Bachelor Thesis

Multi-spectral Mammography using a Medipix3RX Based Detector

Roel G. A. Deckers

Utrecht University

Institute of Subatomic Physics

Faculty of Science

June 17, 2015

Supervisors: Andre Mischke (UU), John Idarraga (UU)

Abstract

In this report we present the performance of a prototype based on Medipix3RX technology for low-dose energy-resolved mammography. We show that current detection methods are inadequate and that there are significant gains to be made by moving to multispectral imaging. We subsequently discuss and demonstrate the capabilities of the Medipix3RX and show how we can use this multispectral information detect cancerous growths in the field of mammography.

The focus of the work is on the computer interface that has been developed and the first images taken with it.

Contents

1 Introduction

At present time, the leading cause of dead from cancer in women worldwide is breast cancer. Breast cancer is the second leading cause of dead from cancer for women in developed countries and the leading cause of dead from cancer in low- and middle-income countries [1]. Early detection is crucial for successful treatment, this is why preventive screening by means of a mammogram is recommended in many places around the world. In an extensive study by Nyström et al. [2] it was shown that preemptive screening can reduce mortality by 13% in women aged 40-49 and by 29% among women ages 50 to 69. In 2010 the U.S. Preventive Services Task Force (USPSTF) recommended biennial screening for women aged 50-59 but not for women under 50 [3].

Unfortunately, at present time a significant amount of women get misdiagnosed. False negative rates are believed to be around 20% [4, 5]. and similar values have been found for false positives [6]. It goes without saying that false negatives can greatly reduce survivability due to late detection, but false positives can also cause a great deal of psychological harm to women and puts an unnecessary strain on medical facilities. Despite these issues, mammography is still the only screening method proven to reduce breast-cancer mortality [1].

Currently, mammograms are taken using either high-resolution full field digitial detectors or by means of screen-film imaging [7]. In an effort to bring down these statistics we have developed a computer interface to the new Medipix3RX chip [8] whose spectroscopic capabilities we hope to leverage to both enhance the visual quality of the mammograms and allow for better computer-aided detection at similar or lower doses as are currently employed.

One of the earliest indicators for breast cancer are what are known as microcalcifications [9], these small deposits of calcium are in some cases easy to miss on mammograms even by experienced medical personal due to the high density and variance of breast tissue. It is believed that by using the know spectral properties of calcium, we should be able to readily mark most if not all microcalcifications. It should furthermore be possible to distinguish different types of tissues based on their spectral characteristics.

In this report we will focus on the work done so far, namely: the development of a computer interface to the Medipix3RX and the first visualization of data obtained with this system.

Moving forward with the project we hope to perform computer aided reconstruction and image segmentation using the newly available spectral information.

2 The Medipix Chip

The Medipix chips are developed in international collaboration by the Medipix Collaboration. The Medipix Collaboration aims to disseminate detector technology developed at CERN for the needs of the LHC to the general public.

The Medipix3RX chip [8], which we leveraged, is a third generation Medipix 12 or 24-bit photon-counting detector capable of thresholding photon energies at up to eight levels simultaneously. Although the Medipix3RX covers a relatively small area (1.4×1.4 cm^2), recent advances in edgeless detectors and alternatives to wire-bonding employing Through-Silicon-via, suggest we can achieve a total area comparable to that of current systems. With a variable pixel pitch of 55μm to 110μm the Medipix3RX also implements high spatial resolution. Current systems have a pixel size in the region of the 80μm. In terms of dose we aim to work with the same type of field used in current screening but the extra spectroscopic information may allow significant reduction.

The Medipix3RX detector has two different spectroscopic modes: cluster mode and charge summing mode. In cluster mode, four pixels are grouped into a single cluster with a spatial resolution of 110μm and eight thresholds. In charge summing mode, four of these neighbouring clusters are connected together and any charge deposited in the 220μm $\times$ 220μm area spanned by these clusters will be assigned only to the cluster with the highest charge deposition, greatly diminishing any artifacts which would otherwise appear as a result of charge drift and diffusion. Using charge summing mode comes at the price of reducing the available thresholds from eight to four while maintaining the same 110μm spatial resolution.

Current work has been performed in charge summing mode using a set of four Medipix3RX chips edgelessly connected in a square on a single board using trough-silicon-via bonding. Unfortunately, due to technical-issues, only three chips were functional. We have good faith that we can leverage the eight different thresholds from cluster mode for mammographic purposes at a later date.

3 Software Interface Development

Our detector setup can be divided into three main parts. First off we have the Medipix3RX chips, which acquire the data. These are attached via a *High Pin Count* (HPC) FMC connector to a *Speedy Pixel Detector Readout module* (SPIDR) [10], which is an FPGA powered daugtherboard currently in development at the NIKHEF for fast and modular processing of data from Medipix-family chips. This in turn, after minimal data-formatting, transfers the data to a standard home-grade PC over ethernet using the *User Datagram Protocol* (UDP). The configuration of both the Medipix3RX chips as wells as the SPIDR board itself can be set using a *Transmission Control Protocol* (TCP) based interface.

The focus of this report is on the processing and visualization of the UDP data, as such we will not go into any more detail on the SPIDR board or its interface with the Medipix3RX chips.

3.1 Requirements

In order to allow for easy development and experimentation several features were required. During the design of the *Graphical User Interface* (GUI) much attention was paid to the usability in a clinical setting, using a modular design to clearly segment data acquisition and visualization from various configurations. At a later stage we plan to add easy to use toggles for switching between direct visualization of the SPIDR-data and reconstructed or otherwise modified data.

It should be noted however, that at the current stage of the project many design decisions have been made to meet our own needs while working on the development and calibration of the detector. The final product might therefore look quite different in order to better facilitate the work-flow of x-ray technicians. In the following, the requirements are discussed.

3.1.1 Easy Deployment

In order to facilitate adoption and limit costs, deployment of the new detector should be made as easy as possible. Therefore the software was made to run on any sufficiently modern hardware and Operating System. Currently, the only requirement for running the software as that the host machine has networking support, supports the Qt framework [11] (this includes all three major operating systems: Windows, Mac and Linux), and OpenGL version 3.1 (this requirement might be brought down at a later date).

3.1.2 Configuration of the detector

As no tools were available, we had to develop our own interface for configuring both the Medipix and the SPIDR Board. This has been incorporated into the GUI, allowing individual configuration of each Medipix chip as well as various SPIDR features. Configurations can be stored to and restored from disk.

3.1.3 Data handling

Facilities needed to be developed to handle the way in which the data arrived from the SPIDR board. Due to the support for multiple Medipix chips on the same detector some peculiarities need to be sorted out with regard to chip orientation and placement.

3.1.4   Failsafe handling of dead chips

Due to technical issues, only three out of four Medipix3RX chips performed as required. This meant the software had to be adapted any number of chips in an intuitive fashion.

3.1.5   Visualization

A visualization module had to be developed in order to present the data to the end-user and allow for intuitive manipulation. Two features were considered crucial: a move- and zoom-able heat map display for the raw data as well as a plot of the histogram of the data.

The heat map was at first this was implemented using premade modules but due to concerns over speed and to allow for fast and easy post-processing of the acquired data we opted to implement the heat map displays ourselves in OpenGL.

3.1.6   Equalization

In order to obtain correct data, an implementation of a proper equalization algorithm had to be written. The details of which fall outside the scope of this report.

3.2 Features implemented

While still in heavy development, we have already reached many major milestones in terms of the development of our GUI. What follows is an incomplete list of features currently supported.

3.2.1 Equalization

The software allows one to run an automated equalization of the detector, the results of which can be stored to disk to be reloaded at a later time.

3.2.2 Configuration

All runtime modifiable parts of the detector are exposed to the user, and can be stored on disk, allowing for easy and rapid reconfiguration. This includes toggling on and off spectroscopic mode, setting the photon-energy thresholds for each individual chip, number and length of the automated triggers and more.

3.2.3 Pixel Masking

In order to perform equalization as well as to get rid of noisy and dead (unresponsive) pixels one has to be able to mask turn-off specific pixels. The Medipix hardware supports the masking of any pixel and this has been implemented in software. Simply right-click on any displayed pixel and select (un)mask pixel from the drop-down menu. It is also possible to save the pixel-mask to file and reuse it at a later date.

3.2.4 Visualization

Acquired data can be displayed to the end-user in an intuitive fashion. A screenshot of the visualization with computer generated data is shown in figure 3. A heat map of the acquired data gets shown to the user in a main window, with a smaller histogram display of to the right. The user can drag the space between the histogram and the main display to resize their relative horizontal sizes (the program defaults to the golden-ratio, with the larger fraction assigned to the heat map). Heat map and histogram are linked in such a fashion that the data-range of the heat map will always correspond to the selected portion of the histogram (indicated by two vertical blue lines).

The heat map display was implemented using OpenGL, data gets transferred as 32-bits integers and stored as-is on the *Graphics Processing Unit* (GPU). Subsequent pseudo-color representation are calculated dynamically on the GPU. This allows us to quickly apply filters and other post-processing methods to the data on the GPU. Work has been started on applying post-processing but, although the initial results have been good, it is not yet presentable.

Several gradients come presupplied with the program (including a colour-blind safe representation) and can be selected in the bottom left corner of the program. Extra care was taken to include a wide variety of gradients but exclude any which might give a false impression of the data, such as the often seen rainbow gradient [12]. Gradients are defined in an easily modifiable human-readable file-format, allowing the end-user to add his own preferred gradients.

The display automatically scales to show all connected chips, and markers along the top and left side of the screen show the coordinates in pixels. In addition, the current position of the mouse-cursor as well as the value of the pixel it is over are displayed in the bottom right corner. Values which fall outside of the currently selected data-range will be clamped to the colours at the appropriate gradient edges, but can also be made transparent by the checkbox "drop out-of-bounds".

To the right of the main display we show a histogram for the current selected threshold (in red, the gray histograms are the different thresholds) as well as various controls. At the top one can select which threshold energy to display. Below the histogram there are buttons to start collecting data, configuration of which is done under the "Config and Monitoring" tab shown at the top, as well as various checkboxs to determine what corrections to apply. Below that we currently have a button for generating fake data (used in debugging the system when no detector is available), and below that are the controls for the histograms.

One can select how many bins the histograms should contain (measured over the whole data-range) and whether to display the full-range, a percentile selection (in terms of area under the histogram curve), or manually select a range. In addition to the controls below the plot, it is also possible to select a data range by clicking the right mouse button and dragging a selection in the plot itself. This allows the user to easily select any data-range of interest.

3.2.5 Flat field correction

The *Flat Field Correction* (FFC) is a technique to increase the image quality, it corrects for artifacts due to pixel-to-pixel sensitivity mismatch as well as different gain factors and any spatial distortions (e.g. surface inconsistencies in the detector protection layer). The FFC requires one to first take an image without a sample present between the detector and the source, then the FFC is given by the formula:

$$p_{\text{corrected}} = p_{\text{measured}}/p_{\text{flatfield}} \times \langle p_{\text{flatfield}} \rangle / \langle p_{\text{measured}} \rangle . \tag{1}$$

Here, $p_{\text{corrected}}$ refers to the corrected pixel, p_{measured} is the uncorrected pixel value, and $p_{\text{flatfield}}$ is the value measured without any sample in front of the detector, and $\langle p \rangle$ means to take the mean over all pixels. This method preserves the mean pixel value. The results of our flat-field measurements at 70 kVp are shown in figure 1 and figure 2. The user can load any flat-field profile and the program will subsequently apply the the correction to any subsequent data.

3.2.6 Summing mode

The program allows the user to integrate over time by means of a simple checkbox. Once turned on, all subsequent data will be summed to the current data instead of overwriting it. Allowing the user to refine his images as much as necessary. This is a feature used for research purposes, and will probably be removed once we move to a clinical setting

3.2.7 Saving and loading of data

In addition to the previously mentioned configurations, it is also possible to save captured data along with the relevant metadata (such as threshold energies) to disk, an reload it at a later data. The data are stored in a custom but straightforward binary format.
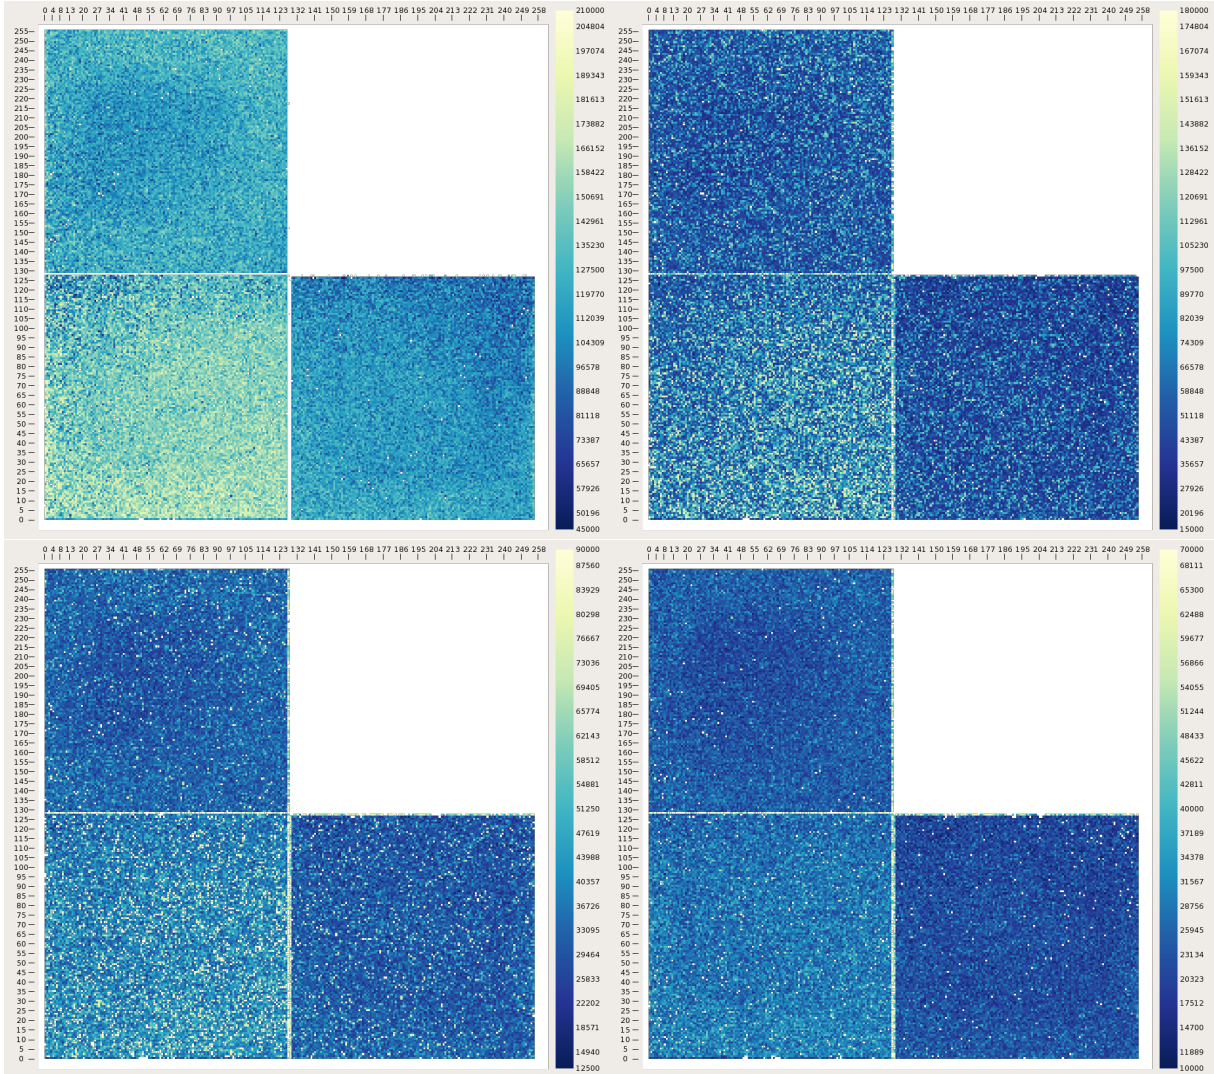
Figure 1: The flat-field results at a 80 kVp at threshold of 50 ADCL (top-left), 60 ADCL (top-right), 70 ADCL (bottom-left), and 80 ADCL (bottom-right).

3.2.8 Scaling

The program scales to multiple or larger detectors of any 2D geometry or even more thresholds. This allows the system to easily be adopted for detectors with a large active-area or new technologies should they become available, as well as down-scale to smaller detectors should they be preferred for other applications due to technical or financial limitations.

3.2.9 Adaptable

Being build using the Qt framework, the program can run natively on any major operating system allowing for easy integration in any IT infrastructure or stand-alone devices. Any major OS (Mac, Windows or Linux) is supported and the only requirements to run the software is OpenGL 3.1+ support, and this might be brought down further with additional research. This significantly lowers the barrier to adoption.
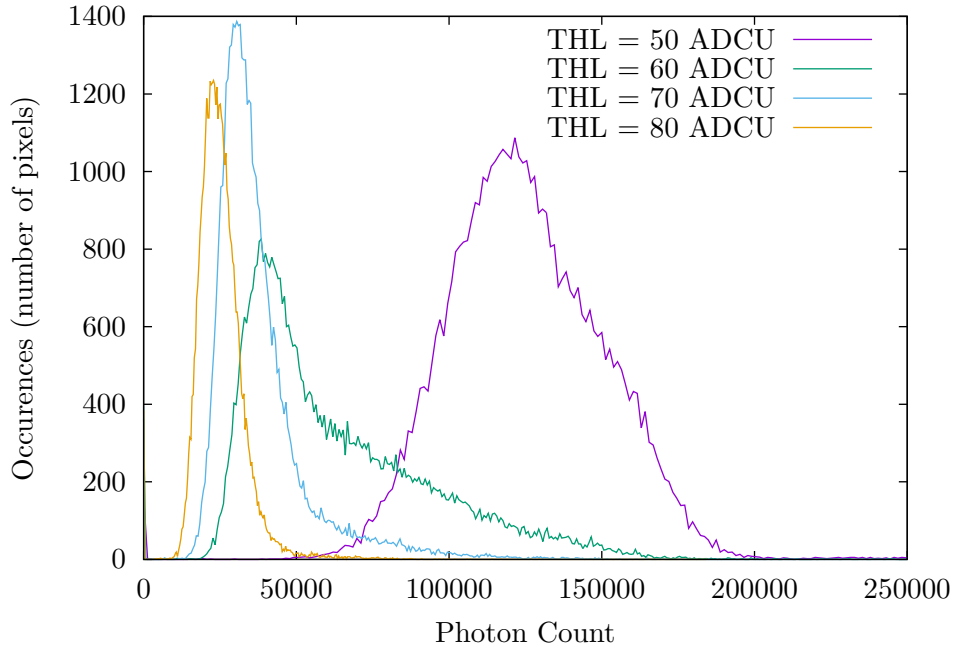
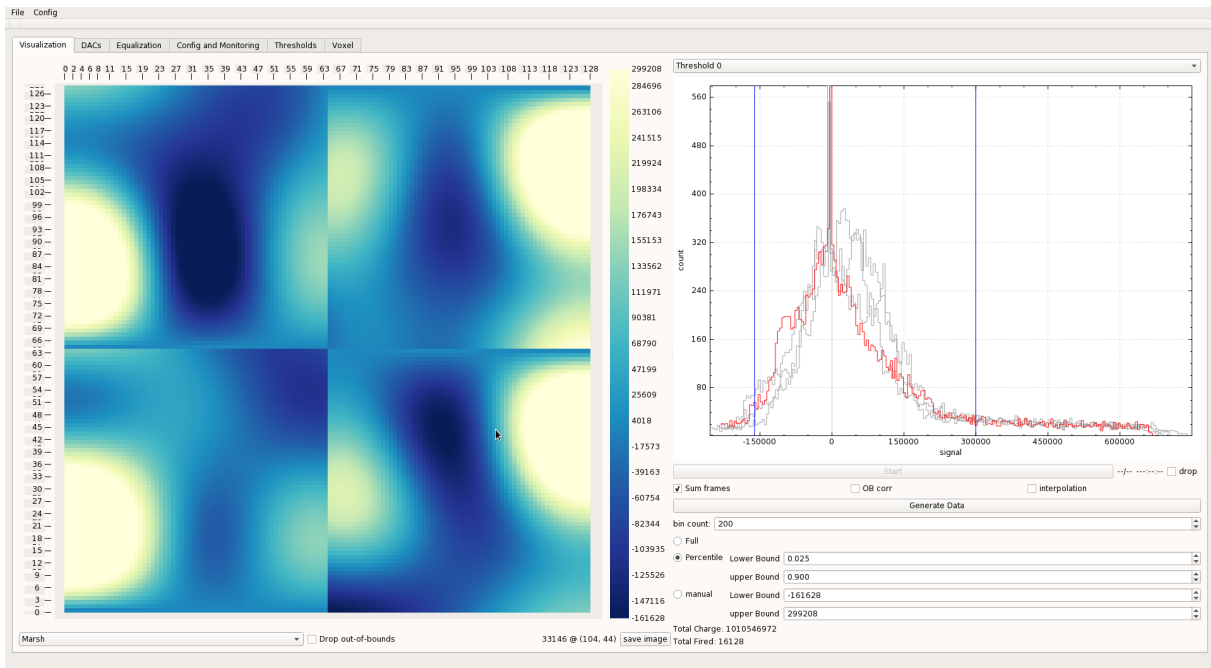Figure 2: The histograms of flatfield at various thresholds

Figure 3: A screenshot of the visualization module we developed, showing computer generated data.

3.3 Future work

In order to move onto a medical setting work still remains to be done on several fronts, including but not limited to:

1. Workflow optimization: Currently the user-interface was designed in such a way as to facilitate rapid development and research. Once we move to a medical setting it is

likely that this interface will not be ideal for radiologists. Due to a modular design, we believe that the software will be easily adoptable to any needs that might presents itself.

2. Compliance with medical standards: The storage and processing of medical data is a heavily regulated field (see for example the DICOM standard). Research remains to be done to make our software compliant.

3. Analysis and Reconstruction: We aim to leverage the spectral information provided by the Medipix chips to improve current state-of-the-art computer-aided detection. Due to the known spectral signatures of cancerous growths and their indicators we believe it should be possible to clearly mark them even in cases where a visual inspection might not have a high rate of success. We will discuss the basic theory behind x-ray analysis in the following subsection.

4. Increased threshold Count: Currently the Medipix chips are run in charge summing mode, which gives us four thresholds to configure. If operated in cluster mode, eight thresholds become available to us, allowing for much greater spectral accuracy at the cost of increased charge diffusion. Work still remains to be done both in the implementation of cluster mode as well as research into the performance w.r.t. charge summing mode.

5. Calibration: The detector is not yet calibrated. We will have to calibrate both the various thresholds to keV as well as determine the best thresholds to use for mammography.

6. Multi-angled imaging: it is believed that the reconstruction and computer aided analysis will greatly benefit from taking multiple images at different angles, similar to CT. This of course comes with a large trade-off in terms on dosage. Research into the optimal balance between dosage and number of images needs to be performed as well as adopting the software to handle multiple angles.

3.4   Reconstruction

The goal is to study (semi-)automated feature recognition. In order to do so we model the interaction between our samples and x-rays by an attenuation based model. X-ray attenuation can be described by

$$I/I_0 = \exp\left(-\mu t\right), \tag{2}$$

where I is the measured beam intensity, I_0 is the incoming beam intensity, t the thickness of our sample, and μ the attenuation coefficient [13]. Both the coefficient and the intensities are functions of the energy E. We can rewrite this equation into a linear equation by taking the log of both sides. Furthermore, if we define $-\log\left(I/I_0\right) = \lambda$ we find that

$$\lambda = \mu t. \tag{3}$$

now t is a function of the type of material and the position (i.e. how much of material M is at location $\vec{x}$) and μ is a function of the type of material and the energy (i.e. what is the spectrum of M at energy E), and finally λ is a function of position and energy. If we restrict ourselves to N_E energy levels and N_M materials we can write the following matrix

expression for each cluster[1]:

$$
\begin{pmatrix} \lambda_0 \\ \vdots \\ \lambda_{N_E} \end{pmatrix} = \begin{pmatrix} \mu_0^0 & \cdots & \mu_0^{N_M} \\ \vdots & \ddots & \vdots \\ \mu_{N_E}^0 & \cdots & \mu_{N_E}^{N_M} \end{pmatrix} \begin{pmatrix} t^0 \\ \vdots \\ t^{N_M} \end{pmatrix},
\tag{4}
$$

or more comprehensively,

$$
\vec{\lambda} = \boldsymbol{\mu} \vec{t}.
\tag{5}
$$

Here, the i'th element of $\vec{\lambda}$ is the λ measured at energy level i. The (i,j)'th element (row major notation) of $\boldsymbol{\mu}$ correspond to the attenuation coefficient of the i'th material at the j'th energy level. The i'th element of $\vec{t}$ rereferso the thickness of the i'th material.

For every pixel one gets such an expression. The matrix is the same for each pixel (assuming we consider the same types of materials and energy levels for each pixel), but the two vectors differ. There are two fundamentally different ways to solve these equations, guided or blind reconstruction (or a hybrid of the two).

3.4.1 Guided reconstruction

When using guided reconstruction one has to know which materials can be present, and what their spectrum looks like. If one knows such a thing, it is possible to prepopulate the matrix $\boldsymbol{\mu}$, leaving one to solve an inverse matrix equation for each pixel.

This method is conceptually and computationally simple, but requires accurate knowledge of the sample. Problems might also arise when spectral profiles become very similar, due to unavoidable noise, the best solution to the matrix equation might be one which does not correspond to reality. We speculate that it might be possible to counteract this by applying a heuristic measure which prefers solutions which are similar to their neighbours'. Unfortunately, this is only applicable when the structures of interest are sufficiently large in comparison to the pixel size.

Another method of guided reconstruction would be to apply a fitting algorithm over the data using the known spectrums. One such method, for which we have high expectations, is the popular Root package RooFit [14] package.

3.4.2 Blind reconstruction

In contrast to guided reconstruction, blind reconstruction does not rely on any a-priori knowledge of the sample. Instead, one dictates a heuristic to quantify the quality of each combination of $\vec{\lambda}, \boldsymbol{\mu}, \vec{t}$ and then applies an optimization algorithm to all pixels' $\vec{\lambda}, \vec{t}$ and $\boldsymbol{\mu}$ simultaneously. There are, of course, infinitely many heuristics possible as well as a multitude of optimization algorithms and it is neigh impossible to say which will be the best performer without running some tests.

Instead of treating this as (solely) an optimization problem one can also apply principal component analysis, k-means clustering, or a adapt a variety of methods from image processing. Unfortunately, a full comparative analysis of these methods goes beyond the scope of the project at present time.

One expected disadvantage of blind reconstruction is that precisely because it has no knowledge of physical properties it will not necessarily converge to a solution with a realistic

---

[1]We remind the reader that in spectroscopic mode, four pixels bond together to for a single cluster.

$\boldsymbol{\mu}$, however we expect the contrast between materials to be preserved in relation to the real distribution and it is applicable to our data.

4   Experimental Setup

Our images were obtained with the help of an X-Ray tube. A variable *Peak kilovoltage* (kVp) was applied across, and 20 mA was applied through, the tube. These properties determine the kinetic energy imparted to electrons in the tube, and subsequently the spectrum of the generated x-rays. The spectrum of the X-Ray tube at 50 kVp is shown in figure 4. Apart from the peaks, the distribution moves as whole along with the value of kVp. The detector
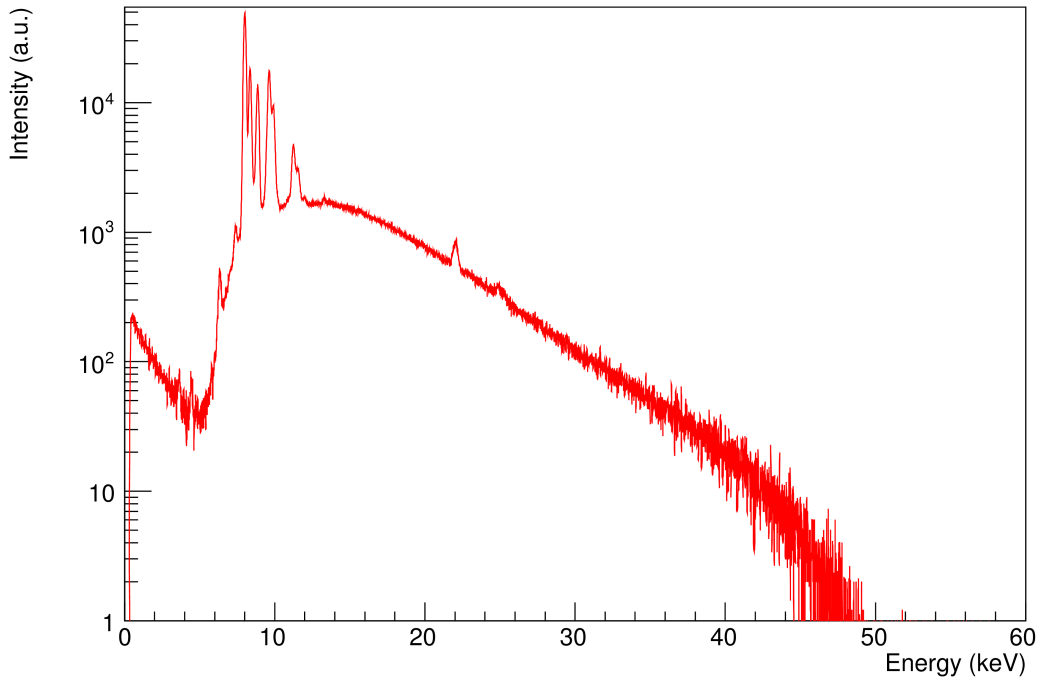
Figure 4: The spectrum of the x-ray tube employed in our experiments. The x-axis indicates the energy in keV of the generated x-rays, and the y-axis the intensity in a.u. at a specific energy.

was placed below the tube, with a sample of interest in between, within 1 cm of the detector. The detector was connected to an FPGA powered board developed at the Nikhef called SPIDR[10], which in turn transmits any data to our software using the *User Datagram Protocol* (UDP). This modular setup allows for rapid prototyping and inexpensive deployment. The detector has not yet been calibrated from internal thresholds to actual units of energy, however a linear relation exists between the internal threshold value and the threshold energy. We will therefore refer to our thresholds in values of *Analog-Digital Converter Units* (ADCU). The thresholds are a lower bound on the energy a photon may have, any photon with an energy lower than the specified threshold will be ignored, a lower threshold therefore means that a larger part of the spectrum will be counted.

In figure 5 a photographic representation is of the whole setup, whereas in figure 6 a detailed view is given of the SPIDR board.

Looking at figure 2 we see that we get a reasonably Gaussian response for a threshold of 50, 70, and 80 and for a threshold of 60 we get what looks like a Landau distribution. Of course, the other distributions are not perfect Gaussians either, but they have a relatively small tail compared to threshold 60. The cause of this discrepancy is currently being investigated.
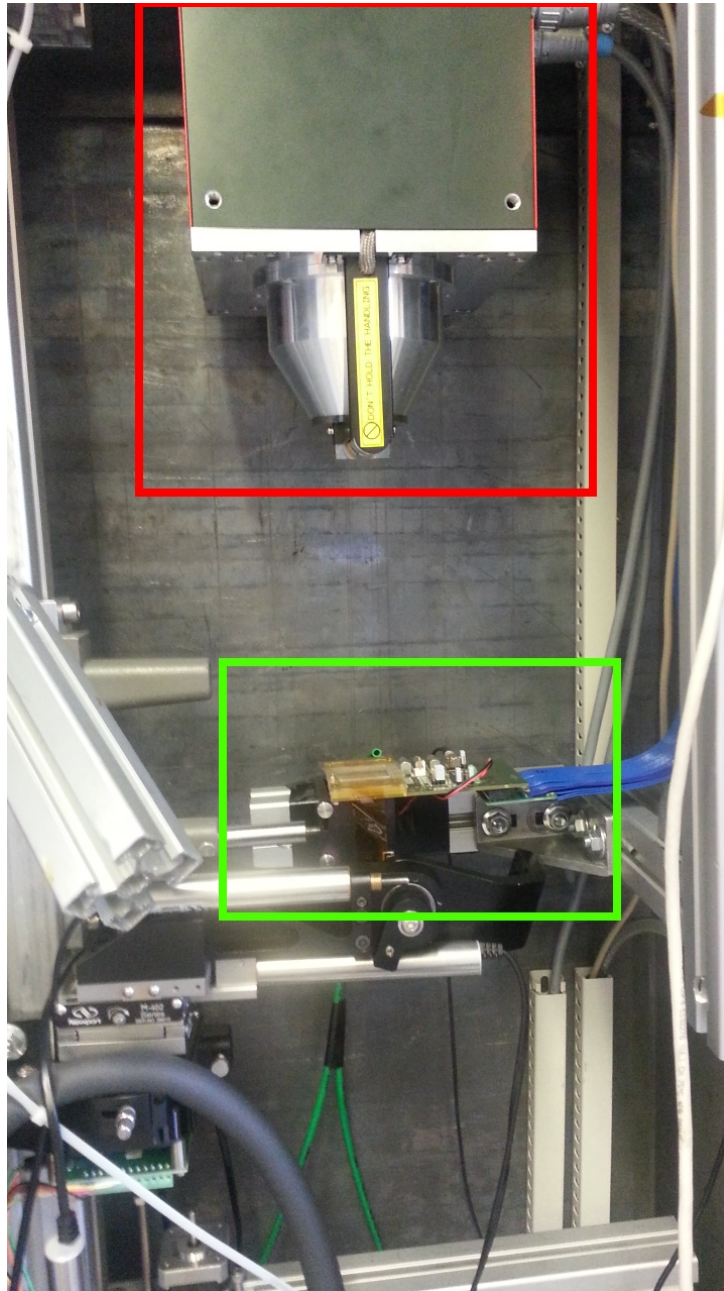

Figure 5: An overview of the experimental setup. We have marked the x-ray tube with a red border, the Medipix3RX is marked with green. The blue cable coming off of the Medipix3RX leads to the SPIDR board which was attached to the wall of the containing x-ray cabinet, next to the x-ray tube.
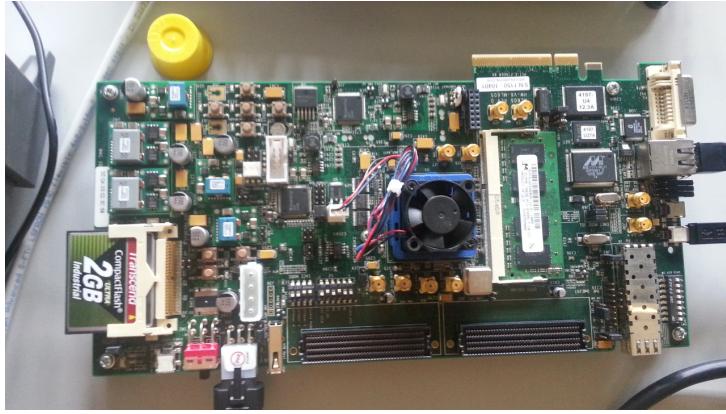
Figure 6: The SPIDR read-out system.
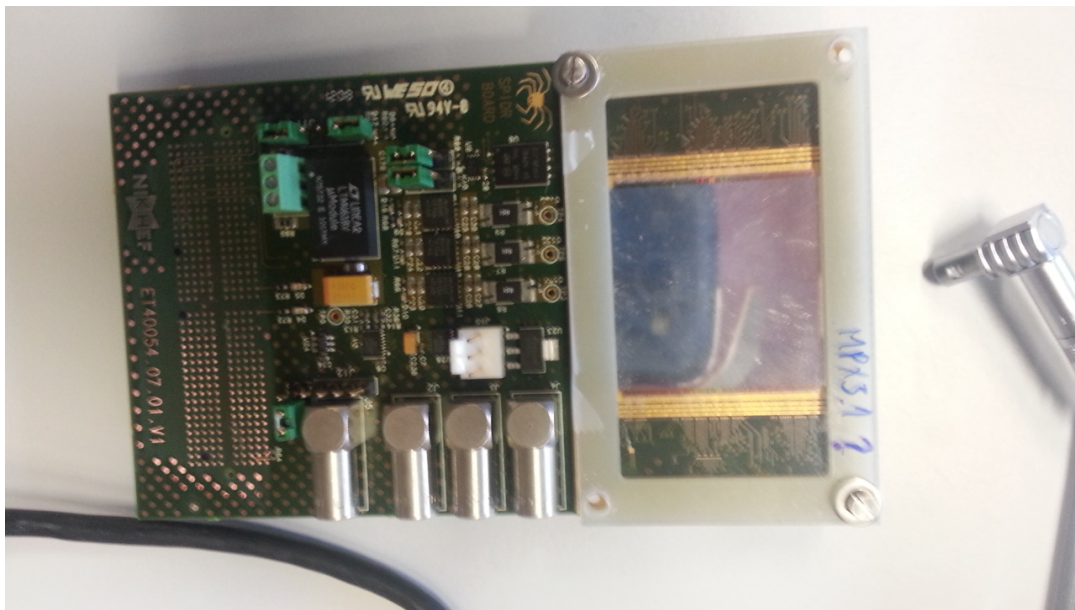
Figure 7: The Medipix detector.

5 Multi-threshold images

5.1 USB Memory Stick

Our first sample was a standard USB memory stick as shown in figure 8. The resulting images are shown in figure 9, the histograms of the resulting images are displayed in 10. The images were taken at a 80 kVp.

Figure 8: Our USB memory stick.

results

Figure 9: A USB memory stick at a threshold of 50 ADCU (top-left), 60 ADCU (top-right), 70 ADCU (bottom-left), and 80 ADCU (bottom-right).
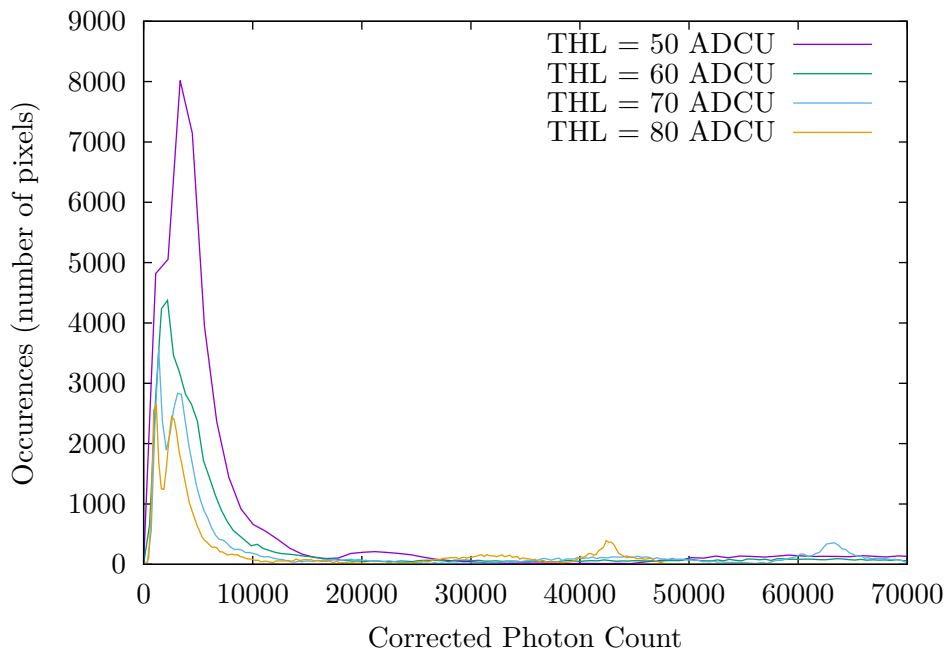
Figure 10: The histograms of the USB memory stick at various thresholds.

Discussion

In these images, one can clearly distinguish the individual copper traces on the circuit board. Showing the spatial capabilities of the detector.

Turning our attention to the spectral information contained within the histograms, we note that depending on the threshold, one can make out two or three peaks. One or two at low photon counts, these correspond to circuit board and its tracing, and a much smaller peak later on. This one is not visualized in the heat map plots but corresponds to the plastic casing, which could be made out along the sides of the circuit board. Should the histogram extend even further, more peaks would become visible, but these correspond not to the USB memory stick but are simply a result of the detector not being fully covered by the stick.

The splitting of the first peak as one increases the threshold could be due to different spectral properties between the board and the copper tracings. However, they could be attributed to or share a common cause with the spread-out spectrum of the flat-field at 60 ADCU. Any relation between these results is currently being investigated.

The bars in the bottom left and right corners of the heat maps are the mounting apparatus.

5.2   Paint Splotches

Our second sample was a collection of undiluted oil-paint swatches in between two pieces of paper, as shown in figure 11. The results can be seen in figure 12 and 13. The images were taken at 80 kVp.
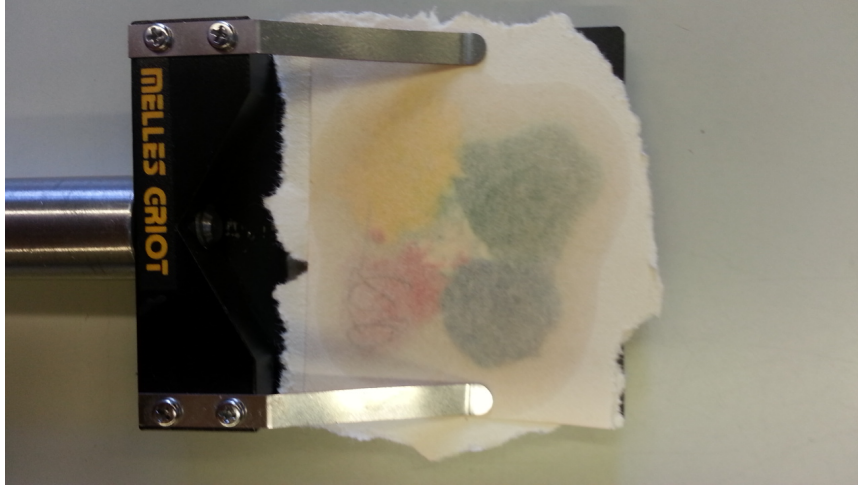

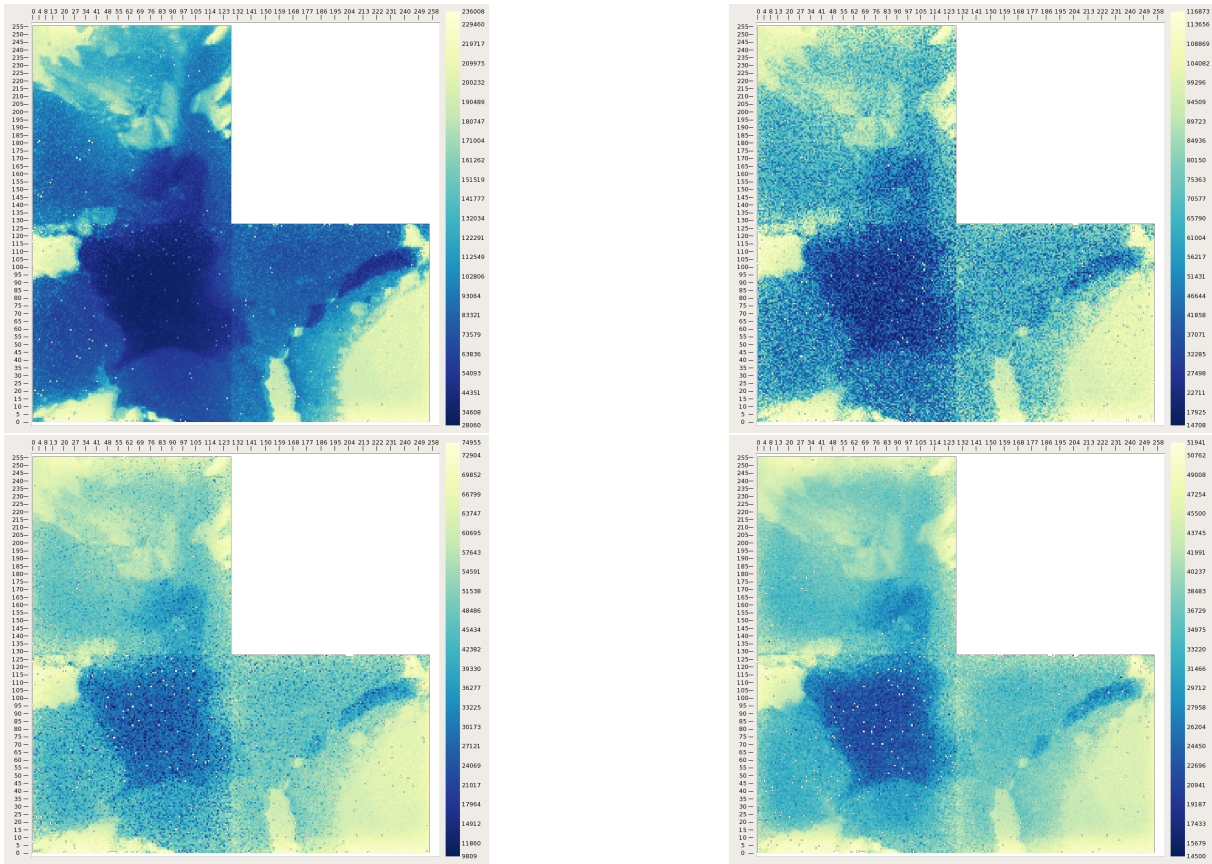Figure 11: Sample of 4 different oil paints in between two pieces of paper.

results

Figure 12: Paint Splotches at a threshold of 50 ADCU (top-left), 60 ADCU (top-right), 70 ADCU (bottom-left), and 80 ADCU (bottom-right).
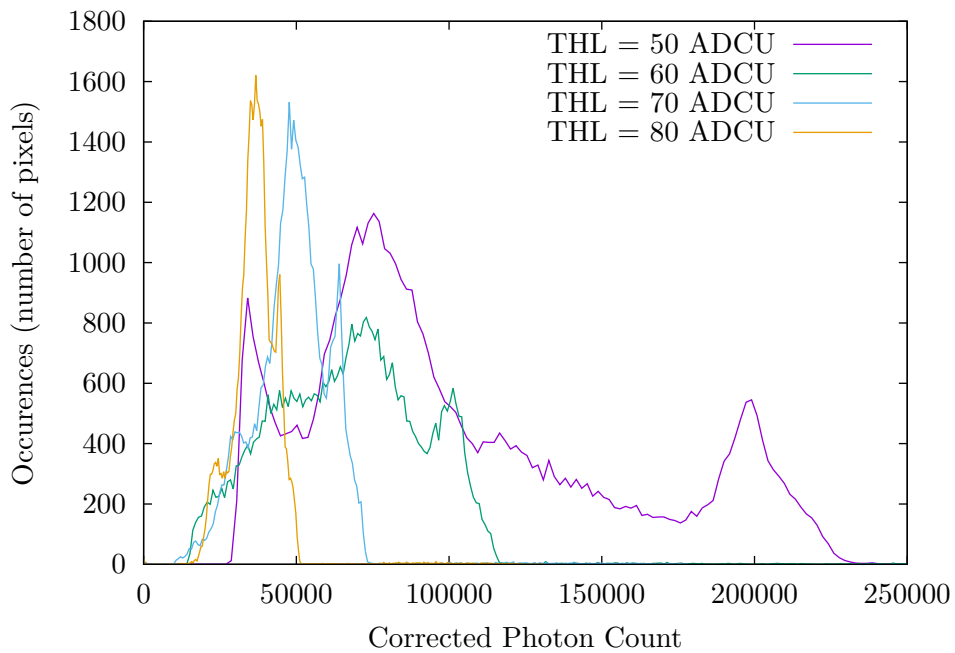

Figure 13: The histograms of the paint blotches at various thresholds

Discussion

In contrast to the USB memory stick, this sample shows little interesting spatial features but has a very peculiar spectral profile. At each threshold one can make out three distinct peaks (at a threshold of 50 ADCU they are at $\sim 200\,000$, $\sim 75\,000$ and $\sim 40\,000$) but the overoverallpe changes dramatically between a threshold of 50 ADCU and 70 ADCU. The leftmost peak can be attributed to the dark spot in the middle of the image (corresponding to a a bit of white paint). The two different peaks seem to be caused by the background (paper and oil), and the different paints. seeing as how it is impossible to make out different colours of paint in the x-ray spectrum by eye, most of this spectrum can probably be attributed to either a binding agent found in all paints or simply the increased thickness w.r.t. the paper and oil.

5.3 Bone

Our third sample was a piece of a raw chicken-leg as shown in figure 14. The results can be seen in figure 15 and 16. This time, there was no discernible difference between the different thresholds (different types of tissue often times have similar spectrums, requiring a great deal of work to find the ideal thresholds). we have opted to show only the best result in order to highlight the image quality. The images were taken at 50 kVp.
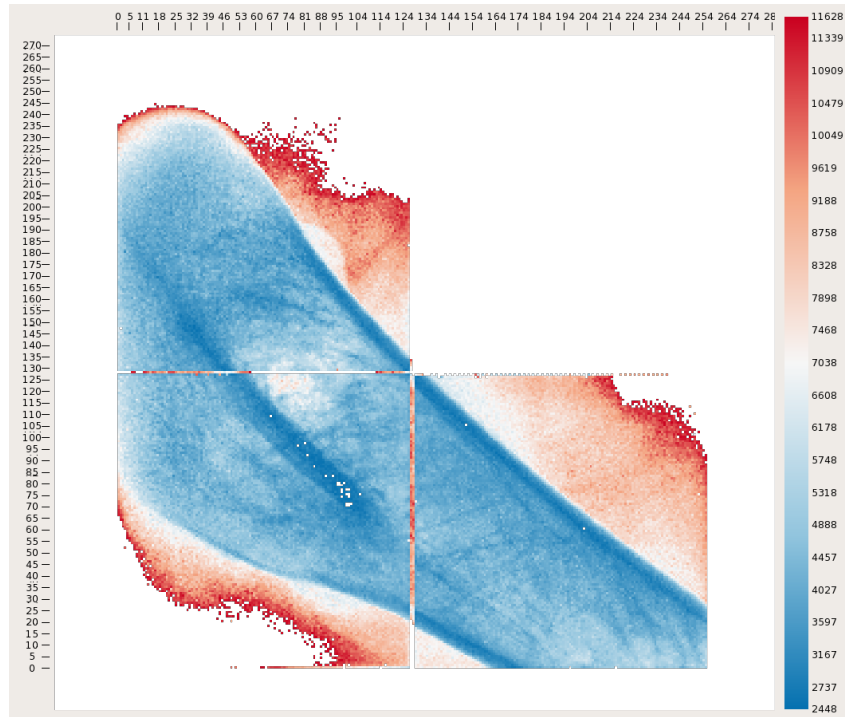

Figure 14: The Chicken leg that was x-rayed.

Results

Figure 15: A chicken leg at a threshold of 50 ADCU. Data range has been clamped to show a clear division between meat and bone.
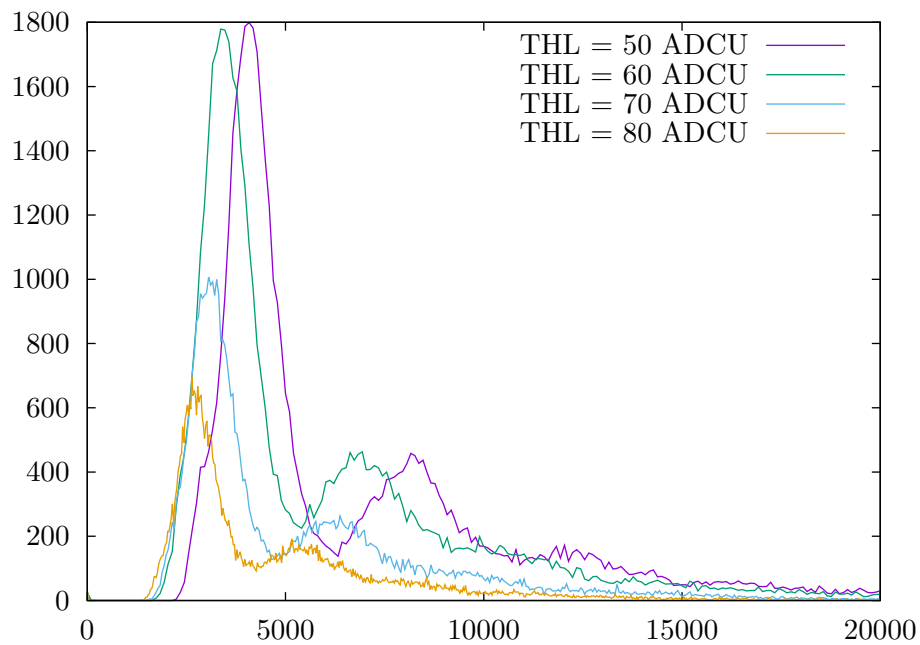

Figure 16: The histograms corresponding to the chicken-leg's x-ray.

Discussion

We can clearly observe two wide peaks in the spectrum. One for the bone and one for the meat. We also see a smaller peak at a corrected photon count of around 3000. This corresponds to the edges of the bone, which are highlighted because the outside of the bone is much more dense than the marrow on the inside. This image clearly shows both the structure of the bone as that of the meat (to a lesser degree).

5.4   Banana peel

Our final sample was a piece of banana peel, as shown in figure 17. The results can be seen in figure 18 and 19. The images were taken at 50 kVp.
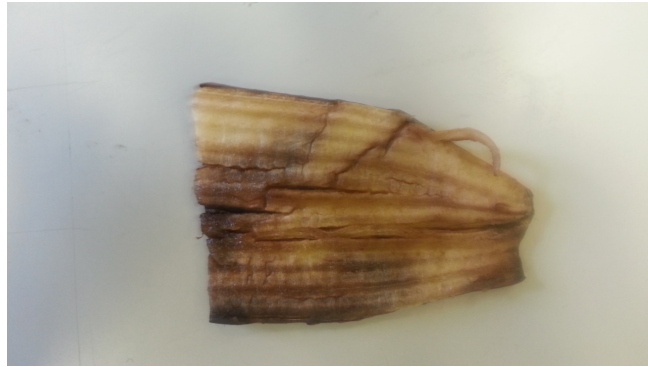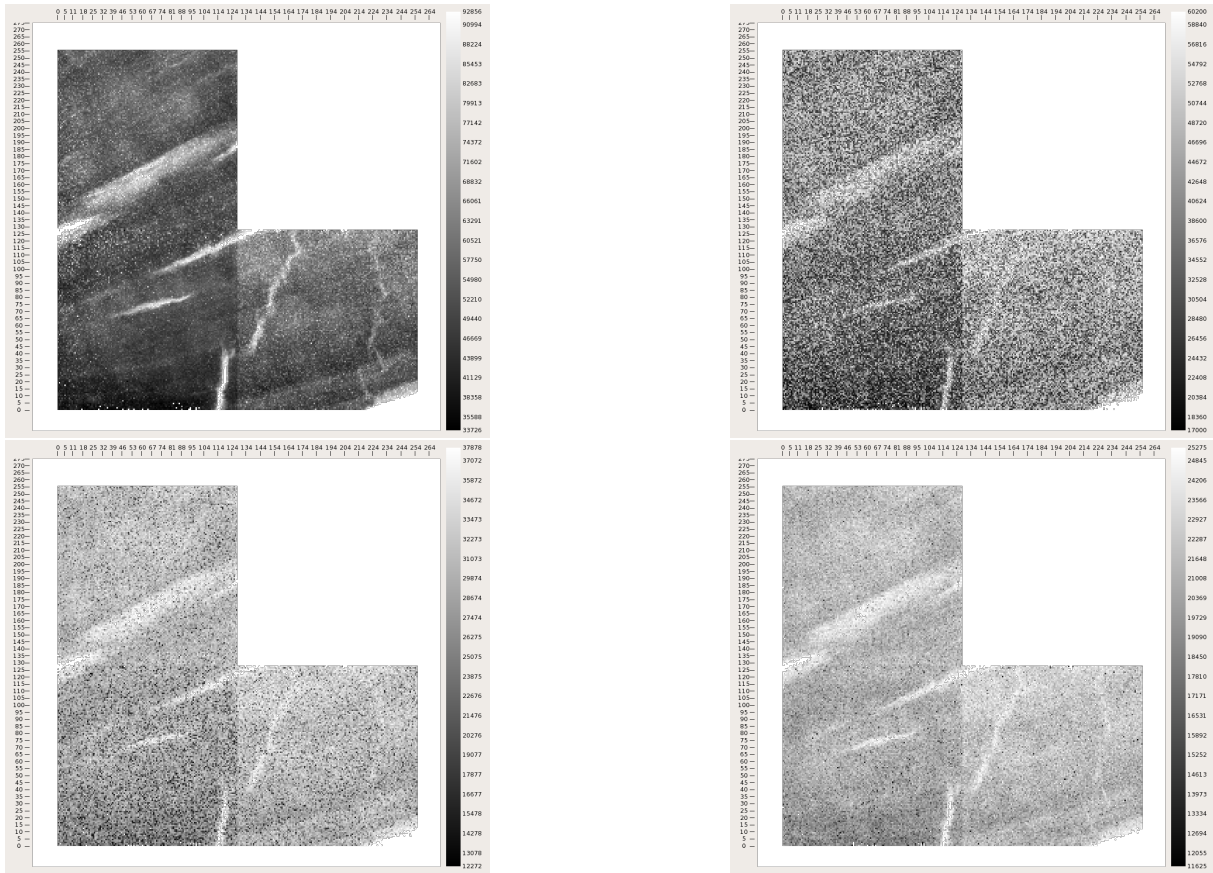

Figure 17: The banana peel sample.

Results

Figure 18: Banana peel at a threshold of 50 ADCU (top-left), 60 ADCU (top-right), 70 ADCU (bottom-left), and 80 ADCU (bottom-right).
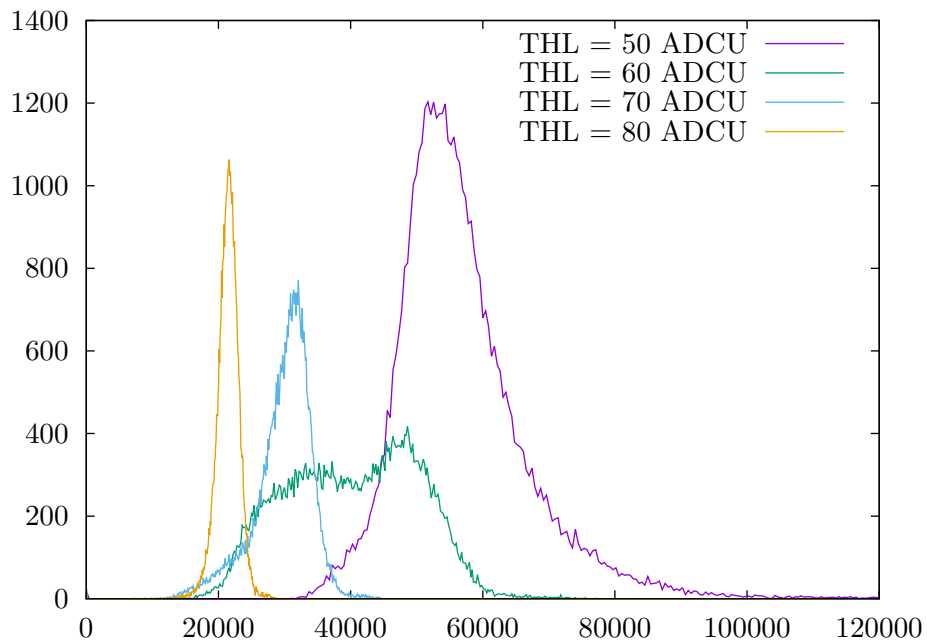

Figure 19: The histograms of the banana peel at various thresholds.

Discussion

As with the paint, we once again see a very characteristic spectrum. As one increases the threshold starting at 50 ADCU and ending at 70 ADCU one starts with a left leaning distribution (that is, it has a long tail to the right of its peak and a sharp drop to the left). after which it will quickly lose its form and change into a much less well-defined distribution. At an ADCU of 60 it is seemingly made up off two very wide peaks at a corrected photon count of roughly 50 000 and 30 000. When one increases the threshold even further, the distribution becomes more well defined again, only this time leaning to the right instead.

While we can not fully exclude the possibility this is in part caused by the distinct spectrum of the x-ray tube at a THL of 60 ADCU, this behaviour is very interesting and might be indicative of a distinct spectral profile between these thresholds.

6 Conclusions

We conclude that studies has shown that there is a need for better breast-cancer detection, and that we have demonstrated that our Medipix3RX based detector performs as promised and can help improve mammographic procedures by means of low-dose spectroscopic analysis.

While, as we discussed, much future work remains to be done before we can move to clinical trials, we have created and present a solid foundation both in theory and software to build upon and are confident that we will be able to successfuly leverage to capabilities of the Medipix3RX to improve upon the current state-of-the-art in mammography technology.

7   Acknowledgements

References

[1] Béatrice Lauby-Secretan, Chiara Scoccianti, Dana Loomis, Lamia Benbrahim-Tallaa, Véronique Bouvard, Franca Bianchini, and Kurt Straif. Breast-cancer screening—viewpoint of the iarc working group. *New England Journal of Medicine*, 2015.

[2] L Nyström, S Wall, LE Rutqvist, A Lindgren, M Lindqvist, S Ryden, J Andersson, N Bjurstam, G Fagerberg, J Frisell, et al. Breast cancer screening with mammography: overview of swedish randomised trials. *The Lancet*, 341(8851):973–978, 1993.

[3] US Preventive Services Task Force et al. Screening for breast cancer: Us preventive services task force recommendation statement. *Annals of internal medicine*, 151(10): 716, 2009.

[4] Richard E Bird, Terry W Wallace, and Bonnie C Yankaskas. Analysis of cancers missed at screening mammography. *Radiology*, 184(3):613–617, 1992.

[5] MG Wallis, MT Walsh, and JR Lee. A review of false negative mammography in a symptomatic population. *Clinical radiology*, 44(1):13–15, 1991.

[6] Joann G Elmore, Mary B Barton, Victoria M Moceri, Sarah Polk, Philip J Arena, and Suzanne W Fletcher. Ten-year risk of false positive screening mammograms and clinical breast examinations. *New England Journal of Medicine*, 338(16):1089–1096, 1998.

[7] John M Lewin, Carl J D'Orsi, R Edward Hendrick, Lawrence J Moss, Pamela K Isaacs, Andrew Karellas, and Gary R Cutter. Clinical comparison of full-field digital mammography and screen-film mammography for detection of breast cancer. *American Journal of Roentgenology*, 179(3):671–677, 2002.

[8] R Ballabriga, W Wong, and X Llopart. Medipix3 manual, 2009.

[9] Heng-Da Cheng, Yui Man Lui, Rita Freimanis, et al. A novel approach to microcalcification detection using fuzzy logic technique. *Medical Imaging, IEEE Transactions on*, 17(3):442–450, 1998.

[10] Martin van Beuzekom. Spidr wiki. URL `https://wiki.nikhef.nl/detector/Main/SpiDr`.

[11] The Qt Foundation. Qt framework. URL `http://www.qt.io/`.

[12] David Borland and Russell M Taylor II. Rainbow color map (still) considered harmful. *IEEE computer graphics and applications*, (2):14–17, 2007.

[13] John H Hubbell and Stephen M Seltzer. Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients 1 kev to 20 mev for elements z= 1 to 92 and 48 additional substances of dosimetric interest. Technical report, National Inst. of Standards and Technology-PL, Gaithersburg, MD (United States). Ionizing Radiation Div., 1995.

[14] The ROOT Team. Roofit. URL `https://root.cern.ch/drupal/content/roofit`.