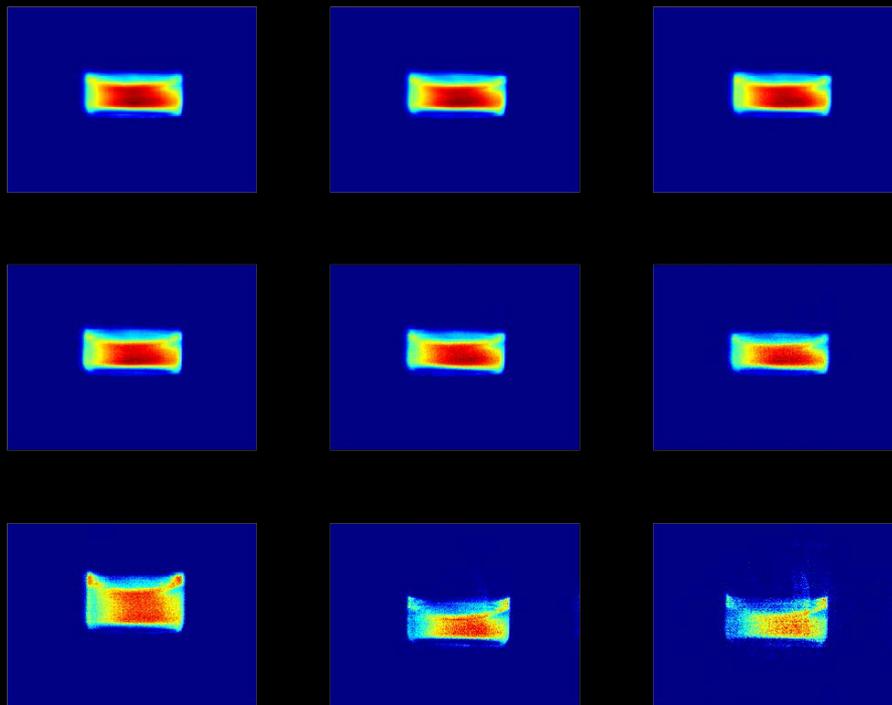




Improving signal-to-noise ratio in fluorescence detection for medical purposes

Sven van Binsbergen & Raimond Frentrop

24 August 2012



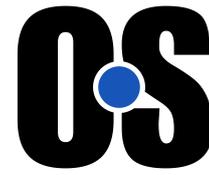
Supervisors:

Dr. Ir. Herman Offerhaus
Ing. Jeroen Kortelak

UNIVERSITY OF TWENTE.

Made possible by

The members of the Optical Sciences Group at the University of Twente and the resources that were available for us in that group. We thank all members of the group for their help and support during the few weeks of our bachelor assignment. Special thanks to our supervisors, Dr. Ir. Herman Offerhaus and Ing. Jeroen Kortterik, for their guidance during these eight weeks.



Also thanks to the Department of Oncology of the Universitair Medisch Centrum Groningen for providing the samples of IRDye[®] 800CW labelled tumors and information regarding the labelling process.

Titlepage figures

The top three images are fluorescence images of three different slices of IRDye[®] 800CW labeled tumors. The images are colormapped and the colormap was altered to remove a very faint diffusive glow around the tumors, caused by reflections in the surrounding paraffin.

The bottom image is a collection of 9 vials containing different concentrations of the IRDye[®] 680RD dye. Each successive vial contains half the concentration of the previous, beginning with 1.0 mg/L in the top-left vial until 3.8 µg/L in the bottom-right vial.

Summary

The main task for this bachelor assignment was to identify the noise of the camera and see if this noise could be reduced to better detect fluorescent dyes used in research on labelling cancer tumors. Labelling tumors with a fluorescent dye enables surgeons to better identify tumors, resulting in a higher percentage of tumor clusters being removed in surgery.

Because of the Flat Field Correction of the PixeLINK PL-B761 camera, there was a neglectable amount of dark current noise and flat field noise. The presence of any hot pixels depends on the exposure time. For short exposure times the images taken show that the hot pixels still have low values, much lower than the signal. This means that for short exposure times the noise mainly depends on the random noise (read/quantization noise and shot noise).

Thus the exposure time should be chosen as short as possible, but high enough to get a detectable amount of signal. A short exposure time is in this case also advantageous because of the possible application of this technique in living tissue: the human body always moves, caused by all kinds of oscillations from for example the heart, muscles and deliberate movement. To get a sharp image the exposure time should be short.

It was also proven that averaging a number of consecutive frames can significantly improve the image quality, by reducing noise and thus increasing the signal-to-noise ratio. Edges of dye-labelled tissue become clearer in this way.

For the chosen camera the noise for exposure times smaller than 50 ms does not differ much and is all well below the quantization level (1 level is 64 counts, the average noise is 36 counts). Considering this, modulation of the signal to subtract noise does not improve the image quality by much. Averaging a lot of images still shows a pattern in the read noise which can be subtracted by using a noise mask, but for single frames this subtraction more often adds noise to the image instead of removing it.

The fluorescent dyes used to label cancer cells can be detected at a short exposure time of 6 ms, even for very low concentrations (3.8 $\mu\text{g/L}$). The concentration of dye present in mouse tumors was estimated to be 155 $\mu\text{g/L}$, meaning that small amounts or thin slices of cancer cells will still be detectable by the camera. Colormapping the resulting image greatly improves the usability.

Contents

Summary	2
Introduction	6
1 Theory	8
1.1 Camera Calibration	8
1.1.1 Noise reduction	8
1.1.2 Signal-to-noise ratio	10
1.2 Fluorescence and dye	11
1.3 Methods for dye detection	12
1.3.1 Filtering of light	12
1.3.2 Gating	13
1.4 Signal Modulation	14
1.5 Usage of colormaps	16
2 Calibration	18
2.1 Setup	18
2.2 Results	19
2.3 Discussion	20
3 Noise Mask	22
3.1 Method	22
3.2 Results	22
3.3 Discussion	24
4 IRDye® 800CW	26
4.1 Equipment	26
4.2 Setup	26
4.3 Results	27
4.4 Discussion	30
5 IRDye® 680RD	32
5.1 Concentration detection limit	32
5.2 Results	32
5.3 Discussion	35
6 Conclusion	36
7 Recommendations	38
Bibliography	38

Introduction

Currently surgeons have to visually interpret during surgery whether certain cells are infected with cancer or not. Sometimes this is no problem, but especially in for example the intestines or ovaries this proves to be very difficult. The result is that many of the patients have to be operated multiple times before all cancer cells have been removed, if that will even happen. A solution to this problem would be an effective way to identify the cancer cells and show the surgeon - either directly or through a camera - where the infected cells are. Current research using a green fluorescent dye already indicates that where surgeons could identify cancer cell clusters of minimal 3 millimeter in size, they can now identify clusters as small as 100 micrometer. [1]

One of the possible ways to identify the cancer cells is using a fluorescent labeling agent that attaches itself to proteins specific for cancer cells. If these labeling agents (dyes) can be made in such a way that they have fluorescent properties, the dye (and thus the cancer cells) could be detected by using a laser or LED (see Figure 1). [9]

The work in this report is based on the work done by Brian van Hoozen but has not been continued where he stopped. Because some of the experimental instruments (camera, dye) were different, the first measurements were redone and the next steps were taken according to those results. [10]

During the project two different methods were considered to detect the fluorescent signal. One is the use of filters to filter out any of the wavelengths coming from the LED or the surrounding environment and ending up with only the wavelengths of the fluorescence. This method was already used by Brian van Hoozen, with some good results. The second method makes use of the lifetime of the dye. If this lifetime would be long enough, the camera could take a picture when the LED is already turned off, but the dye is still fluorescing.

The goal is to detect the fluorescence of the dye using an excitation laser or LED with an intensity low enough to not damage the tissue. The framerate of the imaging should also be high enough to avoid blurring in the image because of movement of the body (caused for example by the heart or by muscles). The resolution should be high enough to also detect very small groups of cancer cells.

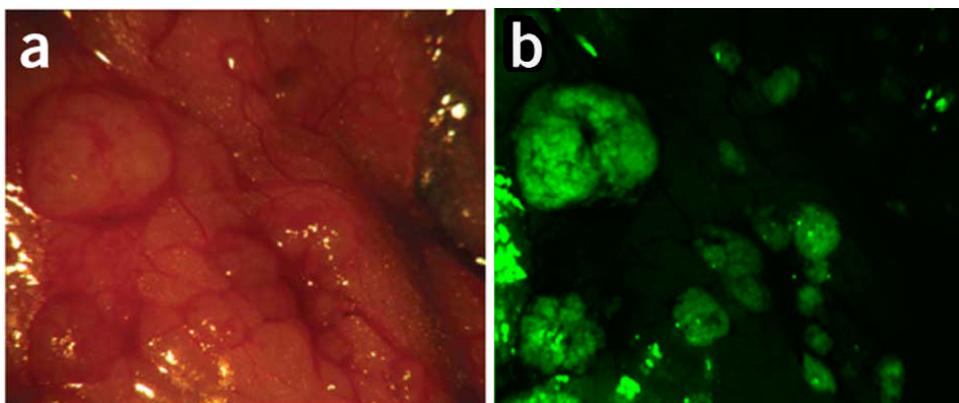


Figure 1: A skilled surgeon can identify about 7 tumors in a), while b) clearly shows there is are a lot more tumors and a lot more smaller clusters as well. Adapted from [9].

Chapter 1

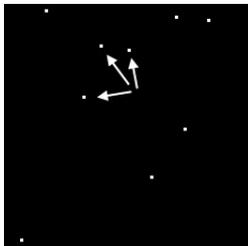
Theory

1.1 Camera Calibration

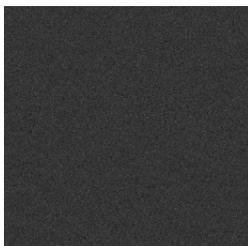
The main reason for calibrating the camera is to get an idea about the kind and intensity of the noise present in an image and to find ways to remove or suppress this noise. Although random noise is - by definition - random and very hard to remove, a part of the noise has a pattern. This can be removed by identifying the different kinds of noise, and create a 'mask' of the constant noise that can be subtracted from all other images.

1.1.1 Noise reduction

The first reason to calibrate the camera is to identify the different kinds of noise to get an idea about the different kinds of noise effects present. Some of these effects are random while others are constant in time. These effects can be divided into five categories [2]:



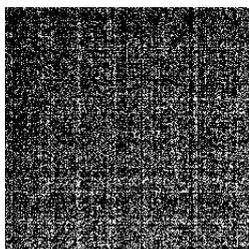
Hot pixels. These are pixels that are more sensitive than the rest of the pixels, and get brighter in time a lot faster than the other pixels. These pixels are mostly consistent (their sensitivity differs somewhat so some are not or less visible at lower exposure times) so they can be measured and subtracted from any taken picture to improve the quality of the image.



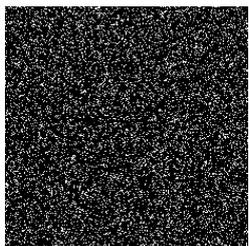
Dark current. When the camera chip is exposed in an optically shielded room, the image will still get brighter with increasing exposure time. This is because the camera operates at a certain temperature, and this causes temperature-dependent dark current (random effects cause electrons to occasionally be excited). Dark current is mostly linear in time, so one dark current frame can be used for different exposure times. Dark current noise does contain a minor amount of random noise as well.



Flat field. Due to mechanical and electronical limitations, it is possible parts of the image are less illuminated than others. Reasons for this are for example the ratio between exposure time and the time it takes for the shutter to open and close. Normally this field creates some kind of gradient on the image that can easily be corrected with a mask. The flat field is usually linear in time.



Read noise. This is noise caused by electronics on the sensor and inside the camera. Ideally it is completely random, but most of the time there is some kind of pattern in it. This pattern can be constant (stripes, waves, spots on the image) or can be moving (like stripes every 100 pixels that are moving to the right or left). The constant read noise can easily be removed, but the moving read noise is very difficult to remove. Another source of read noise is the quantization process.



Shot noise. Every time a picture is taken the amount of photons detected by each pixel is slightly different. This causes a noise both in the entire image as in the pixel over time. Although this noise cannot be removed instantaneously, the significance of this noise can be measured to indicate the precision of the measurement. This noise can normally be removed or at least suppressed by averaging multiple frames.

The effect of hot pixels, dark noise and the flat field on an image is mostly linear in time, so this can be measured once and then extracted from every image, using a multiplier to work with longer or shorter exposures. To measure these constant phenomena, a large number of images is taken by the camera in a completely dark room, with (ideally) an exposure time equal to the exposure time that will be used in the fluorescence measurements. By averaging these images all read and shot noise is removed, leaving only the constant noise patterns, and the resulting image (the constant noise mask) can be subtracted from any image taken during the fluorescence measurements. Some cameras even have a built-in feature (for example, Flat Field Correction) doing exactly this: the camera is calibrated by the manufacturer and has a dark image stored in its memory that represents the flat field and dark noise at certain camera settings. Using an algorithm this dark image can be adjusted for the settings used for each frame that is taken.

Although the read and shot noise are not constant through time, the significance of the noise can be measured, giving the minimal intensity needed for a signal to be detectable. There are two ways to do this. The first is to average a number of noise frames, resulting in an average noise distributed image. One should keep in mind that this gives the average noise intensity and not the exact noise (intensity and location), since the latter is random. Another method is to calculate the noise distribution for each pixel. This takes a bit longer to compute, but should result in a better definition of the noise in the image and will give an indication of any existing pattern in the noise (in a dark room this would mean there is some sort of effect or object in the vicinity contaminating the image).

In both cases the dark mask that was created using the constant noise can be subtracted first, but only if enough images are used or if the exposure time is sufficiently long. Otherwise the subtraction of the mask will not only remove noise, but will also add noise where there is no actual (or a different amount of) noise in the image. This is because the pattern in noise is like Young's double-slit experiment with particles: when the exposure time is short, no pattern is recognizable in the result. But when averaging a lot of short-exposure frames, or when using a long exposure time, the pattern becomes visible. So although there could be a pattern, subtracting this pattern from a short-exposure frame will do more harm than good. So the noise intensity measurement should be using the same approach as the final measurements: the dark image should be subtracted from both the noise and the measurement images when the exposure time is long or many frames are averaged, the dark image should not be used when the exposure time is short. [12]

1.1.2 Signal-to-noise ratio

Using the random noise image obtained in the previous paragraph, the optical power of the noise can be calculated. This can be done by measuring the optical power at the position of the camera chip and take an actual image with the camera for several intensities of the LED. These measurements are then plotted in a graph with the optical power of the beam on the x-axis and the image intensity (expressed in arbitrary camera units for each pixel) on the y-axis. A line can be fit through the measurements, giving a value for "camera units per W optical power". The optical power of the noise can for example be calculated by taking the random noise image and dividing its sum of counts by the "camera units per W optical power". Another method is to determine the standard deviation of the noise. By using this method, any constant noise is not taken into account, but all random noise is. This method should therefore only be used when the camera has a flat-field correction mechanism, or the constant noise is subtracted from the image. [7]

For a more quantitative analysis of the fluorescence images, the signal-to-noise ratio (SNR) can be calculated to indicate the relation between the above mentioned noise (ideally only the random noise is still left) and the signal in the image. In imaging, the SNR is usually defined as follows [8]:

$$SNR = \frac{\mu}{\sigma} \quad (1.1)$$

With μ the mean value of the signal and σ the standard deviation of the signal. Usually not the entire image contains the signal. In that case this formula should be applied to that part of the image which contains the signal. There is not one way to calculate the SNR, so the SNR is still a relative value. One way to do this is take a portion of the image that contains the wanted signal, and calculate both the mean value and the standard deviation of the signal there (in Matlab this can be done using the functions *mean* and *std*). Another way is to calculate the mean value of a portion of the image containing the signal and the mean value of another portion of the image containing only noise, resulting in this formula:

$$SNR = \frac{\overline{signal}}{\overline{noise}} \quad (1.2)$$

Ideally the noise on the signal should be the same as the noise on an area where there is no signal, but in reality the noise on the signal is often a bit bigger because the signal itself is a bit noisy as well. However, because we want to measure the signal against the (as dark as possible) background, we chose to measure the mean of the signal in the portion of the image where a signal is available and the noise in a portion of the image where no signal is available. Even then, there are two 'no-signal' area's in most images: one is almost completely black and contains only camera noise. The other is the portion of the image where paraffin (in which the tumors are embedded) can be seen: the light from the dye bounces through the paraffin, giving an unwanted and therefore 'noisy' signal. In most measurements, we have therefore included two signal-to-noise ratio's: the lower one corresponds to the signal-to-paraffin ratio, while the higher one corresponds to the signal-to-'pure noise' ratio.

1.2 Fluorescence and dye

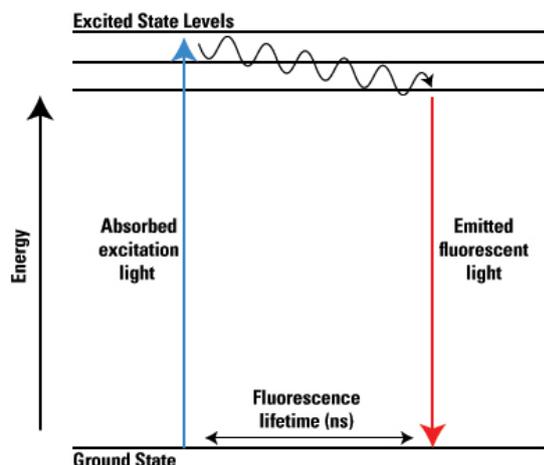


Figure 1.1: Energy scheme of a fluorescent molecule. Adapted from [6]

The second component is the fluorescent dye itself. In this case the NHS-group attaches itself to an antibody before being injected into a mouse. Once injected, the antibody attaches to proteins specific for cancer cells. [5][11]

The absorption and emission spectra for the IRDye[®] 680RD dye can be seen in figure 1.2. When the dye is illuminated with for example light of 660 nm, the resulting fluorescence will hold a broad band of wavelengths, ranging from about 660 nm in the direction of longer wavelengths. The dye used in the rest of the measurements has been tested with an absorption spectrometer and the spectrum was the same as in the dye specifications. For the measurement of different dye concentrations and the minimal resolution, the dye was dissolved in water instead of PBS (Phosphate Buffered Saline). PBS is basically water with added salts and minerals to mimic body fluid, but because the salts and minerals do not affect the dye, there is no reason to use this special dissolvent when water produces the same results.

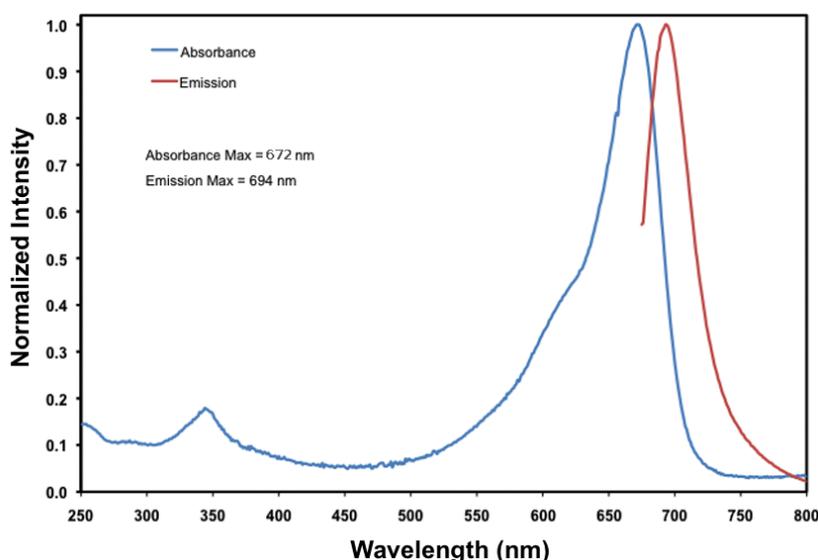


Figure 1.2: Absorbance and emission spectra for the IRDye[®] 680RD NHS Ester dissolved in PBS. Both spectra are normalized separately. Adapted from [3].

1.3 Methods for dye detection

1.3.1 Filtering of light

The main problem of fluorescence is that the light emitted by the dye is often many orders of magnitude smaller than that of the LED, mostly because the light is emitted in all directions from the dye, so only a small part of it will be detected by the camera. So, it is important that the dye has a high fluorescence intensity and that the camera will catch more fluorescence light than excitation light. To minimize the amount of light from the excitation LED falling on the camera chip two filters can be used, separating the LED-wavelengths from the fluorescence wavelengths. The first filter is positioned immediately behind the LED, making sure no wavelength longer than the absorption maximum wavelength are passing through the filter. This can be a lowpass filter, but a bandpass filter works almost as good, because the LED has a limited spectrum of wavelengths it emits. The second filter is placed right in front of the camera, letting all wavelengths caused by fluorescence through, while blocking any unwanted wavelengths from the surroundings. This can be a highpass filter but once again, a bandpass filter works almost just as good.

Since the absorption and emission maximum of the dye are only 22 nm apart, it is difficult to find filters with peaks sharp enough to create a gap between the cut-off wavelengths of the lowpass and highpass filters. In the end, two filters were found that had been pre-made for the Cy5.5 dye, which shares many properties with the 680RD dye. Because the laser in the setup is a 660 nm laser, the first (lowpass) filter should not be needed and only the second filter (690-745 bandpass filter) was ordered. After some measurements it turned out that too much laser light reached the camera. Because of long delivery times, an already available filter was selected to put in front of the LED. Even though this 650SP filter appears to block all LED-light, measurements showed that more than enough light passed through for fluorescence, while blocking virtually all reflection.

In figure 1.3 the spectral characteristics of the dye is combined with the transmission bands of the chosen filters. The LED wavelength is indicated with a red arrow.

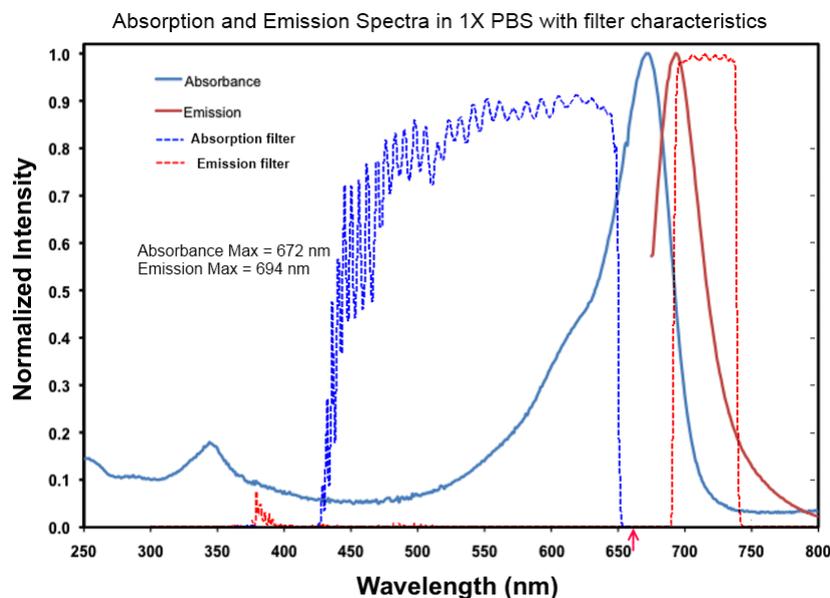


Figure 1.3: Absorption and emission spectra of dye, plotted together with the optical filter bands. The left band belongs to the filter in front of the LED, the right band to the filter in front of the camera. Please note: the spectra of the emission and absorption are normalized separately. Adapted from [3]

1.3.2 Gating

The use of filters can pose some problems. Light rays can reflect and bounce between filters and the sample, resulting in false imaging. Also, the fluorescent light that does come through the filter will be slightly weakened, especially if the wavelength lies near the cut-off wavelength. Finally, using filters makes you throw away a lot of optical power that you could have used. This becomes very clear in figure 1.3: some fluorescent light is blocked by the filter. (Red line under 690 nm and over 745 nm). Also, a lot of the LED power gets lost in the first filter. In the end, a lot of optical power is wasted, resulting in either more advanced cameras or a LED with higher power being needed to compensate for the loss in power.

Although the use of filters gives good results even in setups that lack timing precision, there is another method that could be used if a setup has enough timing precision. This method is called gating. Gating is based on the fact that the source signal (the LED) and the resulting signal (the fluorescence) do not arrive at the camera at exactly the same time. There is a slight delay because the fluorescence is not instantaneous. The delay is typically a few nanoseconds to a few tens of nanoseconds.

The concept is relatively simple: the LED is used to illuminate the sample with fluorescent dye, which will light up. Then, the LED is turned off very fast, in the order of nanoseconds. Due to mechanisms in the dye molecules, their 'turning off' will be delayed a little bit longer (one can think of it as blocking the fuel feed to an engine: it will not stop immediately but first use up the remaining fuel that is present in the engine's own volume). While the LED is turned off and the fluorescent molecule is still emitting, a frame is captured by the camera (see Figure 1.4). Like this, there is no need to filter away LED light because the LED is turned off before the frame is taken.

Despite the effective splitting of the LED and fluorescence signal, there are a couple of problems that make this method very hard to realize. Those problems are connected to the lifetime of the dye (the time it will keep emitting after the LED has been turned off). It turned out to be much smaller than estimated: in the order of tenths of nanoseconds. This makes it nearly impossible to time the exposure of the camera chip synchronously to the LED, simply because of the response time of both the camera and the LED. Also, the minimal exposure time of the camera is approximately 60 nanoseconds, much longer than the lifetime of the dye. Thus, while the camera will record signal for about 0.4 nanoseconds, the remainder of the 60 nanoseconds will only result in much more noise, incredibly decreasing the SNR. At last, the dye is already relatively weak: an emission burst of only tenths of nanoseconds will possibly be much too weak for our camera to detect.

So despite the fact that gating could result in very nice images, this particular dye has a too short lifetime to actually take a picture with good quality. This method will not be further used during the rest of the experiments.

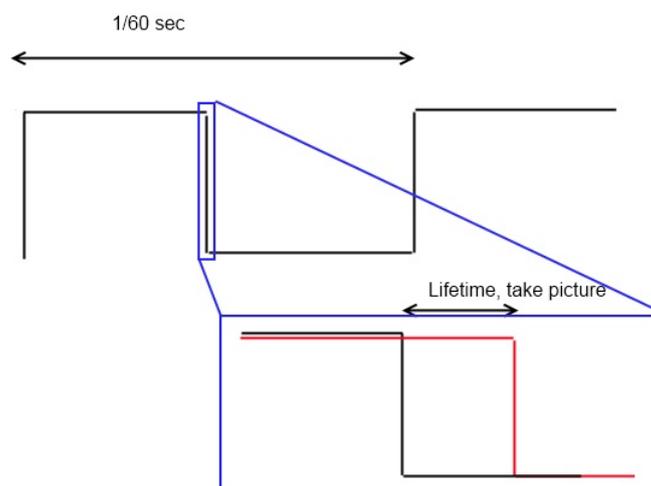


Figure 1.4: The above black square wave represents the signal from the function generator controlling the LED. With gating the camera will only record when the black signal (LED) is already low, while the red signal (fluorescence) is still high.

1.4 Signal Modulation

Detection of fluorescent signals can be a real challenge because of the generally low optical power of fluorescence. A possible technique to improve the detection of a fluorescent signal is lock-in detection. Here, the excitation light is modulated at a certain frequency, resulting in a modulated fluorescence signal. Then, detection is performed at that particular frequency only. As a result, the full signal will be detected, while noise from all frequencies except the used frequency will be weakened or even nearly completely cancelled.

As shown by Brian van Hoozen the intensity of the measured noise is given by:

$$I_{noise} = \int_{-\infty}^{\infty} \left[\frac{4N(\omega)}{i\omega} e^{\frac{i\omega\tau}{2}} \sin^2 \left(\frac{\omega\tau}{4} \frac{\sin\left(\frac{N\omega\tau}{2}\right)}{\sin\left(\frac{\omega\tau}{2}\right)} \right) \right] d\omega \quad (1.3)$$

$N(\omega)$ is an arbitrary function giving the frequency dependence of the noise, τ is the modulation period and N is the number of modulation periods performed.

The strength in this technique lies in the fact that for high values of N the second part of the function

$$\frac{\sin\left(\frac{N\omega\tau}{2}\right)}{\sin\left(\frac{\omega\tau}{2}\right)}$$

approaches the Dirac delta function with the peak at the chosen modulation frequency. Multiplying this part with the rest of the function - containing $N(\omega)$ - thus filters out all frequencies except for the selected frequency.

For the *in-vivo* imaging of tumors, a high value of N is impossible. The used camera can take a maximum of approximately 150 frames per second (fps) at the desired resolution of 320*240. Since two images are needed to get one frame (one image with the fluorescent signal, one with the noise for subtraction), this recording speed drops to 150/2=75 fps. Assuming that any movements of the endoscope can be frozen at a frame rate of 25 fps (which is a big assumptio), this means we can only have 75/25=3 modulation periods within one integration time of $\frac{1}{25}$ second. Even when the movement is minimal and lower frame rates could be achieved, it is desirable to have a fast enough frame rate so the surgeon can see his actions on screen in a smooth movement instead of a lagging image. Therefore, the chosen output frame rate of 25 fps appears to be a nice goal.

For the first measurements (Chapter 2.2) the frames were modulated by taking 60 frames, which would together form one large image, consisting of 10 separate images placed next to each other. Each of those ten images was created by adding and subtracting 6 frames as shown in equation 1.4.

$$Image = \frac{(frame1 - frame2) + (frame3 - frame4) + (frame5 - frame6)}{3} \quad (1.4)$$

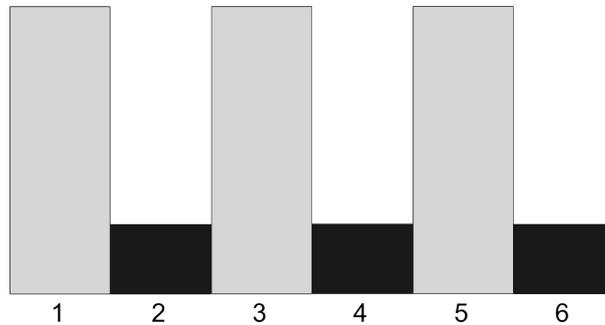


Figure 1.5: Three modulation periods, each consisting of two frames. Together they form one final image, using equation 1.4

Like this, a 2-bin system was simulated which contains three modulation periods within one integration time. Choosing the time between frame 1 and 2 to be $1/75$ second, The modulation frequency became 75 Hz and the output frequency became 25 Hz. This is also displayed in figure 1.5.

Entering $N=3$ (75 fps) into equation 1.3 results in the graph seen in Figure 1.6. As is clear, the figure looks nothing like a Dirac delta peak. However, frequencies close to 75 Hz are still weakened, providing a better signal-to-noise ratio at 75 Hz.

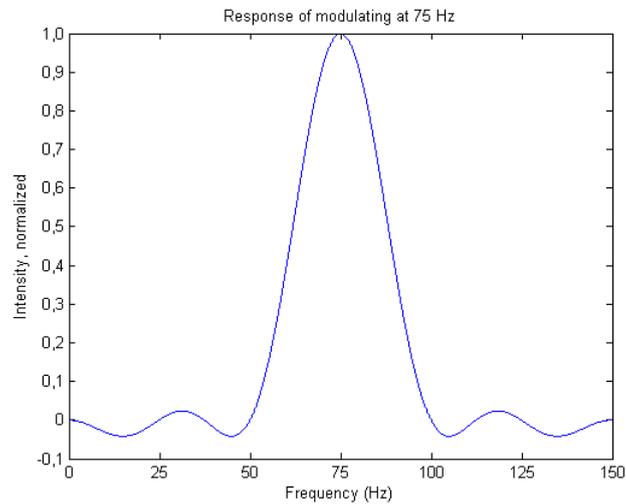


Figure 1.6: The response when modulating at 75 Hz.

1.5 Usage of colormaps

Because of noise, reflections of the material around the dye and other sources of light other than the dye itself, it is sometimes very hard to distinguish the edges of the material labelled with dye from the surrounding material. This is worsened by the fact that a computer is only able to display 256 (2^8) levels of gray. Because data received from the camera is usually more detailed than this (a normal pixel depth is 10 bits, providing $2^{10}=1024$ levels), a lot of information is lost when the image is only shown in gray colors.

To improve the detail of the image, not only the gray levels can be used, but also other colors. The computer defines a color using a combination of red, green and blue. Each of these colors can be defined with 256 levels, resulting in a total of 16 million colors. Gray colors are produced by setting the red, green and blue to the same value.

But even with the use of all available colors it is possible that the detection of especially edges in the image is very difficult. When the difference in number of pixels is low, there can still be a very clear edge, but it is not visible because the different colors still look very much the same to our eye. This can be changed by altering the density of colors per number of counts. By shifting the color density, small differences in colors can be magnified. An example can be seen in Figure 1.7. All three images come from the same raw data. The top image shows the black and white image, while the bottom left image shows a colormapped image. In the bottom right image, the colormap has been edited by shifting the red color more towards the green, which better brings out the shape of the tumor.

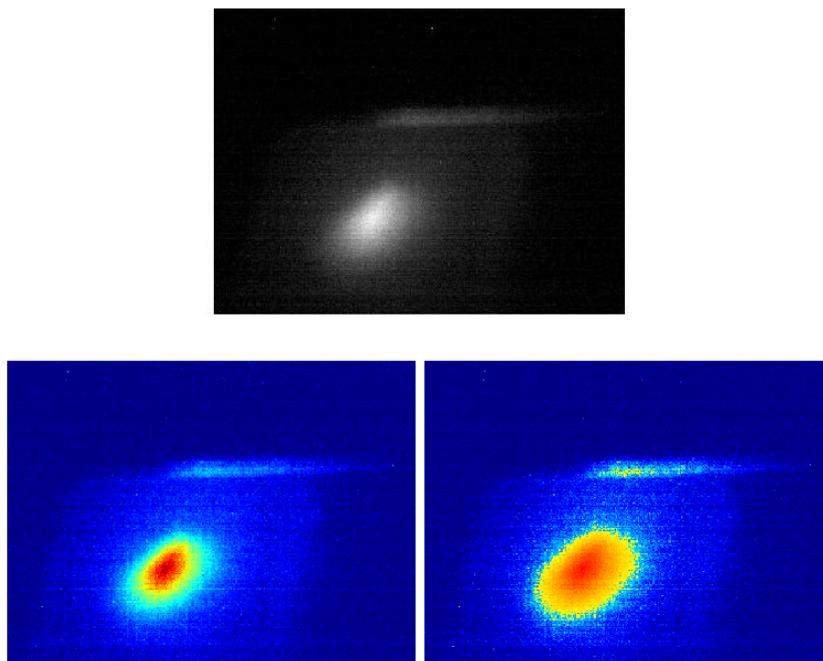


Figure 1.7: Three times the same image, but at the bottom two different colormaps were used. By altering the colormap for the bottom-right image, a small change in value results in a large change in color.

Chapter 2

Calibration

2.1 Setup

As explained in the theory, the camera has to be calibrated both for noise reduction and to get information about the sensitivity of the chip. The camera should be modulated to calibrate for the modulated behaviour of the camera and unmodulated for the unmodulated behaviour.

To measure the noise of the unmodulated camera, both dark and illuminated images are taken. A dark image is taken by closing the aperture of the camera completely. This means no light from the outside falls on the chip, thus the resulting noise is caused by the electronics and any thermal/EM-radiation. This measurement can be taken everywhere inside the optically shielded box, because no other components are needed.

The camera used during the experiments is the PixelINK PL-B761 for USB 2.0. This camera has a Flat Field Correction feature, which already removes a significant amount of noise.

For the calibration of the number of camera counts per optical power, the setup in figure 2.1 is used. A 655 nm laser is used as a light source. Using a neutral density filter and by varying the voltage on the laser, the amount of light falling on the camera chip can be altered. First, the intensity of the light is measured using a photodiode (OSI Optoelectronics PIN-10D). The measured voltage must be multiplied by a factor dependent on the wavelength of the light to obtain the optical power in Watts. Second, the camera chip is placed at the same position as the photodiode, and an image is taken. The resulting optical powers and camera counts can be plotted, and a line can be fitted through these points using best square fitting.

The next step is to determine the standard deviation of the signal of the camera. This was done by taking 10 dark images and joining them together to form one huge 'image'. Of this image, the standard deviation of all pixel values was calculated. The value of the standard deviation has been used throughout the report as a representative value for the noise. This standard deviation is per pixel, and only has to be divided by the slope of the best square fit of the plot to obtain the optical power of the noise per pixel.

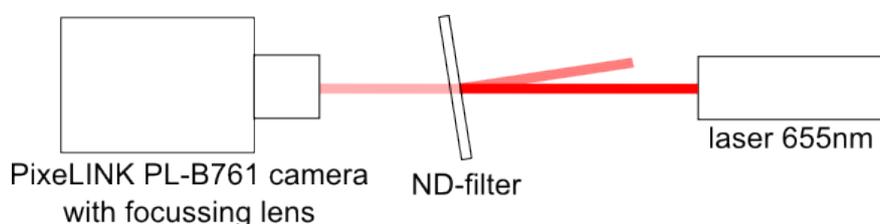


Figure 2.1: Setup for calibrating the camera. With the camera in this setup the number of counts was measured. The camera was replaced by a photodiode at the position of the camera chip to measure the optical power. The neutral density filter was used to prevent the laser from saturating the camera chip.

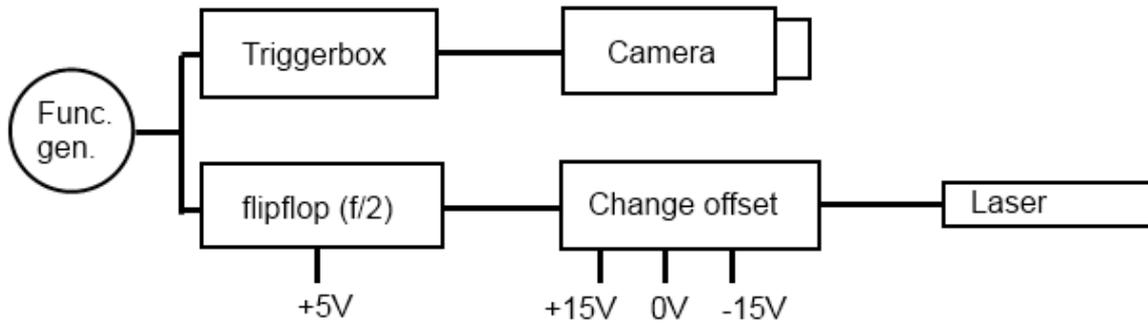


Figure 2.2: Setup for modulating the camera and LED

For the modulated measurements, some parts had to be added to the setup to make the camera and led run in sync at the desired frequency. The position of the camera and laser was the same, but the extra components were connected as shown in figure 2.2. The flipflop was used to halve the frequency of the laser, so that the camera could take a picture when the laser was on and off. The change offset box is a custom built box which let us change the offset in the signal, so that the laser was off when the pulse was low, and on when the pulse was high. The schemes of the flipflop and the "Change Offset"-box can be found in the appendix.

2.2 Results

Using the setup in figure 2.1 and by changing the voltage on the laser LED the optical power on the LED and photodiode were changed and measured. The results can be seen in figure 2.3. The best square fit ($R^2=0.9984$) lies within all the error boxes of the measured values.

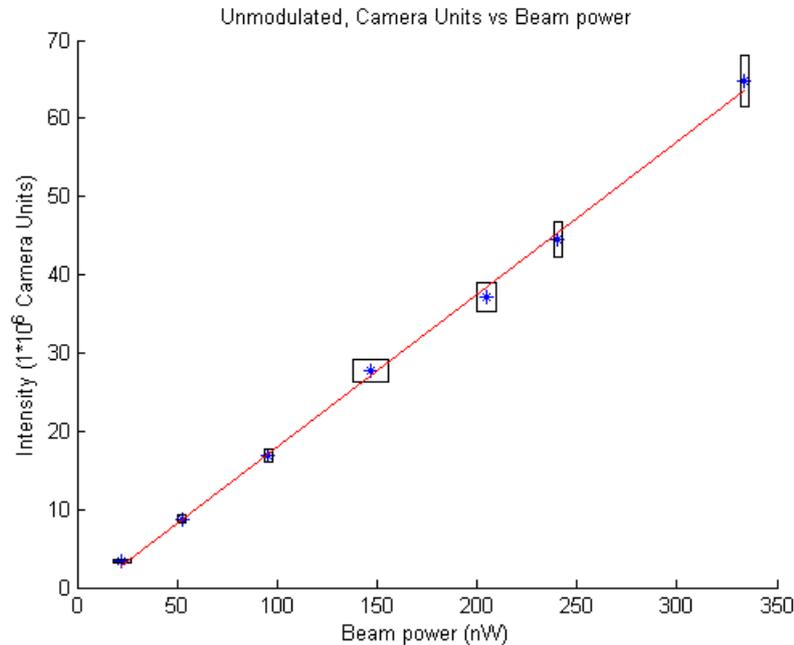


Figure 2.3: The number of camera counts (in total, not per pixel) for different optical powers at 655 nm. The slope of the best square fit can be used to calculate the power of the noise using the number of camera counts of the dark image.

Using the method described in section 1.4, a total of ten dark images was used to calculate the standard devi-

ation of a pixel. By dividing the standard deviation by the slope of the best square fit, the resulting power per pixel was calculated. For the modulating method, the 60 frames were processed into 10 images of which again the standard deviation of a pixel was calculated. This resulted in the following table:

Table 2.1: Initial calibration results

Noise in	Unmodulated	Modulated (75 Hz)
Counts/pixel	36	31
Power/pixel	0.0016158 pW	0.0013869 pW

A second set of measurements was done under noisier circumstances, where both the aperture and the optically shielded box were fully opened.

Table 2.2: Second calibration results (increased noise levels)

Noise in	Unmodulated	Modulated (75 Hz)
Counts/pixel	705	86
Power/pixel	0.031556 pW	0.0038796 pW

As can be seen in figure 2.4, noise levels are clearly quantized. Even though there are many pixels with values 64, 128, 192 and some with 256 or 320, most of them are zero, resulting in the low noise levels shown in table 2.1.

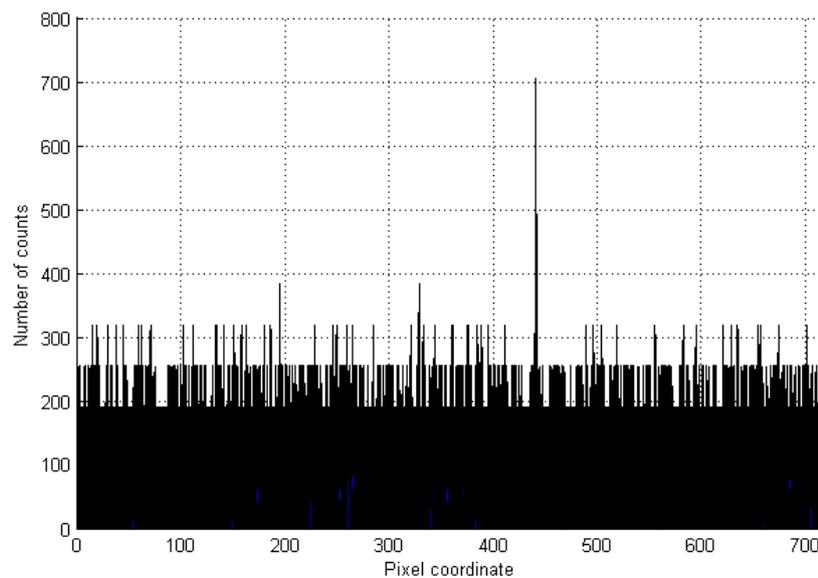


Figure 2.4: Front view of the 3D-surface plot of a dark image

2.3 Discussion

As can be seen from the results, two sets of measurements were performed. The first was under perfect conditions, in an optically shielded box with the aperture closed (Table 2.1). Since noise levels were very low in those conditions and it is very probable that noise reduction methods work less effective in low noise environments, a second set of measurements was taken for which the aperture and box were both opened (Table 2.2).

As can be seen, modulation in the first measurement decreased the noise by 14.3%. In the second measurement, modulation decreased the noise by almost 88%. It can be assumed that at levels with even more noise,

the improvement caused by modulation will be even better. However, since the goal is low noise and not high improvement, it is best to work with the optically shielded box while keeping in mind that with the low levels of noise present modulation will not bring enormous improvements.

The overall noise power seems to be quite small, but another calculation confirmed the order of magnitude of this value. This calculation can be found in the appendix.

Since 64 counts/pixel (10-bit pixel data saved in 16-bit integers) comprise one quantization level and the noise levels are less than half a quantization level in the first measurement, it seems that the random noise intensity is very small. It will be difficult to decrease it by common techniques such as modulation since subtracting from very small values can result in negative values. Averaging multiple images might give a much better result, but will certainly have a negative effect on the frame rate.

This led to the next step: making noise masks. Although the measured noise appeared to be completely random, averaging over many images showed a distinct pattern. The random part of the noise means that dark-frame subtraction will be less efficient than first estimated: clearly defined offset such as an illuminated background will be removed, but random noise will not. In fact, subtracting the random part from another random part will most likely result in almost twice the amount of noise. This is explained in figure 2.5. Because of this effect, it was also tried to set all negative valued pixels to zero after the subtraction. This did decrease the amount of noise on the unilluminated parts a little bit but did not improve the parts of the image with signal.

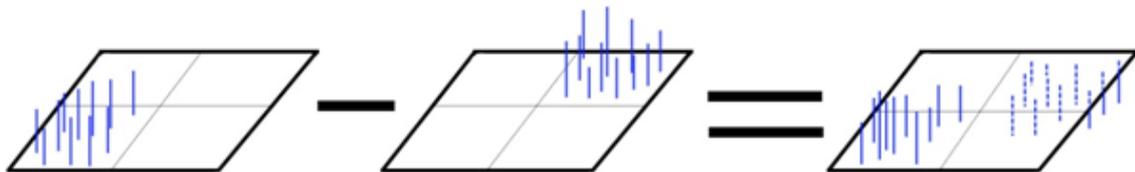


Figure 2.5: Increase of noise when subtracting a second dark image. The location of the "Noise pixels" has been changed for clarity.

Averaging the images will result in better results for two reasons: averaging flattens the random noise and also better brings out the offset noise hidden in the random noise. This offset noise can then be subtracted using a similarly averaged dark image. In the next chapter a noise mask will be created to evaluate the effect of this mask on the image quality.

Chapter 3

Noise Mask

3.1 Method

According to the previous chapter it could be very beneficial to create a noise mask that can be subtracted from a frame to reduce the noise in that frame. In this chapter this hypothesis is tested. As already explained before, the amount of averaged frames (or the exposure time of a frame) can change the effect of a noise mask considerably. Because of the movement of the body and endoscope, the exposure time will be kept short, at 6 ms. The hypothesis will be tested by changing the amount of averaged frames. First a mask will be created using a sufficient amount of frames to make the pattern constant. After that, the noise mask will be subtracted from just 1 frame, an average of 10 frames and an average of 99 frames. By comparing the averaged image and the average image with noise mask subtraction the effect of the noise mask can be analysed. If this effect is still visible and actually improving the result for a small number of averaged frames, it is possible to use the noise mask in the fluorescence measurements.

3.2 Results

As said, a consistent mask should be created using a minimal number of frames, such that the difference in mask between that number of frames and a mask created with one frame less is smaller than a certain percentage. Because there are 1024 quantization levels and the mask can be assumed constant if there is a maximum difference of 1 quantization level, the difference in the mask should be smaller or equal to 0.1%. Figure 3.1 shows the difference in percentage between two noise masks with n and $n-1$ averaged frames. The red line shows the maximum percentual difference needed to create a constant noise mask. So by averaging 1000 dark frames a noise mask can be created. The graph is specifically for an exposure time of 6 ms. Different exposure times give roughly the same graph, indicating that averaging 1000 frames of the same scene (in this case a completely dark room) will give a constant mask.

Figure 3.2 shows the mask resulting from the described method. One can see horizontal and vertical lines forming a pattern that was not visible when

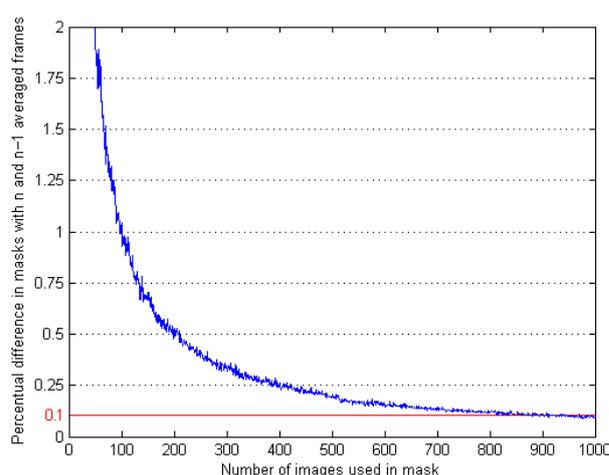


Figure 3.1: Relative difference between a noise mask created using n dark images and a noise mask created using $n-1$ dark images at 6 ms exposure time. For a constant noise mask the difference should be less than 0.1%.

averaging over only a small number of frames.

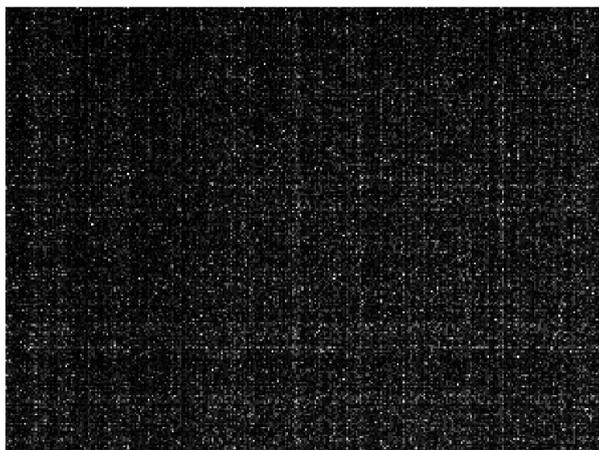


Figure 3.2: The mask resulting from averaging 1000 6ms-frames. For the report, the mask has been edited to have a larger contrast. In reality, differences between pixels were more subtle.

It turned out that for every exposure time, a separate mask should be created. Although one could argue that a 20ms mask is simply the 6 ms mask multiplied by 20/6, figure 3.3 shows this is not the case with this camera for short exposure times under 200 ms. This is most likely caused by the internal (unknown) FFC processing.

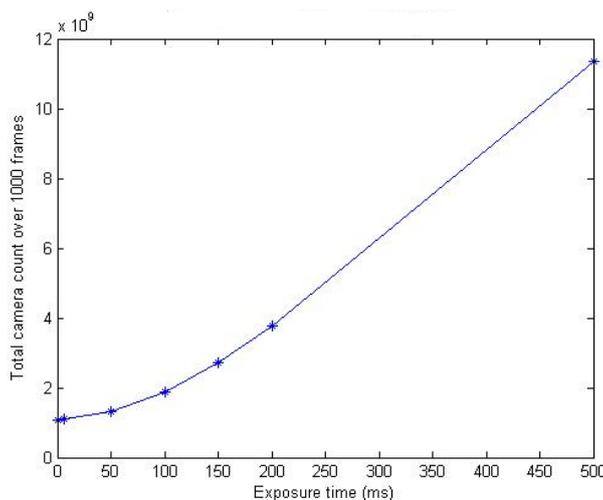


Figure 3.3: The total number of counts as a function of exposure time in a dark image are not linear as one would expect.

Next, the mask is tested. All images are taken at 6 ms exposure time, but the amount of images that is averaged is changed. The results for one frame (Figure 3.4), ten averaged frames (Figure 3.5) and 99 averaged frames (Figure 3.6) are shown. The left subimage shows the resulting image without noise mask subtraction, the right subimage includes the subtraction. The images show what was already known from the calibration measurements: averaging more frames results in an image where any noise pattern is more distinctive, so a noise mask will have much more positive effect on the image quality. For just one frame the quality does not really improve, except for the removal of hot pixels.

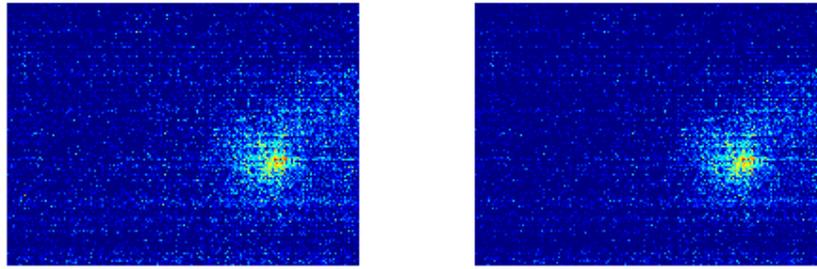


Figure 3.4: No averaging at all. Subtracting the noise mask from one image only removes the hot pixels, no real further improvement.

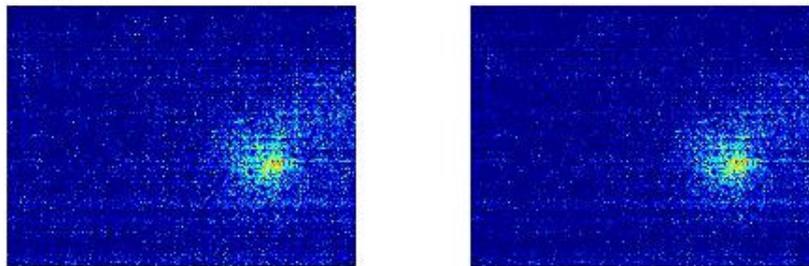


Figure 3.5: The effect of subtracting the dark image from the image is clearly shown: the background is much flatter. Some imagination shows a noise pattern though.

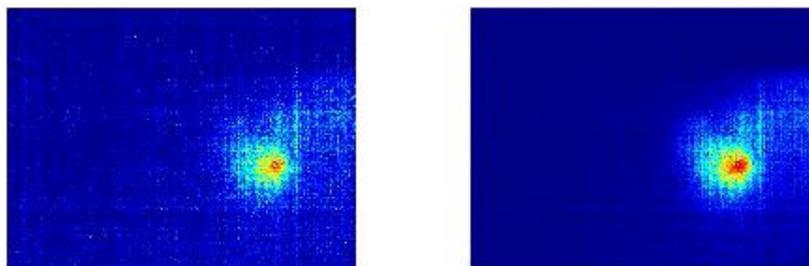


Figure 3.6: In the colormapped image the background is nearly flat when 99 illuminated and dark images are averaged.

3.3 Discussion

The resulting light still visible to the left of the red spot is unavoidable in this case because it is caused by the reflection of the excitation laser. Since the dark images show no laser light, it will not be subtracted from the original image. This problem could be solved by taking a 'dark' image using a laser with the same optical power (compensated for camera sensitivity), focus and focus position as the excitation laser, but with another wavelength that does not excite the dye. This laser should be modulated as well, opposite to the excitation laser so they do not illuminate the image at the same time.

Two things become clear from the acquired measurements. The first is that, as expected and already known, averaging over multiple frames gives better result, since it 'flattens' the random noise. The second is that

the noise mask subtraction gives better results when more frames have been averaged. This can easily be explained: when taking just one frame, the 'random noise' in the image will be very different from the noise pattern in the noise mask. When averaging over multiple images, this 'random noise' slowly transforms into the final pattern, so the noise in the image and the noise in the noise mask look more alike, making subtraction more effective.

For practical reasons it is nice to have a framerate of at least 15 fps, because otherwise it gets hard to work with live images. Taking images with an exposure time of 6 ms, this means that you can use 5 illuminated and 5 dark images for the averaging and mask. This is even less than the amount of frames in Figure 3.5 (which uses 10 illuminated and 10 dark frames) so the choice for the excitation wavelength and the filter is still very important to remove as much unwanted light as possible.

Because images where noise mask subtraction is applied do not always give a realistic view of the measured frame, all images in the following sections will not be averaged and/or cleaned up using a mask. When they are averaged or cleaned up, this will be specified in the figure caption.

Although the original goal of this bachelor's assignment was to find out whether modulation can be applied at much lower frequencies than 1 Mhz as well, other methods to improve responsivity and resolution were also researched because of the conclusions so far: there are two main sources of noise, both of which can not be removed by low-frequency modulation. These sources are:

1. **Low-power random noise**

This noise was very low-power (less than one quantization level) and also random, meaning that modulation (which basically is some fancy dark image subtraction) will have no effect unless we can take averages of a high number of frames. A high number of frames requires a camera with a high fps-value, which beats the purpose of using low-frequency modulation.

2. **Laser reflections**

Some of the laser light reflects very strongly and is visible even though it passes through the filter. Modulation is not going to help in any way.

Chapter 4

IRDye[®] 800CW

4.1 Equipment

Due to long delivery times for the ordered parts (IRDye[®] 680RD dye, corresponding filter and LED), the dye used by Brian van Hoozen (IRDye[®]) was used for the first measurements. Although the goal is to use the IRDye[®] 680RD dye, the setup could already be tested and improved with this old dye. The original dye was already expired, but there were still some usable samples of a tumor labelled with IRDye[®] 800CW, so these were used for the measurements in this chapter.

A number of measurements were taken with both a 785nm laserdiode and a 655 nm laser. The 785nm laserdiode also had a lens in front, which was used for some of the measurements. The 785nm laserdiode was a high-power laser (1W) but its wavelength was very close to the cut-off wavelength of the used filter. The 655nm laser has a much lower power (40mW) and also less absorption in the dye, but was very far away from the cut-off frequency, thus greatly reducing unwanted illumination on the camera's sensor.

Later, when all parts had arrived, measurements were done with the 660nm highpower LED (Roithner LED660-66-60), a 650 nm short pass filter from 3rd Millenium in front of the LED and a 716/40 nm BrightLine[®] filter in front of the camera. Even though the LED was far away from the optimal absorption frequency for the IRDye[®] 800CW, the results are still very detailed because of the high power of the LED.

In order to test not only responsivity but also resolution, a narrow equilateral triangle (base of 1.5mm, both sides 22.7 mm) was cut out of a black paper and then stuck to the tumor sample. The smallest tumor detectable was then deduced from the smallest point of the triangle that could still be seen. In order to improve the resolution of the camera for a better result, a 200mm positive lens was put between the camera and sample, increasing the size of the image on the camera chip so a larger part of the chip could be used (752x480 pixels instead of 320x240 pixels used for imaging most of the tumor images).

4.2 Setup

As stated before, modulation of the frames does not give any improvement of the image for low number of frames, so all components needed for the modulation (square wave generator, frequency divider, camera trigger hardware, etc.) can be left out of the setup. The resulting setup is shown in Figure 4.1. The focussing lens on the camera was a Cosmocar Television Lens (16mm 1:1.4) with a Tiffen Sky 1-A filter (UV-filter) attached to it. Figures 4.3 and 4.4 were taken with this setup. Figure 4.5 was created by replacing the 785 nm laser by a 655nm laser. The 785 nm laser performed at 11,6 V and 101 mA, while the 655 nm laser performed at 3.3V and 0.45A.

The results in figure 4.6 and further on are the result of the setup as shown in figure 4.2. Here the 660 nm powerLED was used at 12V and 0.42A, together with a 650 nm lowpass filter to create a gap between the wavelength

range from the LED and the emission range from the sample (the latter is done with the already present high-pass filter). Since the powerLED's beam was extremely divergent, a 30mm positive lens was used to focus the beam on the sample so that the entire sample was illuminated. Please note that two different tumor samples were used for both setups.

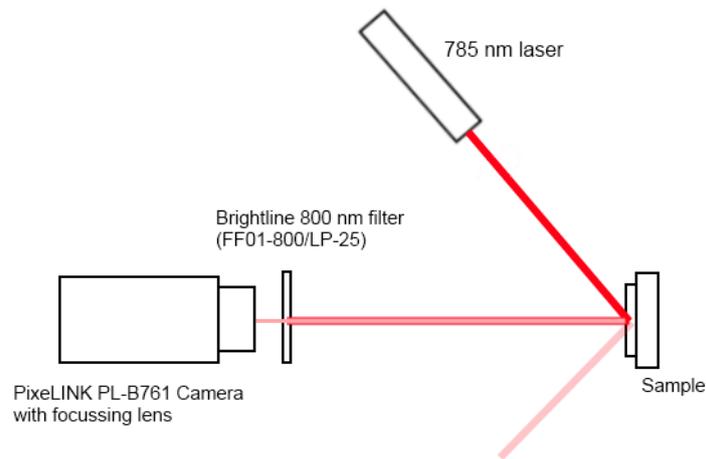


Figure 4.1: The used setup for earlier measurements of the IRDye[®] 800CW

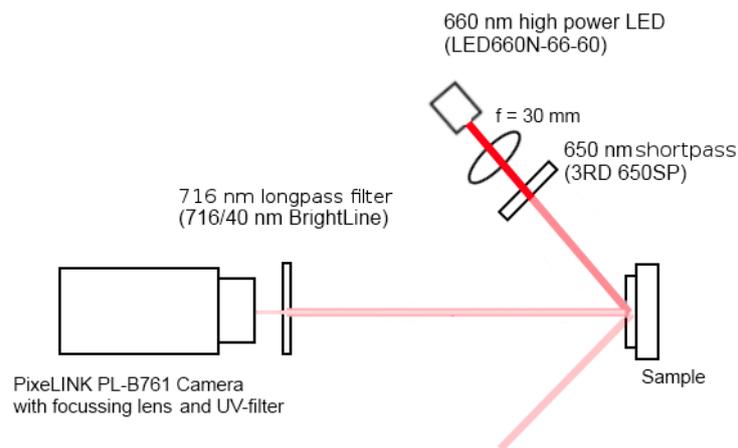


Figure 4.2: The used setup for later measurement of the IRDye[®] 800CW

4.3 Results

Figures 4.3 and 4.4 are both images taken using the 785 nm laser. The improvement when increasing the exposure time is very clear, including the fact that the noise is much more smooth than with short exposure times. Using this result switching the laser with the 655 nm laser was expected to improve the quality even further. But the problem with the 655 nm laser was that the beam had too much of a Gaussian spreading: the beam was very strong in the center, but quickly decreased in intensity toward the outside. This meant that not the entire tumor could be illuminated evenly, resulting in more of a spot than the actual shape of the tumor.

The 660 nm laser in the setup in figure 4.2 did improve the quality though, as can be seen in Figure 4.6. Even with one image taken at 6 ms the shape of the tumor is very distinct, and averaging of six frames flattens a lot of the random noise creating an even better result. One can even distinguish details on the tumor itself.

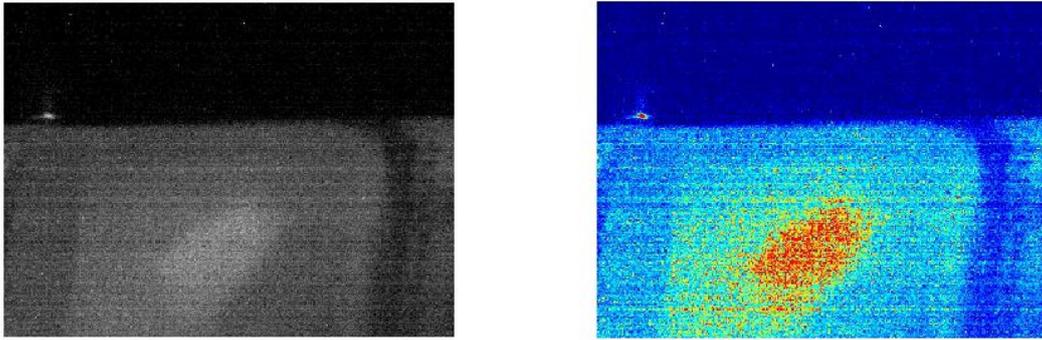


Figure 4.3: A tumor sample embedded in paraffin labelled with the dye. The camera is exposed for 50 ms. The image roughly shows the shape of the tumor, but still there is a lot of background light. The calculated signal-to-noise ratio was 7.9 or 17.8.

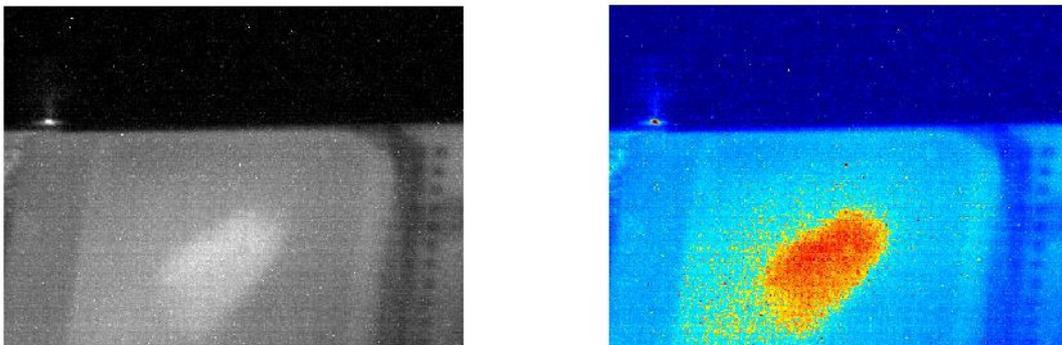


Figure 4.4: Extending the exposure time to 200 ms gives a higher signal-to-noise ratio, so the contours of the tumor are much more defined. The calculated signal-to-noise ratio was 17.4 or 23.7.

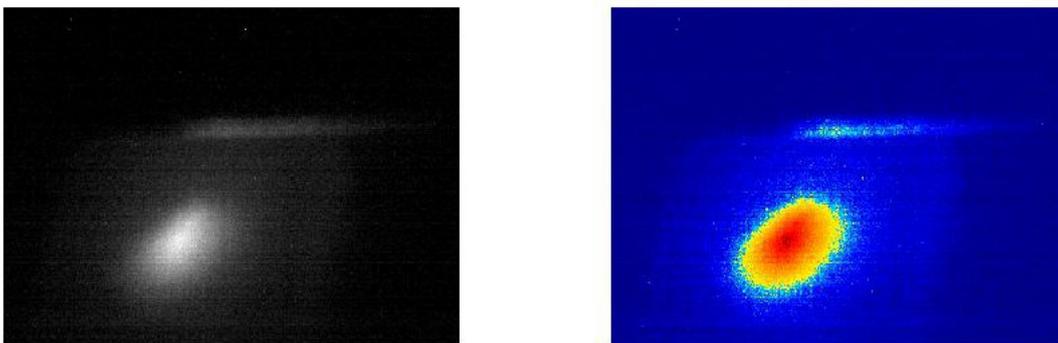


Figure 4.5: The tumor illuminated for 50 ms with the 655 nm laser. The tumor is visible but the shape is not detailed enough when compared to figure 4.4. Due to the lower wavelength of the laser, noise levels are much lower. The calculated signal-to-noise ratio was 23.7 or 58.5.

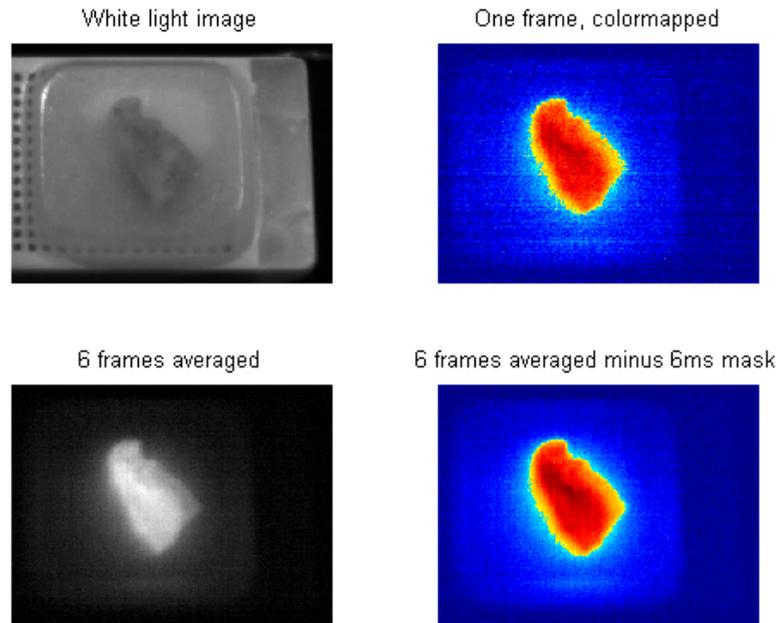


Figure 4.6: Using the powerLED, the tumor is very well visible. The signal-to-noise levels for the single frame, 6 averaged frames and 6 averaged frames minus the 6 ms mask are respectively 24.6 & 35.9, 36.8 & 67.1, and 36.4 & 62.5.

Figure 4.7 shows the result from the measurement with the paper wig. Calculations with a reference image having a ruler in the image show that in the current setup, 3.15 camera pixels correspond to 100 micrometer

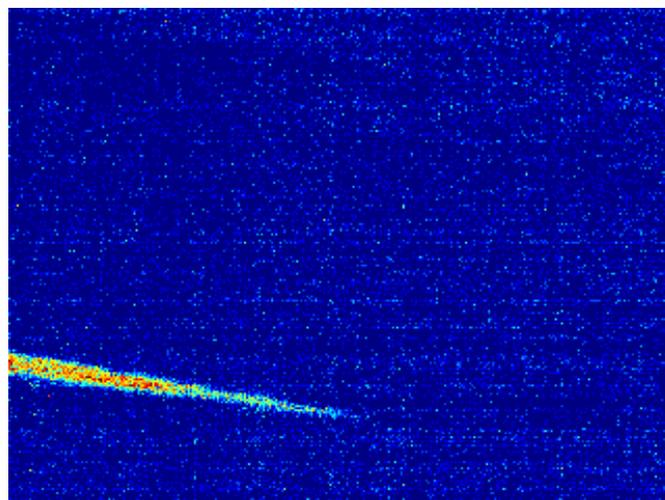


Figure 4.7: The tip of the V-shaped cut.

in length. Zooming in on the tip (figure 4.8) gives us a clearer view on the resolution. Opinions may vary on when the tip is no longer distinguishable from the noise, but 4-5 pixels appears to be a realistic estimate. This means the camera can detect tumors as small as 130-160 micrometer.

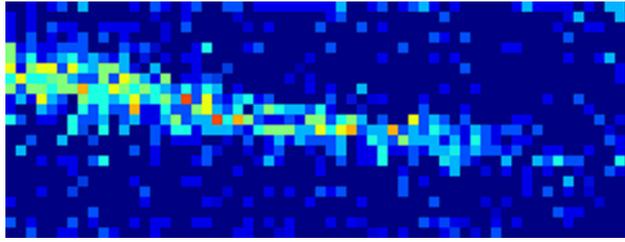


Figure 4.8: The tip of the V-shaped cut, zoomed in. One pixel corresponds to approximately $32\ \mu\text{m}$

4.4 Discussion

After some experimenting it seemed the lack of proper illumination was caused by a lens that was attached to the 785 nm laser. This caused too much focussing and also spherical aberration, so only a part of the sample could be illuminated. By removing the lens and moving the laser closer to the sample (because the laser is very divergent) the quality of the images greatly improved, thanks to the more homogeneous illumination (see figure 4.3). A longer exposure time removes more noise and gives a very clear image of the tumor, in which the contours are very defined (see Figure 4.4).

As can be seen in figures 4.3 and 4.4 there is still a lot of background light: the paraffin and sample holder are still clearly visible, but they do not have any fluorescent dye. The cause for this is the 785 nm laser itself, because it is very broadband. For medical purposes this is not a very bad thing, because it provides an orientation and shows the contours of the surrounding tissue. But for analysis it is a challenge because it creates a lot more noise. Some of the light of the laser has a wavelength greater than 808 nm (the cutoff wavelength of the longpass filter) so it is recorded by the camera. Switching to a 655nm LED drastically reduced this background lighting, but at a cost: more optical power is needed to get the same amount of fluorescence, because the dye is less sensitive at this wavelength (approximately one tenth of the sensitivity at 785 nm [4]). The two options to compensate for this are a more powerful LED or a more focussed beam. The first option is preferred because the LED could still illuminate the entire sample instead of a small part of it. The best result we got from the 'low-power' (40mW) 655nm LED is shown in Figure 4.5.

Figure 4.6 shows the results of the setup in figure 4.2 with IRDye[®] 800CW samples. Surprisingly, we got pretty good results. The top left figure shows the shape of the tumor, as seen under normal white-light illumination. The top right figure shows the colormapped readout at 6 ms. The bottom left picture shows the average of 6 frames. The bottom right picture shows the average over 6 frames with the 6ms dark image subtracted. As can be seen, this does not improve the image, because the noise in the image and mask are still too different: the only way to obtain a better result is averaging over more frames, so noise mask subtraction has more effect.

But even for the single frame at 6 ms, it is clear that without colormapping (let alone editing the colormap) the contours of the tumor are still clearly visible.

In terms of signal-to-noise ratio, all images were surprisingly clear. Even figure 4.3 shows a ratio of at least 7.9. Better measurements such as the 6-frames-averaged image in figure 4.6 have a signal-to-noise ratio of almost 37. With high ratio's as these, it is clear that lower illumination powers or lower dye concentrations should still yield clear results.

The images also show that the signal-to-noise ratio is not the only important aspect. With a ratio of 7.9 one would expect a nice image, clearly showing the location and shape of the tumor. However, when looking at figure 4.3 it can be seen that although the ratio is decent, the resolution is a bit disappointing: The tumor is shown as a blob and no real shape can be deducted from the image. Figure 4.6 shows a much clearer shape of the tumor, but this is not only because of the higher signal-to-noise ratio. The fact that a stronger and more evenly spread light was used will have played a role as well.

With the setup used for the resolution measurement, it is shown that tumors with a minimum size of about 160 micrometer can be detected. Averaging and subtraction of an (averaged) dark frame increases resolution even more, but was not applied in this case because it would be unrealistic. Averaging over 30 frames yielded great results, but would require a frame rate much higher than possible with this camera. It is suspected that

when using the IRDye[®] 680RD results will be even better, due to a better match between dye and setup.

Chapter 5

IRDye[®] 680RD

5.1 Concentration detection limit

For the IRDye[®] 680RD the first thing to measure was the lowest detectable concentration of the dye. Especially because no tumor samples were available. In order to measure this we prepared a set of 19 samples with concentrations ranging from 500 mg/L to only 1.9 µg/L, halving the concentration with each sample. Although the samples could be prepared in even lower concentrations, 3.8 µg/L was thought to be sufficiently low when compared to the expected concentration in mice. The estimation for this lowest value can be seen in the next section.

For all measurements in this chapter the setup in Figure 4.2 was used, which was the same setup that produced the best result in the previous chapter.

Concentration in mice

To make a comparison to the dye concentration that would be found in mice (more specifically: inside the tumor tissue), an estimation was made for this concentration using information from the UMCG in Groningen. This calculation is based on the earlier used IRDye[®] 800CW but is expected to be roughly the same for the IRDye[®] 680RD.

Values estimated beforehand are:

- A mouse receiving a subcutane injection of cancer cells develops 1 tumor.
- The size of this tumor is about 1 cm³
- A maximum of 10% of the injected antibodies plus dye attaches to the tumor.

Mice are injected with 100 µg of antibodies. On average, each antibody is attached to 2.5 dye molecules. With a molar mass of 150.000 g/mole for the used antibodies and 1166 g/mole for the IRDye[®] 800CW, this means 1.554 µg of dye molecules is injected. Ten percent is attached to the tumor, resulting in 0.155 µg of dye to detect in one tumor. Combined with its volume, the estimated concentration of dye in a tumor is 155 µg/L.

5.2 Results

In figure 5.1, nine images from samples with different concentrations of IRDye[®] 680RD can be found. All samples hold approximately 0.5 mL of solution. Variations in this amount are not caused by errors in pipetting, but rather because some vials were used in earlier experiments trying to find a minimum tumor size. Each image is normalized to itself, so sample 18 was nowhere near as bright as sample 10, but it was still very distinguishable against the background. Samples 1-9 are not included because the signal is in fact saturating the camera.

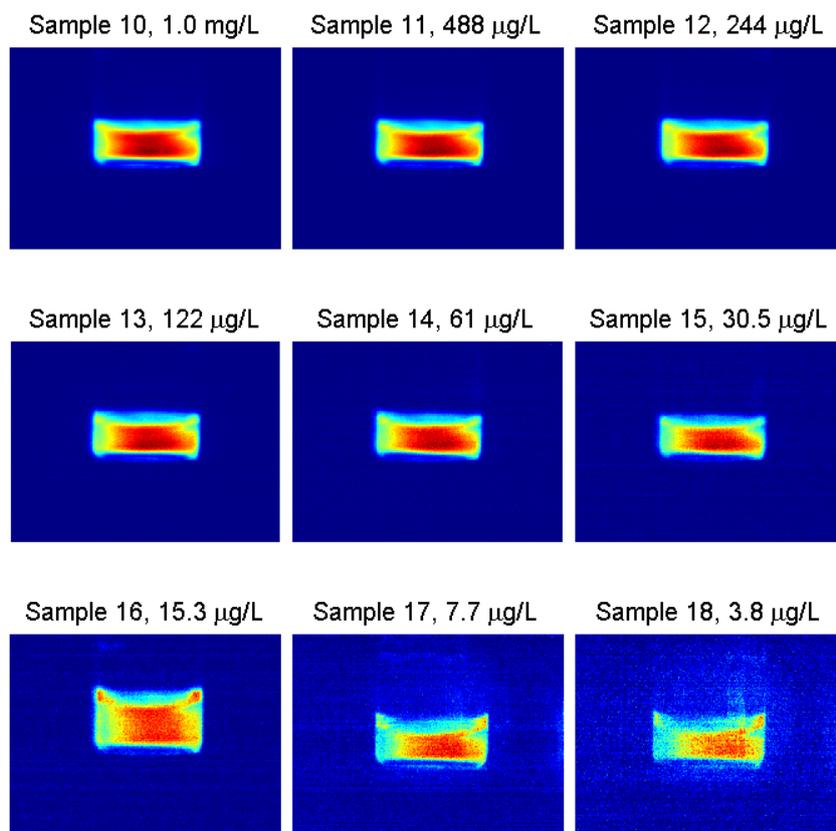


Figure 5.1: 9 samples of IRDye® 680RD

Sample 18, 3.8 $\mu\text{g/L}$

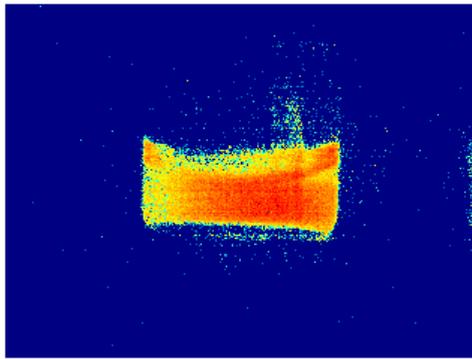


Figure 5.2: Sample 18 with an edited colormap. The signal-to-noise ratio for this data was 11.4 or 14.1, depending on whether you take the full surface of the dye as 'signal', or only the area where intensity is highest.

The results from figure 5.1 are unedited, except for the colormap. When editing the colormap itself, sample 18 can be made even clearer, as can be seen in figure 5.2.

As can be seen, samples 12 and 13, where the concentrations are 244 $\mu\text{g/L}$ and 122 $\mu\text{g/L}$, are very well detected. These concentrations are comparable to the estimated concentration in a tumor, as explained above. Even sample 18, with a concentration of only 3.8 $\mu\text{g/L}$ can still be easily detected. When taking into account that in a solution, *all* dye can be excited because the water is transparent (so the center shows more illumination because the vials are cylindrical), whereas in a tumor it will most likely be only the dye closest to the surface, it is very handy that a much lower concentration can still be detected.

5.3 Discussion

The measurement on the low concentrations of dye shows that a concentration of 3.8 $\mu\text{g/L}$ is still detectable (colormapped in Figure 5.2), which is about 40 times lower than the expected concentration in tumors. Some experimenting with the colormapping shows that it is even very easy to colormap all images on-the-fly. At a constant concentration the image has a maximum amount of noise that is always approximately the same. So from an image at a certain concentration a constant value can be subtracted to remove this noise, and the resulting image still has a lot of noise but can be successfully colormapped.

Samples 17 and 18 show a small amount of reflection in the vial, making the vial visible in the image. This reflection is most likely caused by the fluorescence itself, because the other concentrations already show that there are absolutely no reflections from the laser itself.

Chapter 6

Conclusion

Based on our findings, several conclusions can be drawn on the noise characteristics of the camera used and on the detection of fluorescent dye.

The camera bought for this research, a PL-B761, has very low noise levels. In 10-bit mode it can distinguish 1024 different intensity levels. Each level consists of 64 'counts' due to the data being written in 16-bit files. Calibration measurements showed that in dark images, the standard deviation of the signal (and thus the noise) was only 36 'counts', which is less than one quantization level. Calculations showed that one quantization level corresponds to approximately 57 photons. Looking at the absolute noise, no pixels (except for very few hot pixels) had values over 4 quantization levels.

Even though the noise from the camera was so low, it still was the main source of noise in the used setup. As it appears, the environment has little effect on the images. This is of course caused by the fact that the measurements were performed in an optically shielded box. Since the final application of the researched methods, endoscopic tumor detection, will happen inside a human body, using the box is justified.

As a result, the detected noise is not the measurement of unwanted signals (for example, the light of TL-lamps), but really the (random) noise generated and detected by the camera when reading the image (read noise and dark counts).

Only when averaging dark images over a significant amount of frames (at least 100 are advised) the random noise will slowly start to form a fixed pattern.

Due to the fact that most noise is random noise, the effectiveness of the modulation technique is greatly limited. It is based on specific timed subtractions of frames resulting in a signal with a specific frequency being maintained while all other frequencies (partially) cancel themselves. In the case of random noise, like here, this will not work. Only when averaging over many images will the fixed pattern appear. Unfortunately, this brings along a new problem: the 100 averaged frames need to be taken within a certain time span, say 20 ms, in order to 'freeze' the image and any movement in the body. This will require a highspeed camera, which beats the purpose of low-frequency modulation.

Unfortunately, it was not possible to do measurements on tumors containing IRDye[®] 680RD. As a result some measurements were done on the IRDye[®] 800CW using the LED and filter setup that was prepared for the IRDye[®] 680RD. Even though a lot of LED-light was blocked (the LED was 660nm, followed by a 650nm low-passfilter) and the wavelengths were approximately 100 nm lower than ideal, the results were great and were even better than that of earlier dedicated setups. The best signal-to-noise ratio achieved was 36.8. The power of the LED, combined with an even illumination and camera with a high responsivity were responsible for those great results. The lowest concentration of IRDye[®] 680RD measured was 3.8 µg/L, much lower than the expected dye concentration in mice. The minimal resolution when measuring the samples with IRDye[®] was approximately 160 µm, providing a vast improvement over regular visual detection by eye. A right choice for the filters made sure that reflections did not ruin the image quality.

Chapter 7

Recommendations

The current measurements on the samples containing IRDye[®] 800CW were done using the 'mismatched' equipment. Even so the results were surprisingly good: very little noise was detected and the dye could be detected at very low concentrations. Therefore, we assume that when measuring samples containing IRDye[®] 680RD, results will be even better. At exposure times of 6 ms, the signal might be so strong that the noise is of almost no concern. As a result, lower concentrations of dye could be used. Therefore, we strongly advise to continue measurements as soon as the samples containing IRDye[®] 680RD become available.

The current method of detecting tumors separates the detected dye very well from any background signal. As a result, the tumor stands out against a virtually black (or in the case of the used colorscheme, blue) background. Although the tumor can now be easily detected, surgeons might find it difficult to exactly locate it since they have no reference: all other light is blocked. This problem could be circumvented by adding a second bright LED which illuminates the surroundings of the tumor. This LED would have to be modulated with the original LED, making sure that only one LED is turned on at a given moment. These two images can be added together.

When considering higher framerates, another improvement can be found in averaging the image over multiple frames. First, the random noise is averaged, resulting in a lower overall noise level. Second, the patterned noise that builds up over time (as described in section 1.1.1) will become clearer, making it easier to compensate for with dark image subtraction. Figure 7.1 shows two versions of the V-shaped cut as it appears in figure 4.7. The first version is the original frame where only 2 hot pixels have been removed manually. The second version is averaged over 30 frames, after which the mask produced in chapter 3 was subtracted. At 30 frames and a framerate of 150 fps, the output framerate would be only 5 fps. Since this is very low, it might be an idea to get back to the original idea of using a 1MHz Time-of-Flight camera: this time not to use its modulation properties but to acquire a high number of frames, resulting in a decent framerate when averaging over many frames..

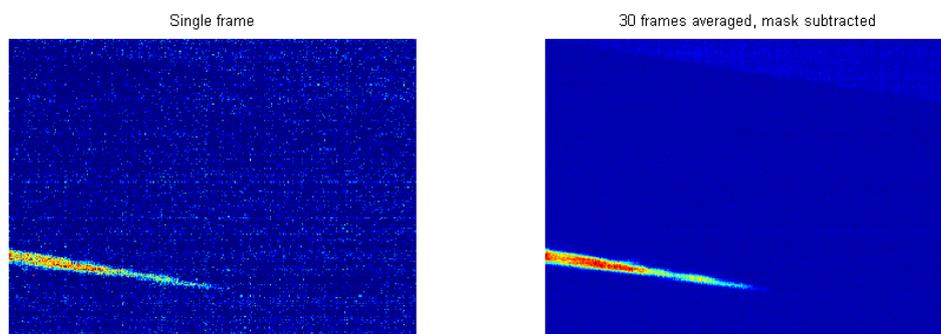


Figure 7.1: Improvement achieved when averaging over multiple (30) frames and then subtracting a dark image. The noise levels decreased by more than a factor 7. (The standard deviation of the noise signal was 35.2 in the left figure, and only 4.87 in the right)

Bibliography

- [1] Elizabeth K. Gardner.
Purdue technology used in first fluorescence-guided ovarian cancer surgery.
<http://www.purdue.edu/newsroom/research/2011/110918LowSurgery.html>, September 2011.
- [2] R. Healey, G.E.; Kondepudy.
Radiometric CCD Camera Calibration and Noise Estimation.
IEEE Transactions on Pattern Analysis and Machine Intelligence, 16(3), 1994.
- [3] Li-Cor.
IRDye[®] 680RD NHS Ester product overview.
http://www.licor.com/bio/products/reagents/irdye/irdye680rd_nhsEster.jsp.
- [4] Li-Cor.
IRDye[®] 800CW NHS Ester product overview.
- [5] Li-Cor.
IRDye[®] Infrared Dyes.
http://biosupport.licor.com/docs/IRDye_brochure_2010.pdf.
- [6] Piercenet.
Fluorescent Probes.
<http://www.piercenet.com/browse.cfm?fldID=4DD9D52E-5056-8A76-4E6E-E217FAD0D86B>.
- [7] Craig Stark.
Signal to Noise: Understanding it, Measuring it, and Improving it (Part 1), 2009.
- [8] Li Tan.
Digital Signal Processing Fundamentals and Applications.
Elsevier, 1 edition, 2007.
ISBN: 9780123740908.
- [9] G. van Dam, G. M.; Themelis.
Intraoperative tumor-specific fluorescence imaging in ovarian cancer by folate receptor-alpha targeting:
first in-human results.
Nat. Med., 17(1315-U202), 2011.
- [10] B.L. Jr. Van Hoozen.
Sensitive Fluorescence Detection using a Camera from the Gaming Industry.
Lawrence University Honors Projects, 25, 2012.
- [11] A. G. T. Terwisscha van Scheltinga.
E-mail conversation about labelling of cancer cells.
- [12] T. Kanade Y. Tsin, V. Ramesh.
Statistical Calibration of CCD Imaging Process.
Proceedings Eight IEEE International Conference of Computer Vision, 1(480-487), 2001.

Appendix

Camera specifications

The noise we measured was $1.6158 \cdot 10^{-15}$ W per pixel. Combined with an exposure time of 6 ms this amounts to a total energy of $9.6948 \cdot 10^{-18}$ J per pixel.

Since the single photon energy of 655nm light is $3.033 \cdot 10^{-19}$ J, this means the noise corresponds to 32 photons.

The optical power of the measured noise was calculated from the number of camera counts, which was 36. This means that one quantization level (containing 64 counts) corresponds to $64/36 * 32 = 57$ photons.

Flip-flop

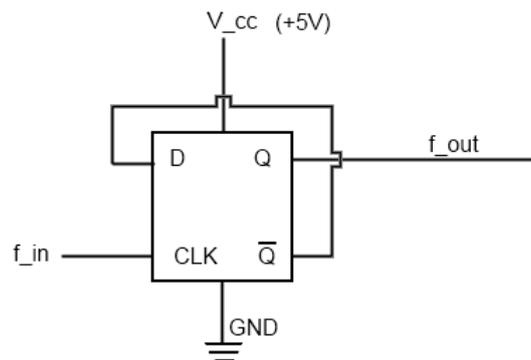


Figure 7.2: Schematic of the used frequency divider.

Modulator

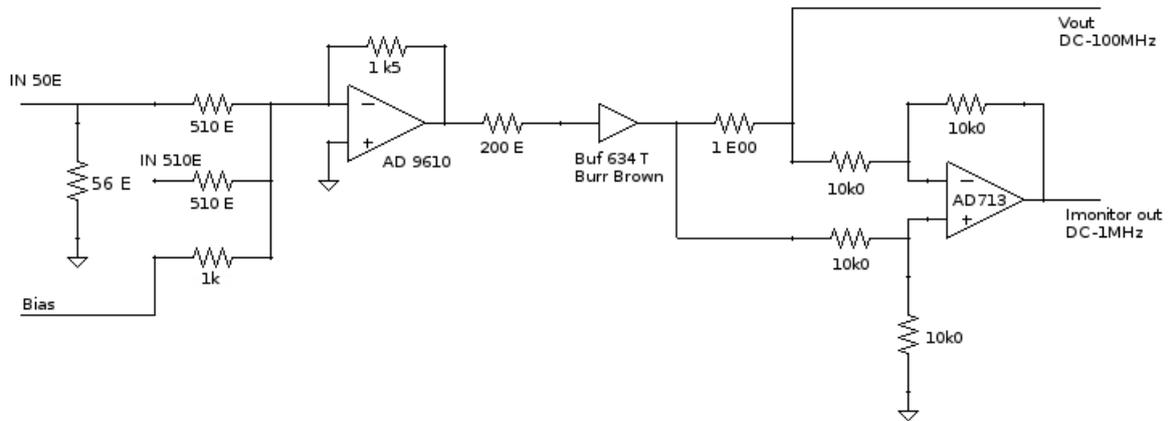


Figure 7.3: Schematic of the modulator device. The basic function is to add an offset to the signal so the high peak of a square wave is high enough for a LED to emit light, and the low peak is below this minimum voltage.

