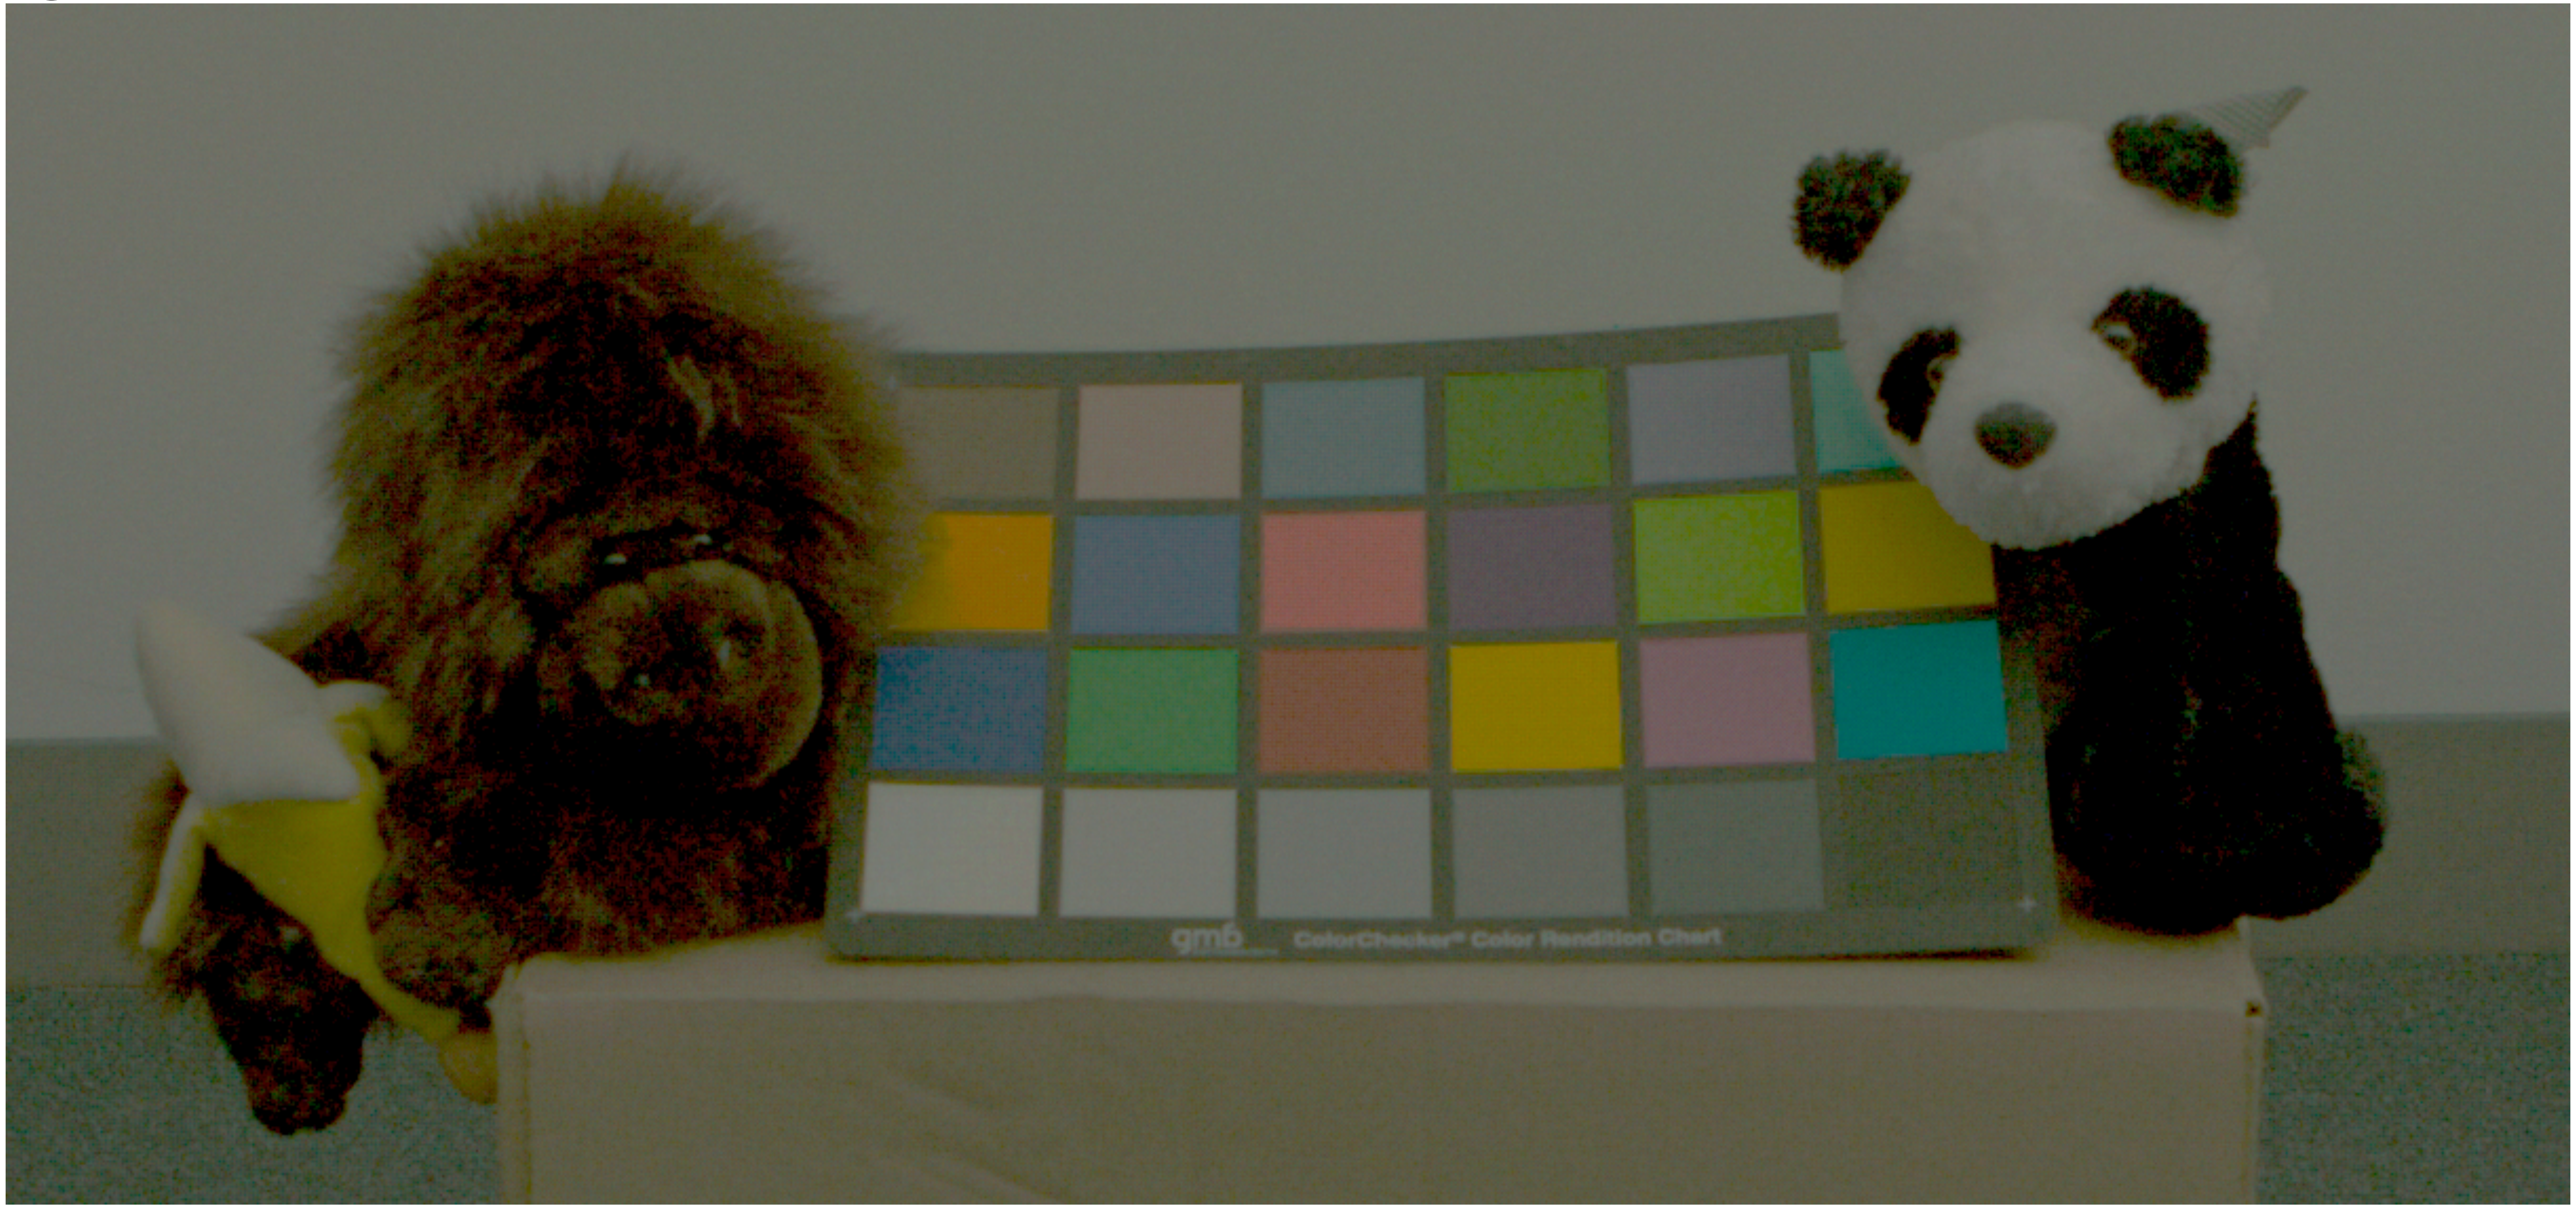

After a slide by Frédo Durand

Canon 1D IIN at ISO 400

log SNR

After a slide by Frédo Durand

Nikon D3s at 1600 ISO

After a slide by Frédo Durand

Nikon D3s at 1600 ISO

After a slide by Frédo Durand

Canon 1D Mark IIN at 1600 ISO

After a slide by Frédo Durand

Recap and questions?

Noise level depends on

- pixel intensity
- ISO
- color channel
- camera

After a slide by Frédo Durand

Sources of noise

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

On average you might get 100 photons, but sometimes it will be 98, sometimes 103, etc.

This phenomenon is governed physics and the value follows the Poisson distribution.

After a slide by Marc Levoy

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

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

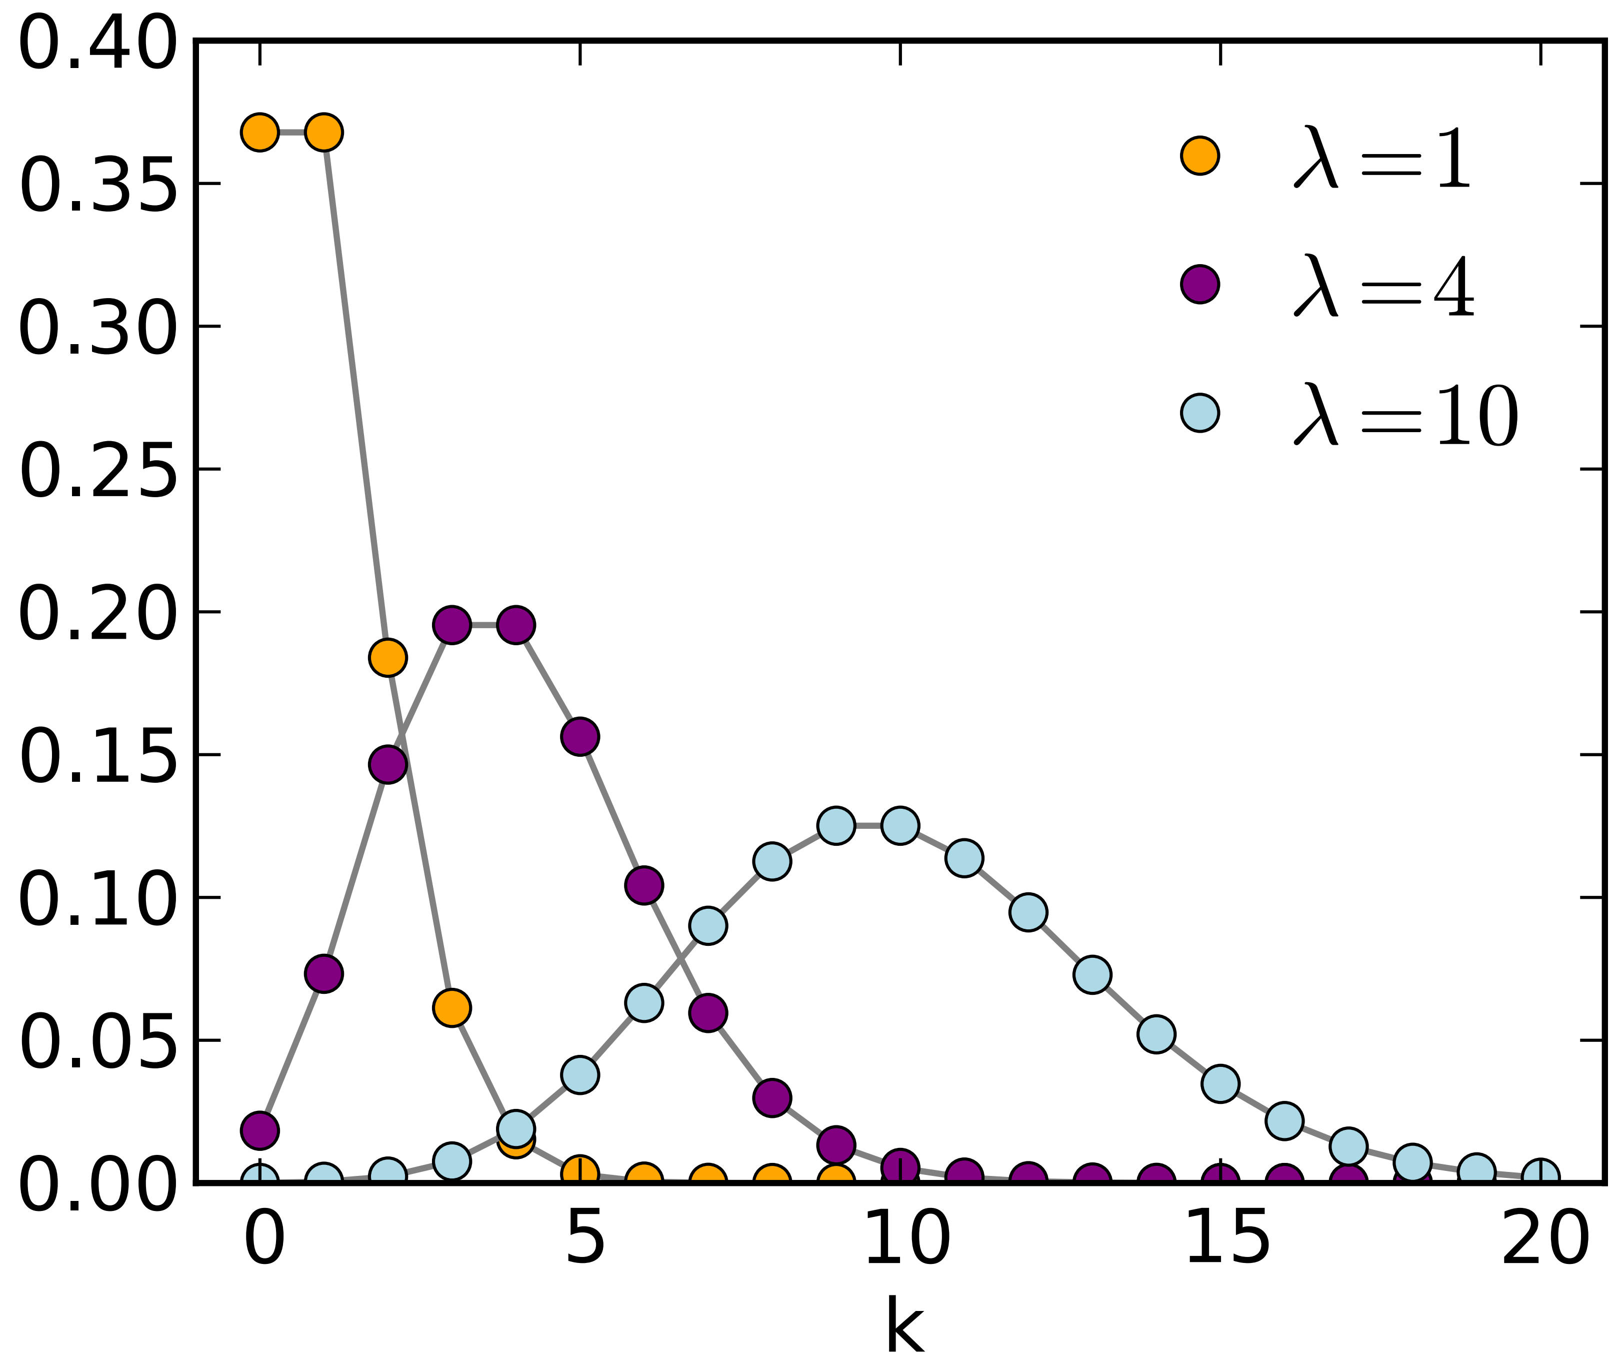

After a slide by Marc Levoy

Poisson distribution

The mean and variance of the Poisson distribution are

$$\mu = \lambda \qquad \sigma^2 = \lambda$$

The standard deviation is

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

Deviation grows slower than average.

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

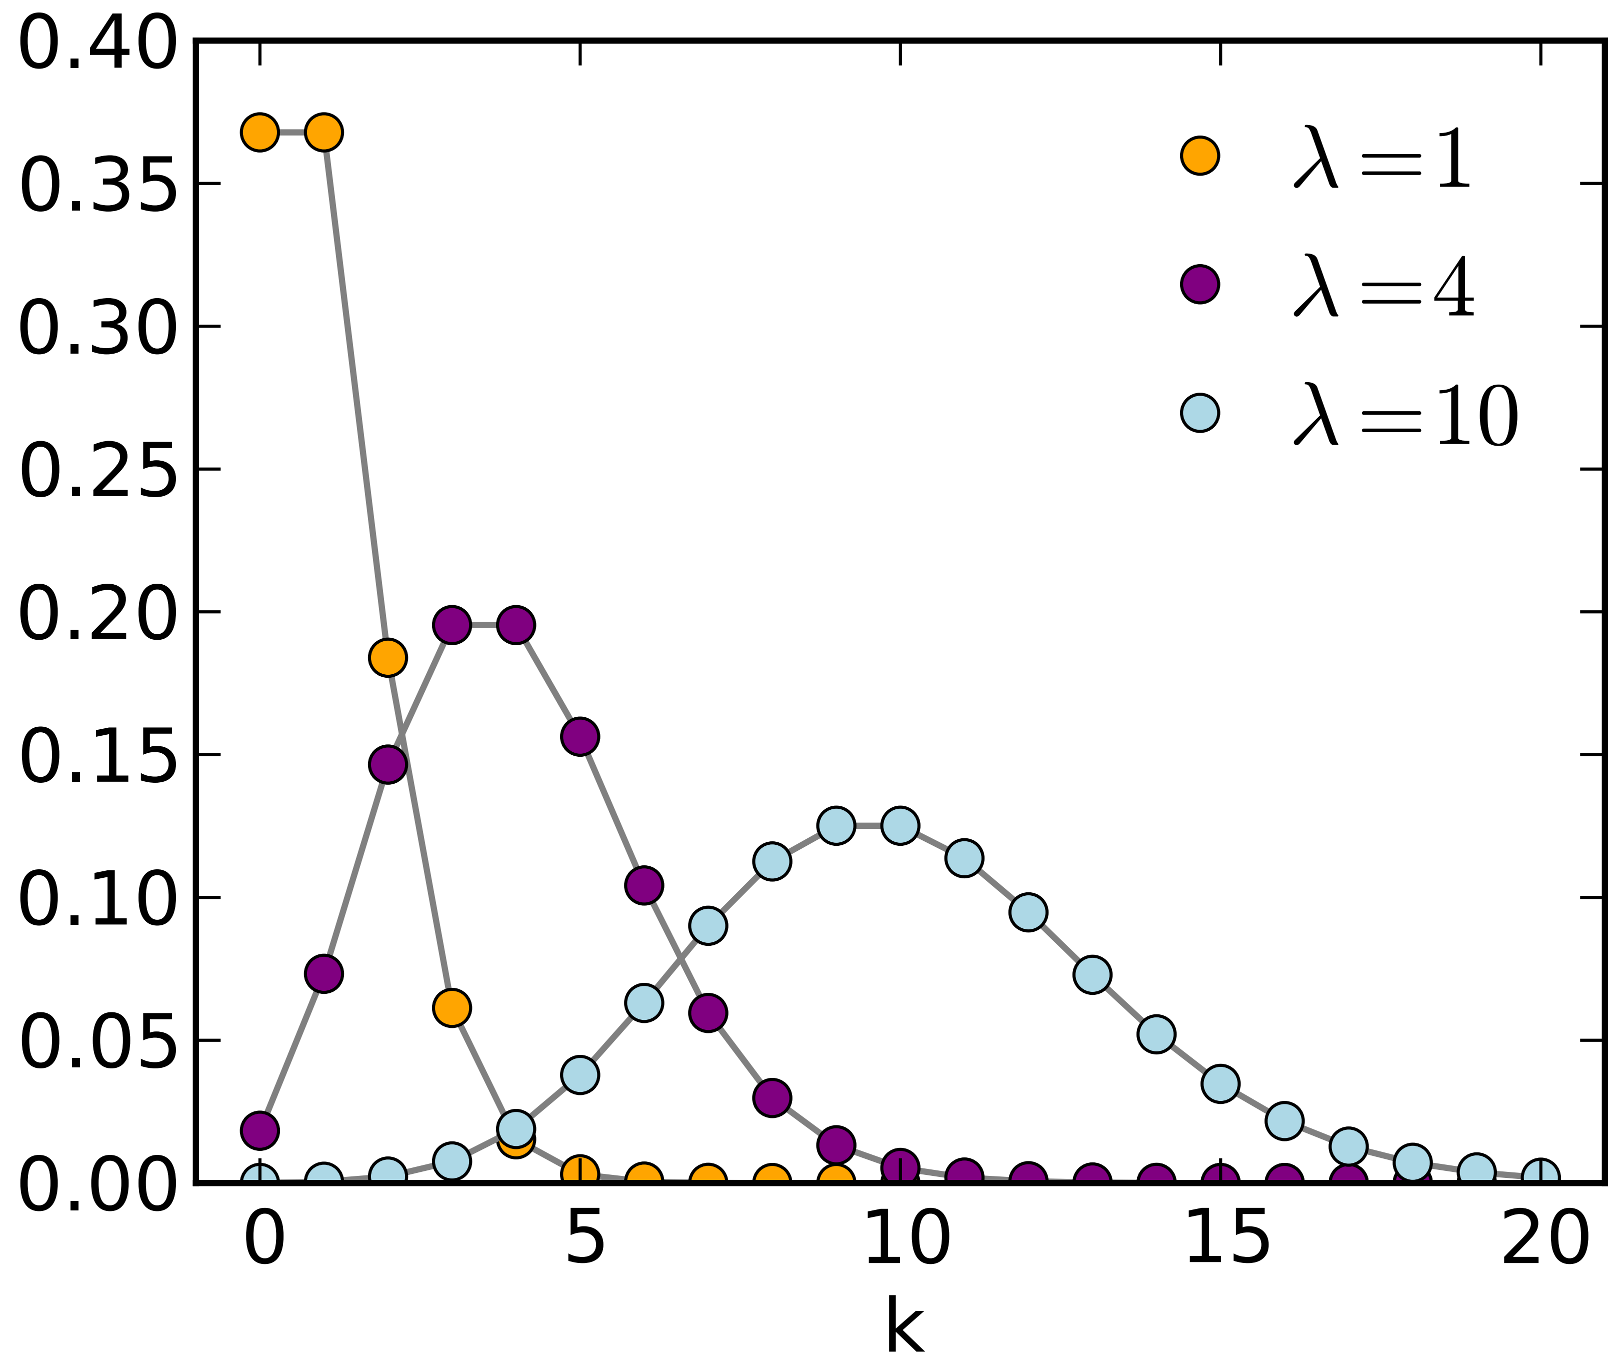

After a slide by Marc Levoy

Photon shot noise

Photons arrive in a Poisson distribution

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

so

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

Shot noise scales as square root of number of photons

After a slide by Marc Levoy

Canon 1D IIN at ISO 3200

log SNR – dominated by Poisson, ~sqrt(image)

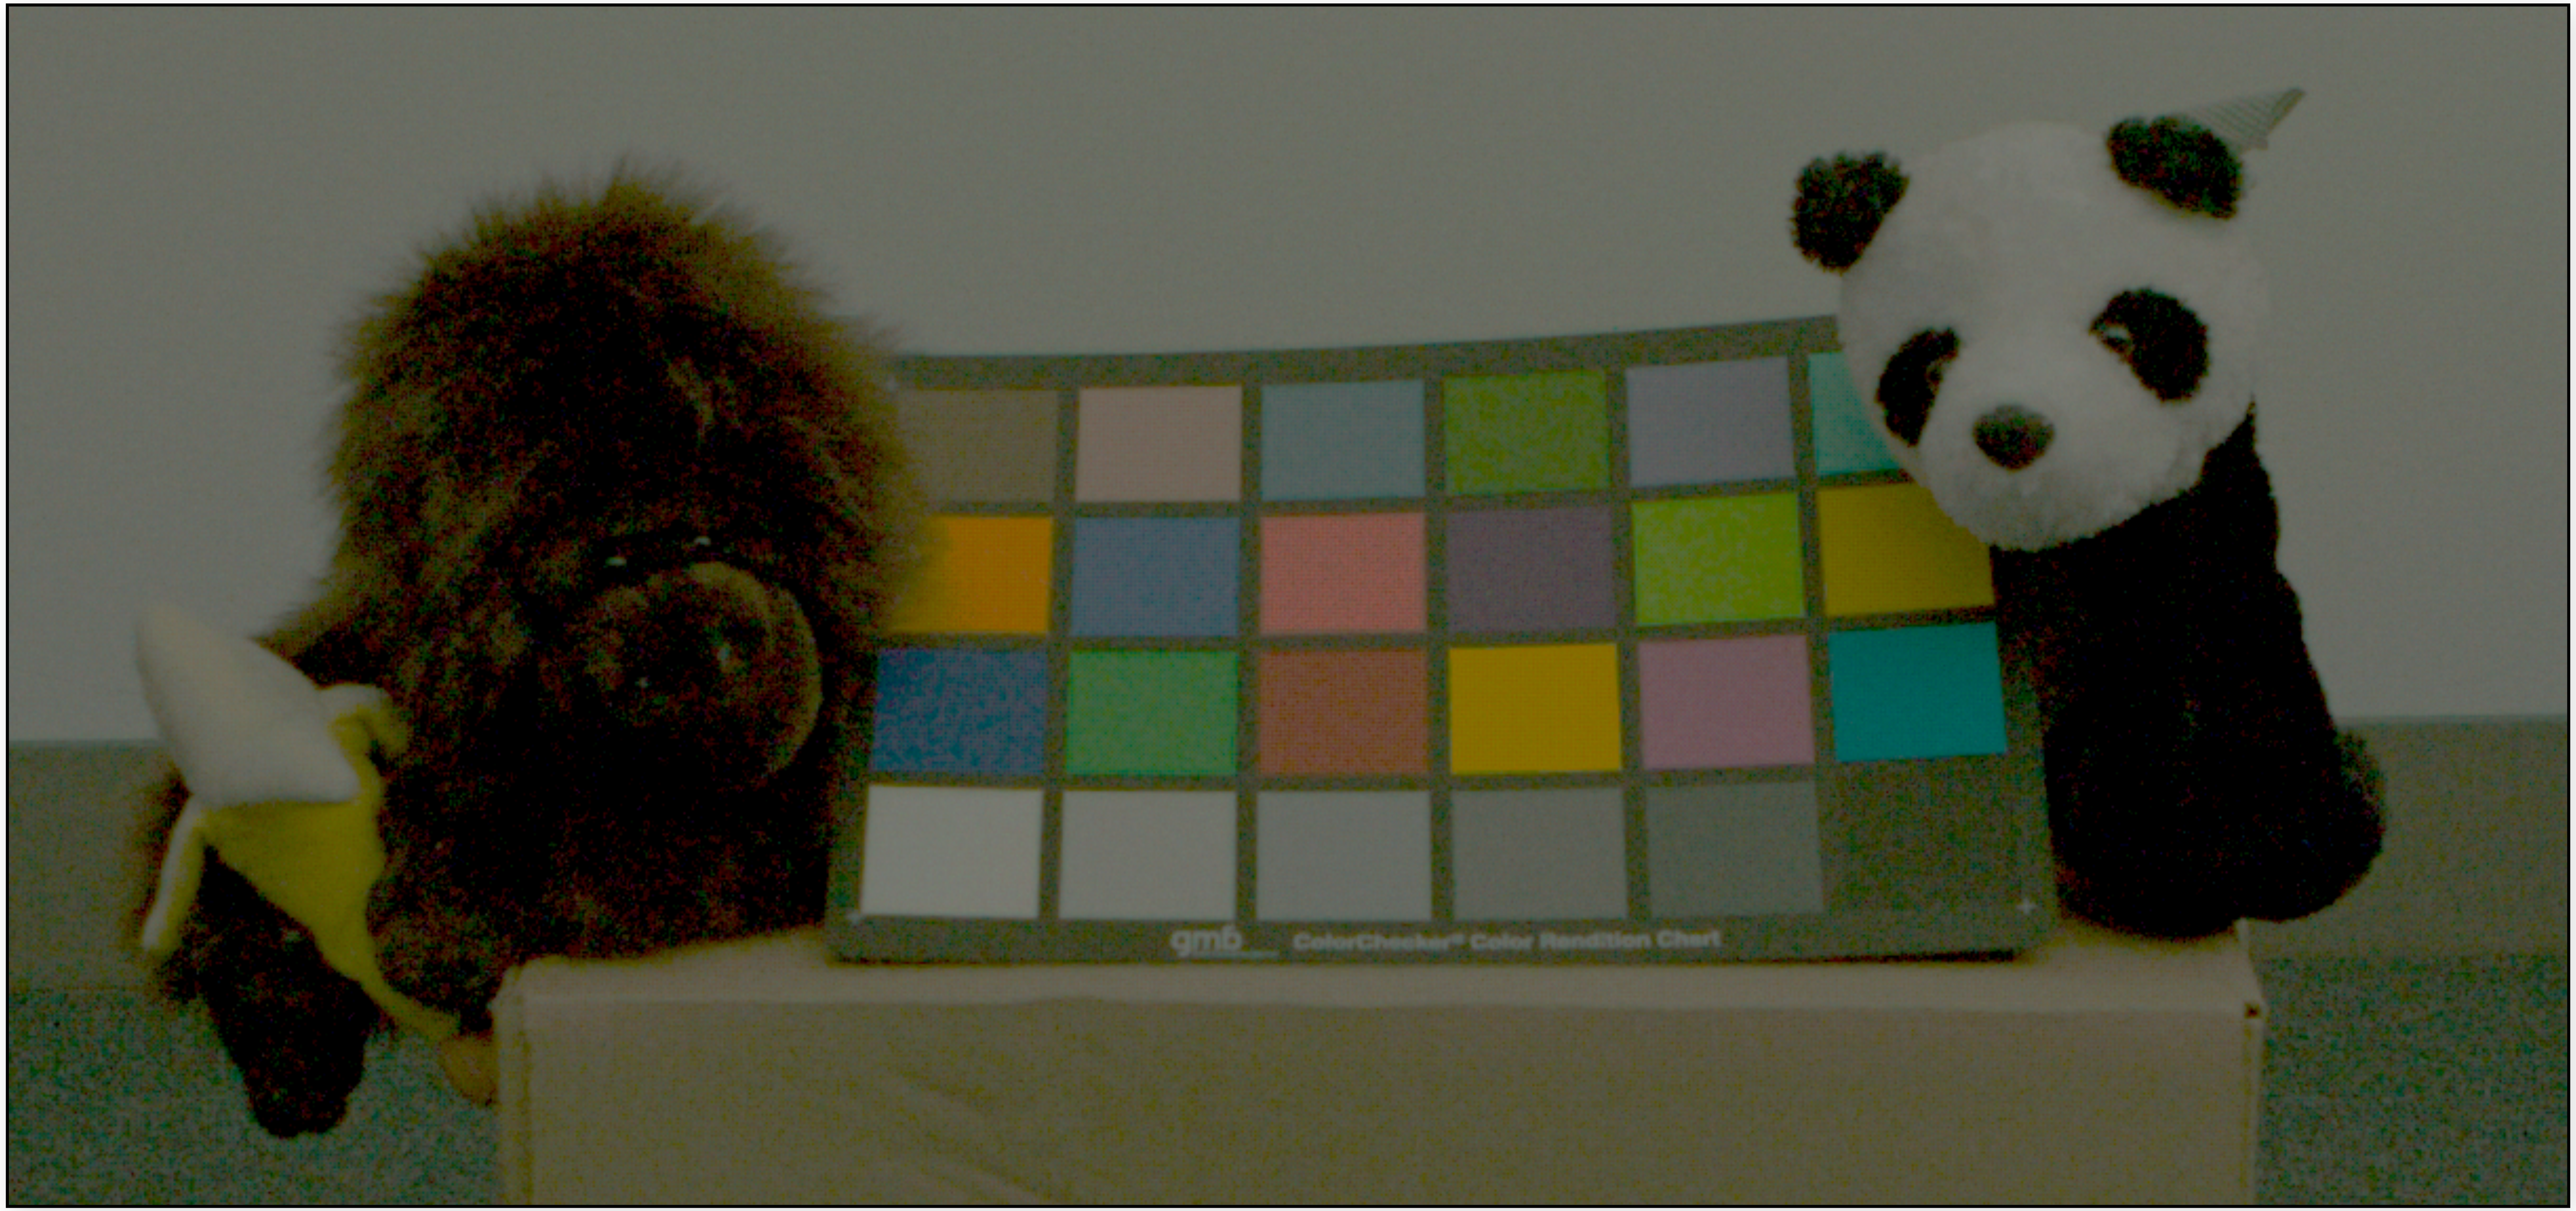

After a slide by Frédo Durand

Dark current

Electrons dislodged by random thermal activity

Increases linearly with exposure time

Increases exponentially with temperature

Varies across sensor

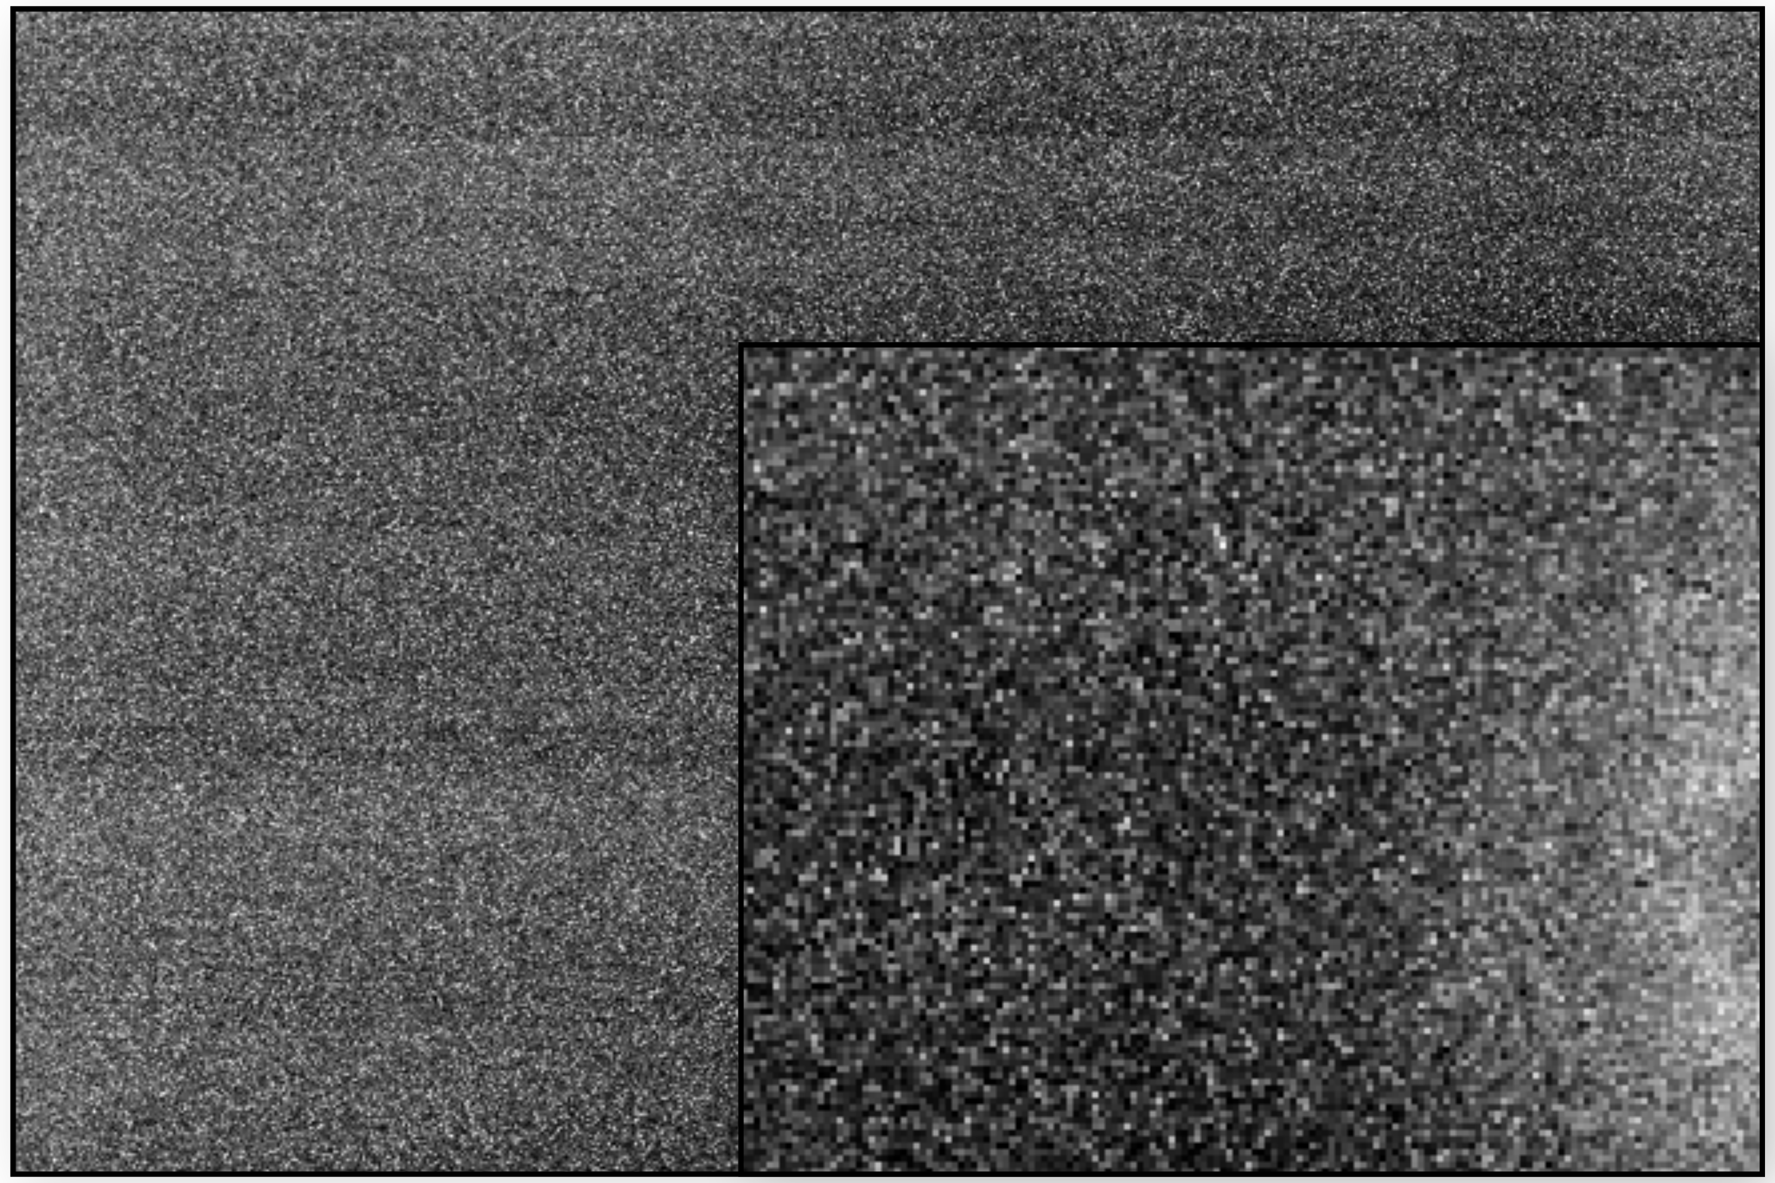

(http://theory.uchicago.edu/~ejm/pix/20d/tests/noise/)

Canon 20D, 612 sec exposure

After a slide by Marc Levoy

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

After a slide by Marc Levoy

Fixing dark current and hot pixels

Solution #1: chill the sensor

Solution #2: dark frame subtraction

- available on high-end SLRs
- compensates for average dark current
- also compensates for hot pixels and FPN

After a slide by Marc Levoy

Fixed pattern noise (FPN)

Manufacturing variations across pixels, columns, blocks

Mainly in CMOS sensors

Doesn’t change over time, so read once and subtract

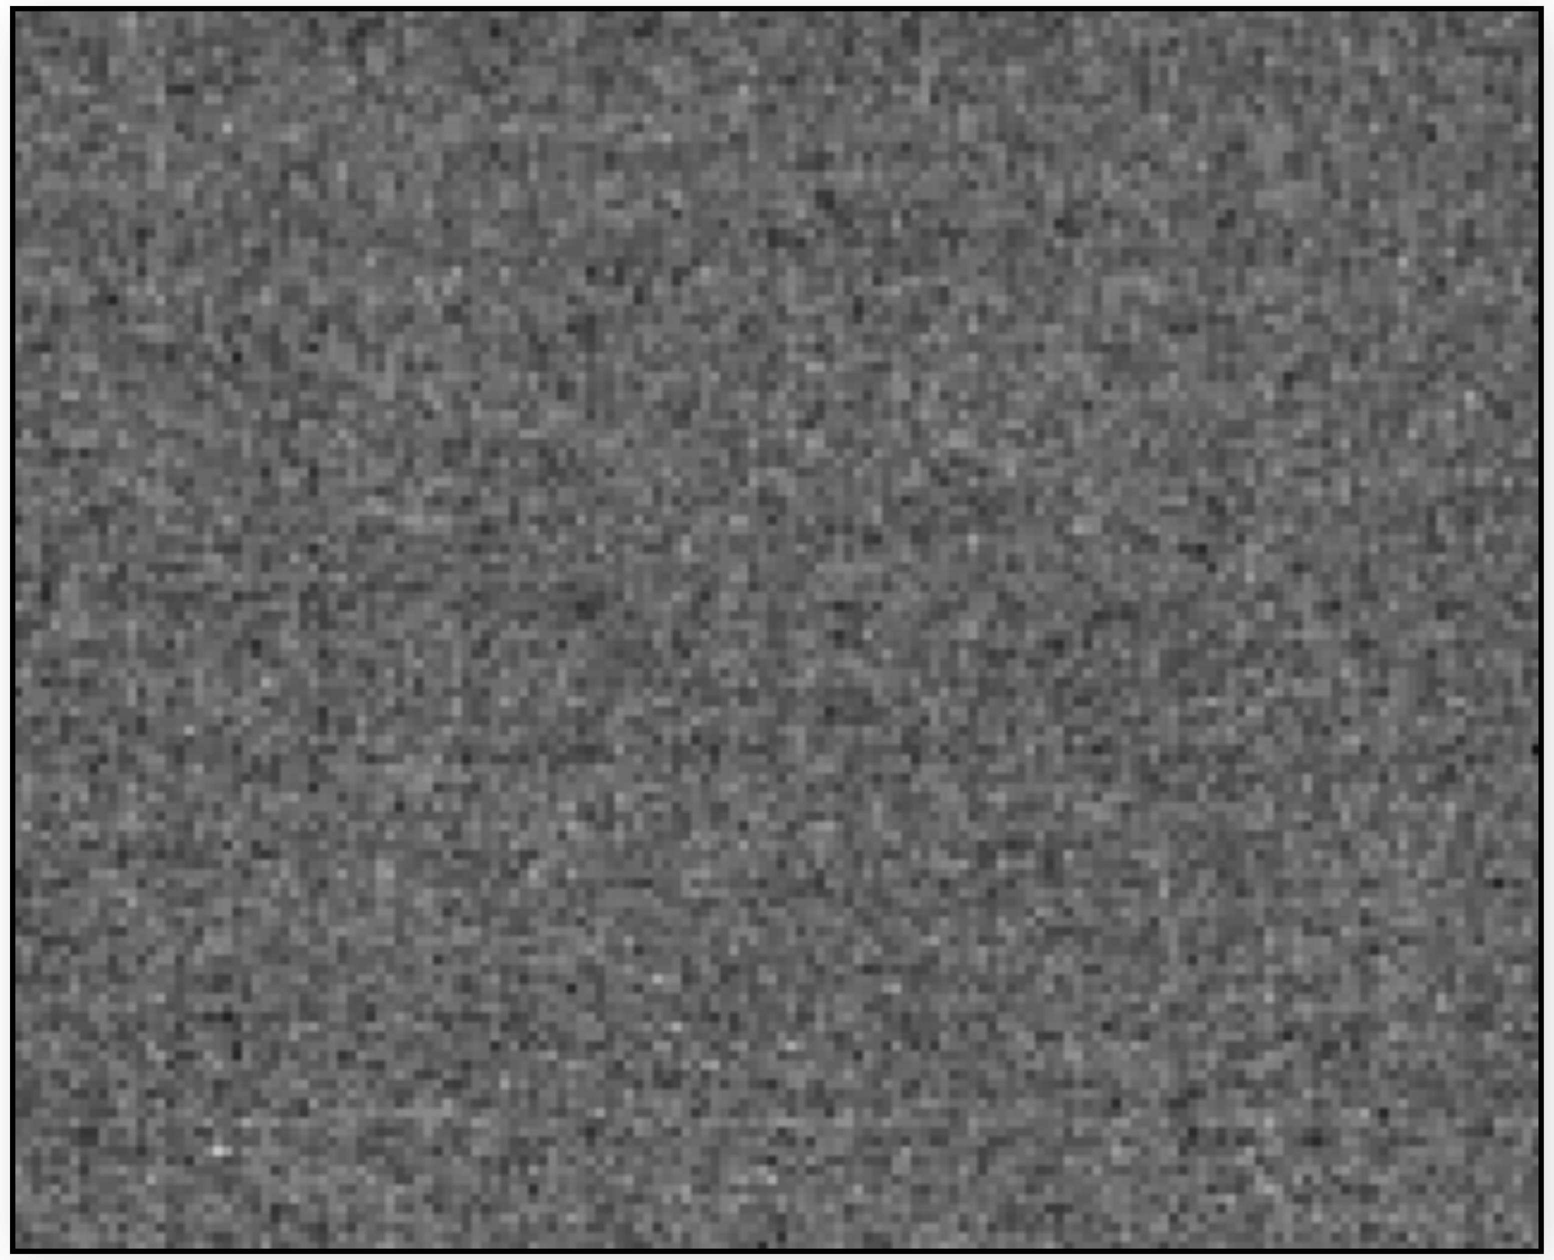

Canon 20D, ISO 800, cropped

After a slide by Marc Levoy

Read noise

Thermal noise in readout circuitry

Again, mainly in CMOS sensors

Not fixed patterns, so only solution is cooling

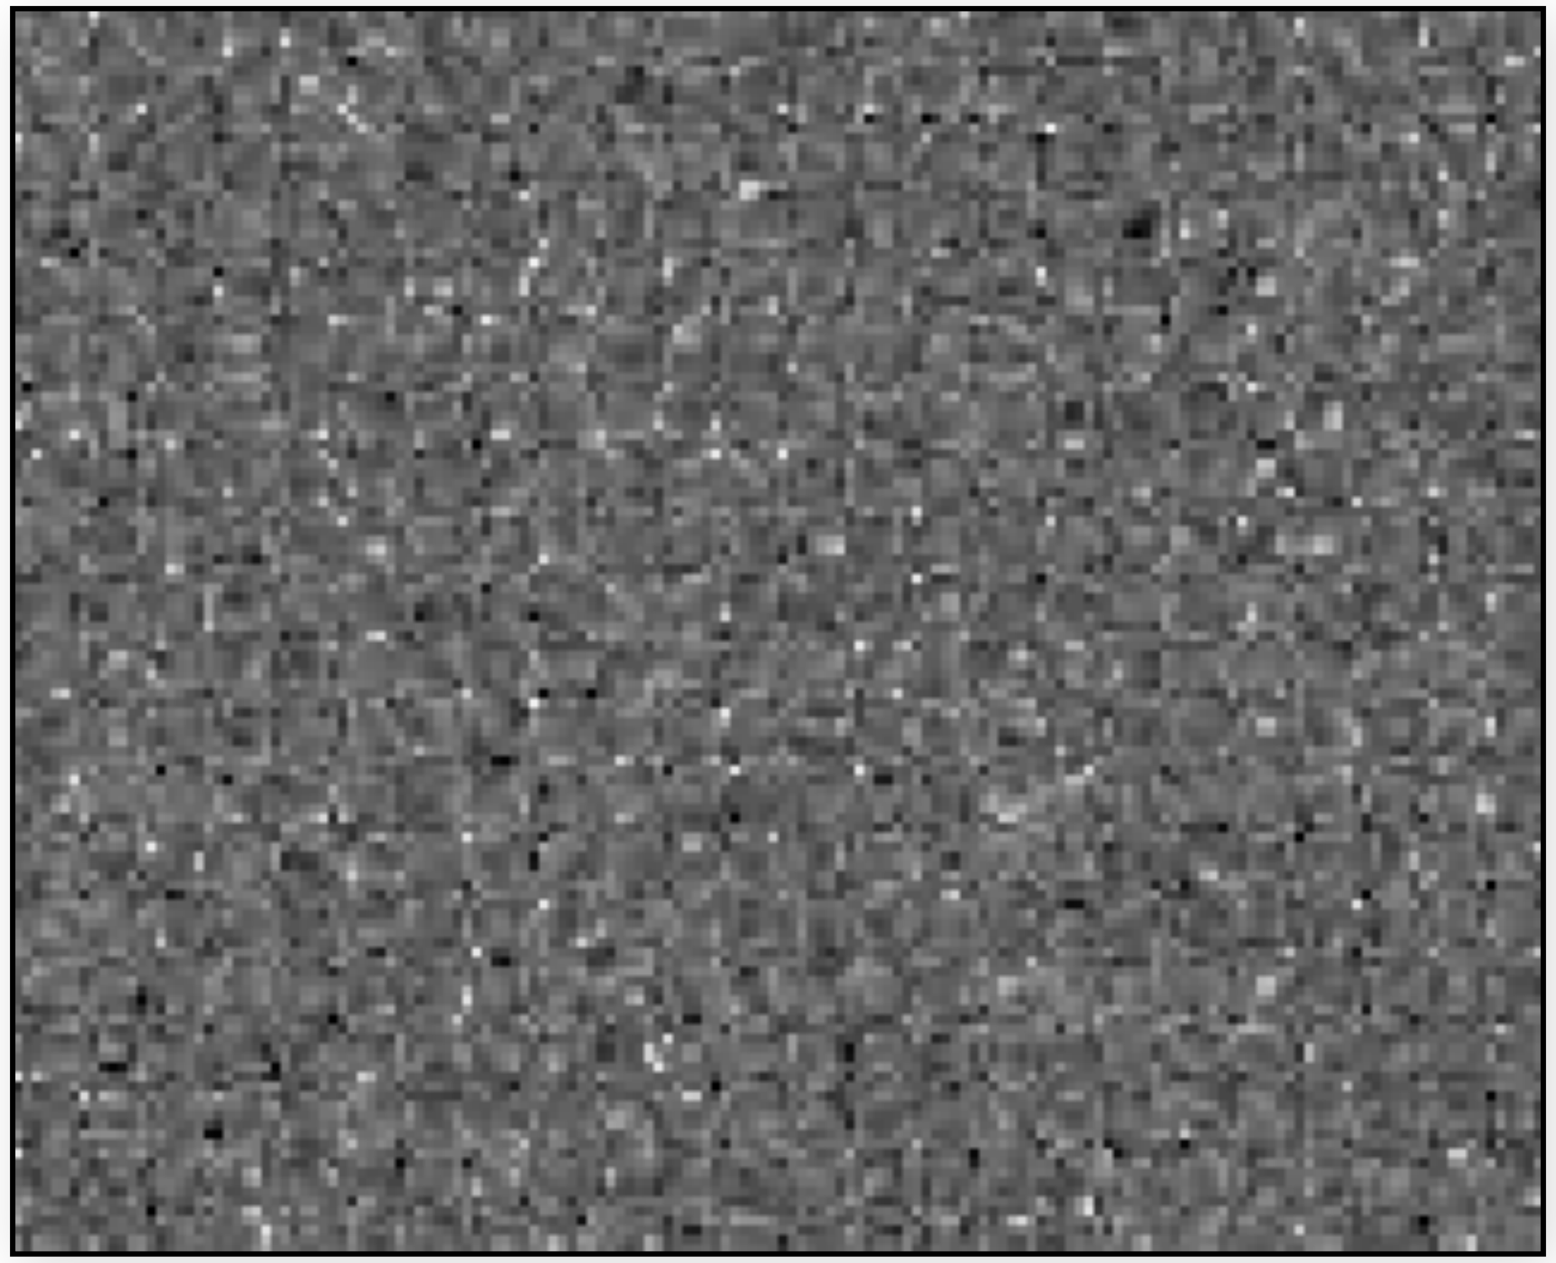

Canon 1Ds Mark III, cropped

After a slide by Marc Levoy

Recap

Photon shot noise

- **unavoidable** randomness in number of photons arriving
- grows as the sqrt 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**

After a slide by Marc Levoy

Digital pipeline

Photosites transform photons into charge (electrons)

- The sensor itself is *linear*

Gets amplified (depending on ISO setting)

Then goes through analog-to-digital converter

- up to 14 bits/channel these days

Stop here when shooting RAW

Then demosaicing, denoising, white balance, a response curve, gamma encoding are applied

Quantized and recorded as 8-bit JPEG

After a slide by Frédo Durand

Pipeline & noise

This is a conceptual diagram, don't take it too literally

- e.g. the A-to-D converter is a serious source of noise, but usually electronic noise, not quantization artifacts

Noise = (photon noise + readout noise) * amplification + post-amplification noise

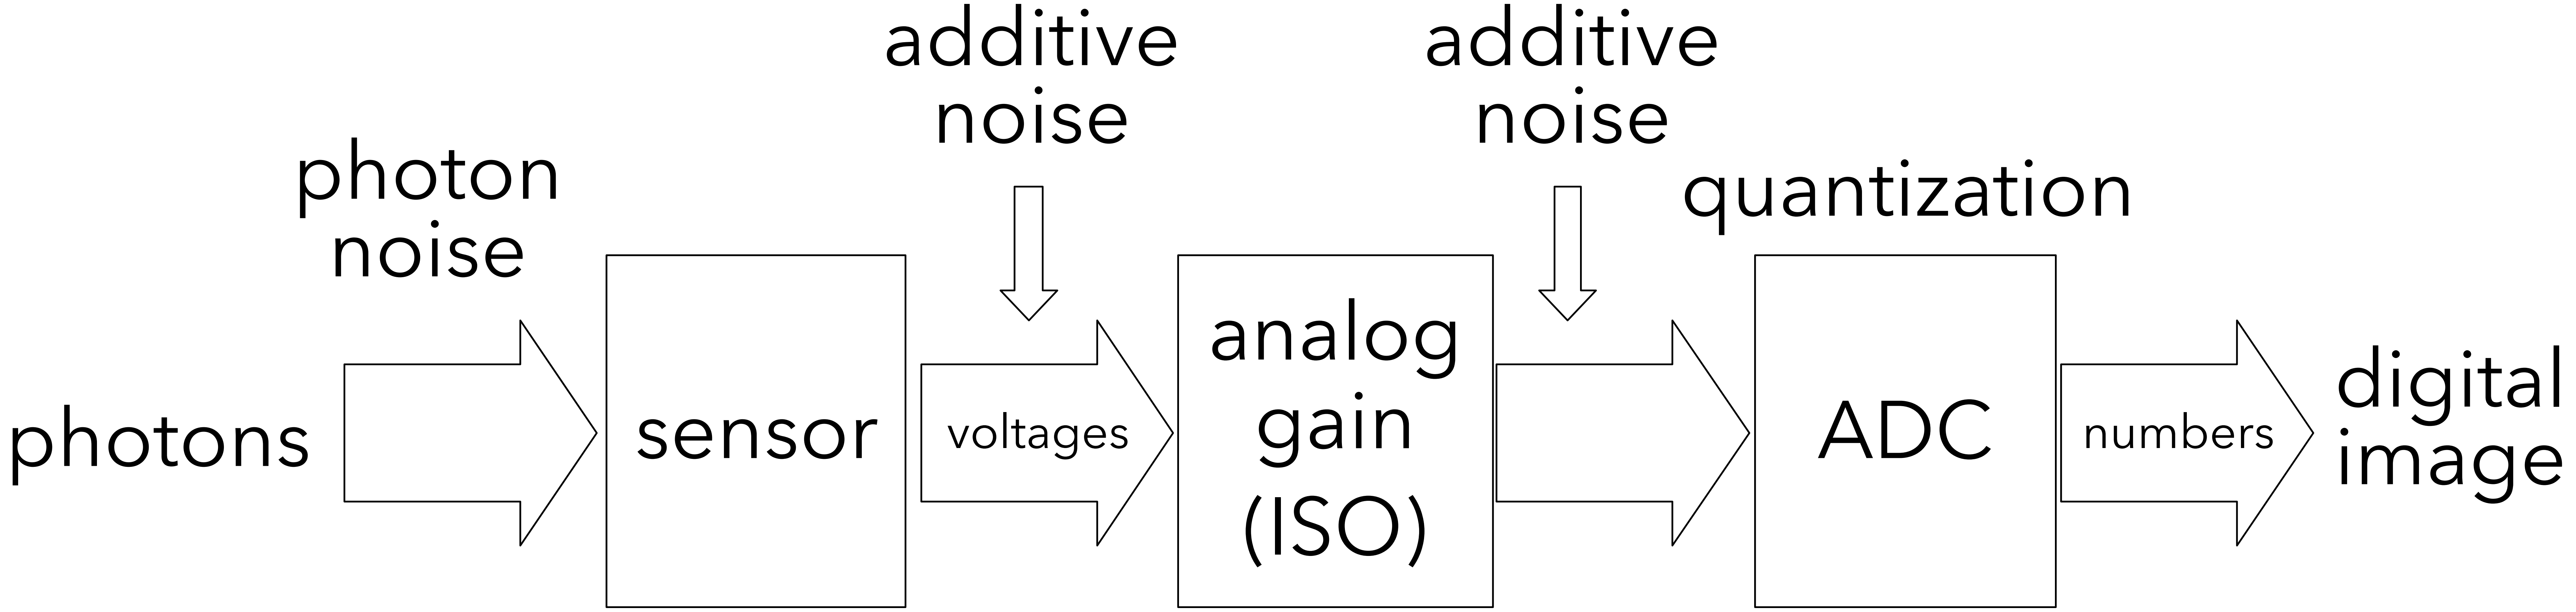

After a slide by Frédo Durand

ISO amplifies

e.g. going from ISO 100 to ISO 400 amplifies by 4

both noise & signal

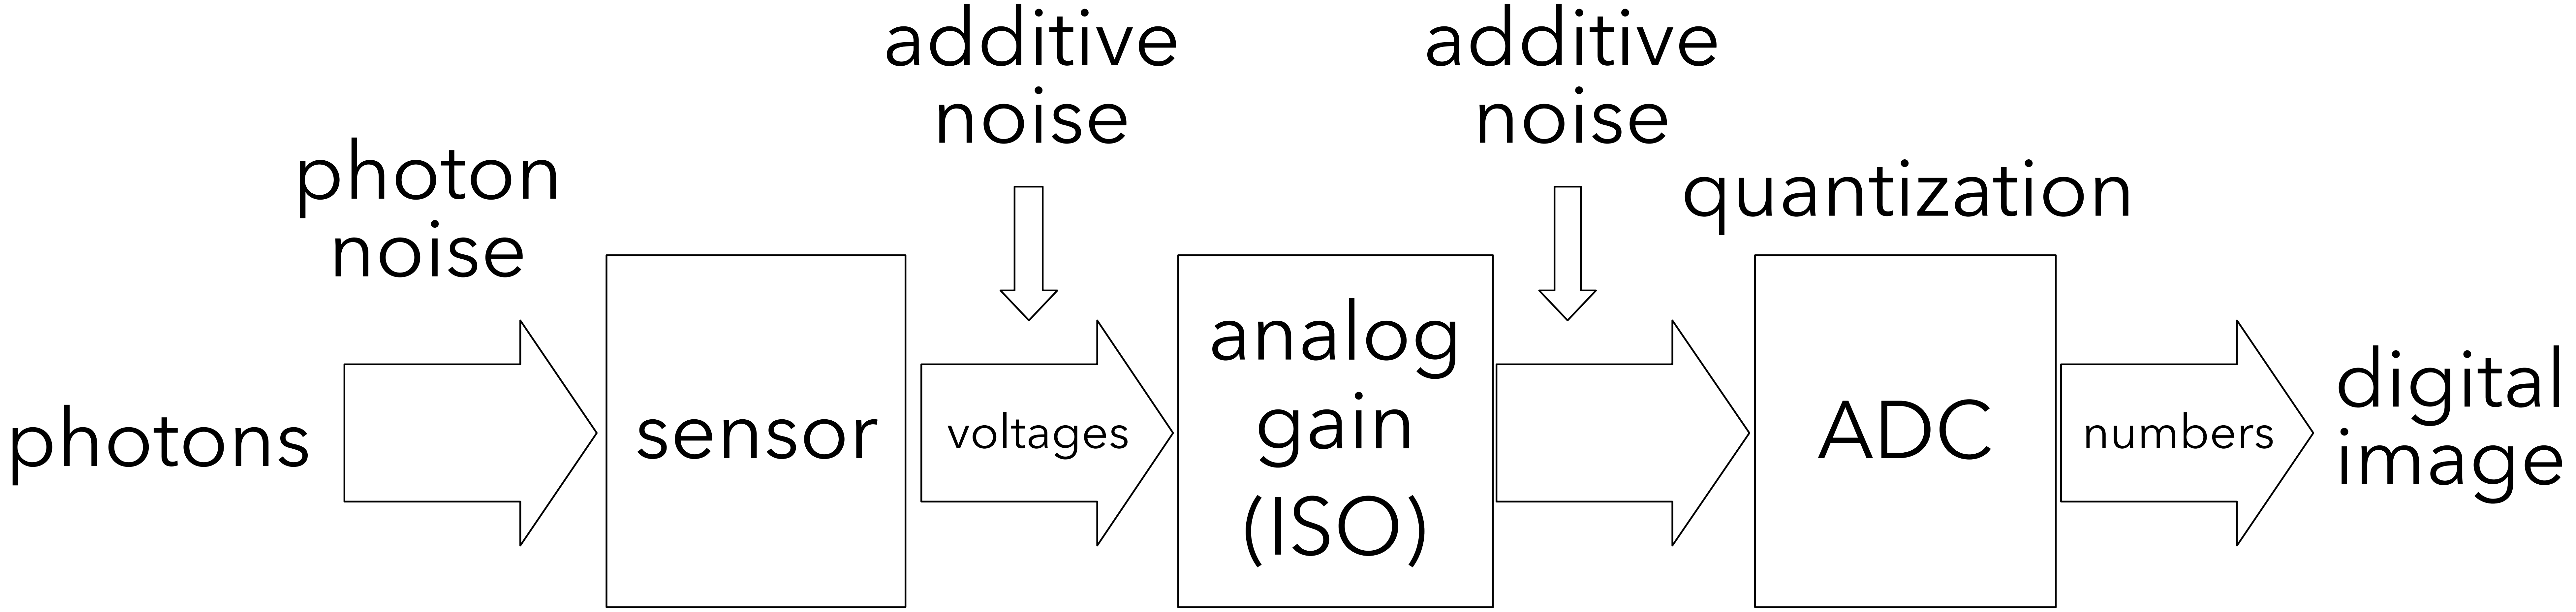

After a slide by Frédo Durand

ISO 3200

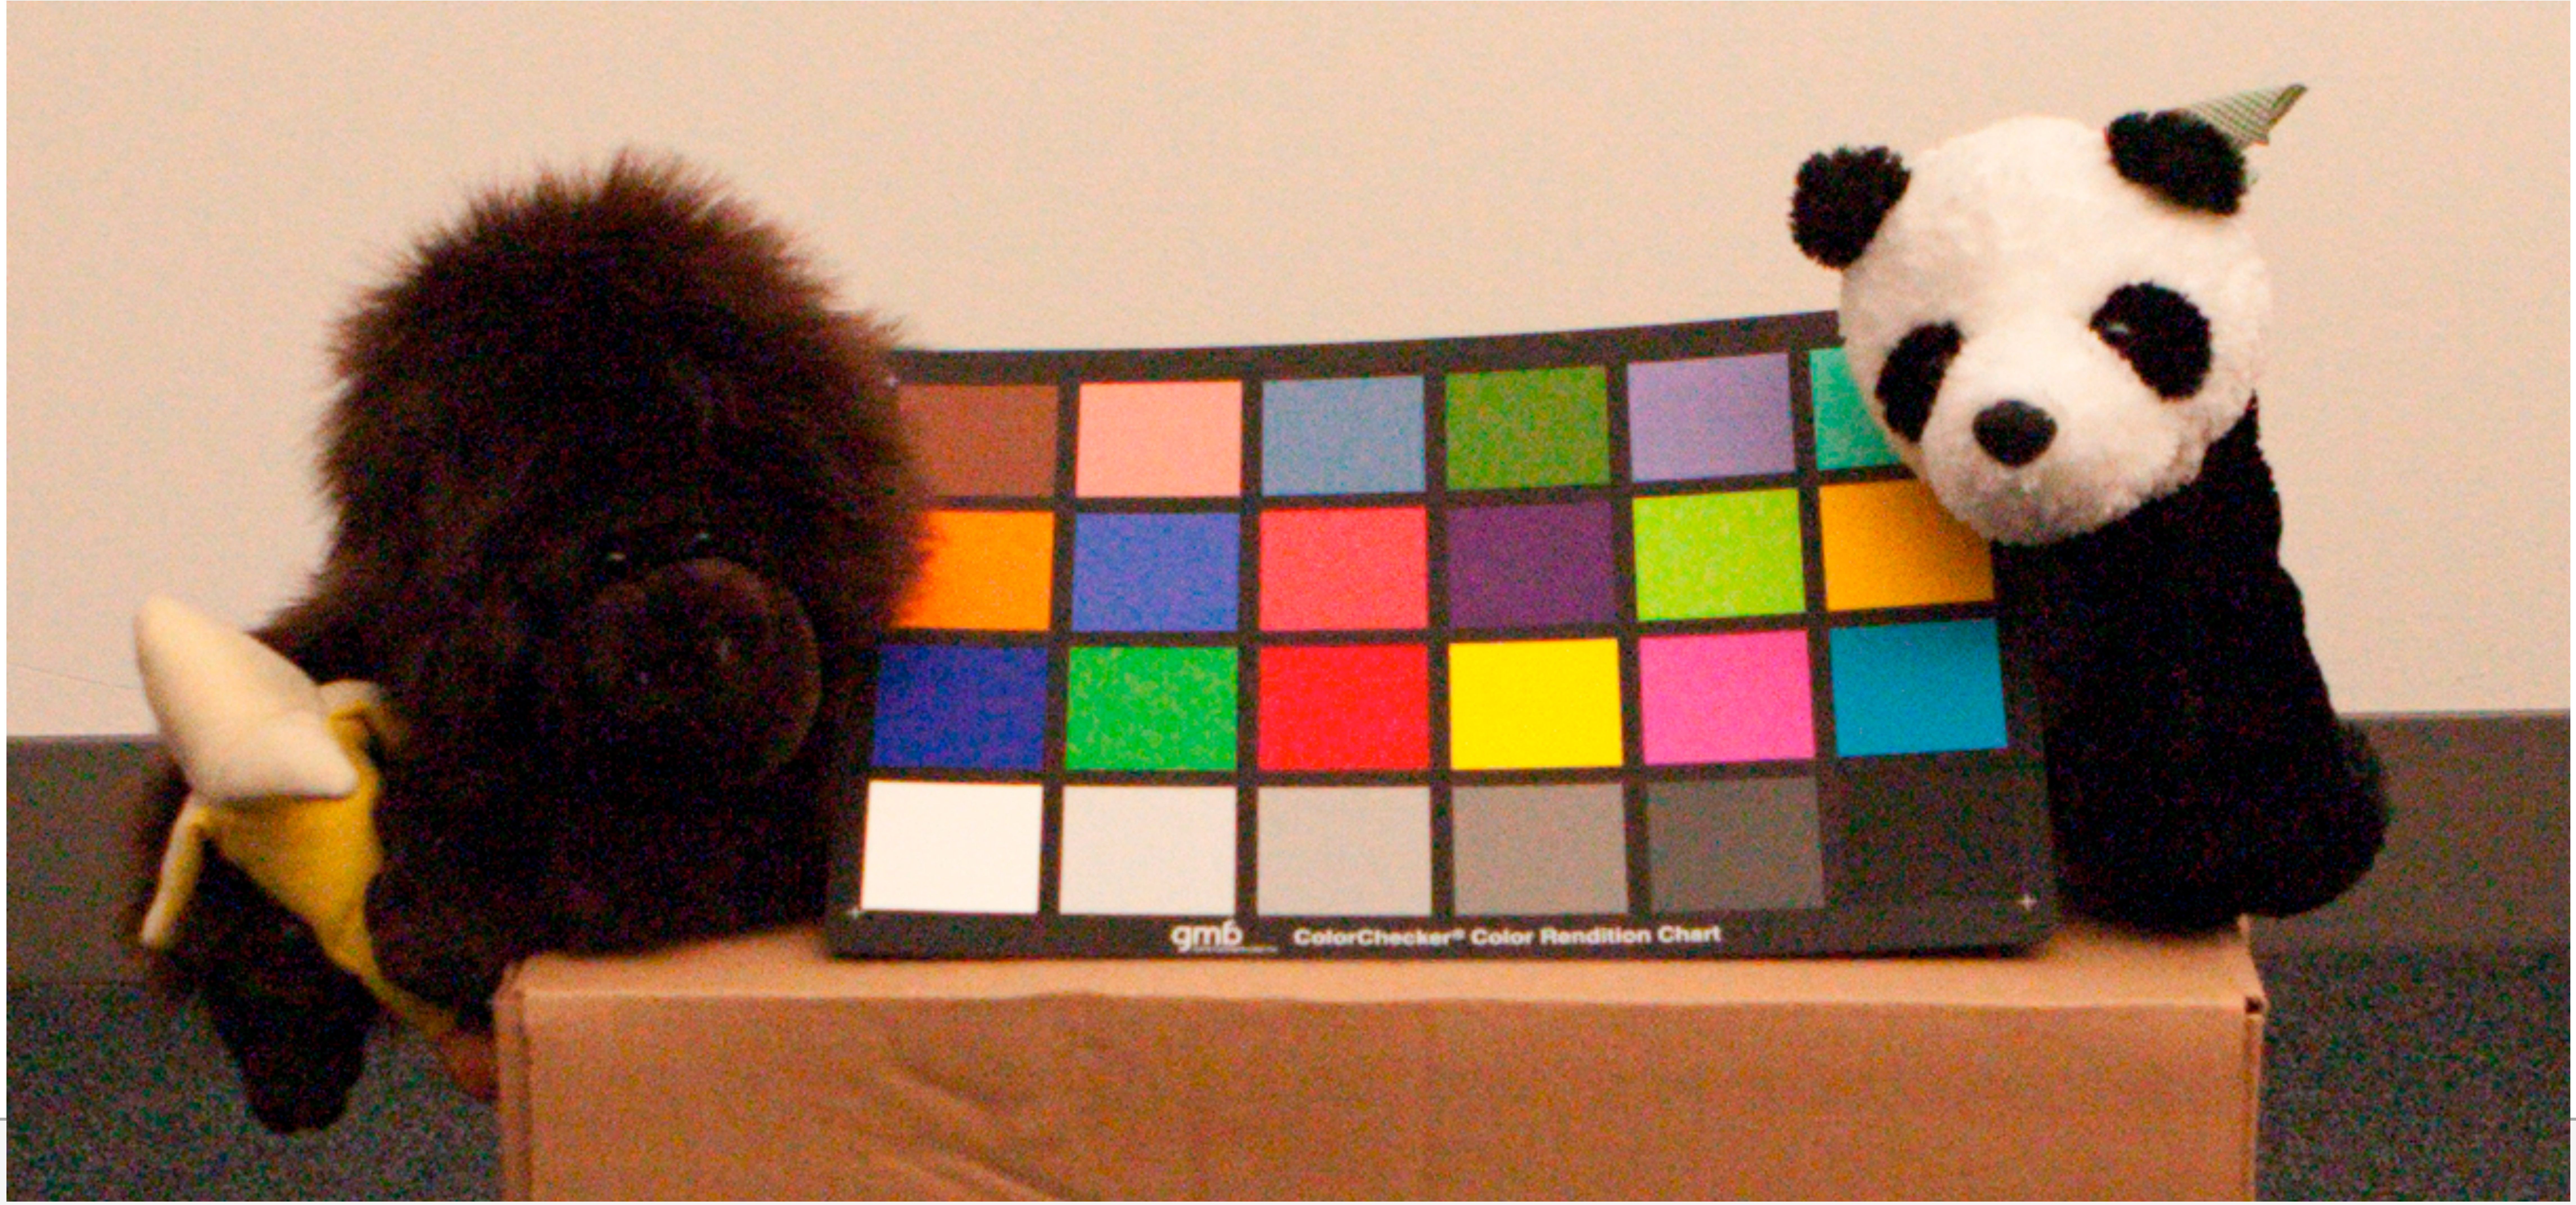

After a slide by Frédo Durand

ISO 100

A lot less noisy!

After a slide by Frédo Durand

Pipeline & noise

For a given signal level, and a given desired image brightness…

Two alternatives

- use low ISO and brighten digitally
- use high ISO to get brightness directly

The latter gives less noise because you don't amplify post-gain noise

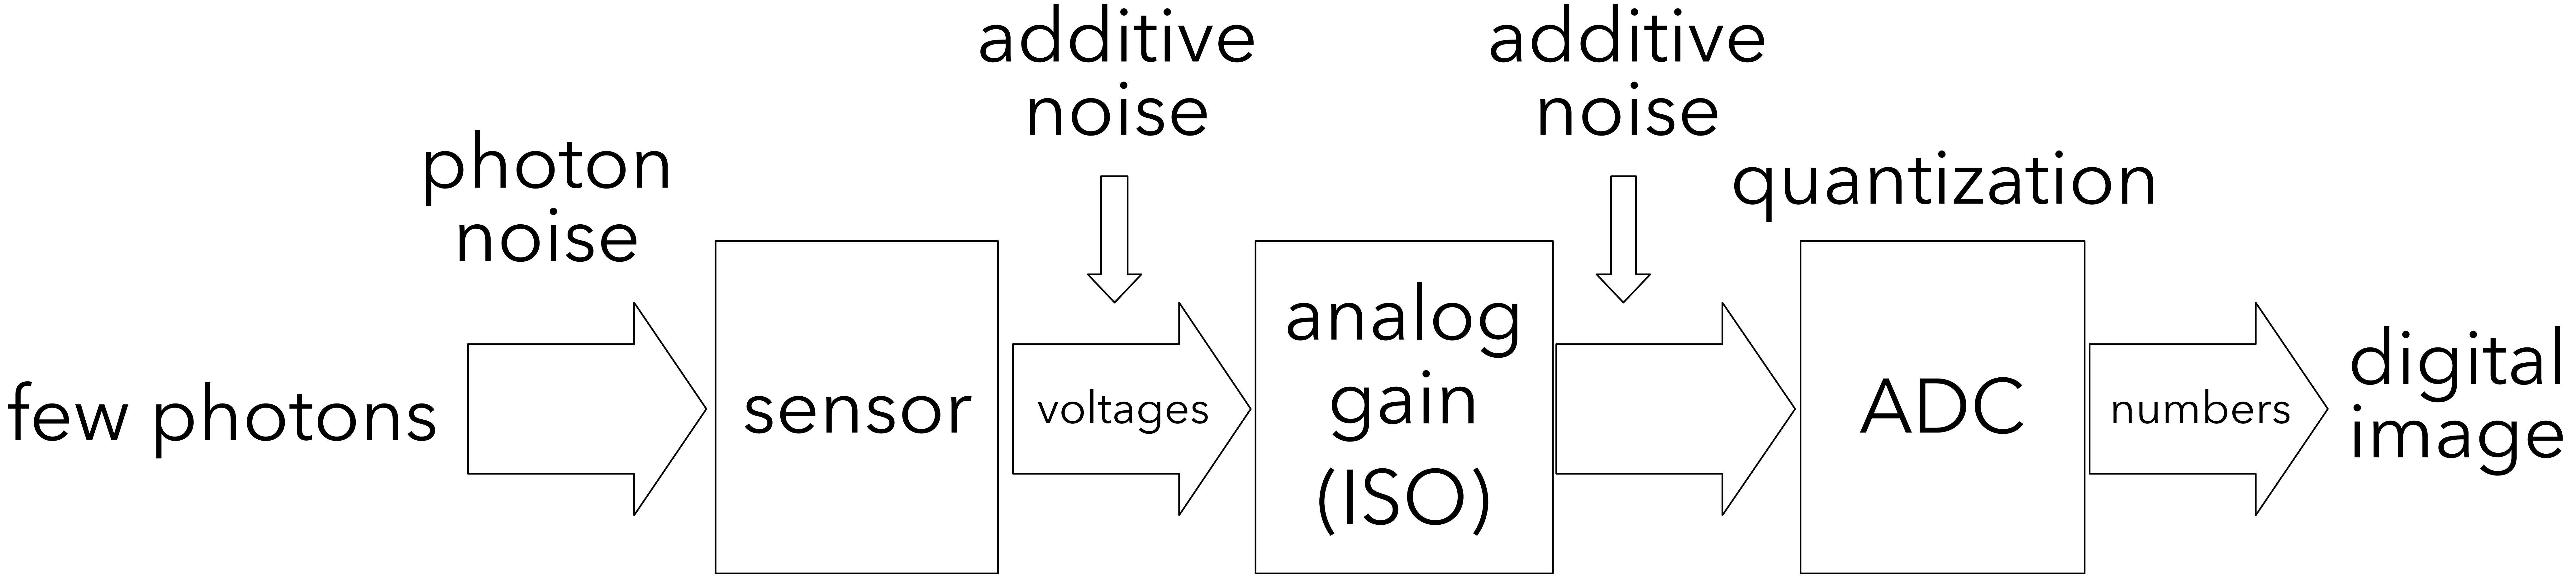

After a slide by Frédo Durand

ISO recap

ISO is a simple gain

- amplifies noise as well

But when the signal is low, it’s better to amplify as early as possible (ISO rather than digitally)

Ideally, you make sure the signal is high by use a slower exposure or larger aperture

After a slide by Frédo Durand

Brain teaser

For the same light level and electronic (per photosite read noise), and same total sensor size, is it better to:

- have a 16 Mpixel sensor and average groups of 4 pixels to yield a 4 Mpixel image
- have a 4Mpixel sensor (with bigger photosites)

Analyze photon noise and read noise

- careful about adding vs. averaging

After a slide by Frédo Durand

Different regimes

For bright pixels (in fact, most pixels), photon noise dominates

For dark pixels: electronic (read) noise dominates

For long exposures, thermal noise kicks in

After a slide by Frédo Durand

Questions?

Slide credits

Frédo Durand

Marc Levoy