Noise in cameras and photodiodes
Laboratory Instructions

Version: 2018-02-19

Preparations for the lab

Read the following (can be found on the lecture notes web page):

- Chapter 5 from from Image Acquisition and Preprocessing for Machine Vision Systems by P. K. Sinha (Sections 5.7, 5.9.5 and 5.9.6 are not included; sections 5.1-5.3 and the appendix may be omitted since they overlap with other parts of the course)
- Photon Transfer Noise (Section 3.3 not to be included)
- Browse pp 777-791 of chapter 18.6 in the course book (Saleh & Teich)
- This lab instruction

Do the following exercises:

1) A CCD camera has a read noise of 5 electrons, a fixed pattern noise quality factor of 3% and a dark current of 2 electrons/s/pixel. When the pixels are exposed to light, charges accumulate at a rate of 100 counts/s . Answer the following questions for an exposure time of 1s and for an exposure time of 10s.

   a) Give the RMS values for the different noise sources.
   b) Give the RMS value for the total noise.
   c) Give the signal to noise ratio.
   d) If you could make a perfect camera which didn't suffer from any technical noise. What would the signal to noise ratio be?

2) Consider the circuit diagram in Figure 1 (left) showing a reverse biased photodiode in series with a load resistor R_L.

   a) Given that the light impinging onto the photodiode generates a photocurrent $i = 2$ µA, how big is the shot-noise current assuming a 1 Hz bandwidth?
   b) Given that the resistance $R_L = 100$ kΩ, how large is the thermal noise current in a 1 Hz bandwith?
   c) At what photocurrent i is the shot-noise equal to the thermal noise? What voltage drop across the resistance R_L will this photocurrent give rise to?

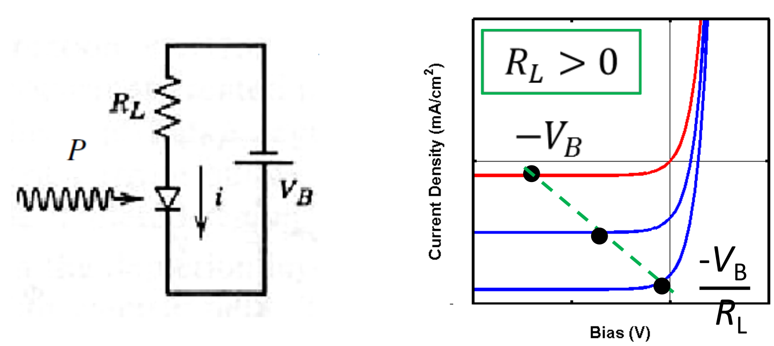

Figure 1. *Left*: Circuit diagram of reverse biased photodiode in series with a resistor R_L. *Right*: Current-voltage curves for different light intensities. The dashed line is the load line describing how the voltage across the diode decreases with increasing light intensity.

Purpose

In this lab we will study the noise performances of a CCD camera and of a photodiode, both with and without electron multiplication. We all know that images taken in low light conditions are noisy. We will here look at how we can quantify the noise and the performance of the camera so that different cameras can be compared. We will also investigate the fundamental noise limit: photon shot noise. Furthermore, we will investigate how the noise performance of a CCD camera can be improved through binning and electron multiplication gain.

Background

The CCD camera

The CCD was invented in 1969 by Willard Boyle and George Smith, at AT&T Labs. In 2009 they received the Nobel Prize in Physics for their discovery. In a CCD camera the entire area of the pixel is used for collecting light. The accumulated pixel charge is then sequentially shifted out of the CCD array to a sense node where charge is converted to voltage which is then amplified. The read out signal represents how many electrons were accumulated on each pixel. The obvious advantage with the CCD is that the entire pixel area is used for light collection. The architecture is also very versatile in low light conditions where the charge to voltage conversion is very noisy. If the light collected by each pixel is too low to give a reasonable signal to noise ratio (SNR) several pixels can be merged into a super-pixel during the readout. This process is called *binning* and is a way of sacrificing spatial resolution for improved SNR. Another possibility to improve low light performance is to include a gain register before the charge to voltage conversion. The gain register multiplies the number of electrons in each pixel via impact ionization and the device is therefore called electron multiplying CCD (*EMCCD*).

Noise in imaging

Noise is present in any image. Pixels which on average receive the same number of photons read out different values according to a Gaussian distribution. How well a camera performs is not measured by how large signal it outputs but rather by the signal to noise ratio (SNR).

$$SNR = \frac{signal}{noise} = \frac{\overline{S}}{\sigma}$$

Where $\overline{S}$ is the mean pixel value and σ is the standard deviation of the pixel values. The standard deviation is here expressed as a root mean square (RMS) value.

The main noise source is ***photon shot noise*** which comes from the fact that the photons are randomly distributed in a stream of light. It can be shown that the standard deviation for an average photon number $\overline{N}_{photons}$ is $\sqrt{\overline{N}_{photons}}$ (see Photon streams Ch. 12.2 pp 458-465). This means that the SNR is ultimately limited by the photon shot noise limit:

$$SNR_{Shot} = \frac{\overline{N}_{photons}}{\sigma_{photons}} = \frac{\overline{N}_{photons}}{\sqrt{\overline{N}_{photons}}} = \sqrt{\overline{N}_{photons}}$$

In a similar manner the standard deviation for the signal is

$$\sigma_{shot} = \sqrt{\overline{S}}$$

So that in terms of the signal the shot noise limit can be expressed as

$$SNR_{Shot} = \frac{\overline{S}}{\sqrt{\overline{S}}} = \sqrt{\overline{S}}$$

The second noise source is the ***read noise*** of the camera. It is generated by the electronics as the charge on a pixel is read out. Before each readout, the pixel should be reset to a well-defined number of electrons. This process isn't perfect and there is some variance to the reset value. This noise will be independent of the signal.

$$\sigma_{read} = constant$$

The third noise source is ***fixed pattern noise*** (FPN). It is called "fixed" because it is not random, it is the same from one image to another. FPN can come from e.g. dust specs on the sensor but this is generally easy to clean. In a CMOS camera FPN originates from fabrication inconsistencies since each pixel has its own amplifier, i.e. not all pixels have the same gain. Fixed pattern noise therefore scales with the signal.

$$\sigma_{FPN} = P_N * S$$

Where P_N is the FPN quality factor.

For a CCD FPN is generally not a problem since all pixel values are read out through the same output node. In the early days if digital cameras FPN was a big problem for CMOS cameras and therefore CCDs dominated the market. As fabrication methods have improved, FPN can nowadays generally be ignored (at least for consumer electronics).

The fourth noise source is ***dark current noise***. It originates from thermally generated electrons that accumulate on the pixels during long integration times. If all pixels are equally susceptible to dark electrons the mean number of dark electrons is $\overline{m} = D * t$ (where D is a constant and t is time) then

$$\sigma_{dark} = \sqrt{\overline{m}} = \sqrt{D * t}$$

Nowadays the dark current is so small that dark current noise isn't really relevant until you reach exposure times of several seconds.

Noise in a photodiode circuit

The signal-to-noise ratio in electric circuits is usually defined in terms of power. In the photodiode circuit in fig. 1, the electric power in the load resistor is $P_{\text{electric}} = R_L \cdot i^2$. The signal-to-noise ratio is, with this definition, therefore given by

$$SNR = \frac{i^2}{{\sigma_i}^2}$$

where *i* is the mean value of the current and σ_i its standard deviation. Since the current is a flow of discrete particles, electrons, there is shot-noise in a similar way as for photons. The number of electrons passing a cross-section during a time *t* is it/e leading to a standard deviation in the number of electrons of $\sqrt{it/e}$. The standard deviation of the current during the time *t* is therefore given by

$$\sigma_{shot} = \frac{e}{t}\sqrt{it/e} = \sqrt{2eiB}$$

where $B = 2/t$ is the bandwidth. The dependence on bandwidth can be understood by considering that large bandwidth requires a short measurement time leading to few electrons and thus high uncertainty.

A resistor like the one in the circuit in fig. 1 will introduce additional noise. Due to the thermally driven random motion of electrons in the resistor there will be charge fluctuations that give rise to what is called thermal noise or Johnson-Nyquist noise. Expressed as current, this noise is given by

$$\sigma_{th} = \sqrt{\frac{4kTB}{R}}$$

where *k* is the Boltzmann constant and *T* is the absolute temperature.

The gain *G* in an avalanche photodiode (APD) will enhance both the signal and the shot noise but not the thermal noise and adding gain may therefore improve the SNR at low light levels. However, variations in the gain will introduce additional noise described by the excess noise factor *F*. The shot-noise can in this case be written as

$$\sigma_{shot} = \sqrt{2eGiFB}$$

Combining shot-noise and thermal noise leads to the signal-to-noise ratio

$$SNR = \frac{i^2}{\sigma_{shot}^2 + \sigma_{th}^2} = \frac{(eG\eta\phi)^2}{2e^2G^2\eta\phi FB + \sigma_{th}^2}$$

For the last equality we used that the photocurrent *i* is given by the photon flux ϕ, the quantum efficiency η and the gain *G* such that $i = eG\eta\phi$. Without gain we can set *G*=*F*=1 and the SNR becomes shot-noise limited for sufficiently large photon flux.

Laboratory Assignment

Noise in a CCD camera

The CCD camera is an Andor Luca camera with the following specifications:

Active pixels:	658x496	
Pixel size:	10x10 µm	
Pixel well depth:	16000	(how many electron each pixel can hold before saturating)
Read Noise	1-15 electrons (RMS)	

The CCD is controlled via the LabView program "Luca Control" on the desktop of the computer. Start by familiarizing yourself with how the program works. DO NOT TAKE ANY IMAGES UNTIL YOU ARE SURE YOU WON'T SATURATE THE CAMERA.

Familiarizing yourself with the CCD controls

Figure 2 (at the end of the instructions) shows the user interface for the camera. To the left there are three control panels. In the middle the image is displayed. To the right you control the cursors and read out statistical data from the image.

In control panel 1 you initialize the camera by pressing the button "Init". Do this and wait until the "Init ok" light comes on. In control panel 2 you can control which pixels are read out, how the pixels are binned together during readout and which electron multiplication gain is used. Let's leave the settings at their default values to start with (Fig. 1 shows the default settings). In control panel 3 you set the exposure time and take the images. Every time you make a change to the camera settings (exposure time, binning, EM gain) you should press the "Set" button. You can either take a single image ("Get Single Image") or a series of images ("Get Image Series"). Start by setting the exposure time to its minimum value and take a single image.

An image is now displayed at the center of the screen and statistical data of the Selected Area (area within the yellow cursors) is seen on the right. To the right of the image is also a color bar giving you the intensity scale for the image. Above the intensity scale there is an indicator that reads “Max pixel value”, during the exercises you should keep an eye on this so that it doesn’t exceed 16000.

Familiarizing yourself with the Matlab interface

When an image is taken a Matlab script is called from Labview and a Matlab plot showing a histogram of the pixel values in the Selected Area pops up. Labview opens the Matlab command window but nothing else. During the measurements all data will be passed from Labview to Matlab. To access the data you need to see the Matlab workspace, to open it run the command “workspace” in the command window. You will also need to plot the data in various ways which is done by writing a script in the matlab editor, to open the editor run the command “edit” in the command window. You should now have 3 Matlab windows open (command, workspace and editor).

When you take an image series Matlab will plot the histogram for the pixel values within the Selected Area for each image. The statistics will be stored in the following vectors:

StatsMean	(A vector with the pixel mean value, the signal)
StatsStd	(A vector with the pixel standard deviation, the noise)
StatsTime	(A vector with the exposure times)

Once you take a new image series these vectors will be replaced. To save them for later just rename them in the Matlab workspace. It makes sense to save the workspace from time to time with Ctrl+S. Make sure you save all the vectors listed above for later analysis. You can also save specific Matlab figures to use in your report.

Read noise

The first task in evaluating the camera performance is to measure the read noise. Turn off the light completely and take a dark image. Take an image with “Get Single Image” and look at the histogram. What is the signal and what is the standard deviation? Clearly the camera reads out a signal which doesn’t originate from light exposure. This must be subtracted from further images. On Labview control panel 3 you can record a dark image by pressing “Get Dark Image”. The mean pixel value of the dark image will now be subtracted from all the following images. Take a new image (still in the dark) with “Get Single Image”. Have a look at the histogram, what happens to the mean pixel value and what happens to the standard deviation?

Questions:

1. What is the measured read noise and how does it compare to the specification for the camera?

Investigate the SNR as a function of exposure time

On the Labview control panel create a new vector “Exposure times, series” with exposure times between your minimum and maximum exposure time. Choose exposure times which are appropriate to display on a logarithmic scale. Take an image series and then process the data in Matlab.

Questions:

1. How many photons (on average) hit each pixel per second? What is the photon flux density impinging the sensor?
2. Plot the SNR as a function of the signal. Use a log-log plot since both SNR and the signal span several orders of magnitude.

3. Mark the photon shot noise limit in the plot. At what point does the photon shot noise equal the read noise?
4. Do the measured data points match what you expect based on the read noise?
5. Plot the SNR as a function of exposure time. Mark the photon shot noise limit in the plot. At what exposure time does the photon shot noise equal the read noise?
6. Can the SNR ever exceed the photon shot noise limit?

Binning

Change the binning settings in Labview control panel 2. You need to press "Set" for the change to take effect. You also need to take a new Dark Image as well as rescale the Selected Area, ask your supervisor about this.

Before you take an entire image series you should take a few single images. Start at a low exposure time and increase it slowly to find when the image starts to saturate. It will happen at a shorter exposure time than before. Change the "Exposure time, series" vector on the Labview control panel 3 to reflect the new maximum exposure time. If needed ask your supervisor how to delete elements in a Labview vector (don't just set them to zero). Once you have made sure you've changed the names of the image statistics vectors in Matlab you can take a new image series with the new binning settings.

Questions:

1. Is the increase in signal proportional to the binning factor?
2. Compare the $SNR_{with\ binning}$ with $SNR_{without\ binning}$in a plot. Do this for at least one binning setting, more if you have time.
3. Include the photon shot noise limit in your plots. When do the binned images reach the shot noise limit compared to the unbinned images?
4. Compare the resolution of a few images with different binning settings. In the report you should visualize the change in resolution in a suitable way.
5. When does it make sense to use binning?

Electron multiplication gain

IN ORDER NOT TO DAMAGE THE CAMERA; MAKE SURE YOU DON'T SATURATE THE PIXELS!

We will now investigate how EM gain can be used to improve the SNR under some circumstances. The EM gain value in Labview control panel 1 has to this point been set to 0. This doesn't mean zero gain just that the EM gain function is disabled. We will now take a few image series with different EM gain settings and evaluate the SNR as a function of exposure time. For these measurements you can return to the 1x1 binning used originally. Just like before you need to make sure that the camera doesn't saturate. The more gain you introduce, the earlier it will saturate. Change the EM gain value and then take a few single images starting from the minimum exposure time. Increase until the maximum pixel value at the top right corner of the image on the Labview screen approaches 16000. For every new EM gain value you need to take a new Dark Image before you take the exposure time series. Update the "Exposure time, series" vector (REMEMBER TO PRESS SET IN ORDER FOR THE CHANGES TO TAKE EFFECT) and then acquire the image series. Do this with a few different EM gain value settings (find the maximum allowed exposure time before each image series acquisition). Use EM gain values up to 100 but not more.

Questions:

1. What is the signal gain in your measurements?
2. Does the EM gain value represent the signal gain, i.e. is the gain 100 if you set the EM gain value to 100?
3. Compare the SNR as a function of exposure time for different EM gain settings with the measurements without EM gain. Also compare with the photon shot noise limit calculated in the case of no EM gain.

4. Compare a few images with different EM gain (and no EM gain) at the minimum exposure time.
5. Also do the comparison for an exposure time when the image is at the shot noise limit even without EM gain.
6. When does it make sense to use EM gain?
7. The EM multiplication register makes 536 transfers. Based on the signal gain you observe what is the ionization probability at each transfer?

Can fixed pattern noise and the dark current noise be ignored? Sanity check.
Whenever you make an assumption you should try to check if it is valid. Take an image with long exposure time and look for fixed pattern noise. You should also take an image in the dark with the longest exposure time you have used and check if the noise level is higher than just the read noise. Use a piece of cloth to cover the camera objective to make sure no light effects the measurement. Do a very long exposure and measure the dark current (expressed in dark electrons accumulated per second).
Questions:

1. Is there any fixed pattern noise that might have influenced the measurements?
2. What is the dark current? Is the assumption that dark current noise can be ignored valid?

Noise in a photodiode

During this part of the laboratory exercise, you will measure the signal-to-noise ratio of a reverse biased photodiode at different light intensities and compare this with what you would expect from theory.

With and without reverse bias
Start by applying zero bias to the photodiode and measure the voltage drop over the 100 kΩ load resistor using a voltmeter while varying the LED current. Plot the voltage as a function of the LED current in a log-log plot. Note: because of the log-log scale, it is suitable to have a non-linear distribution of your data points. Start for instance at 50 µA LED current and double the current for each successive data point.
Now connect the battery to reverse bias the photodiode and repeat the measurement. Plot the data in the same graph as for the unbiased photodiode. Discuss the results. The graph to the right in Fig. 1 can be useful in understanding what is going on when applying a reverse bias: The dashed line is called the load line and shows how the voltage over the diode decreases as the increasing photocurrent results in an increased voltage drop across R_L thus leaving a smaller fraction of V_B to reverse bias the photodiode. How would you draw the load line for the case of V_B=0?

Thermal noise in resistors
Disconnect both the voltmeter and the photodiode from the 100 kΩ resistor and measure the voltage noise that appear across the resistor by using the lock-in amplifier. In these experiments, the lock-in amplifier is used as a very sensitive AC voltmeter which measures the RMS deviations (in $\mathrm{V}/\sqrt{\mathrm{Hz}}$) from the mean voltage. In this first measurement, the mean voltage is of course zero since you measure the voltage across a resistor which is not connected to any voltage or current source. Repeat the measurement with a 1 kΩ resistor. Do the values for the voltage noise that you get for the two resistors agree with what you expect?

Noise in the photodiode signal

Now re-connect the photodiode to the $100\ k\Omega$ resistor and measure the noise with the lock-in amplifier as you vary the light intensity of the LED. To be able to calculate the signal-to-noise ratio (SNR) you also need to measure the "signal", i.e. the (normal) mean voltage using the voltmeter. However, the voltmeter is very noisy and should not be connected during the noise measurements. You can either connect and disconnect the voltmeter between each data point or you set the LED current to the same values as you used earlier when you evaluated the effect of a reverse bias and use the voltages that you obtained then. Plot the noise as a function of the voltage in a log-log plot. In the same graph, plot the shot noise that you would expect from the photocurrent. Discuss and compare with theory. Finally calculate and plot the SNR in a log-log plot. In the same graph, include the shot-noise limited SNR. Is your data shot-noise limited for any light intensities? Discuss the deviations.

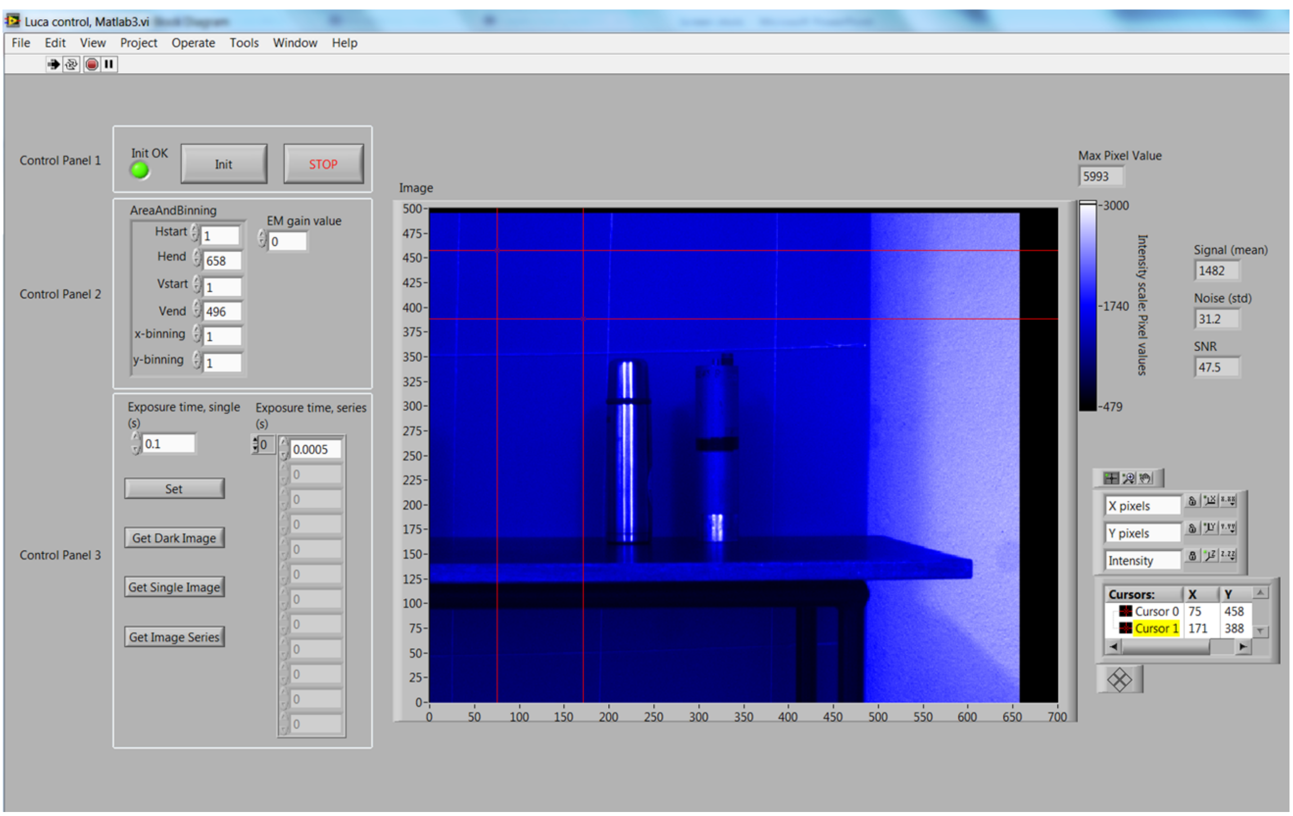

Figure 2