Fundamentals of Imaging
Image acquisition: digital

Prof. Dr. Charles A. Wüthrich,
Fakultät Medien, Medieninformatik
Bauhaus-Universität Weimar
caw AT medien.uni-weimar.de

Solid state sensors

- A sensor places photo-sensitive elements in an array.
- Use photoelectric effect: generate electro-hole pairs as result of photons coming in, and measure charge
- Two main types:
  - CCD: charge coupled devices
  - CMOS: complementary metal-oxide semiconductor
- Typical characteristics of sensors:
  - *Pixel count*: there are limits to it due to diffraction (airy disk in diffraction through a pinhole)
    - The intensity of the first ring is 1.75% that of the center disc and is located at a radius *r* of $r=1.22\lambda F$ (wavelength, aperture)
    - *r* measures resolving power of lens, and indicates minimum spacing of 2 points (Rayleigh limit)
  - *Angular response*: light does not come straight into the sensor

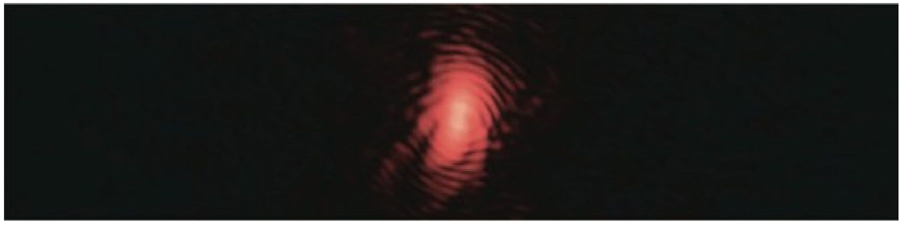

Solid state sensors

- *S/N ratio*: charge may accumulate or be lost
  - *noise-equivalent power*, radiant power that produces a signal-to-noise ratio of 1.
  - *detectivity D* : reciprocal of noise-equivalent power
- *Dynamic range*: range of irradiance detectable
- *Responsivity*: amount of signal generated per unit of image irradiance, determined by quantum efficiency of each pixel and its fill factor
- *Linearity:* both
  - collection of charge in response to incident photons,
  - conversion to a digital signal

  achieve a linear relationship between the number of photons and the value of the resulting signal
- This is true for most sensors (not film)
- *Pixel uniformity*: differences in pixels due to manufacturing. Sometimes manufacturers correct this in firmware

CCD sensors

- In a CCD sensor, photosensitive elements are photodiodes arranged in an array, basically capacitors
- A photodiode can absorb photons, attracting electrons which reduce the voltage across the diode proportionally to the amount of incident power.
- When exposure starts, photodiodes collect charge until filled (*full-well capacity*).
- At end of exposure, charges are sent to A/D converter.
- Image is read one pixel at a time (slow)

- Thus 3 architectures:

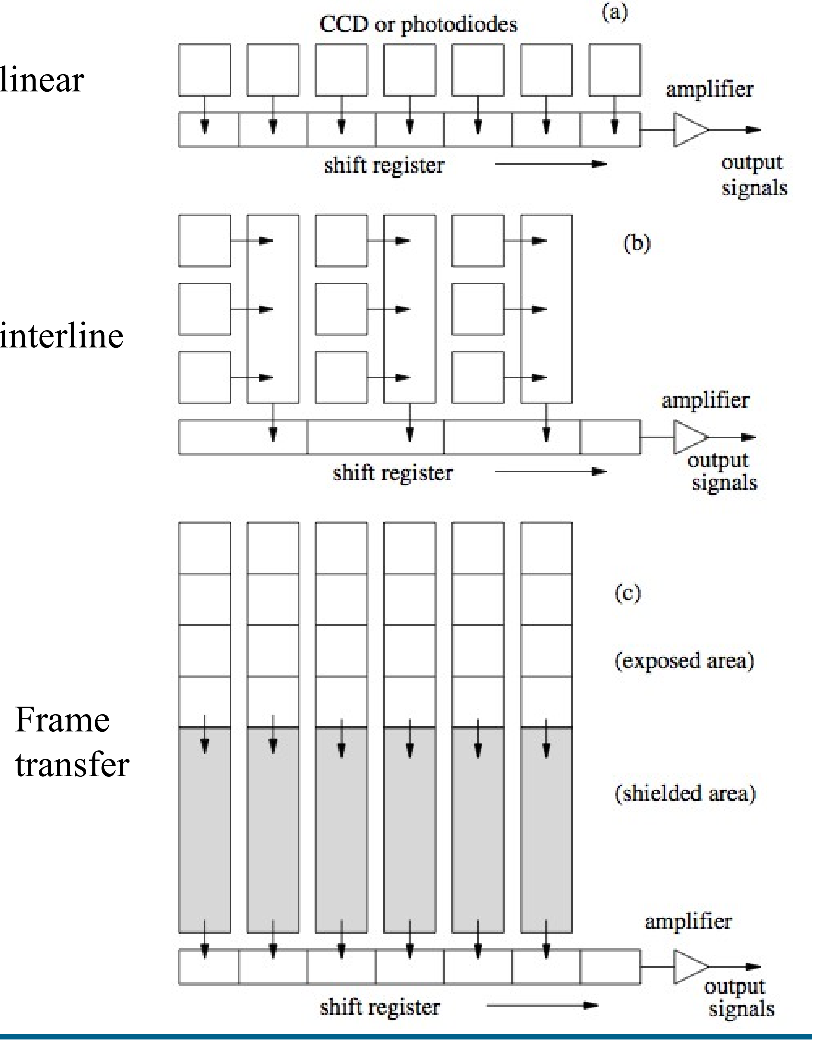

CCD sensors

- Noise in CCD sensor comes from:
  - photon statistics,
  - the CCD array itself:
    - transfer noise +
    - dark current (during exposure, thermal agitation generates electrons) +
    - manufacturing imperfections
    - Cosmic noise hitting sensor
  - on-chip amplifier,
  - off-chip amplifier(s),
  - A/D converter,
  - electrical interference,
  - signal processing steps

- Thus 3 architectures:

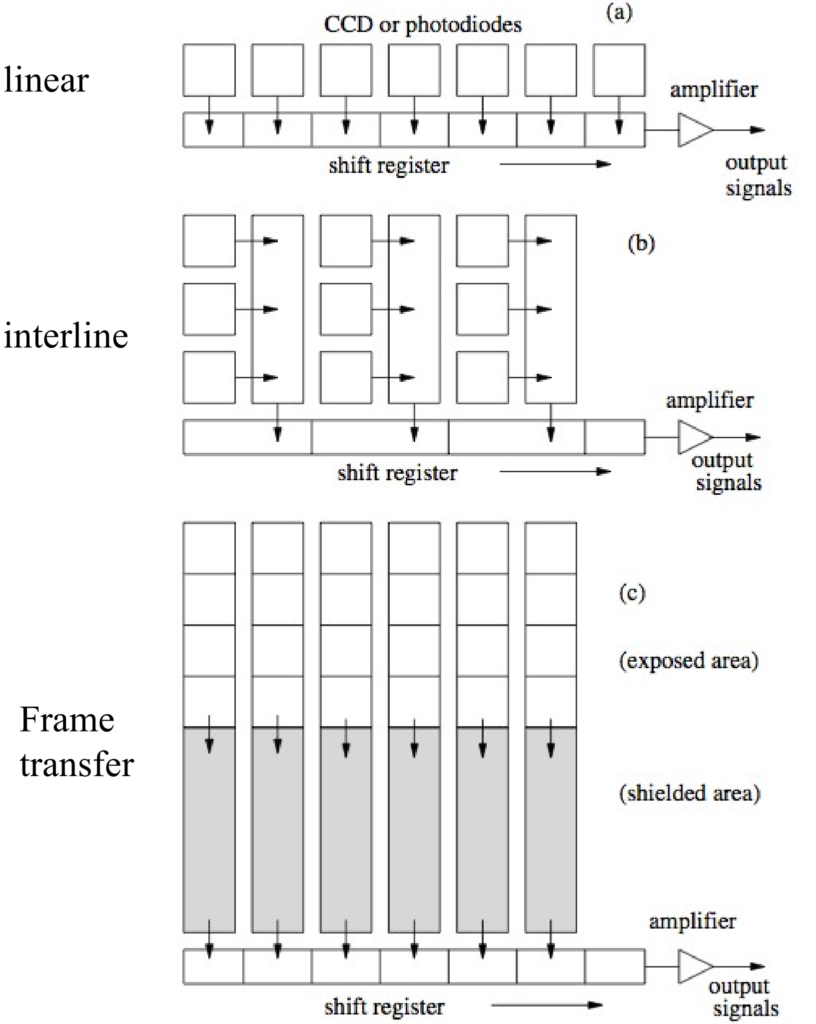

CMOS sensors

- CMOS is used for chips, but can be adapted for sensing light
- In CMOS technology, it is possible to integrate the sensor array, the control logic, and potentially analog-to-digital conversion on the same chip.
- Building blocks:
  - a pixel array,
  - analog signal processors,
  - row and column selectors,
  - timing and control block.
- Nowadays, the photosensor at each pixel is augmented with additional circuitry, such as a buffer/amplifier, yielding an *active pixel sensor* (APS)
- APS sensors allow a high frame rate
- Analog-to-digital (A/D) conversion circuitry may be included for each pixel: *digital pixel sensor* (DPS) (high-speed).
- Drawback:
  - additional circuitry takes up space
  - Heat of A/D converter generates temperature differences, which increase noise

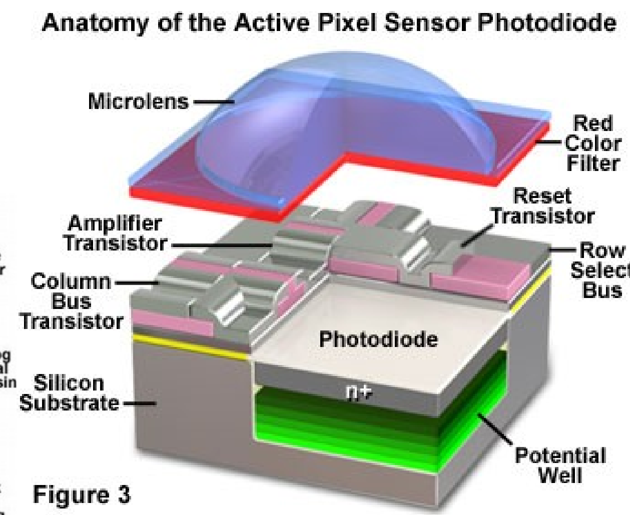

Figure 3

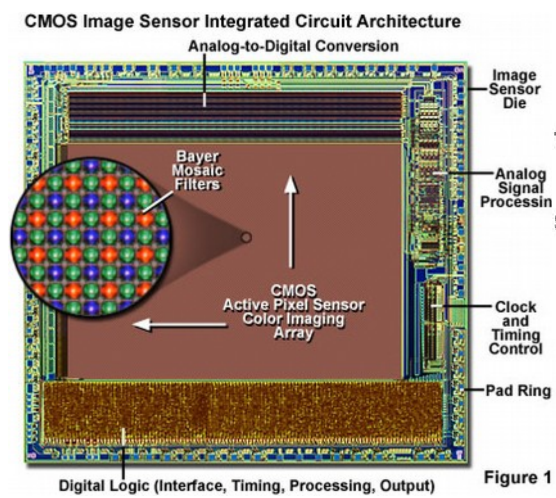

Figure 1

Colour in digital sensors

- Sensors discussed were just B/W
- To achieve colour, thin coloured micro-lenses are molded on top of each sensor.
- Most used: Bayer pattern
- Notice that pattern sense non-contiguous locations for a single colour

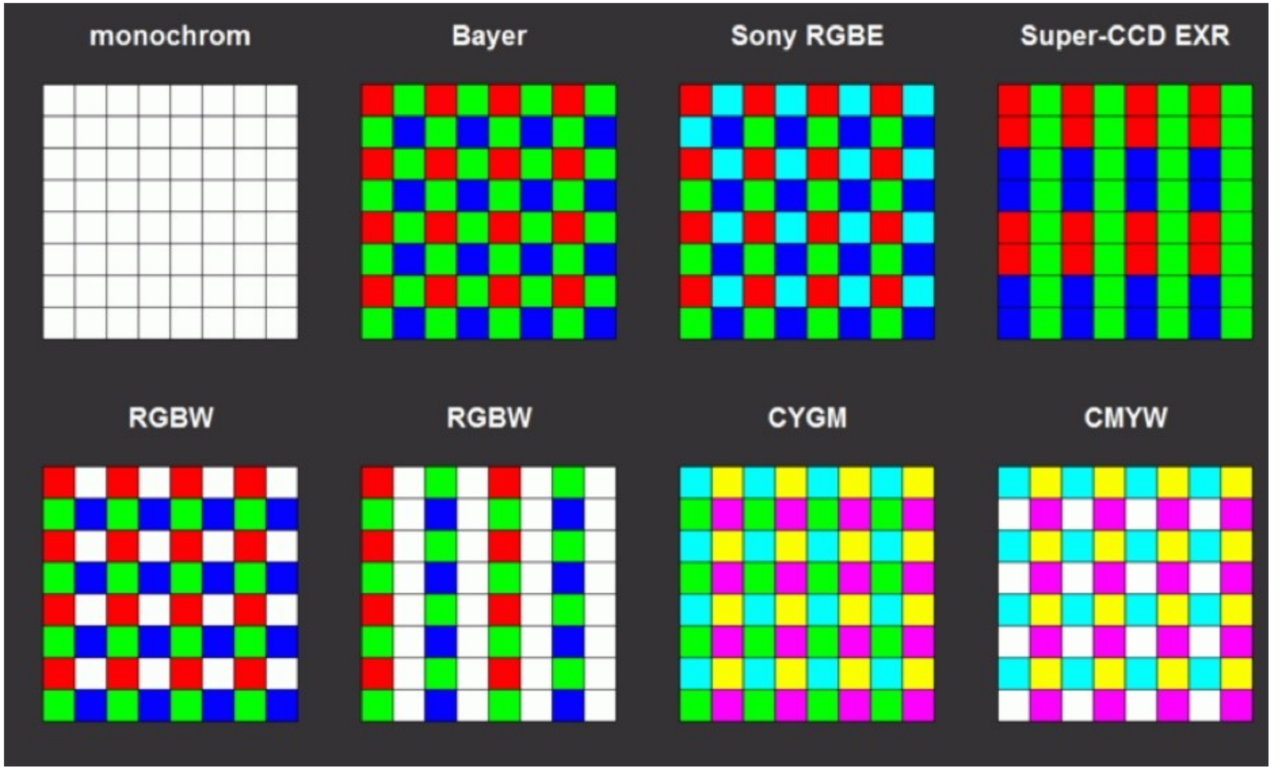

ISO in digital sensors

- Digital sensors have a parameter called ISO sensitivity
- Higher ISO values are obtained by modifying amplifying gain before A/D conversion
- This of course amplifies also noise, resulting in more grainy images
- To avoid this, manufacturers implement noise reduction algorithms, which in turn degrade the original image

ISO 200

ISO 6400

Color Filter Arrays

- Sensors are monochrome, covered by RGB color filters (CFA: *color filter array*)
  - mostly arranged on a Bayer pattern (2x green)
- How to reconstruct image?
- Naïve approach: combine 4 neighbouring pixels to obtain pixel color:
  - poor spatial resolution
- Better approach: reconstruct image same resolution as sensor by interpolation (*demosaicing*)

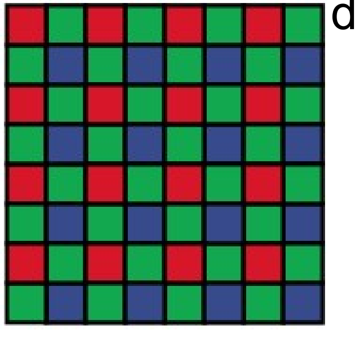

- Luminance values are estimated from green values
- If luminance for non-green pixels computed through interpolation: blur
- Non-linear adaptive average is then used:
  - Edge detection
  - Ensure object edges not blurred
- If one assumes RGB correlated in a local image region, then edge detection will give more detailed values

Color Filter Arrays

- In a camera, the luminance channel (G) may be augmented by two chrominance channels:
  - R-G (C_R)
  - B-G (C_B).
- Note that the green value was computed using the CFA interpolation scheme.
- This very simple solution is intended to minimize firmware processing times.
- To compute missing chrominance values, linear interpolation is employed between neighboring chrominance values
  - sometimes neighbors used for interpolaton may be located diagonally.
- Luminance and chrominance values are then converted to RGB

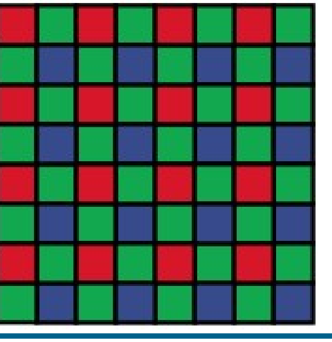

White balancing

- Can be automatic, or user specified
- In auto, camera has to infer the illuminant fast
- One way to do it is making gray-world assumption: average colour is 18% gray:
  - Compute average colour, and correct accordingly
- Second approach: assume pixel with highest intensity as white
  - Fails if non-white light sources are in picture, e.g. traffic light
- Third approach: analyze color gamut of picture
  - Use statistical assumptions on average surface reflectance and emission spectra of light sources
  - Color gamut is compared to pre-defined database
- Usual approach: combine the 3 methods
- Pro photographers prefer to shoot raw and correct in image processing software

A mathematical camera model

- Math models of camera help understand error sources accumulating in acquisition
- Model by Healey and Kondepudy [94]
- Assume # of electrons collected at a photosite (x, y) as integral over surface area of the pixel (u, v) and over all wavelengths λ to which the sensor is sensitive.
- Because each electron is charge unit, total charge at photosite is

$$Q(x,y) = \Delta t \int_{\lambda} \int_{x}^{x+u} \int_{y}^{y+v} E_e'(x,y,\lambda)\, S_r(x,y)\, c_k(\lambda)\, dx\, dy\, d\lambda.$$

where

- Δt: integration time in secs
- E'_e: spectral irradiance incident onto sensor (W/m^2)
- S_r: spatial response of pixel
- $c_k(\lambda)$: ratio of charge collected to the light energy incident during integration (C/J)

- An ideal system has no noise and no losses ⇒ all charge converted to voltage and amplified with gain g, leading to voltage

$$V(x,y)=Q(x,y)g$$

- The A/D converter converts this into a number n ST by b bits, $0 \leq n \leq 2^b-1$.
- If quantization step is s, voltage is rounded to $D=ns$, so that $(n(x, y)-0.5)s < V(x, y) \leq (n(x, y)+0.5)s$
- Camera firmware processes image n(x,y)

Sensor Noise Characteristics (SNC) - Reset noise

- Noise can be defined as any component of the output signal not derived from irradiance onto the sensor
- Some such noise will vary in time (temporal noise) depending on
  - Picture
  - Time
- Other noise depends on sensor imperfection and will induce same imperfection for all pictures

- Before image can be taken, potential wells are reset
- This takes time: reset-time constant
- For high speed applications, there might not be time to reset fully
- Therefore, some potential wells may still be charged:
  reset noise.
- This can be eliminated by doing *correlated dual sampling*
  - Resetting pixel
  - reading pixel once to know reset charge left
  - Read pixel for final image

SNC - Fixed pattern noise

- When sensor is uniformly lit, charge collected at every pixel will vary due to fabrication errors (Dark Signal Non-Uniformity)
- This is corrected by taking picture with lens cap on
- Resulting image measures fixed-pattern noise
- In the camera model, fixed pattern noise shows up as:
  - Variation in sensor response S_r
  - Variation in the quantum efficiency c_k (λ).
- These 2 sensor characteristics are each scaled by a constant which is fixed (but different) for each pixel.
- Call the constants $k_1(x,y)$ and $k_2(x,y)$, and, because they are constant, they can be taken out of integral.
- Set $k(x,y) = k_1(x,y)\ k_2(x,y)$. the camera model produces a charge Q_n for charge collection site (x,y):

  $$Q_n(x, y) = Q(x, y)\ k(x, y)$$

  where n indicates that Q includes noise
- The fixed pattern noise $k(x,y)$ is taken to have a mean of 1 and a variance of σ_k^2, which depends on the quality of the camera design
- This model works if one assumes that neighbouring pixels do not interact

SNC - Dark Current and Dark Current Shot Noise

- In absence of light, electrons might still reach potential wells collecting charge
  - for example because of thermal vibrations
- *Dark current* is the uniform distribution of electrons over the sensor collected in the dark
  - Can be corrected in postprocessing
- *Dark current shot noise* is instead a non-uniformly distributed source of noise which accumulates in the dark
  - Cannot be corrected in postprocessing
  - Should be minimized
- For extreme applications, cooling the sensor helps

- Dark current is independent of the number of photoelectrons generated.
- The charge associated with the number of dark electrons Q_{dc} produced by dark current is added to the signal
- This gives a total charge Q_n:

$$Q_n = Q(x, y)k(x,y) + Q_{dc}(x, y).$$

SNC - Photon shot noise

- Photons arrive at sensor randomly, with Poisson distribution
- If the mean increases, so does the variance in signal
- Thus, given a threshold, eventually photon noise dominates other noises
- However, below this threshold, the other noises dominate
- Photon shot noise can be reduced by collecting a higher number of photo-electrons
  - widen the aperture
  - increase exposure time up to the limit imposed by the full-well capacity of the sensor.
- Photons do Poisson distribution, and if sensor linear
  ⇒ photoelectrons also follow Poisson distribution
- The uncertainty in the number of collected electrons can be modeled with a zero mean random variable Q_s, and its variance depends
  - on number of photoelectrons,
  - the number of dark electrons.
- Thus,

$$Q_n = Q(x,y)k(x,y) + Q_{dc}(x,y) + Q_s(x,y).$$

Transfer noise

- Transferring accumulated charges to the output amplifier may cause errors, ⇒ transfer noise
- Can be neglected, because read-out efficiency of modern CCD devices can be greater than 0.99999.
- The amplifier of a CCD device generates additional noise with zero mean.
- This *amplifier noise* is independent of the amount of charge collected and therefore determines the noise floor of the camera.
- The amplifier applies a gain *g* to the signal, but also introduces noise and applies low-pass filtering to minimize the effects of aliasing.
- The amplifier noise (called read noise) is indicated with Q^{-}_{r} ,
  ⇒ output voltage of the amp is:

$$V(x,y) = (Q(x,y)k(x,y) + Q_{dc}(x,y)+Q_{s}(x,y)+Q_{r}(x,y))\ g$$

- Note that
  - fixed pattern noise is multiplied with the desired signal, whereas the remaining noise sources are added to the signal.
  - combined signal and noise is then amplified by a factor of *g*

Quantization noise

- The voltage V produced by the amplifier is subsequently
  - sampled and
  - digitized by the A/D converter.
- This leads to a further additive noise source Q_q, which is independent of V
- This gives:

$$D(x,y) = (Q(x,y)\, k(x,y)+Q_{dc}(x,y)+Q_s(x,y)+Q_r(x,y))\, g+Q_q(x,y).$$

- Quantization noise Q_q is a zero-mean random variable which is uniformly distributed over the range $[-0.5q, 0.5q]$ and has a variance $q^2/12$

Implications of noise

- Most noise sources are independent of the image irradiance except photon shot noise.
- Thus, all noise sources become increasingly important when photographing dark environments
  - Longer exposures required
- Illustration shows log-log plot of photon electrons increasing linearly with irradiance
- Look at picture:
  - # of photon electrons increases linearly till well capacity reached
  - After this, Q cannot increase
  - All other noise sources are independent from incoming photons,
  - except photon shot noise
- But standard deviation of photon shot noise increases as the square root of the signal level:
  - the amount of photon shot noise is a straight line w/ slope 0.5
- So sensors should have lots of full well capacity, but this depends on surface area

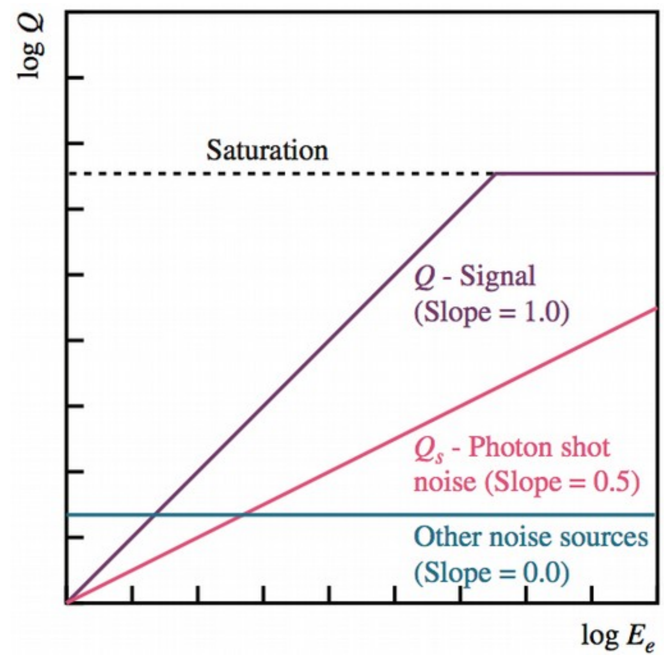

Implications of noise

- Except high-end digital cameras (FF sensors), sensors get smaller, due to lower cost of smaller opticals.
- Useful for cell phones and web cams.
- Additionally, we see more pixels per sensor.
- The surface area of each photosite is reduced, leading to a lower full-well capacity, and therefore saturation occurs at lower image irradiances.
- In addition, the dynamic range of a sensor is related to both the full-well capacity and the noise floor.

- For high dynamic range cameras, it would be desirable to have a high full-well capacity:
  - even in the case that all other noise sources are eliminated as much as possible, dynamic range is limited by photon shot noise

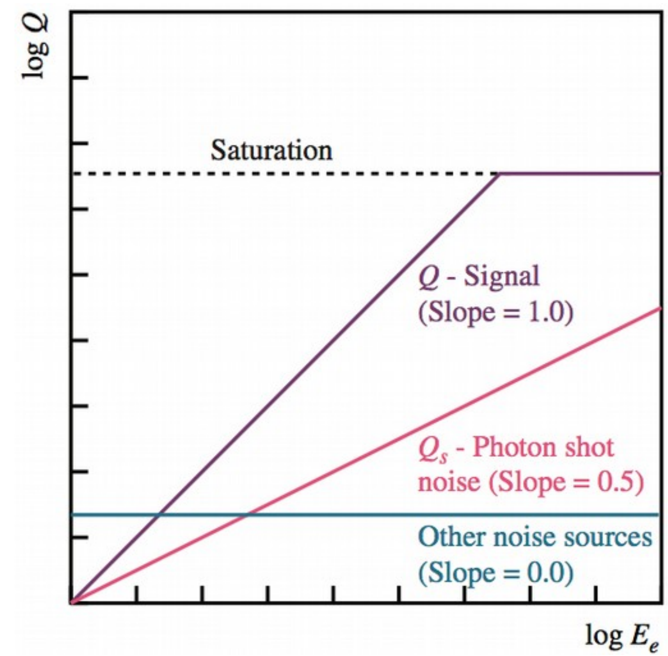

Measuring camera noise

- To measure the noise introduced by a digital camera, one can take test targets and compare the sensor output with known values
  - For instance, images can be taken with the lens cap fitted.
- Or, uniformly illuminate a Lambertian surface and take images of it.
- If sensor characteristics are to be estimated, it will be good to defocus such an image.
- Variation in sensor output will be determined by
  - camera noise,
  - (small) variations in illumination of the surface,
  - (small) variations in reflectance properties of the surface.

- We'll assume test surface
  - to face camera
  - nearly uniformly illuminated with a light source
- Reflected radiance L_e is then

  $L_e=(\rho_d/\pi)L$

  where
  - light source strength=L
  - ρ_d/π: fraction of light reflected to camera
- The lens in front of the sensor then focuses the light onto the sensor so that the irradiance incident upon a pixel is given by

$$E'_e = \frac{\rho_d}{2} L \left(\frac{n'}{n}\right)^2 G$$

  which we rewrite as

  $E'_e = k_s \rho_d L$

Measuring camera noise

- To model the (small) variation of the illumination and reflectance as function of position on the image plane, the factors in this equation are parameterized as function of wavelength λ and position (x, y) on the sensor:
  $$E'_e(x,y,\lambda) = k_s\, \rho_d(x,y,\lambda) L(x,y,\lambda)$$
- The average illumination on the test target is $\overline{L}(\lambda)$ and average reflectance is $\overline{\rho}_d(\lambda)$
- The illumination onto the test card surface that is ultimately projected onto pixel (x,y) is then
  $$L(x,y,\lambda) = \overline{L}(\lambda)+L_r(x,y,\lambda)$$
  where:
  - L_r deviation from the average $\overline{L}(\lambda)$ for this pixel
- Expected value is then
  $$\mathrm{E}(L(x,y,\lambda)) = L(\lambda)$$
  and expected value of residual
  $$L_r(x,y,\lambda)=0$$
- Similarly, reflectance of the test card at position projected onto sensor location (x,y) can be split into
  - average ρ_d and
  - zero mean deviation $\rho_r(x,y,\lambda)$:
  $$\rho_d(x,y,\lambda) = \overline{\rho}_d(\lambda)+\rho_r(x,y,\lambda)$$
- Where
  - $\mathrm{E}(\rho_d(x,y,\lambda)) = \overline{\rho}d(\lambda)$
  - $\mathrm{E}(\rho_r(x,y,\lambda)) = 0$
- It is reasonable to assume there is no correlation between illumination and reflection of test card: thus
  $$E'_e(x,y,\lambda) = k_s(\overline{L}(\lambda)\overline{\rho}_d(\lambda)+\varepsilon(x,y,\lambda))$$
  with
  $$\varepsilon(x,y,\lambda) = \rho_r(x,y,\lambda)\overline{L}(\lambda) + \overline{\rho}_d(\lambda)L_r(x,y,\lambda)+\rho_r(x,y,\lambda)L_r(x,y,\lambda)$$

Measuring camera noise

- The expected value of ε is then 0
- The charge collected by the sensor can now be split into a constant component Q_c, and a spatially varying component $Q_v(x,y)$

  $$Q(x, y) = Q_c + Q_v(x, y)$$

  where

$$Q_c = k_s \Delta t \int_\lambda \int_x^{x+u} \int_y^{y+v} \bar{L}(\lambda) \bar{\rho}_d(\lambda) S_r(x,y) c_k(\lambda)\, dx\, dy\, d\lambda$$

$$Q_v(x,y) = k_s \Delta t \int_\lambda \int_x^{x+u} \int_y^{y+v} \varepsilon(x,y,\lambda) S_r(x,y) c_k(\lambda)\, dx\, dy\, d\lambda.$$

- As $E(\varepsilon(x,y,\lambda)) = 0$, we have that $Q_v(x,y)$ has a zero mean and a variance that depends on the variance in $\bar{L}(\lambda)$ and $\bar{\rho}_d(\lambda)$.

- During calibration procedure, image capture is important to control illumination of test card to achieve illumination that is as uniform as possible.
- Similarly, the test card should have as uniform a reflectance as possible to maximally reduce the variance of $Q_v(x,y)$.

Most of the variance in the resulting signal will then be due to the sensor noise, rather than to non-uniformities in the test set-up.

Noise variance

- The variance of the system consists of several components.
- We first discuss these components, leading to an expression of the total variance of the sensor.
- Then we can estimate its value by photographing a test card multiple times
- Quantized values $D(x, y)$ can be modeled as random variables as follows:
  $D(x,y) = \mu(x,y)+N(x,y)$
- Expected value $E(D(x,y))$ is $\mu(x,y)$:
  $\mu(x,y)=Q(x,y)k(x,y)g+E(Q_{dc}(x,y)g)$
- zero-mean noise is modeled by $N(x,y)$:
  $N(x,y) = Q_s(x,y)g+Q_r(x,y)g+Q_q(x,y)$

- The noise sources can be split into
  - a component that does not depend on the level of image irradiance
  - a component that does.
- Photon shot noise, modeled as a Poisson process, increases with irradiance:
  $Q_s(x, y)\,g$
- Accounting for the gain factor g, the variance associated with this Poisson process given by
  $g^2(Q(x, y)k(x, y)+ E(Q_{dc}(x, y)))$
- signal-independent noise sources, amplifier noise and quantization noise, are given by
  $Q_r(x,y)g+Q_q(x,y)$
  and have combined variance of
  $g^2\, \sigma_r^2/12$
  where σ_r^2 variance of amplifier noise

Noise variance

- Total variance σ^2 in noise introduced by the sensor is then sum of these two variances:

$$\sigma^2 = g^2 \left(Q(x,y)\,k(x,y) + \mathrm{E}(Q_{\mathrm{dc}}(x,y))\right) + g^2\sigma_r^2 + \frac{q^2}{12}.$$

- Expected value of dark current for given pixel can be replaced by
  - sum of the average expected value over the whole sensor +
  - deviation from this expected value for a given pixel: if we put $Q_{\mathrm{E(dc)}}(x,y) = \mathrm{E}(Q_{\mathrm{dc}}(x,y))$, then

$$Q_{\mathrm{E(dc)}}(x,y) = \bar{Q}_{\mathrm{E(dc)}} + Q_{d\mathrm{E(dc)}}(x,y)$$

  where
    - $Q_{\mathrm{E(dc)}}$ average expected value of the dark current
    - $Q_{d\mathrm{E(dc)}}(x,y)$ deviation from the expected value

- Now we can rewrite:

$$\sigma^2 = g^2 \left(Q(x,y) + \bar{Q}_{\mathrm{E(dc)}}\right) + g^2 \left(k(x,y) - 1\right) Q(x,y) + g^2 \left(Q_{d\mathrm{E(dc)}}(x,y) - \bar{Q}_{\mathrm{E(dc)}}\right) + g^2\sigma_r^2 + \frac{q^2}{12}.$$

and if we assume $|k(x,y)-1| \ll 1$ and $|Q_{d\mathrm{E(dc)}}(x,y) - \bar{Q}_{\mathrm{E(dc)}}| \ll \bar{Q}_{\mathrm{E(dc)}}$ then the variance of sensor noise can be approximated as:

$$\sigma^2 \approx g^2 \left(Q(x,y) + \bar{Q}_{\mathrm{E(dc)}}\right) + g^2\sigma_r^2 + \frac{q^2}{12}$$

Estimating total variance

- By photographing a uniformly illuminated test card twice one gets two pictures with same subject (=expected pixel value $\mu(x,y)$) but different noises:
  $D_1(x,y) = \mu(x,y)+N_1(x,y)$
  $D_2(x,y) = \mu(x,y)+N_2(x,y)$
- Subtracting the images one gets an image with 0 mean and spatial variance $2\sigma^2$
- The expected value of D_1 or D_2:
  $\mu=Q(x,y)\, g + \overline{Q}_{E(dc)}\, g$
- If spatial variance minimized, one can replace Q(x,y) with spatial mean $\overline{Q}$:
  $\mu=\overline{Q}\, g + \overline{Q}_{E(dc)}\, g$
- So the variance in terms of μ is
  $$\sigma^2 = g\mu + g^2\sigma_r^2 + \frac{q^2}{12}$$
- For two images D_1,D_2, μ can be estimated as
  $$\hat{\mu} = \frac{1}{XY}\sum_{x=0}^{X}\sum_{y=0}^{Y} D_1(x,y) + \frac{1}{XY}\sum_{x=0}^{X}\sum_{y=0}^{Y} D_2(x,y)$$
  and the variance can be estimated as
  Mean of difference image
  $$\hat{\sigma}^2 = \frac{1}{XY-1}\sum_{x=0}^{X}\sum_{y=0}^{Y}\left(D_1(x,y) - D_2(x,y) - \hat{\mu}_d(x,y)\right)^2$$
- We have not considered illumination yet: to estimate the amp gain g and variance of signal-dependent noise sources $\sigma_c^2 = g^2\sigma_r^2 + q^2/12$ one can vary illumination

Estimating total variance and dark current

- We can then find $\hat{g}$ and $\hat{\sigma}_c^2$ with a line fitting technique

$$\sum_i \frac{\left(\hat{\sigma}_i^2 - \left(\hat{g}\hat{\mu}_i + \hat{\sigma}_c^2\right)\right)^2}{\mathrm{var}\left(\hat{\sigma}_i^2\right)}$$

  where the sum is over a set of image pairs, indexed with an *i*
- Finally the variance can be estimated:

$$\mathrm{var}\left(\hat{\sigma}_i^2\right) \approx \frac{2\left(\hat{\sigma}_i^2\right)^2}{XY-1}$$

- In the absence of light, a camera will still output a signal as a result of dark current.
- By fitting the lens cap, the number of photo-generated electrons will be zero: quantized value output by the sensor is given by

$$D(x,y) = \left(Q_{\mathrm{dc}}(x,y) + Q_{\mathrm{s}}(x,y) + Q_{\mathrm{r}}(x,y)\right) g + Q_{\mathrm{q}}(x,y)$$

- The mean value of signal is then $Q_{E(dc)}(x,y)g$, and because all noises have 0 mean, its variance is $\sigma^2(x,y)$.
- Taking a number *n* of pictures, and averaging them, variance will be reduced to $\sigma^2(x, y)/n$.
- If *n* large enough, pixel values will converge to accurate estimate of dark current $Q_{E(dc)}(x,y)g$.

Estimating fixed pattern noise

- We now have estimates of
  - dark current,
  - amplifier gain, and
  - total variance in the noise
- To estimate fixed pattern noise, it is not enough to vary lighting uniformly:
  - non-uniform lighting variations have to be taken
  - Additionally, the orientation of the test card has to vary
- We can average out the variation due to illumination and retain the fixed pattern noise $k(x, y)$, here called *photo response non-uniformity* (PRNU)
- Suppose that we create n_1 illumination conditions and capture n_2 images for each of these conditions
- Total noise $N(x,y)$ has 0 mean: we average frames for imaging condition i

  $k(x,y)Q_i(x,y)g+Q_{E(dc)}(x,y)g$

  with variance $\sigma_i^2(x,y)/n^2$

- Subtracting estimate of dark current yelds

  $d(x,y) \approx k(x,y)Q_i(x,y)g$

- This estimate varies due to fixed pattern noise and differences in illumination and reflectance
- For small pixel neighbourhoods, variations in illumination and reflectance are small: one can compute the mean $\bar{d}(x,y)$ over small windows, usually 9x9 windows

  $\bar{d}(x, y) \approx Q_i(x, y)g$

- Ratio between a single pixel estimate and windowed estimate is a rough approximation of the fixed pattern noise $k_e(x,y)$

$$\frac{d(x,y)}{\bar{d}(x,y)} \approx \frac{k(x,y)\,Q_i(x,y)\,g}{Q_i(x,y)\,g} \approx k_e(x,y)$$

- To refine the approximation, average ratio over n_1 imaging conditions is computed:

$$k_e(x,y) \approx \frac{1}{n_1}\sum_{i=0}^{n_1}\frac{d_i(x,y)}{\bar{d}_i(x,y)}$$

Calibrating cameras

- To recover a linear relationship between the irradiance values and the pixel encoding produced by the camera, we need to model the non-linearities introduced by in-camera processing.
- The process of recovering this relationship is known as camera characterization
- Typically, this is done in two ways:
  - Using spectral sensitivity, which needs expensive equipment such as a monochromator
  - Using predefined targets
- Target-based techniques make use of a set of differently colored samples that can be measured with a spectrophotometer
- Captured colors and the target values are then matched

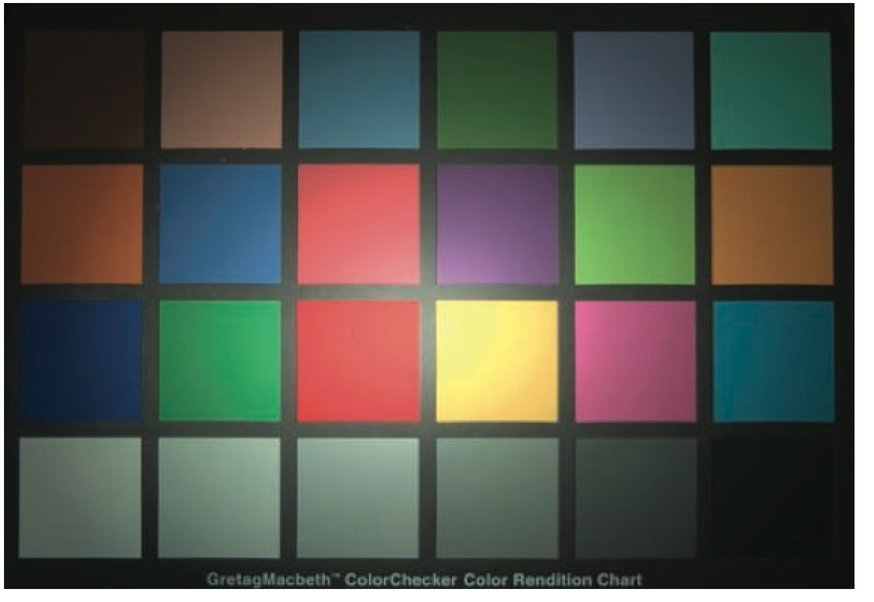

Calibrating cameras

- When few shots are made, the measured data can be seen as color differences of the device (error) and ideal color values
- A transformation can be computed such that the difference between the transformed device output and the ideal response is minimized
- The function

$$\underset{f_k}{\operatorname{argmin}} \sum_{n=1}^{N} ||f_k(p_n) - P_n||^2$$

  has to be minimized.
- Here ||.|| is CIELAB ΔE^*_{ab} color difference metric
  - p_n pixel value recorded by the camera for the *n*th stimulus,
  - P_n corresponding measured response,
  - N total number of samples,
  - f_k transformation being estimated for the *k*th color channel

- Typical techniques for finding the mapping from known data include linear and polynomial regression, as well as neural networks
- Once the camera is calibrated, the recovered function f_k can be applied to all images captured
- Images obtained this way are called *device independent*