

Noise



15-463, 15-663, 15-862
Computational Photography
Fall 2020, Lecture 7

Course announcements

- Homework 2 is out.
 - Due October 2nd.
 - Requires camera *and* tripod.
 - Start early! *Substantially* larger programming and imaging components than in Homework 1.
 - Generous bonus component, up to 50% extra credit.
 - Homework 2 (and many later homeworks) requires submitting large files, so be mindful of that.
 - Do not leave uploading your Homework 2 solution for the last minute!
 - No really: start early!

Overview of today's lecture

- Leftover from lecture 6: other aspects of HDR imaging.
- Leftover from lecture 6: tonemapping.
- Leftover from lecture 6: Some notes about HDR imaging and tonemapping.
- A few motivating examples.
- Sensor noise.
- Noise calibration.
- Optimal weights for HDR.

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).

A few motivating questions from things we've seen

Side-effects of increasing ISO

Image becomes very grainy because noise is amplified.

- Why does increasing ISO increase noise?



ISO 80



ISO 800




ISO 1600

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.

- Why is this not a problem with multi-exposure HDR?

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.



Note: many possible weighting schemes

Many possible weighting schemes

“Confidence” that pixel is noisy/clipped

- What are the optimal weights for merging an exposure stack?

$$w_{ij} = \exp\left(-4 \frac{(I_{lin_{ij}} - 0.5)^2}{0.5^2}\right)$$




Sensor noise

A quick note

- We will only consider per-pixel noise.
- We will not consider cross-pixel noise effects (blooming, smearing, cross-talk, and so on).


Noise in images

Results in “grainy” appearance.



The (in-camera) image processing pipeline


Which part introduces noise?




The (in-camera) image processing pipeline

Which part introduces noise?

- Noise is introduced in the green part.

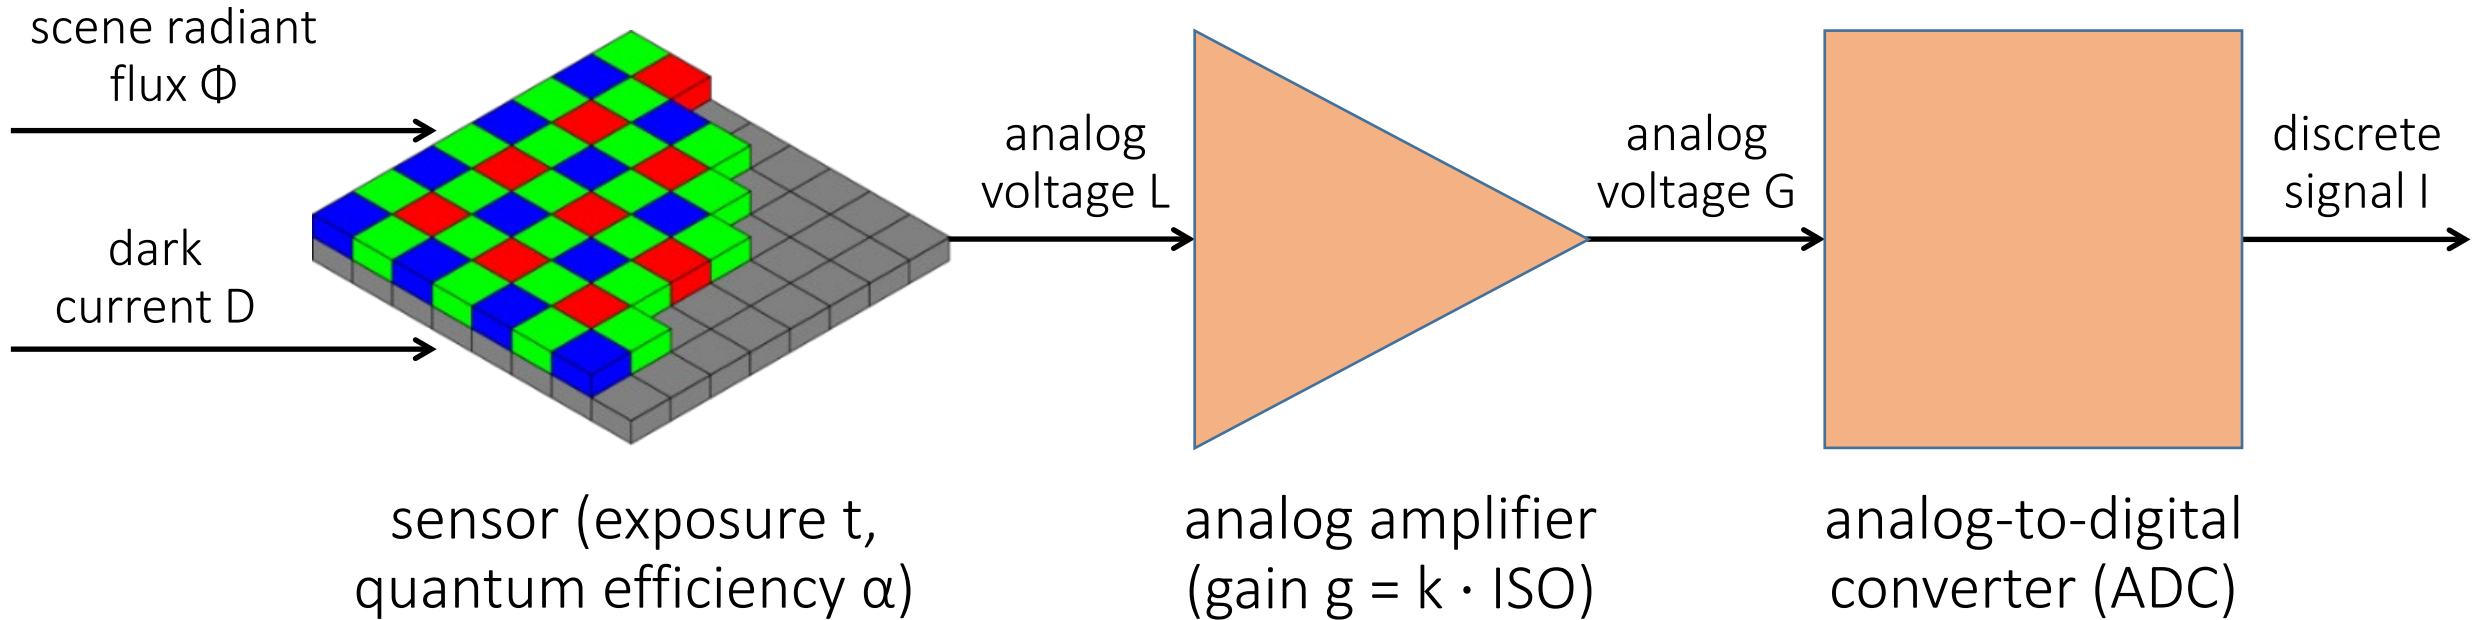


The noisy image formation process

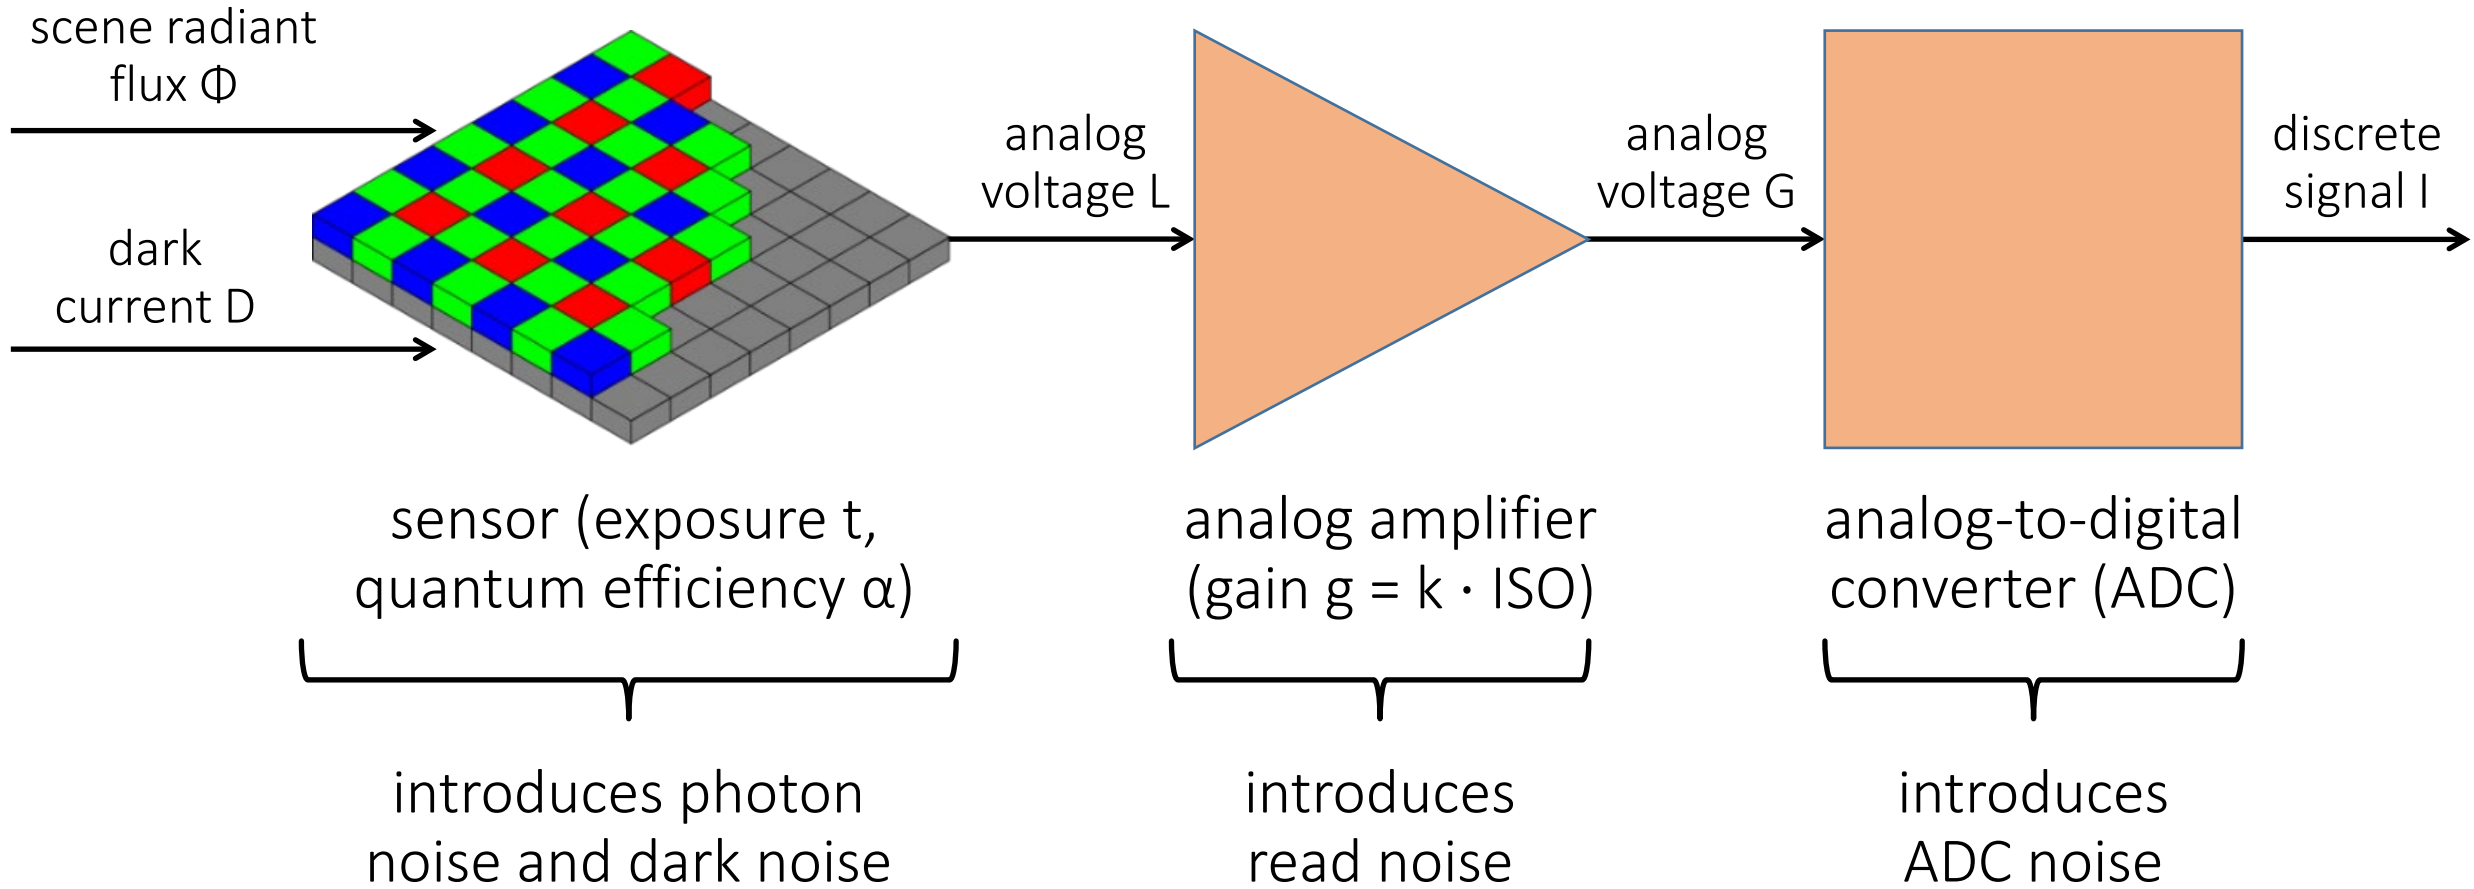


What are the various parts?

The noisy image formation process



The noisy image formation process



- We will be ignoring saturation, but it can be modeled using a clipping operation.

Poisson distribution

Is it a continuous or discrete probability distribution?

Poisson distribution

Is it a continuous or discrete probability distribution?

- It is discrete.

How many parameters does it depend on?

Poisson distribution

Is it a continuous or discrete probability distribution?

- It is discrete.

How many parameters does it depend on?

- One parameter, the *rate* λ .

What is its probability mass function?

Poisson distribution

Is it a continuous or discrete probability distribution?

- It is discrete.


How many parameters does it depend on?

- One parameter, the *rate* λ .

What is its probability mass function?

$$N \sim \text{Poisson}(\lambda) \Leftrightarrow P(N = k; \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$

What is its mean and variance?



Poisson distribution

Is it a continuous or discrete probability distribution?

- It is discrete.

How many parameters does it depend on?

- One parameter, the *rate* λ .


What is its probability mass function?

$$N \sim \text{Poisson}(\lambda) \Leftrightarrow P(N = k; \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$

What is its mean and variance?

- Mean: $\mu(N) = \lambda$
 - Variance: $\sigma(N)^2 = \lambda$
- } The mean and variance of a Poisson random variable both equal the rate λ .

What is the distribution of the sum of two Poisson random variables?



Poisson distribution

Is it a continuous or discrete probability distribution?

- It is discrete.

How many parameters does it depend on?

- One parameter, the *rate* λ .

What is its probability mass function?


$$N \sim \text{Poisson}(\lambda) \Leftrightarrow P(N = k; \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$

What is its mean and variance?

- Mean: $\mu(N) = \lambda$
 - Variance: $\sigma(N)^2 = \lambda$
- } The mean and variance of a Poisson random variable both equal the rate λ .

What is the distribution of the sum of two Poisson random variables?

$$N_1 \sim \text{Poisson}(\lambda_1), N_2 \sim \text{Poisson}(\lambda_2) \Rightarrow N_1 + N_2 \sim \text{Poisson}(\lambda_1 + \lambda_2)$$



Photon noise


A consequence of the discrete (quantum) nature of light.

- Photon detections are independent random events.
- Total number of detections is Poisson distributed.
- Also known as shot noise and Schott noise.

photon noise depends on
scene flux and exposure

$$N_{\text{detections}} \sim \text{Poisson}[t \cdot \alpha \cdot \Phi]$$

simulated mean
#photons/pixel



Dark noise

A consequence of “phantom detections” by the sensor.

- Electrons are randomly released without any photons.
- Total number of detections is Poisson distributed.
- Increases *exponentially* with sensor temperature (+6°C ≈ doubling).

dark noise depends on exposure but *not* on scene

$$N_{\text{detections}} \sim \text{Poisson}(t \cdot D)$$

Can you think of examples when dark noise is important?

Dark noise

A consequence of “phantom detections” by the sensor.

- Electrons are randomly released without any photons.
- Total number of detections is Poisson distributed.
- Increases *exponentially* with sensor temperature (+6°C ≈ doubling).

dark noise depends on exposure but *not* on scene

$$N_{\text{detections}} \sim \text{Poisson}[t \cdot D]$$

Can you think of examples when dark noise is important?

- Very long exposures (astrophotography, pinhole camera).

Can you think of ways to mitigate dark noise?



Dark noise

A consequence of “phantom detections” by the sensor.

- Electrons are randomly released without any photons.
- Total number of detections is Poisson distributed.
- Increases *exponentially* with sensor temperature (+6°C ≈ doubling).

dark noise depends on exposure but *not* on scene

$$N_{\text{detections}} \sim \text{Poisson}[t \cdot D]$$

Can you think of examples when dark noise is important?

- Very long exposures (astrophotography, pinhole camera).

Can you think of ways to mitigate dark noise?


- Cool the sensor.



Fundamental question

Why are photon noise and dark noise Poisson random variables?

The noisy image formation process



- What is the distribution of the sensor readout L ?

The distribution of the sensor readout

We know that the sensor readout is the sum of all released electrons:

$$L = N_{\text{photon_detections}} + N_{\text{phantom_detections}}$$

What is the distribution of photon detections?

The distribution of the sensor readout

We know that the sensor readout is the sum of all released electrons:

$$L = N_{\text{photon_detections}} + N_{\text{phantom_detections}}$$

What is the distribution of photon detections?

$$N_{\text{photon_detections}} \sim \text{Poisson}(t \cdot \alpha \cdot \Phi)$$

What is the distribution of phantom detections?

The distribution of the sensor readout

We know that the sensor readout is the sum of all released electrons:

$$L = N_{\text{photon_detections}} + N_{\text{phantom_detections}}$$

What is the distribution of photon detections?

$$N_{\text{photon_detections}} \sim \text{Poisson}(t \cdot \alpha \cdot \Phi)$$

What is the distribution of phantom detections?

$$N_{\text{phantom_detections}} \sim \text{Poisson}(t \cdot D)$$

What is the distribution of the sensor readout?

The distribution of the sensor readout

We know that the sensor readout is the sum of all released electrons:

$$L = N_{\text{photon_detections}} + N_{\text{phantom_detections}}$$

What is the distribution of photon detections?

$$N_{\text{photon_detections}} \sim \text{Poisson}(t \cdot \alpha \cdot \Phi)$$


What is the distribution of phantom detections?

$$N_{\text{phantom_detections}} \sim \text{Poisson}(t \cdot D)$$

What is the distribution of the sensor readout?

$$L \sim \text{Poisson}(t \cdot (\alpha \cdot \Phi + D))$$

The noisy image formation process



Read and ADC noise

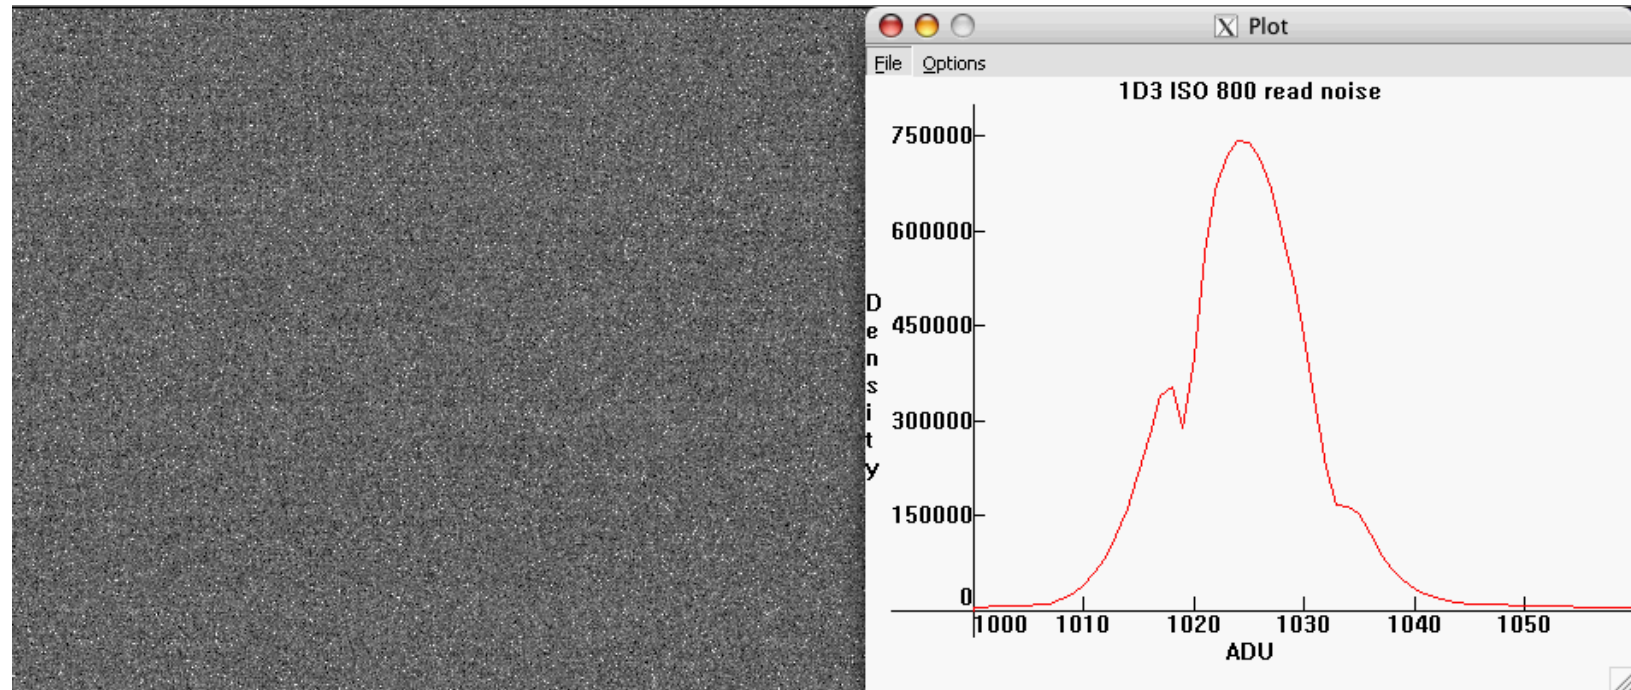
A consequence of random voltage fluctuations before and after amplifier.

- Both are independent of scene and exposure.
- Both are normally (zero-mean Gaussian) distributed.
- ADC noise includes quantization errors.

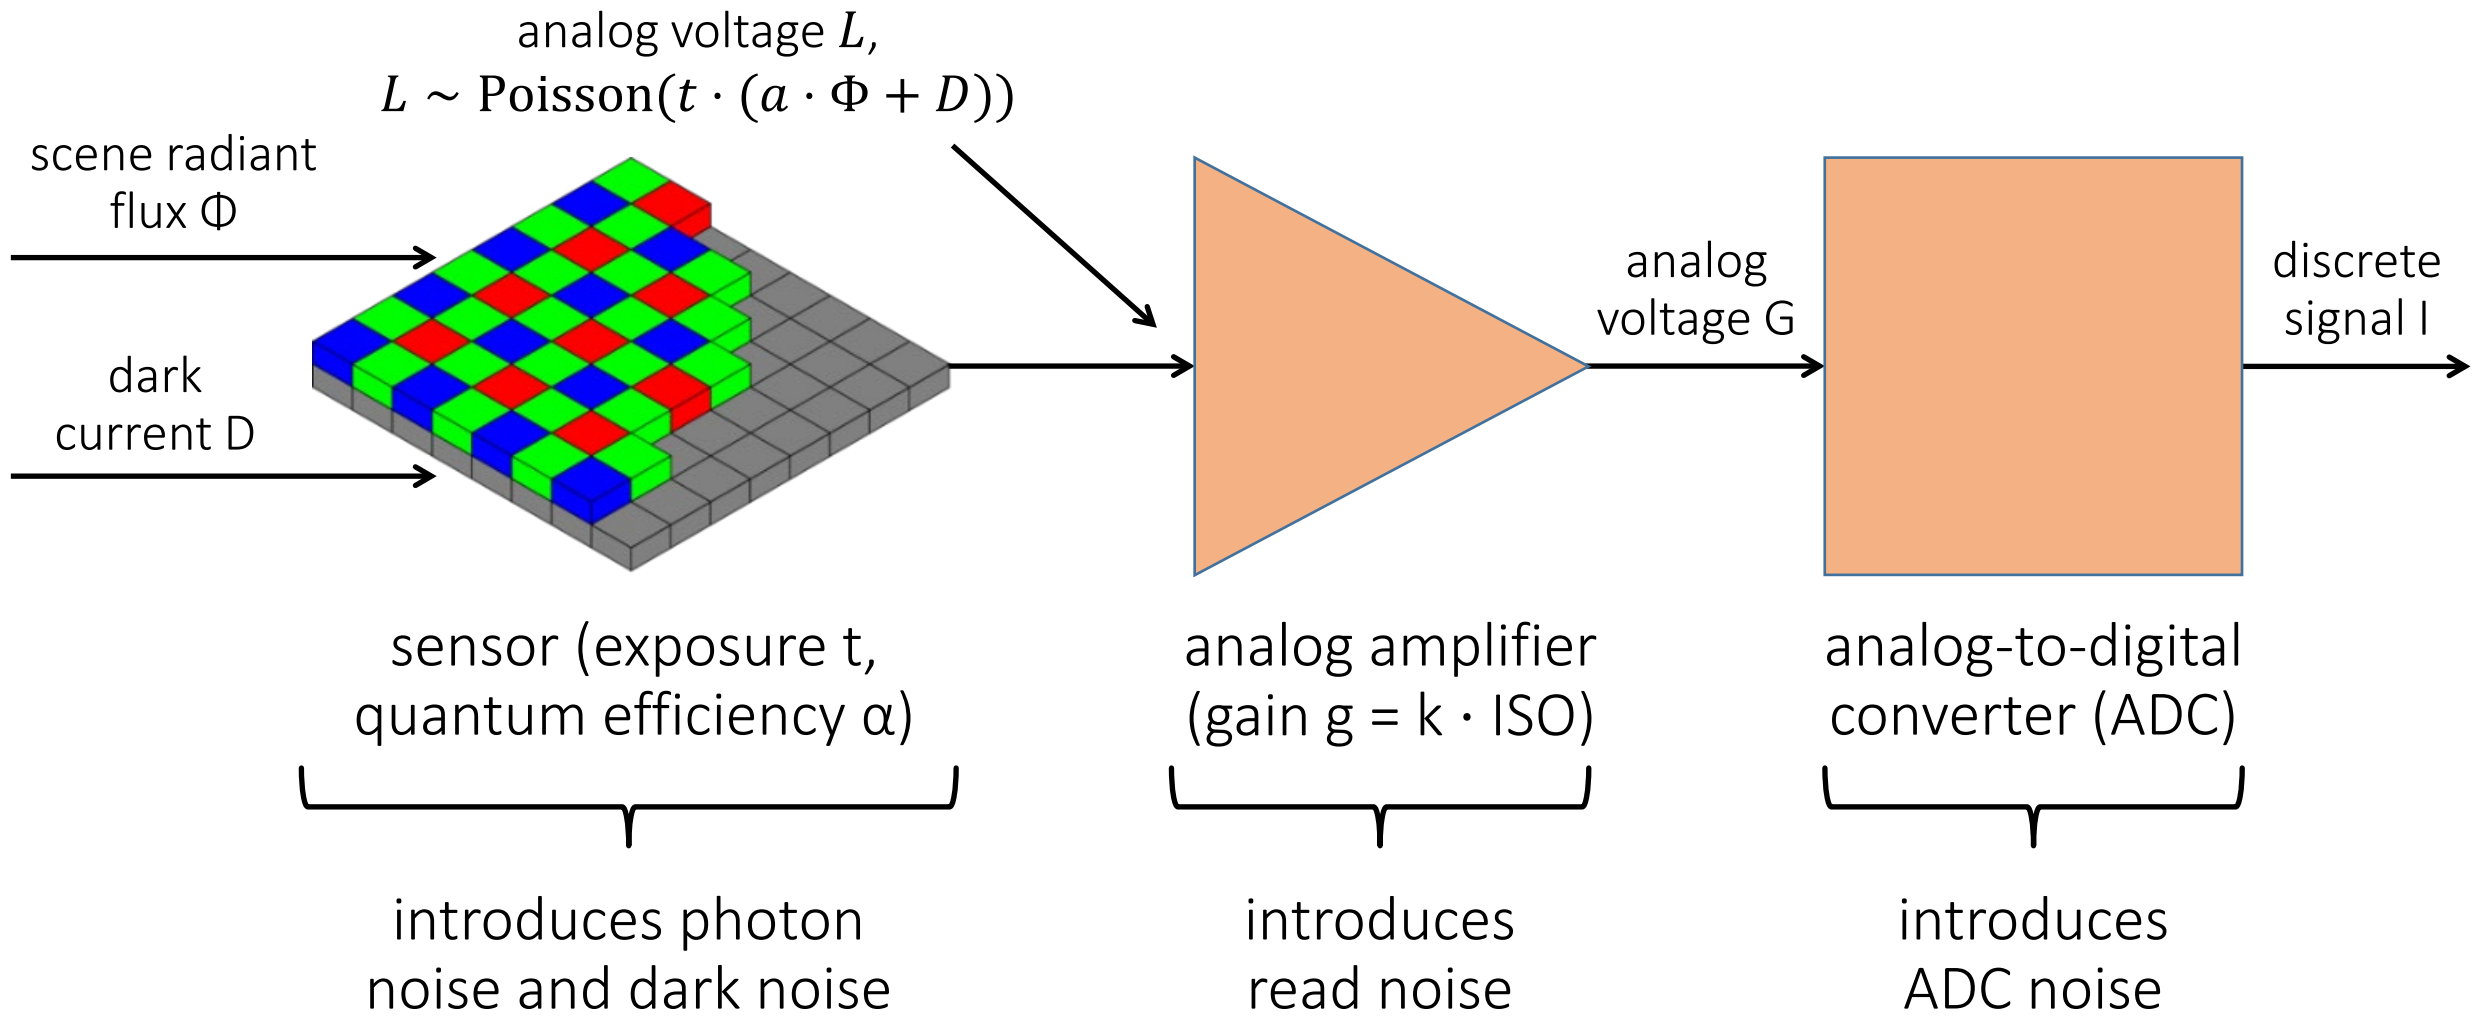
$$n_{\text{read}} \sim \text{Normal}(0, \sigma_{\text{read}})$$

$$n_{\text{ADC}} \sim \text{Normal}(0, \sigma_{\text{ADC}})$$

Very important for dark pixels.



The noisy image formation process



- How can we express the voltage G and discrete intensity I ?

Expressions for the amplifier and ADC outputs

Both read noise and ADC noise are *additive* and *zero-mean*.

- How can we express the output of the amplifier?

Expressions for the amplifier and ADC outputs

Both read noise and ADC noise are *additive* and *zero-mean*.

- How can we express the output of the amplifier?

$$G = L \cdot g + n_{\text{read}} \cdot g$$

← don't forget to account for the ISO-dependent gain

- How can we express the output of the ADC?

Expressions for the amplifier and ADC outputs

Both read noise and ADC noise are *additive* and *zero-mean*.

- How can we express the output of the amplifier?

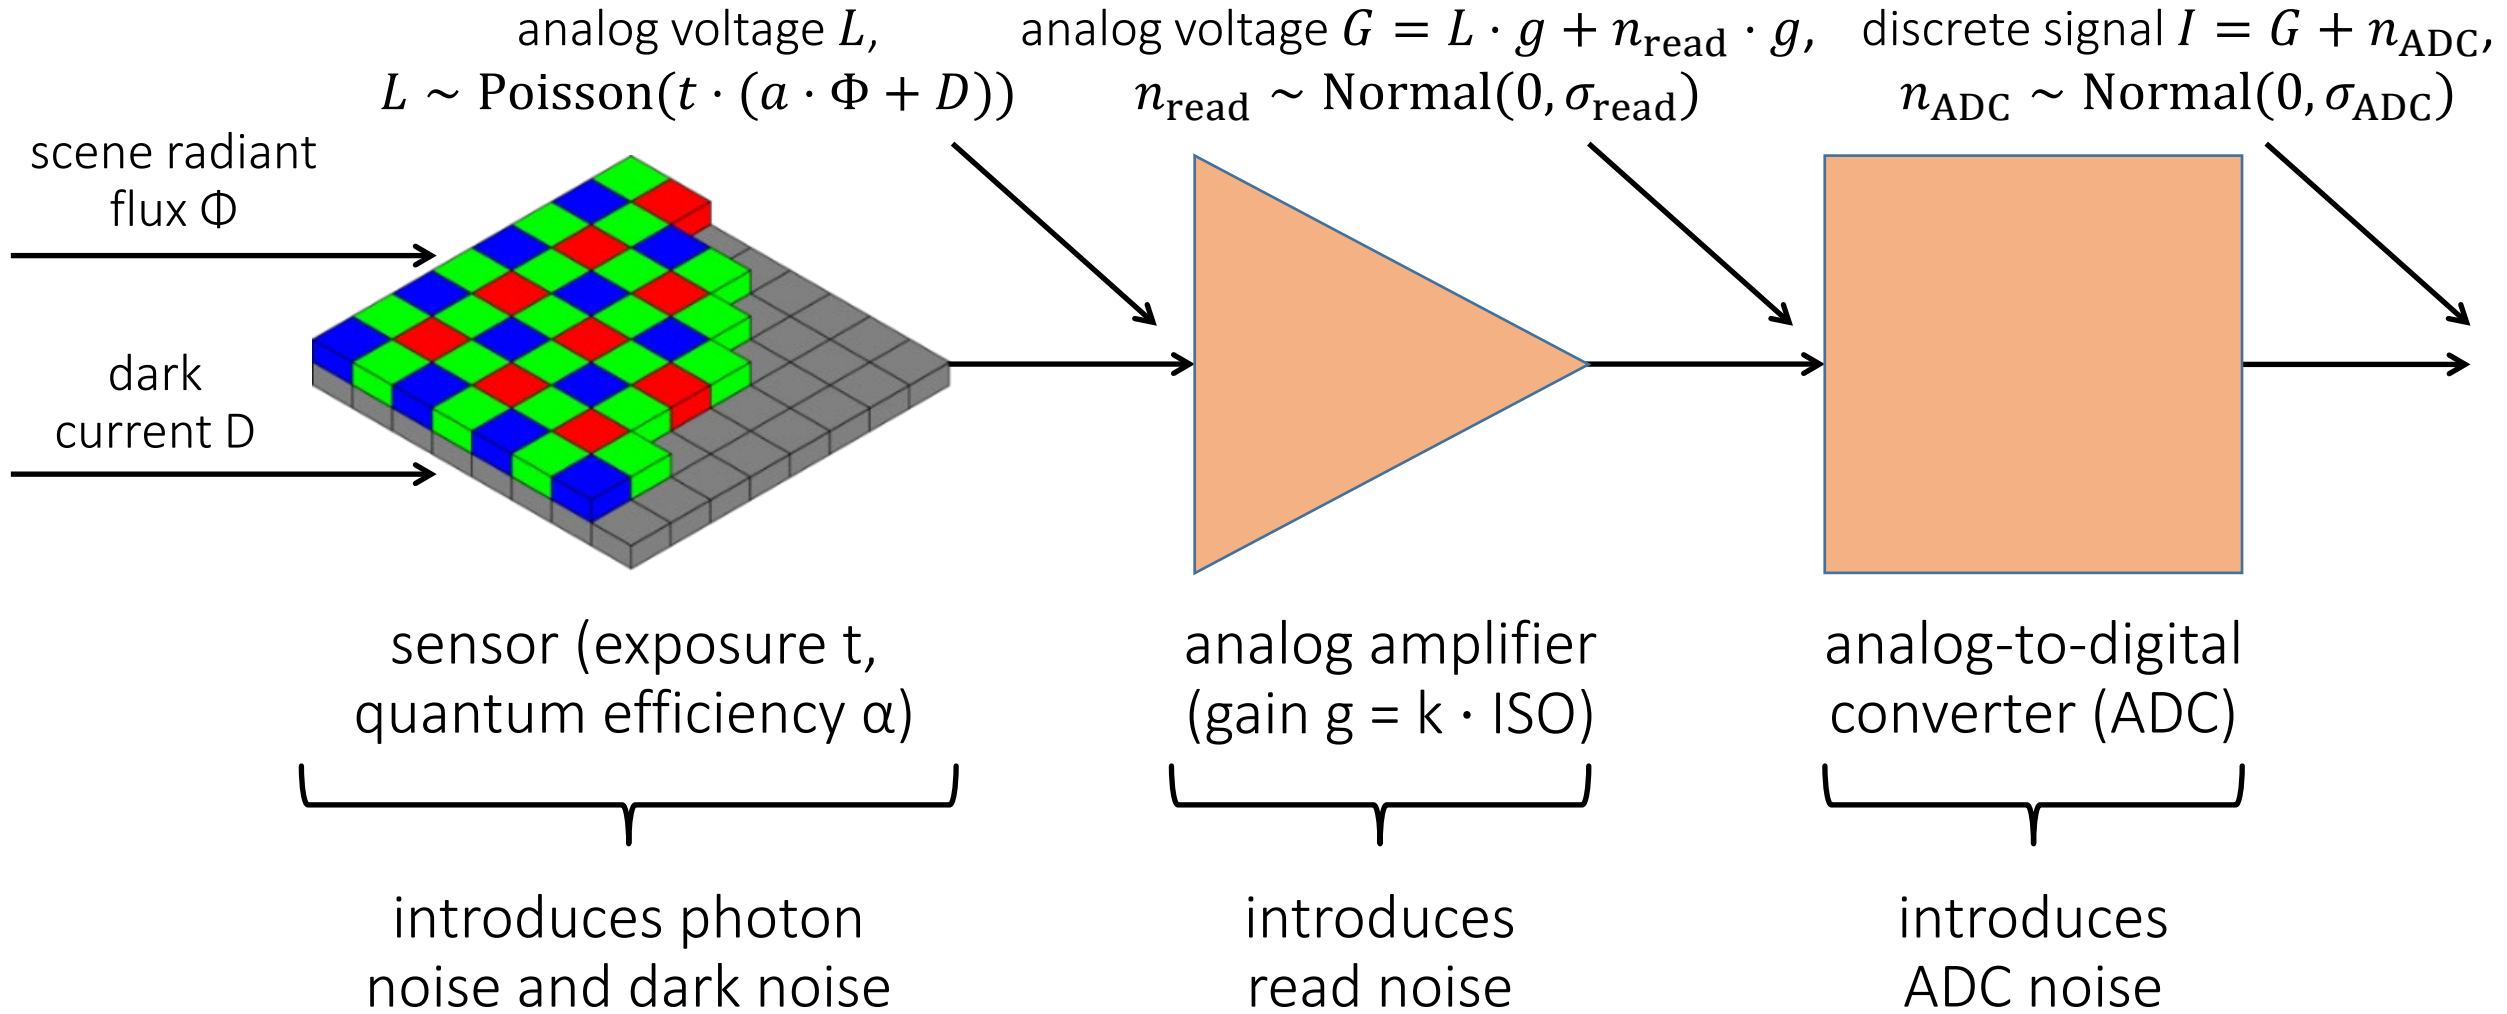
$$G = L \cdot g + n_{\text{read}} \cdot g$$

← don't forget to account for the ISO-dependent gain

- How can we express the output of the ADC?

$$I = G + n_{\text{ADC}}$$

The noisy image formation process



Putting it all together

Without saturation, the digital intensity equals:

$$I = L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}} \quad \text{where}$$

$$L \sim \text{Poisson}(t \cdot (a \cdot \Phi + D))$$

$$n_{\text{read}} \sim \text{Normal}(0, \sigma_{\text{read}})$$

$$n_{\text{ADC}} \sim \text{Normal}(0, \sigma_{\text{ADC}})$$

What is the mean of the digital intensity (assuming no saturation)?

$$E(I) =$$

Putting it all together

Without saturation, the digital intensity equals:

$$I = L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}} \quad \text{where}$$

$$L \sim \text{Poisson}(t \cdot (a \cdot \Phi + D))$$

$$n_{\text{read}} \sim \text{Normal}(0, \sigma_{\text{read}})$$

$$n_{\text{ADC}} \sim \text{Normal}(0, \sigma_{\text{ADC}})$$

What is the mean of the digital intensity (assuming no saturation)?

$$\begin{aligned} E(I) &= E(L \cdot g) + E(n_{\text{read}} \cdot g) + E(n_{\text{ADC}}) \\ &= \end{aligned}$$

Putting it all together

Without saturation, the digital intensity equals:

$$I = L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}} \quad \text{where}$$

$$L \sim \text{Poisson}(t \cdot (a \cdot \Phi + D))$$

$$n_{\text{read}} \sim \text{Normal}(0, \sigma_{\text{read}})$$

$$n_{\text{ADC}} \sim \text{Normal}(0, \sigma_{\text{ADC}})$$

What is the mean of the digital intensity (assuming no saturation)?

$$\begin{aligned} E(I) &= E(L \cdot g) + E(n_{\text{read}} \cdot g) + E(n_{\text{ADC}}) \\ &= t \cdot (a \cdot \Phi + D) \cdot g \end{aligned}$$

What is the variance of the digital intensity (assuming no saturation)?

$$\sigma(I)^2 =$$

Putting it all together

Without saturation, the digital intensity equals:

$$I = L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}} \quad \text{where}$$

$$L \sim \text{Poisson}(t \cdot (a \cdot \Phi + D))$$

$$n_{\text{read}} \sim \text{Normal}(0, \sigma_{\text{read}})$$

$$n_{\text{ADC}} \sim \text{Normal}(0, \sigma_{\text{ADC}})$$

What is the mean of the digital intensity (assuming no saturation)?

$$\begin{aligned} E(I) &= E(L \cdot g) + E(n_{\text{read}} \cdot g) + E(n_{\text{ADC}}) \\ &= t \cdot (a \cdot \Phi + D) \cdot g \end{aligned}$$

What is the variance of the digital intensity (assuming no saturation)?

$$\begin{aligned} \sigma(I)^2 &= \sigma(L \cdot g)^2 + \sigma(n_{\text{read}} \cdot g)^2 + \sigma(n_{\text{ADC}})^2 \\ &= \end{aligned}$$

Putting it all together

Without saturation, the digital intensity equals:

$$I = L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}} \quad \text{where}$$

$$L \sim \text{Poisson}(t \cdot (a \cdot \Phi + D))$$

$$n_{\text{read}} \sim \text{Normal}(0, \sigma_{\text{read}})$$

$$n_{\text{ADC}} \sim \text{Normal}(0, \sigma_{\text{ADC}})$$

What is the mean of the digital intensity (assuming no saturation)?

$$\begin{aligned} E(I) &= E(L \cdot g) + E(n_{\text{read}} \cdot g) + E(n_{\text{ADC}}) \\ &= t \cdot (a \cdot \Phi + D) \cdot g \end{aligned}$$

What is the variance of the digital intensity (assuming no saturation)?

$$\begin{aligned} \sigma(I)^2 &= \sigma(L \cdot g)^2 + \sigma(n_{\text{read}} \cdot g)^2 + \sigma(n_{\text{ADC}})^2 \\ &= t \cdot (a \cdot \Phi + D) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \end{aligned}$$

How do we compute mean and variance in practice?

How do we compute mean and variance in practice?

Mean: capture multiple *linear* images with identical settings and average.

$$\bar{I} = \frac{1}{N} \sum_{n=1}^N I_n \xrightarrow{N \rightarrow \infty} E(I)$$

How do we compute mean and variance in practice?

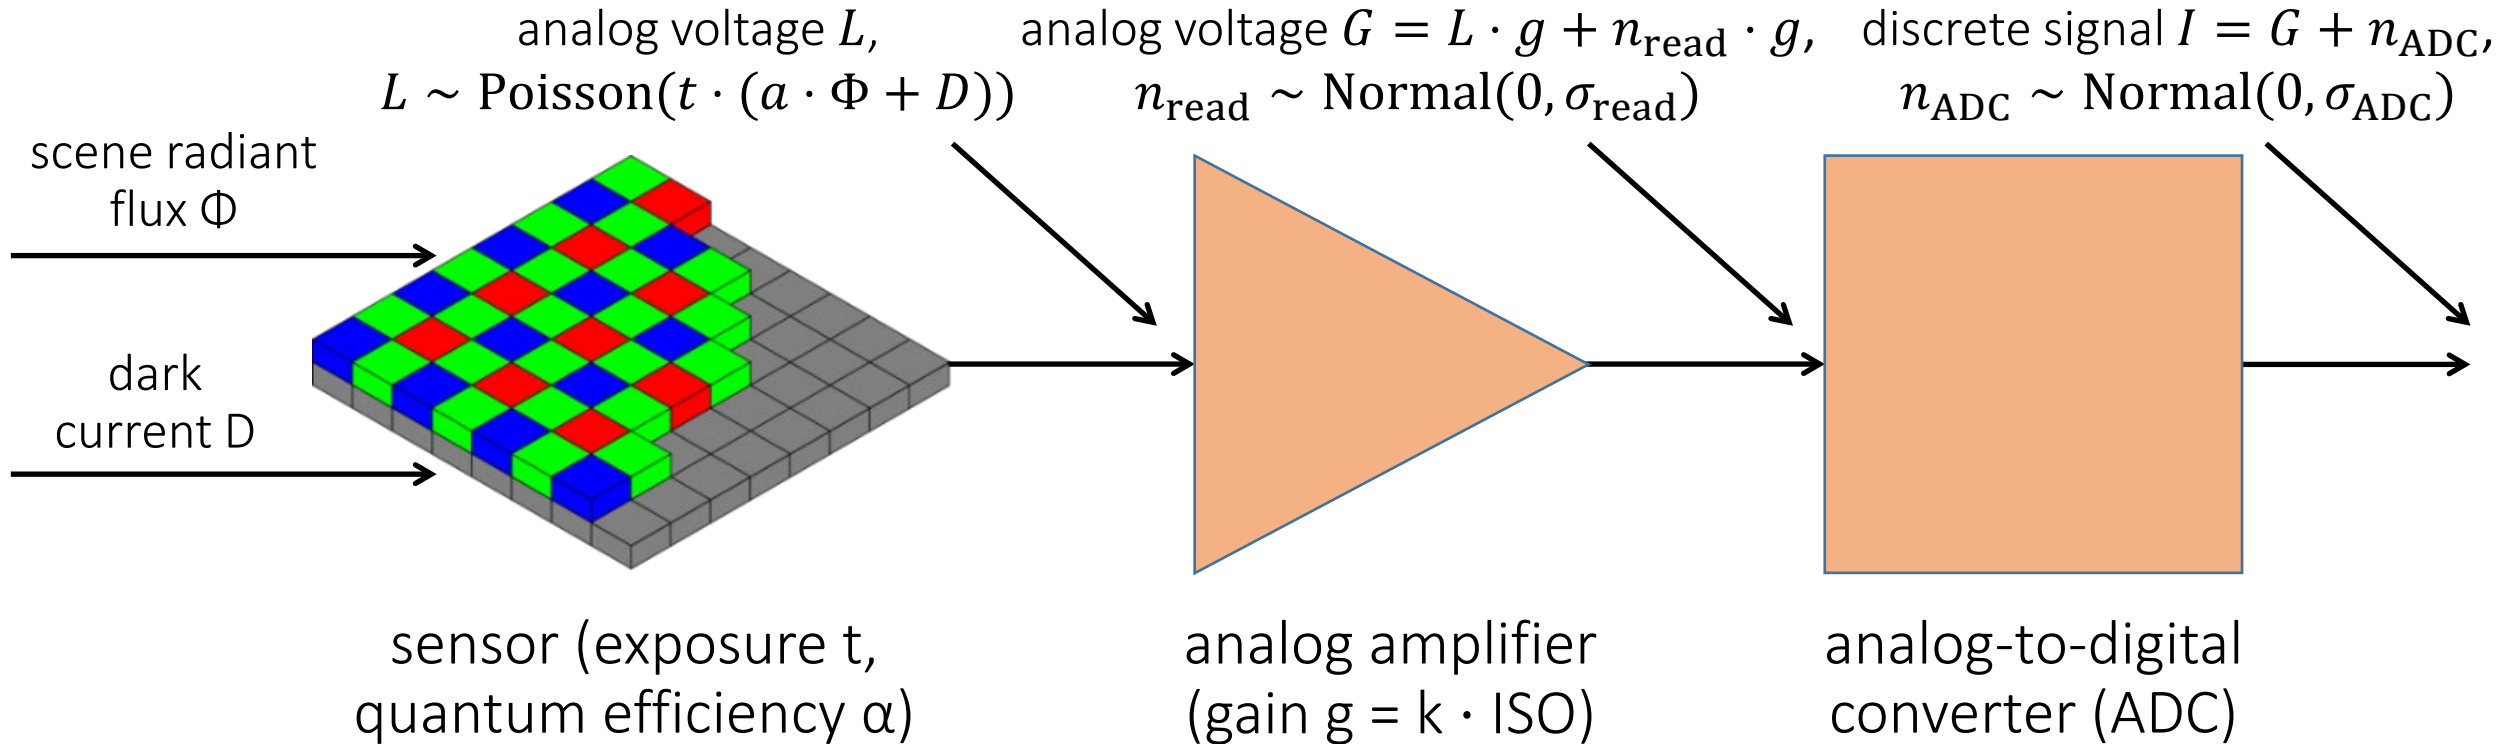
Mean: capture multiple *linear* images with identical settings and average.

$$\bar{I} = \frac{1}{N} \sum_{n=1}^N I_n \xrightarrow{N \rightarrow \infty} E(I)$$

Variance: capture multiple *linear* images with identical settings and form variance estimator.

$$\bar{\Sigma} = \frac{1}{N-1} \sum_{n=1}^N (I_n - \bar{I})^2 \xrightarrow{N \rightarrow \infty} \sigma(I)^2$$

The noisy image formation process



discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}}, I_{\text{max}})$$

saturation level \nearrow

intensity mean and variance (without saturation):

$$E(I) = t \cdot (\alpha \cdot \Phi + D) \cdot g$$

$$\sigma(I)^2 = t \cdot (\alpha \cdot \Phi + D) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2$$

Affine noise model

Combine read and ADC noise into a single *additive* noise term:

$$I = L \cdot g + n_{\text{add}} \quad \text{where} \quad n_{\text{add}} = n_{\text{read}} \cdot g + n_{\text{ADC}}$$

What is the distribution of the additive noise term?

Affine noise model

Combine read and ADC noise into a single *additive* noise term:


$$I = L \cdot g + n_{\text{add}} \quad \text{where} \quad n_{\text{add}} = n_{\text{read}} \cdot g + n_{\text{ADC}}$$

What is the distribution of the additive noise term?

- Sum of two independent, normal random variables.

$$n_{\text{add}} \sim \text{Normal}(0, \sqrt{\sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2})$$

Affine noise model



discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{add}}, I_{\text{max}})$$

intensity mean and variance (without saturation):

$$E(I) = t \cdot (a \cdot \Phi + D) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi + D) \cdot g^2 + \sigma_{\text{add}}^2$$

Some observations

Is image intensity an *unbiased* estimator of (scaled) scene radiant flux?

Some observations

Is image intensity an *unbiased* estimator of (scaled) scene radiant flux?

- No, because of dark noise (term $t \cdot D \cdot g$ in the mean).
- Averaging multiple images cancels out read and ADC noise, but *not* dark noise.

When are photon noise and additive noise dominant?

Some observations

Is image intensity an *unbiased* estimator of (scaled) scene radiant flux?

- No, because of dark noise (term $t \cdot D \cdot g$ in the mean).
- Averaging multiple images cancels out read and ADC noise, but *not* dark noise.

When are photon noise and additive noise dominant?

- Photon noise is dominant in very bright scenes.
- Additive noise is dominant in very dark scenes.

Can we ever completely remove noise?

Some observations

Is image intensity an *unbiased* estimator of (scaled) scene radiant flux?

- No, because of dark noise (term $t \cdot D \cdot g$ in the mean).
- Averaging multiple images cancels out read and ADC noise, but *not* dark noise.

When are photon noise and additive noise dominant?

- Photon noise is dominant in very bright scenes.
- Additive noise is dominant in very dark scenes.

Can we ever completely remove noise?

- We cannot eliminate photon noise.
- Super-sensitive detectors have pure Poisson photon noise.

single-photon avalanche photodiode (SPAD)



Summary: noise regimes

<u>regime</u>	<u>dominant noise</u>	<u>notes</u>
bright pixels	photon noise	scene-dependent
dark pixels	read and ADC noise	scene-independent
low ISO	ADC noise	post-gain
high ISO	photon and read noise	pre-gain
long exposures	dark noise	thermal dependence

discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}}, I_{\text{max}})$$

intensity mean and variance (without saturation):

$$E(I) = t \cdot (a \cdot \Phi + D) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi + D) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2$$

Summary: noise regimes

<u>regime</u>	<u>dominant noise</u>	<u>notes</u>
bright pixels	photon noise	scene-dependent
dark pixels	read and ADC noise	scene-independent
low ISO	ADC noise	post-gain
high ISO	photon and read noise	pre-gain
long exposures	dark noise	thermal dependence

Does this mean that using high exposure makes images more “noisy”?

discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}}, I_{\text{max}})$$

intensity mean and variance (without saturation):

$$E(I) = t \cdot (a \cdot \Phi + D) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi + D) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2$$

Signal-to-noise ratio

Variance versus signal-to-noise ratio

Variance?

Variance versus signal-to-noise ratio

Variance is an *absolute* measure of the (squared) magnitude of noise:

$$\sigma(I)^2 = E \left((I - E(I))^2 \right) = E(I^2) - E(I)^2$$

Signal-to-noise ratio (SNR)?

Variance versus signal-to-noise ratio

Variance is an *absolute* measure of the (squared) magnitude of noise:

$$\sigma(I)^2 = E \left((I - E(I))^2 \right) = E(I^2) - E(I)^2$$

Signal-to-noise ratio (SNR) is a *relative* measure of the (inverse squared) magnitude of noise:

$$\text{SNR} = \frac{E(I)^2}{\sigma(I)^2}$$

When noise *decreases*:

- The variance...
- The SNR...

Variance versus signal-to-noise ratio

Variance is an *absolute* measure of the (squared) magnitude of noise:

$$\sigma(I)^2 = E \left((I - E(I))^2 \right) = E(I^2) - E(I)^2$$

Signal-to-noise ratio (SNR) is a *relative* measure of the (inverse squared) magnitude of noise:

$$\text{SNR} = \frac{E(I)^2}{\sigma(I)^2}$$

When noise *decreases*:

- The variance decreases.
- The SNR increases.

The case of sensor noise

Assuming for simplicity that *there is no dark current*:

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \qquad SNR = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2}$$

What happens when the exposure time or flux (or both) are very large?

The case of sensor noise

Assuming for simplicity that *there is no dark current*:

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \qquad SNR = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2}$$

What happens when the exposure time or flux (or both) are very large?

- We can ignore additive (read and ADC) noise terms.

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 \qquad SNR = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2} = t \cdot a \cdot \Phi$$

What happens when the exposure time or flux (or both) are very small?

The case of sensor noise

Assuming for simplicity that *there is no dark current*:

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \quad \text{SNR} = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2}$$

What happens when the flux or exposure time (or both) are very large?

- We can ignore additive (read and ADC) noise terms.

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 \quad \text{SNR} = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2} = t \cdot a \cdot \Phi$$

What happens when the flux or exposure time (or both) are very small?

- We can ignore scene-dependent noise terms.

$$\sigma(I)^2 = \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \quad \text{SNR} = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{\sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2}$$

The case of sensor noise

Assuming for simplicity that *there is no dark current*:

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \qquad SNR = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2}$$

What happens when the exposure time or flux (or both) are very large?

photon-noise-limited case

- We can ignore additive (read and ADC) noise terms.

$$\sigma(I)^2 = t \cdot a \cdot \Phi \cdot g^2 \qquad SNR = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{t \cdot a \cdot \Phi \cdot g^2} = t \cdot a \cdot \Phi$$

What happens when the exposure time or flux (or both) are very small?

additive-noise-limited case

- We can ignore scene-dependent noise terms.

$$\sigma(I)^2 = \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2 \qquad SNR = \frac{(t \cdot a \cdot \Phi \cdot g)^2}{\sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2}$$

The case of sensor noise

As flux or exposure time (or both) increase:

- The noise variance increases.
- The SNR also increases.

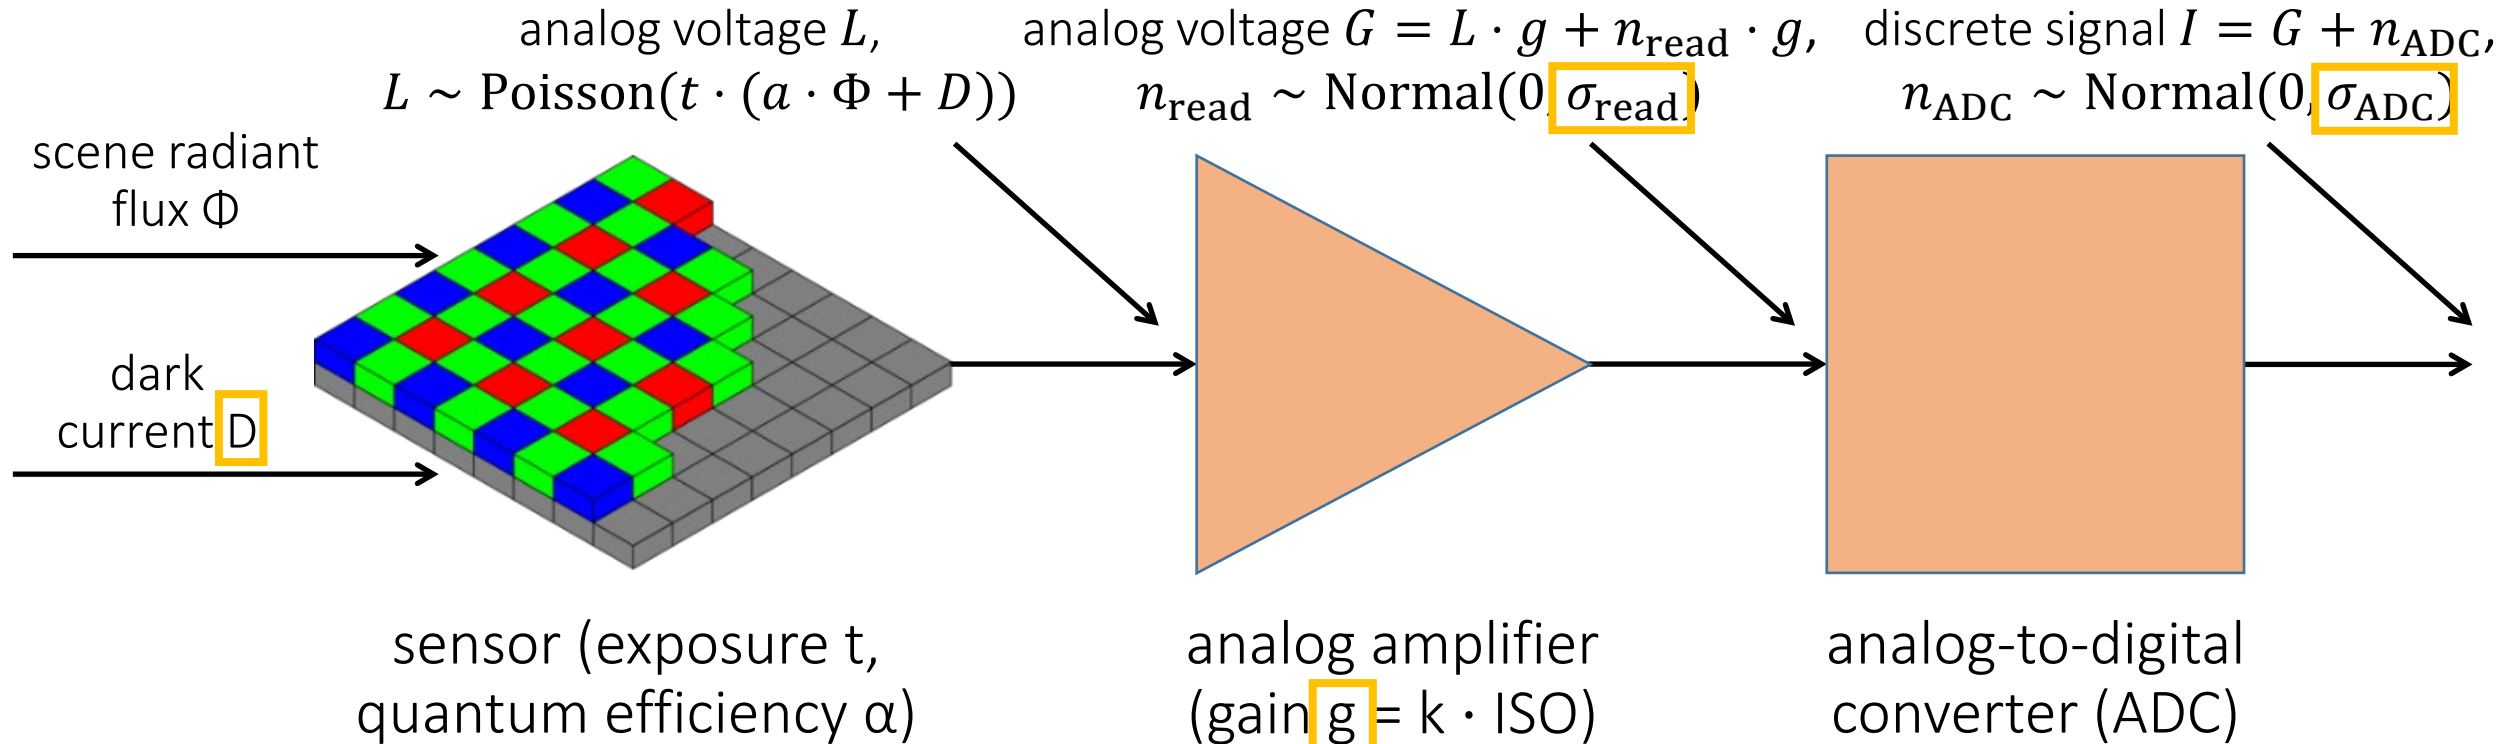
Even though the absolute magnitude of noise increases, its relative magnitude compared to the signal we are measuring decreases.

→ Our measurements become *less noisy* as flux or exposure time (or both) increase.

(For the case of exposure time, we need to be careful to also take into account dark noise.)

Noise calibration

How can we estimate the various parameters?



discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}}, I_{\text{max}})$$

saturation level \nearrow

intensity mean and variance:

$$E(I) = t \cdot (a \cdot \Phi + D) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi + D) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2$$

Estimating the dark current

Can you think of a procedure for estimating the dark current D ?

Estimating the dark current

Can you think of a procedure for estimating the dark current D ?

- Capture multiple images with the sensor completely blocked and average to form the *dark frame*.

Why is the dark frame a valid estimator of the dark current D ?

Estimating the dark current

Can you think of a procedure for estimating the dark current D ?

- Capture multiple images with the sensor completely blocked and average to form the *dark frame*.

Why is the dark frame a valid estimator of the dark current D ?

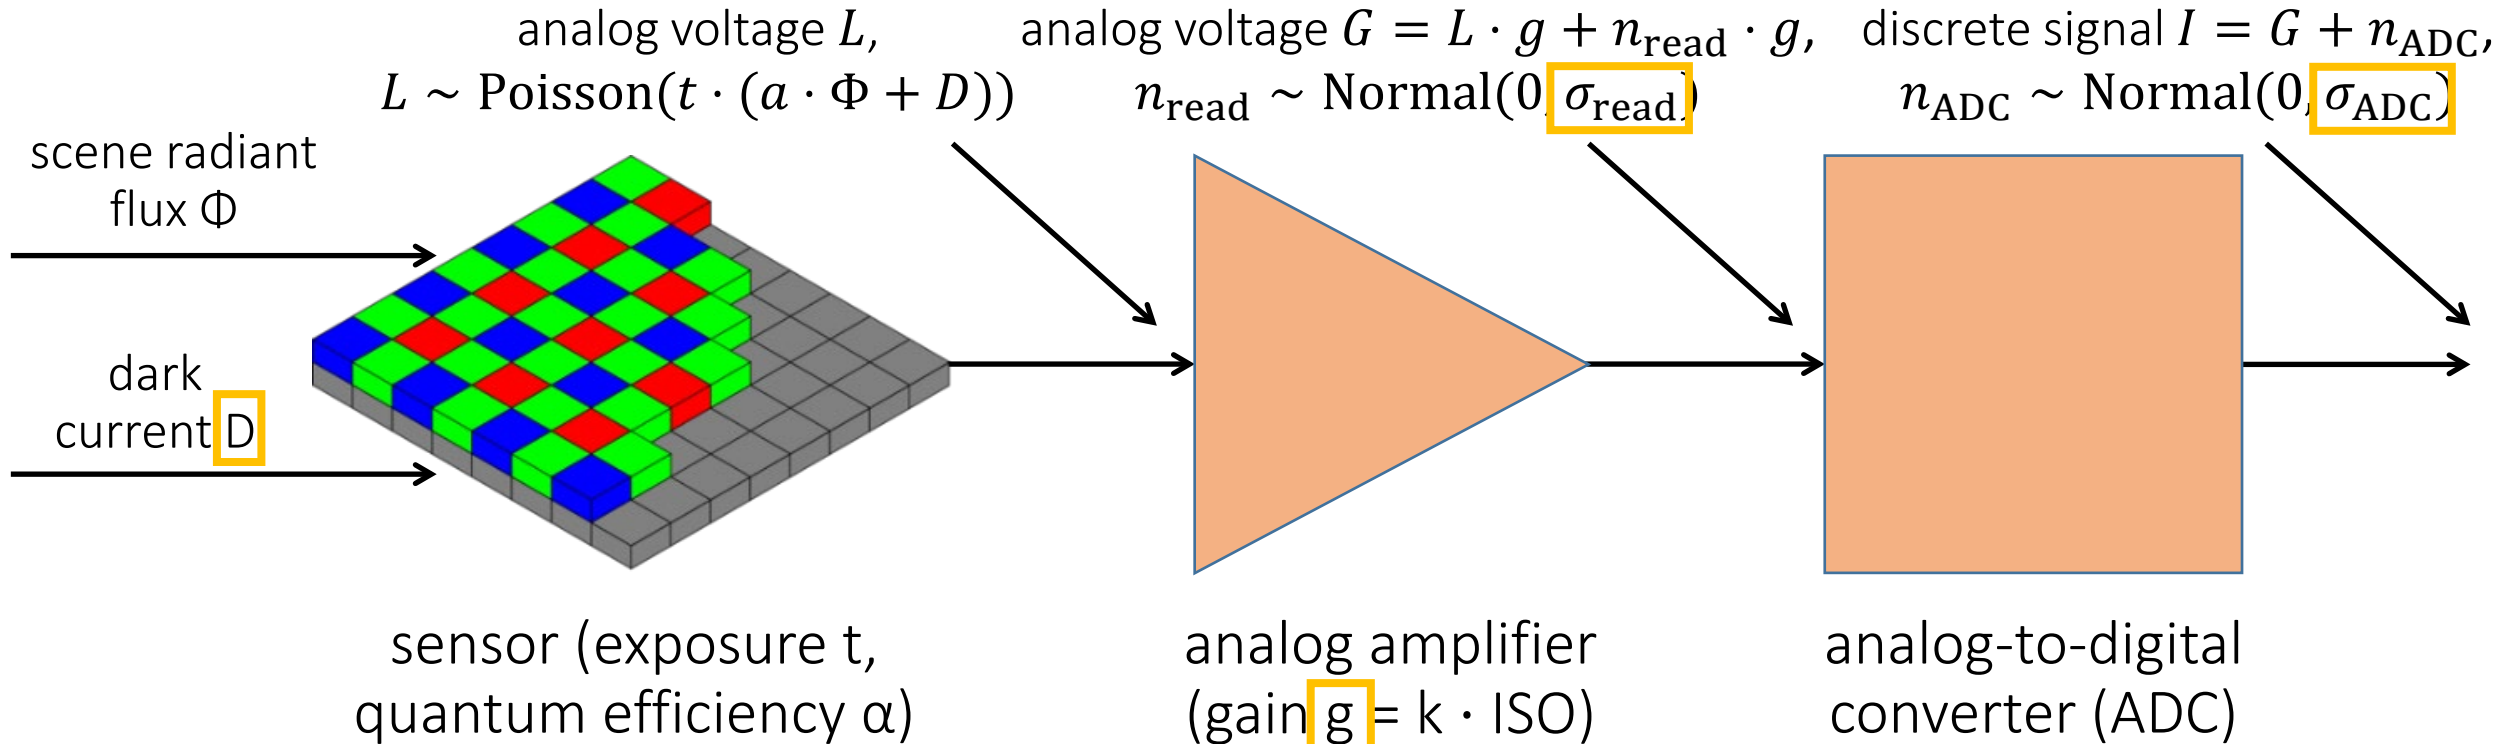
- By blocking the sensor, we effectively set $\Phi = 0$.
- Average intensity becomes:

$$E(I) = t \cdot (a \cdot 0 + D) \cdot g = t \cdot D \cdot g$$

- The dark frame needs to be computed separately for each ISO setting, unless we can also calibrate the gain g .

For the rest of these slides, we assume that we have calibrated D and removed it from captured images (by subtracting from them the dark frame).

Noise model before dark frame subtraction



discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}}, I_{\text{max}})$$

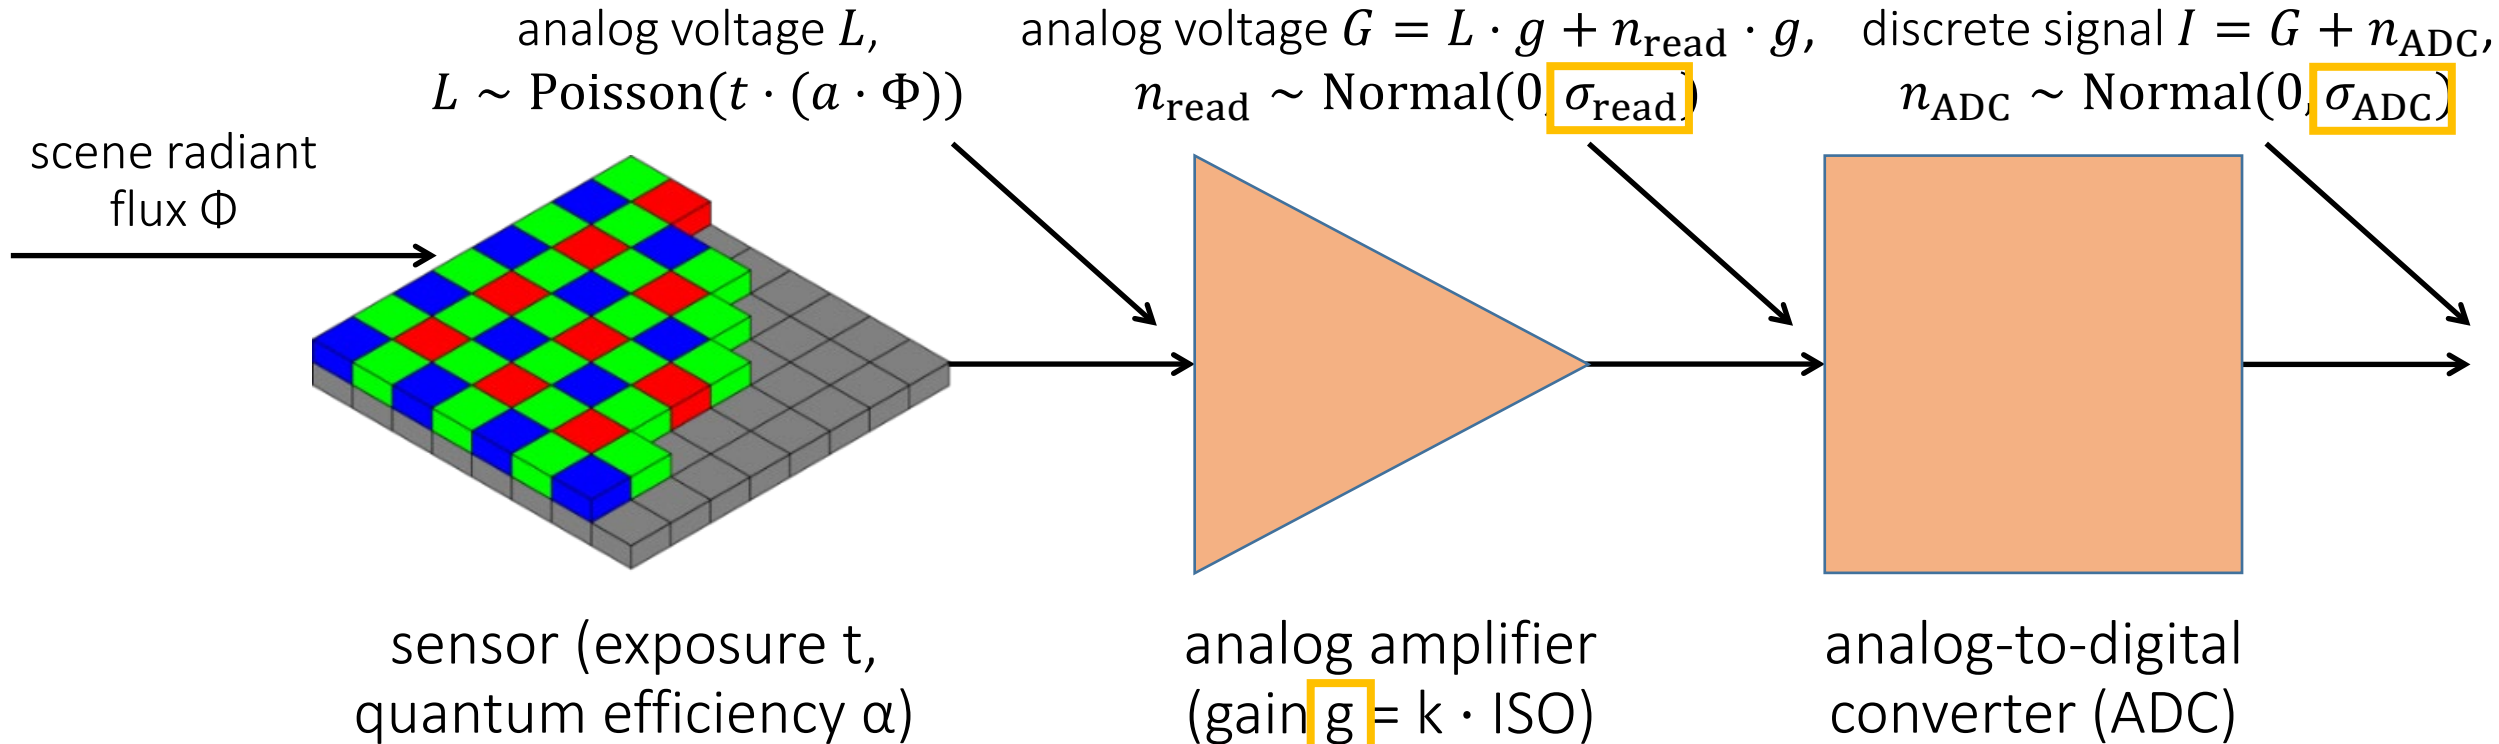
saturation level \nearrow

intensity mean and variance:

$$E(I) = t \cdot (a \cdot \Phi + D) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi + D) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2$$

Noise model after dark frame subtraction



discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{read}} \cdot g + n_{\text{ADC}}, I_{\text{max}})$$

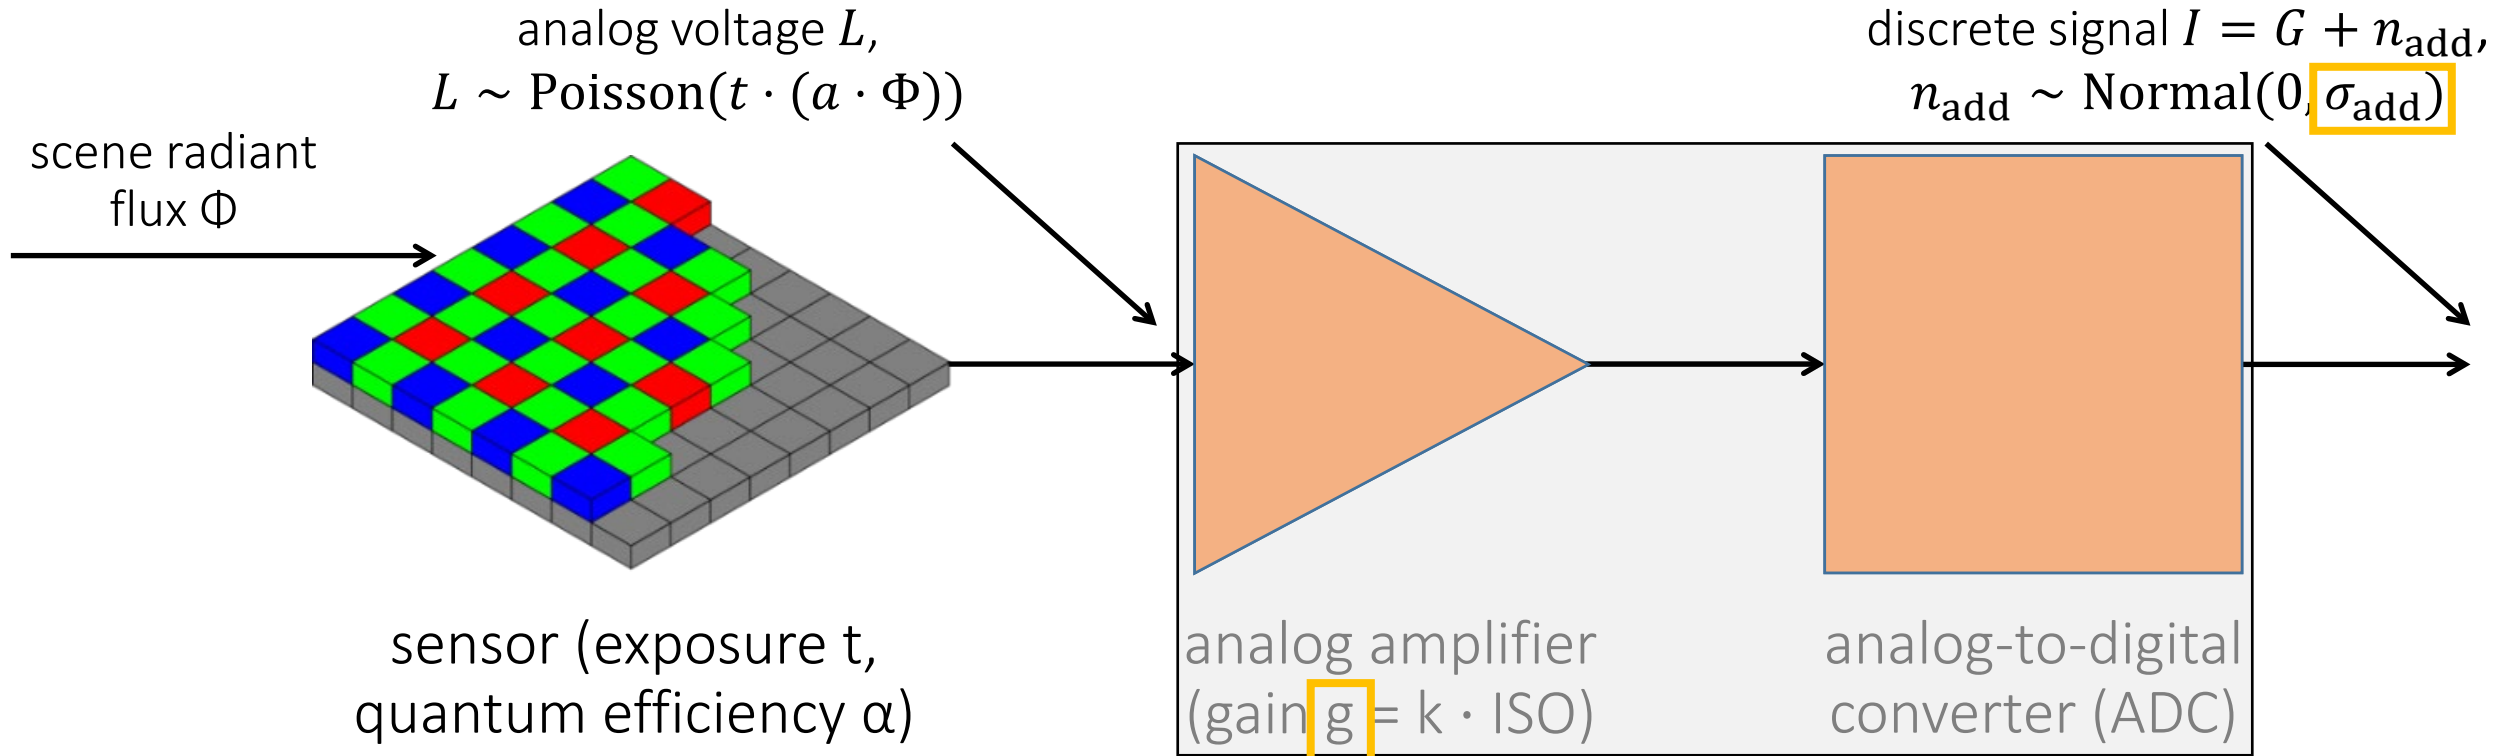
saturation level \nearrow

intensity mean and variance:

$$E(I) = t \cdot (a \cdot \Phi) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi) \cdot g^2 + \sigma_{\text{read}}^2 \cdot g^2 + \sigma_{\text{ADC}}^2$$

Affine noise model after dark frame subtraction



discrete image intensity (with saturation):

$$I = \min(L \cdot g + n_{\text{add}}, I_{\text{max}})$$

intensity mean and variance:

$$E(I) = t \cdot (a \cdot \Phi) \cdot g$$


$$\sigma(I)^2 = t \cdot (a \cdot \Phi) \cdot g^2 + \sigma_{\text{add}}^2$$

Estimating the gain and additive noise variance

Can you think of a procedure for estimating these quantities?



Estimating the gain and additive noise variance

1. Capture a large number of images of a grayscale target.





Estimating the gain and additive noise variance

1. Capture a large number of images of a grayscale target.
2. Compute the empirical mean and variance for each pixel, then form a mean-variance plot.



Estimating the gain and additive noise variance

1. Capture a large number of images of a grayscale target.
2. Compute the empirical mean and variance for each pixel, then form a mean-variance plot.





$$E(I) = t \cdot (a \cdot \Phi) \cdot g$$

$$\sigma(I)^2 = t \cdot (a \cdot \Phi) \cdot g^2 + \sigma_{\text{add}}^2$$

$$\Rightarrow \sigma(I)^2 = E(I) \cdot g + \sigma_{\text{add}}^2$$

Estimating the gain and additive noise variance

1. Capture a large number of images of a grayscale target.
2. Compute the empirical mean and variance for each pixel, then form a mean-variance plot.
3. Fit a line and use slope and intercept to estimate the gain and variance.



equal to line slope



$$\sigma(I)^2 = E(I) \cdot g + \sigma_{\text{add}}^2$$

equal to line intercept

How would you modify this procedure to separately estimate read and ADC noise?

Estimating the gain and additive noise variance

1. Capture a large number of images of a grayscale target.
2. Compute the empirical mean and variance for each pixel, then form a mean-variance plot.
3. Fit a line and use slope and intercept to estimate the gain and variance.



equal to line slope

$$\sigma(I)^2 = E(I) \cdot g + \sigma_{\text{add}}^2$$

equal to line intercept

How would you modify this procedure to separately estimate read and ADC noise?

- Perform it for a few different ISO settings (i.e., gains g).

Important note

Noise calibration should be performed with RAW images!

Optimal weights for HDR merging

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 ← (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.



Note: many possible weighting schemes

Many possible weighting schemes

“Confidence” that pixel is noisy/clipped

- What are the **optimal** weights for merging an exposure stack?

$$w_{ij} = \exp\left(-4 \frac{(I_{lin_{ij}} - 0.5)^2}{0.5^2}\right)$$



RAW (linear) image formation model

(Weighted) radiant flux for image pixel (x,y) : $\alpha \cdot \Phi(x, y)$

Exposure time:


t_5



t_4



t_3



t_2



t_1



What weights should we use to merge these images, so that the resulting HDR image is an optimal estimator of the weighted radiant flux?

Different images in the exposure stack will have different noise characteristics

Simple estimation example

We have two noisy *unbiased* estimators x and y of the same quantity (e.g., a pixel's intensity).

- The two estimators have variance $\sigma[x]^2$ and $\sigma[y]^2$.

Simple estimation example

We have two noisy *unbiased* estimators x and y of the same quantity (e.g., a pixel's intensity).

- The two estimators have variance $\sigma[x]^2$ and $\sigma[y]^2$.

Assume we form a new estimator from the *convex* combination of the other two:

$$z = a \cdot x + (1 - a) \cdot y$$

Simple estimation example

We have two noisy *unbiased* estimators x and y of the same quantity (e.g., a pixel's intensity).

- The two estimators have variance $\sigma[x]^2$ and $\sigma[y]^2$.

Assume we form a new estimator from the *convex* combination of the other two:

$$z = a \cdot x + (1 - a) \cdot y$$

What criterion should we use in selecting the mixing weight a ?

Simple estimation example

We have two noisy *unbiased* estimators x and y of the same quantity (e.g., a pixel's intensity).

- The two estimators have variance $\sigma[x]^2$ and $\sigma[y]^2$.

Assume we form a new estimator from the *convex* combination of the other two:

$$z = a \cdot x + (1 - a) \cdot y$$

What criterion should we use in selecting the mixing weight a ?

- We should select a to minimize the variance of estimator z . (Why?)

What is the variance of z as a function of a ?

Simple estimation example

Assume we have two noisy estimators x and y of the same quantity (e.g., intensity at a pixel).

- The two estimators have variance $\sigma[x]^2$ and $\sigma[y]^2$.

Assume we form a new estimator from the *convex* combination of the other two:

$$z = a \cdot x + (1 - a) \cdot y$$

What criterion should we use in selecting the mixing weight a ?

- We should select a to minimize the variance of estimator z .

What is the variance of z as a function of a ?

$$\sigma[z]^2 = a^2 \cdot \sigma[x]^2 + (1 - a)^2 \cdot \sigma[y]^2$$

What value of a minimizes $\sigma[z]^2$?

Simple estimation example

Simple optimization problem:

$$\frac{\partial \sigma[z]^2}{\partial a} = 0$$

$$\Rightarrow \frac{\partial (a^2 \cdot \sigma[x]^2 + (1-a)^2 \cdot \sigma[y]^2)}{\partial a} = 0$$

$$\Rightarrow 2 \cdot a \cdot \sigma[x]^2 - 2 \cdot (1-a) \cdot \sigma[y]^2 = 0$$

$$\Rightarrow a = \frac{\sigma[y]^2}{\sigma[x]^2 + \sigma[y]^2} \quad \text{and} \quad 1 - a = \frac{\sigma[x]^2}{\sigma[x]^2 + \sigma[y]^2}$$

Simple estimation example

Putting it all together, the optimal linear combination of the two estimators is

$$z = \underbrace{\frac{\sigma[x]^2 \sigma[y]^2}{\sigma[x]^2 + \sigma[y]^2}}_{\text{normalization factor}} \cdot \underbrace{\left(\frac{1}{\sigma[x]^2} x + \frac{1}{\sigma[y]^2} y \right)}_{\text{weights inversely proportional to variance}}$$

Simple estimation example

Putting it all together, the optimal linear combination of the two estimators is

$$z = \underbrace{\frac{\sigma[x]^2 \sigma[y]^2}{\sigma[x]^2 + \sigma[y]^2}}_{\text{normalization factor}} \cdot \underbrace{\left(\frac{1}{\sigma[x]^2} x + \frac{1}{\sigma[y]^2} y \right)}_{\text{weights inversely proportional to variance}}$$

More generally, for more than two estimators,

$$z = \frac{1}{\sum_{i=1}^N \frac{1}{\sigma[x_i]^2}} \cdot \sum_{i=1}^N \frac{1}{\sigma[x_i]^2} x_i$$

This weighting scheme is called Fisher weighting and is a BLUE estimator.

Back to HDR

Given *unclipped* and *dark-frame-corrected* intensity measurements $I_i[x, y]$ at pixel $[x, y]$ and exposures t_i , we can merge them optimally into a single HDR intensity $I[x, y]$ as

$$I[x, y] = \frac{1}{\sum_{i=1}^N w_i[x, y]} \cdot \sum_{i=1}^N w_i[x, y] \frac{1}{t_i} I_i[x, y]$$

The per-pixel weights $w_i[x, y]$ should be selected to be inversely proportional to that pixel's variance at each image in the exposure stack.

- How do we compute this variance?

Back to HDR

Given *unclipped* and *dark-frame-corrected* intensity measurements $I_i[x, y]$ at pixel $[x, y]$ and exposures t_i , we can merge them optimally into a single HDR intensity $I[x, y]$ as

$$I[x, y] = \frac{1}{\sum_{i=1}^N w_i[x, y]} \cdot \sum_{i=1}^N w_i[x, y] \frac{1}{t_i} I_i[x, y]$$

The per-pixel weights $w_i[x, y]$ should be selected to be inversely proportional to that pixel's variance at each image in the exposure stack.

- How do we compute this variance? → Use affine noise model.

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = ?$$

Back to HDR

Given *unclipped* and *dark-frame-corrected* intensity measurements $I_i[x, y]$ at pixel $[x, y]$ and exposures t_i , we can merge them optimally into a single HDR intensity $I[x, y]$ as

$$I[x, y] = \frac{1}{\sum_{i=1}^N w_i[x, y]} \cdot \sum_{i=1}^N w_i[x, y] \frac{1}{t_i} I_i[x, y]$$

The per-pixel weights $w_i[x, y]$ should be selected to be inversely proportional to that pixel's variance at each image in the exposure stack.

- How do we compute this variance? → Use affine noise model.

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} \sigma[I_i[x, y]]^2$$

$$\Rightarrow \sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = ?$$

Back to HDR

Given *unclipped* and *dark-frame-corrected* intensity measurements $I_i[x, y]$ at pixel $[x, y]$ and exposures t_i , we can merge them optimally into a single HDR intensity $I[x, y]$ as

$$I[x, y] = \frac{1}{\sum_{i=1}^N w_i[x, y]} \cdot \sum_{i=1}^N w_i[x, y] \frac{1}{t_i} I_i[x, y]$$

The per-pixel weights $w_i[x, y]$ should be selected to be inversely proportional to that pixel's variance at each image in the exposure stack.

- How do we compute this variance? → Use affine noise model.

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} \sigma[I_i[x, y]]^2$$

$$\Rightarrow \sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} (t_i \cdot \alpha \cdot \Phi[x, y] \cdot g^2 + \sigma_{\text{add}}^2)$$

Computing the optimal weights requires:

1. calibrated noise characteristics.
2. knowing the radiant flux $\alpha \cdot \Phi[x, y]$.

This is what we wanted to estimate!

Simplification: only photon noise

If we assume that our measurements are dominated by photon noise, the variance becomes:

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} \left(t_i \cdot \alpha \cdot \Phi[x, y] \cdot g^2 + \sigma_{\text{add}}^2 \right) \simeq ?$$

Simplification: only photon noise

If we assume that our measurements are dominated by photon noise, the variance becomes:

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} \left(t_i \cdot \alpha \cdot \Phi[x, y] \cdot g^2 + \sigma_{\text{add}}^2 \right) \simeq \frac{1}{t_i} \alpha \cdot \Phi[x, y] \cdot g^2$$

By replacing in the HDR merging formula and *assuming only valid pixels*, the HDR estimate simplifies to:

$$I[x, y] = \frac{1}{\sum_{i=1}^N t_i} \cdot \sum_{i=1}^N I_i[x, y]$$

Notice that we no longer weight each image in the exposure stack by its exposure time!


Some comparisons



original weights



optimal weights assuming
only photon noise



Simplification: only photon noise

When is this a good assumption?

More general case

If we cannot assume that our measurements are dominated by photon noise, we can approximate the variance as:

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} \left(t_i \cdot \alpha \cdot \Phi[x, y] \cdot g^2 + \sigma_{\text{add}}^2 \right) \simeq \frac{1}{t_i^2} \left(I_i[x, y] \cdot g + \sigma_{\text{add}}^2 \right)$$

Where does this approximation come from?

More general case

If we cannot assume that our measurements are dominated by photon noise, we can approximate the variance as:

$$\sigma\left[\frac{1}{t_i} I_i[x, y]\right]^2 = \frac{1}{t_i^2} \left(t_i \cdot \alpha \cdot \Phi[x, y] \cdot g^2 + \sigma_{\text{add}}^2 \right) \simeq \frac{1}{t_i^2} \left(I_i[x, y] \cdot g + \sigma_{\text{add}}^2 \right)$$

Where does this approximation come from?



- We use the fact that each pixel intensity (after dark frame subtraction) is an unbiased estimate of the radiant flux, weighted by exposure and gain:

$$E[I_i[x, y]] = t_i \cdot \alpha \cdot \Phi[x, y] \cdot g$$


Some comparisons

tone-mapped merged HDR



standard weights



optimal weights



ground-truth



What about ISO?

Noise-Optimal Capture for High Dynamic Range Photography

Samuel W. Hasinoff Frédo Durand William T. Freeman
 Massachusetts Institute of Technology
 Computer Science and Artificial Intelligence Laboratory

Abstract

Taking multiple exposures is a well-established approach both for capturing high dynamic range (HDR) scenes and for noise reduction. But what is the optimal set of photos to capture? The typical approach to HDR capture uses a set of photos with geometrically-spaced exposure times, at a fixed ISO setting (typically ISO 100 or 200). By contrast, we show that the capture sequence with optimal worst-case performance, in general, uses much higher and variable ISO settings, and spends longer capturing the dark parts of the scene. Based on a detailed model of noise, we show that optimal capture can be formulated as a mixed integer programming problem. Compared to typical HDR capture, our method lets us achieve higher worst-case SNR in the same capture time (for some cameras, up to 19 dB improvement in the darkest regions), or much faster capture for the same minimum acceptable level of SNR. Our experiments demonstrate this advantage for both real and synthetic scenes.

rameters of an exposure sequence, and we show that this reduces to solving a mixed integer programming problem. In particular, we show that, contrary to suggested practice (e.g., [5]), using high ISO values is desirable and can enable significant gains in signal-to-noise ratio.

The most important feature of our noise model is its explicit decomposition of additive noise into pre- and post-amplifier sources (Fig. 1), which constitutes the basis for the high ISO advantage. The same model has been used in several unpublished studies characterizing the noise performance of digital SLR cameras [7, 20], supported by extensive empirical validation. Although all the components in our model are well-established, previous treatments of noise in the vision literature [13, 18] do not model the dependence of noise on ISO setting (i.e., sensor gain).

To the best of our knowledge, varying the ISO setting has not previously been exploited to optimize SNR for high dynamic range capture. However, in the much simpler context of single-shot photography, the *expose to the right* tech-

- We need to separately account for read and ADC noise, as read noise is gain-dependent.
- We can optimize our exposure bracket by varying both shutter speed and ISO

Bonus part of Homework 2 (+ 50%!)

Real capture results

	input	tone-mapped merged HDR	bright region	dark region
standard exposure bracketing	<p>1/800 s ISO 100</p> <p>1/400 s ISO 100</p> <p>1/200 s ISO 100</p>			<p>5.9 dB</p>
SNR-optimal sequence	<p>1/6400 s ISO 800</p> <p>1/1600 s ISO 6400</p> <p>1/125 s ISO 6400</p>			<p>16.2 dB</p>
“ground truth”	<p>+15 more ...</p>			

References

Basic reading:

- Szeliski textbook, Sections 10.1, 10.2.
- Hasinoff et al., “Noise-Optimal Capture for High Dynamic Range Photography,” CVPR 2010.
A paper on weighting different exposures based on a very detailed sensor noise model, additionally discussing combining shutter speed and ISO changes.
- Healey and Kondepudy, “Radiometric CCD camera calibration and noise estimation,” PAMI 1994.
A detailed paper on radiometric and noise calibration based on the noise model we discussed.
- Martinec, “Noise, Dynamic Range and Bit Depth in Digital SLRs,” 2008, <http://theory.uchicago.edu/~ejm/pix/20d/tests/noise/index.html>
A very detailed discussion of noise characteristics and other performance aspects of digital sensors.

Additional reading:

- Kirk and Andersen, “Noise characterization of weighting schemes for combination of multiple exposures,” BMVC 2006.
A great paper on the variance characteristics of most common HDR weighting schemes.
- Granados et al., “Optimal HDR Reconstruction with Linear Digital Cameras,” CVPR 2010.
This paper extends the analysis of optimal HDR weights to consider spatially-varying noise effects.
- Hasinoff, “Fundamentals of Computational Photography: Sensors and Noise,” ICCP 2010 tutorial, <https://people.csail.mit.edu/hasinoff/hdrnoise/hasinoff-sensornoise-tutorial-iccp10.pptx>
A detailed tutorial on sensors and noise.
- Hasinoff et al., “Time-constrained photography,” ICCV 2009.
- Hasinoff and Kutulakos, “Light-efficient photography,” PAMI 2011.
These two papers examine noise-optimal acquisition and merging schemes for *focal* and *aperture* stacks, rather than exposure stacks.
- Ratner et al., “Optimal multiplexed sensing: bounds, conditions and a graph theory link,” Optics Express 2007.
- Ratner and Schechner, “Illumination Multiplexing within Fundamental Limits,” CVPR 2007.
These two papers discuss the effect of different types of noise when fusing multiple images in the context of illumination multiplexing.
- Gupta et al., “Photon-Flooded Single-Photon 3D Cameras,” CVPR 2019.
A paper on the noise characteristics of single-photon-sensitive cameras.