SUB-PIXEL REGISTRATION FOR LOW COST, LOW DOSAGE, X-RAY PHASE CONTRAST IMAGING

A thesis presented in partial fulfilment of the requirements for the degree of

Master of Engineering

in

Electronic and Computer Engineering

at Massey University, Palmerston North,

New Zealand.

Hamish Bradley

2021

Contents

List of Tables

List of Figures

Abstract

X-ray phase contrast imaging is an imaging modality that measures the phase shift of the X-ray wavefronts as they travel through different materials. This gives a higher contrast between regions of similar X-ray attenuation, in a medical sense this corresponds to a higher contrast of soft tissues. A new area of research for X-ray phase contrast imaging is to use shadow-based intensity modulation to generate these images. This thesis explores a range of different registration techniques, and their suitability for phase contrast imaging using shadow based intensity modulation.

Image registration is a key step in generating the phase contrast images as it is related to the x and y differential phase. These are then integrated to generate the phase contrast image. Therefor a high accuracy sub-pixel registration technique will provide high quality phase contrast images. The registration techniques explored are 1D curve 2-D surface fitting to a correlation map, phase registration, Newton-Raphson method, and optimal interpolation filtering. These registration techniques were tested with images that are noise free, as well as images corrupted by Poisson noise. The Newton-Raphson, and the optimal interpolation filters show the most promise due to low errors in the noise free environment. In the presence of noise, the Newton-Raphson method performs poorly, and hence requires a good denoising method, while the optimal interpolation filters do not get any improvement from any denoising techniques.

Currently the Newton-Raphson based method are used widely in digital image correlation, however the optimal interpolation filtering has the benefit of not being limited by the choice of interpolation technique, and it removes the iterative process, and depending on the size of the optimal interpolation filter it performs better than, or only marginally worse than the Newton-Raphson method.

Acknowledgements

Firstly, I would like to thank my supervisors Professor Donald Bailey and Dr Steven Le Moan, who were always there to keep me on track and support me whenever I needed it. Also the members of our research group, the Centre of Research in Image and Signal Processing, which during many meetings assisted in providing a fresh set of eyes.

I would also like to thank Peter Gänz and Professor Sven Simon, who were the working on the project alongside us in Germany. A huge thanks to VolkswagenStiftung whose financial support allowed us to partake in this research.

This work would not have been completed without the emotional support from my parents, who have always been there cheering me on all the way.

Chapter 1

Introduction

1.1 X-ray Phase Contrast Imaging

When X-ray radiation was first discovered in 1895[6], Wilhelm Röntgen dubbed them 'X' for undefined, and they have maintained that name even though much more is known about them now. Shortly after their discovery in December 1895, Rontgen realised the medical possibilities when he took the first X-ray "shadow" image of his wife's hand where you could clearly see the bone structure, and by May 1896, the first mass produced medical X-ray imaging machine (fluoroscope) was produced.

These "traditional" X-ray images measure the attenuation of the X-ray photons through an object. For medical imaging this is strongly related to the density of tissues. As bones have a significantly higher density than muscles and other soft tissues, there is a high attenuation difference resulting in a very high contrast. This is beneficial for identifying imperfections within bones e.g. broken bones, fractures and quality of metallic implants. While this is useful for identifying issues related to bones, there are many medical situations where the issues are within the soft tissues exclusively. Soft tissues do not always have a large density difference, making X-ray attenuation imaging unfeasible. This is where X-ray phase contrast imaging is beneficial.

X-ray phase contrast imaging is an effective imaging modality that increases the contrast of materials with a similar X-ray attenuation. This modality measures the phase shift of the X-rays as they travel through the different materials. This phase shift is introduced due to the different refractive indices inherent within materials, which affects the speed of the X-rays giving different phase shifts through different materials.

These phase shifts cannot be measured directly due to the sensors measuring the intensity of the X-rays. What can be measured are lateral shifts at the sensor that

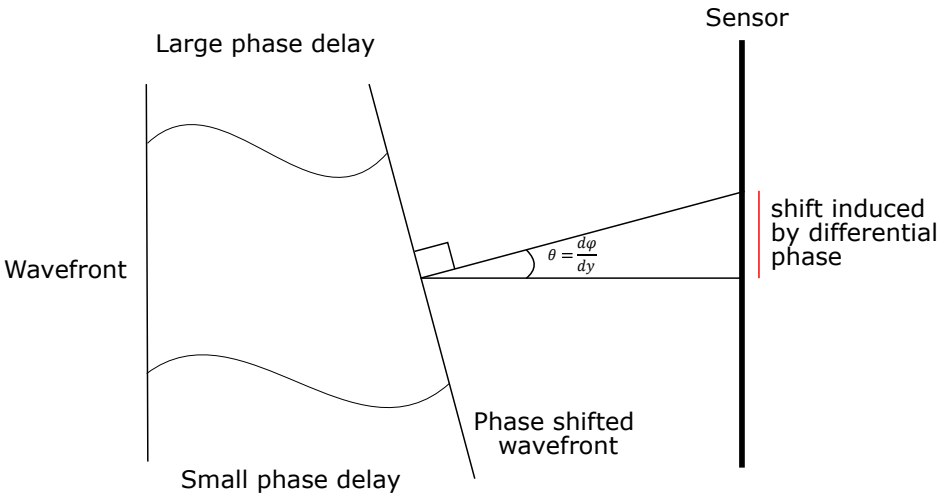

Figure 1.1: Example of lateral shift due to refraction

are caused by the refraction. Such measurement is achieved by imposing a pattern on the X-ray beam, this differential phase will induce a shift that can be measured by the movement of this pattern at the sensor. Figure 1.1 shows how the refraction angle is proportional to the differential phase. Currently the two common commercial and industrial techniques to spatially modulate the X-ray beam involve using a fine diffraction grating, or speckle imaging. Both create a known pattern which enables any fine shifts to be detected.

1.1.1 Grating based imaging

Phase contrast grating based imaging can be achieved by utilising two or more precision engineered diffraction gratings[7][1]. There are many slight variants on the grating based imaging, but the two main methods utilise X-ray Talbot interferometry [8] and retrieval of X-ray shifts by Moiré fringes [9]. The gratings used in these methods need to be precision engineered to achieve a grating period in the order of micrometres.

Though the techniques for phase shift retrieval are fundamentally different, the gratings used are often the same. The first “phase” grating should have a negligible X-ray attenuation but a small period between the fins to spread the incoming beam via diffraction. This diffraction creates an interference pattern with constructive and destructive regions. The second “analyser” grating needs to have a much higher X-ray attenuation to minimise the number of X-rays where the destructive interference is to increase the contrast of the interference pattern as seen in fig 1.2.

The gratings are etched onto a silicon wafer which has a relatively low X-ray attenuation, and then the analyser grating is electrochemically plated with gold to increase the X-ray attenuation, however this makes the analyser gratings very expensive. The

change in intensity values between a reference image containing only the shadow pattern, and the objective image with the object of interest, contains the phase information.

Figure 1.2: Grating phase contrast imaging setup, image from [1]

Talbot interferometry

X-ray Talbot interferometry [8] uses the Talbot effect of near field diffractions. When a wave travels through a diffraction grating, the image of this grating is repeated identically, at a certain distance, known as the Talbot length (Z_T) such that:

$$Z_T = \frac{2a^2}{\lambda} \tag{1.1}$$

where λ is the wavelength of the light, and a is the period of the grating. At distances other than the Talbot length, the interference pattern has a higher frequency. At $\frac{Z_T}{2}$ the pattern has the same frequency but is 180 out of phase. The relationship between the Talbot length and the fringe frequency is inversely proportional, such that $\frac{Z_T}{4}$ has double the fringe frequency, and $\frac{Z_T}{8}$ has 4 times the original frequency. this can be seen in fig 1.3. For phase contrast imaging, the distance between the grating and the sensor is often the Talbot length or half of the Talbot length as it the grating period is already very small (in the micrometre range). The other distances all produce fringe patterns that are smaller than the period of the grating, and that is often not desired due to the pixel size of the sensor. Positioning the grating and sensor at Talbot lengths also gives the fringes the highest contrast as the constructive and destructive interference is at its maxima.

This method is used in phase contrast imaging by placing the object of interest behind the first diffraction grating. The object then affects the X-rays wave fronts that approach the grating due to the refraction through the object. This affects the Talbot

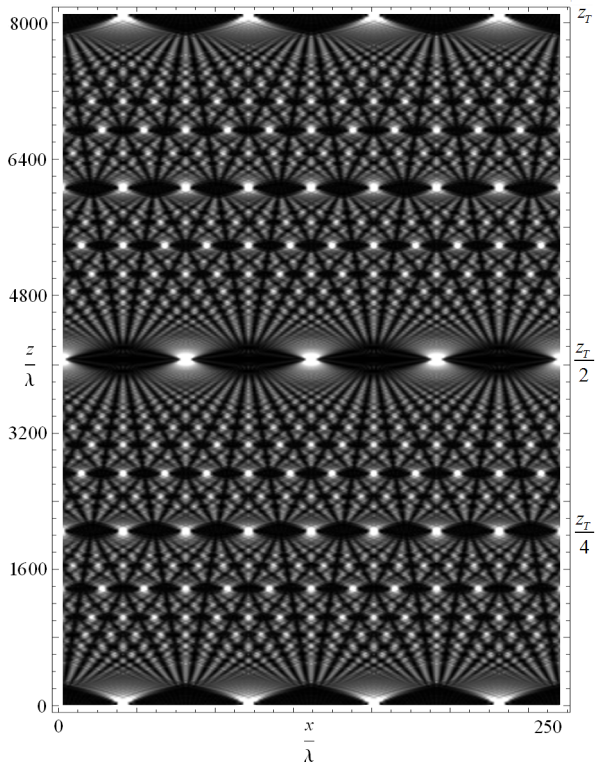

Figure 1.3: Example of Talbot carpet. Z shows distance to sensor and x horizontal distance parallel to the grating. Here the grating period is $a = 64\lambda$, from [2]

pattern as the additional phase shift, resulting from the object of interest, causes the patterns to no longer be aligned. In order to separate the phase information from the absorption and imperfections of the gratings, phase stepping is used. Phase stepping is when the phase grating is stepped across one period of the grating and at every step the new image is taken. The intensity oscillations at each pixel, can be represented by a Fourier series and hence the phase and attenuation can be easily separated.

Moiré fringes

Moiré fringes occur when there are two similar but not identical diffraction gratings are in the same beam of light[10][11]. These differences can be in frequency of the grating, angle of grating or a displacement of the grating. These differences result in an overlap of gratings in different areas, which creates lower frequency fringes as seen in fig 1.4.

These moiré fringes can then be treated the same as the fringes from the Talbot interferometry as it is the distortion of these fringes that contain the phase information. By comparing the uninterrupted moire fringe pattern, and the fringe pattern that has been shifted, the displacement can be calculated, which is related to the differential phase as in fig 1.1. Due to the lower frequency of the interference fringes it has a

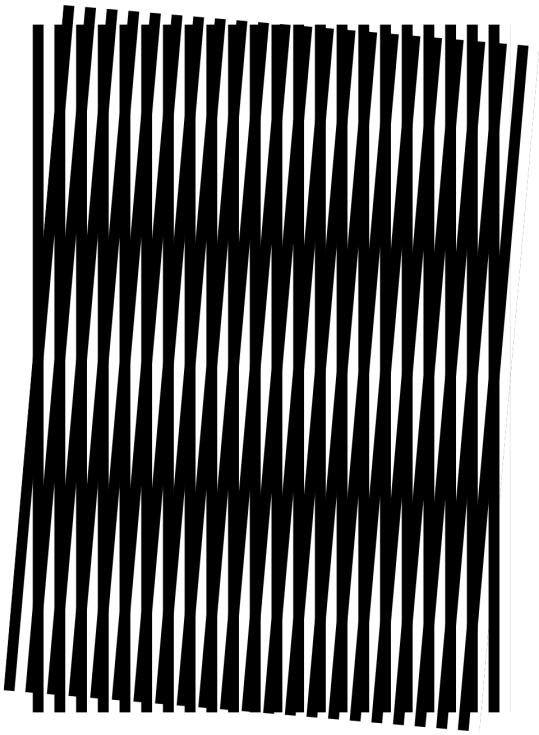

Figure 1.4: Example of moiré fringes with two identical period gratings offset by 5 degrees, from [3]

lower differential phase resolution. However, this lower resolution is countered by not requiring the phase stepping technique, this results in only requiring a single image rather than multiple images. This reduces the amount of storage required and time to take the images.

1.1.2 Speckle based imaging

X-ray speckle is caused by the constructive and destructive interference of a coherent monochromatic light source scattered from a rough surface or object. A synchrotron based X-ray source is the main way of producing monochromatic and coherent X-ray waves. Any disturbance to the light source causes random scattering of the photons, resulting in the interference to create speckle. This speckle can be split into two different regimes, near field and far field speckle. Far field speckle is reliant on the wave and its properties, while the near field speckle is reliant on the scattering material only. For X-ray phase contrast imaging, the near field speckle is much more accessible due to the speckle pattern being reliant on the scattering, and is independent of the X-ray beam

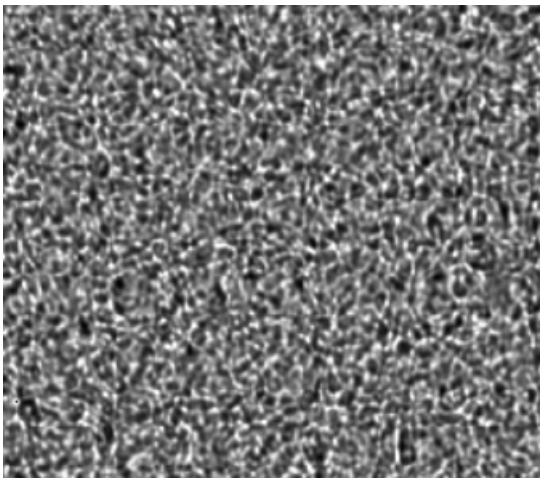

Figure 1.5: Example of X-ray speckle, from [4]

[4].

When the sample is introduced into this speckle image, it distorts the speckle pattern due to refraction within the sample. This refraction causes a local shift of the speckle pattern which, when the shifted image is registered with the original, unshifted, speckle image, gives the lateral displacement. This displacement is proportional to the differential phase shifts.

Speckle based imaging requires a coherent monochromatic X-ray source to generate high contrast speckle patterns, which is most commonly achieved with a synchrotron generator. However there are only a few number of these ($\approx$ 60 world-wide) which makes it essentially unavailable for routine medical uses.

A recently proposed alternative is to replace the speckle pattern with a spatially modulated intensity pattern due to attenuation. This enables polychromatic X-rays to be utilised. For example, commercial sandpaper produces a shadow pattern seen in fig 1.6. The attenuation is caused by the grains that mimics true speckle pattern generated by the synchrotron radiation.

This technique shares the same principles as the real speckle and some grating methods. A *reference image* of the uninterrupted sandpaper "shadow" pattern is needed. As the X-ray wave fronts travel through the sample, a phase shift is induced due to the different refractive indices which distort the shadow pattern, generating the *objective image*. This distortion in the objective image can be measured via local image registration, which is normally achieved within small windows across the image. This registration finds the displacement of the shadow pattern introduced by the object of interest. As with the grating based methods, the lateral shifts at each pixel relates to

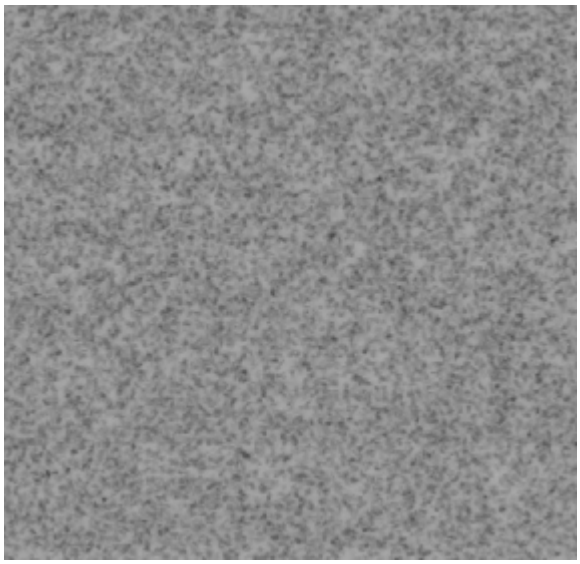

Figure 1.6: Example of X-ray sandpaper shadow, a sheet of 100 grit sandpaper

the differential phase and hence the registration generates a vector-valued differential phase image. The differential phase image is then integrated to generate the phase contrast image.

The main contributors to the effectiveness of shadow based phase contrast imaging are

- Average shadow particle size – small particles are optimal as these determine the spatial resolution
- Shadow particle visibility – high contrast helps with the reconstruction algorithm as the edges are more defined, smaller particles can be utilised.
- Randomness – needs to be non-periodic, non-repetitive, and non-isotropic to eliminate accidental matching.

To achieve these requirements, a material with micrometre sized pores or lumps is needed. A material such as sandpaper gives control over the grain size, and grain material which in turn gives control over the shadow size, and contrast, while maintaining a random distribution.

This thesis examines the feasibility of this method with low X-ray dosage. As X-rays have a high energy, the refractive indices are very close to 1, and the differential phase shifts are very small. This implies that a good sub-pixel registration algorithm is essential to give the most accurate differential phase shifts and hence the phase contrast image. With low X-ray dosage, noise is an issue to overcome, specifically Poisson noise due to the random numerical difference in photon counts per pixel.

The scope of this project is to investigate:

- The effects of different sub-pixel registration techniques.
- The effects of noise on differential phase and phase contrast images.
- The effects of denoising techniques on differential phase and phase contrast images.
- the relationship between particle grain size, registration window, and quality of phase contrast images.

Chapter 2

Sub-pixel Registration

Image registration is the process of aligning two or more images to each other [12][13]. These images can be from different sensors, different locations, or different times. If there is mutual information, that is they share the same scene, one image can be distorted to match with the other image on a shared co-ordinate system. This is widely used in remote sensing[14][15], super-resolution imaging[16][17], mechanical displacement metrology[18][19], and medical CT scans[20][21]. For many of these methods, registering the two images to the nearest pixel is not sufficient; this has resulted in a lot of research into extending registration techniques to a sub-pixel accuracy.

What is meant by the term sub-pixel? Within digital cameras, there is spatial quantization due to physical distance between discrete sensing elements. When performing registration, the two or more images are aligned to the nearest pixel. Sub-pixel registration extends this alignment to an accuracy less than a pixel.

In terms of X-ray phase contrast imaging, image registration is used to measure the offset caused by refraction (which relates to the differential phase) induced in each pixel by the object. As the X-rays travel through the object, the wave fronts will be refracted slightly due to the X-rays travelling through different materials with different refractive indices. As X-rays have a high energy, the refractive indices of materials are close to 1. This refraction results in a shift at the detector level which is identifiable. This shift due to the refraction angle is magnified over the distance between the sample and the sensor [22]. As the distance between X-ray source and sensor in commercial, hospital based, X-ray systems is relatively small, pixel-level registration is not suitable even with significant magnification. Therefore, a high accuracy sub-pixel registration method is essential to achieve good results.

2.1 Problem definition

Phase contrast imaging requires measuring the differential phase at each pixel. The *reference image* provides a known pattern. The *objective image* is a distorted reference image due to the object of interest introducing different phase shifts to the X-ray wave fronts due to changes within the object (different material, thickness, etc). The distortions introduced by the object of interest are found by mapping the offset from the reference image to the corresponding location in the objective image. This offset is proportional to the *differential phase*, which can be integrated to retrieve the final phase contrast image.

This requires the sub-pixel offset to be estimated for every pixel. This is commonly achieved by registering the images within small localised windows. The window size will affect the precision of the sub-pixel registration as a larger window will effectively average the differential phase shift within the window. The window size therefore affects the accuracy of the differential phase images. The window will blur the differential phase images, and in turn the phase contrast image.

The reference and objective image windows are first normalised via zero mean contrast normalisaton (ZMCN).

$$ZMCN = \frac{I(x,y) - \mu}{\sigma} \qquad \text{for } (x,y) \in \text{ window} \tag{2.1}$$

where I is the pixel intensity at each location within the window, μ and σ are the mean and standard deviation of all the pixels within the window respectively. This mitigates the luminance difference between the objective and reference images caused by the attenuation from the sample and software normalisation performed in the X-ray images. Figure 2.1 shows the effects of normalisation on the distribution of pixel values. The distribution is shifted and scaled to have a mean value of 0 and a standard deviation of 1.

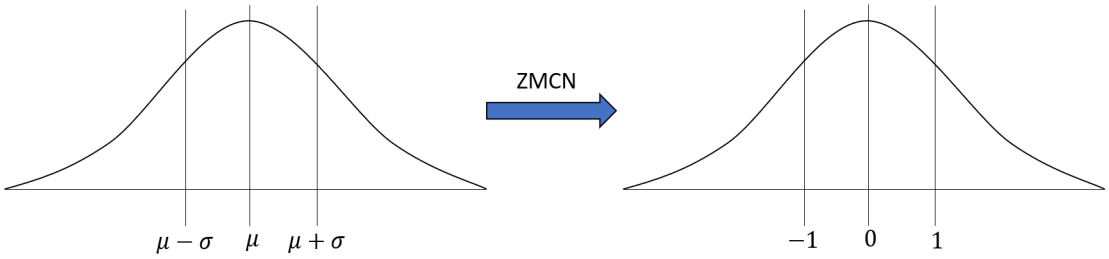

Figure 2.1: How ZMCN affects the distribution of pixel values

This chapter investigates the effectiveness of several sub-pixel image registration techniques and their feasibility for phase contrast imaging. The sub-pixel registration

algorithms evaluated here are:

- Correlation based methods
- Phase correlation methods
- Newton-Raphson root finding algorithm methods
- Optimal interpolation filtering

2.2 Correlation Based Methods

Perhaps the simplest subpixel registration method is fitting a 1 dimensional curve, or 2-dimensional surface to a correlation map. Square windows with a size of $w \times w$ centred around the same spatial position (x, y) in both the normalised reference and objective images are used to calculate the correlation coefficient. Shifting the window within the objective image creates a map of integer offset correlation values for both x and y directions.

When applied to ZMCN windows (reference = R, objective = I), zero normalised cross correlation gives the correlation co-efficient in the range of -1 and 1, where -1 implies the images are exact opposites, 1 implies the windows are exactly the same and a ZNCC values of 0 implies the windows have no similarity.

$$C(i,j) = \frac{\sum_{m=-\frac{w}{2}}^{\frac{w}{2}} \sum_{n=-\frac{w}{2}}^{\frac{w}{2}} R(x+m, y+n) \times I(x+m+i, y+n+j)}{\sqrt{\sum_{m=-\frac{w}{2}}^{\frac{w}{2}} \sum_{n=-\frac{w}{2}}^{\frac{w}{2}} R(x+m, y+n)^2 \times \sum_{m=-\frac{w}{2}}^{\frac{w}{2}} \sum_{n=-\frac{w}{2}}^{\frac{w}{2}} I(x+m+i, y+n+j)^2}} \tag{2.2}$$

2.2.1 1 Dimensional Curve Fitting

This correlation map gives the pixel-level displacement because the image only has pixel values at integer samples. To estimate the offset to sub-pixel accuracy it is necessary to estimate the true peak location. By assuming the peak is separable, the location of the peak can be found in both x and y directions independently of each other. The shape of the peak can be modelled using any number of models [23]. For phase contrast imaging using shadow patterns, the correlation map, as seen in figure 2.2 appears to be a paraboloid, so a parabola is fitted to the correlation map. The values chosen for calculating the sub-pixel peak, in the x direction, are $C_0 = C(i_0, j_0)$,

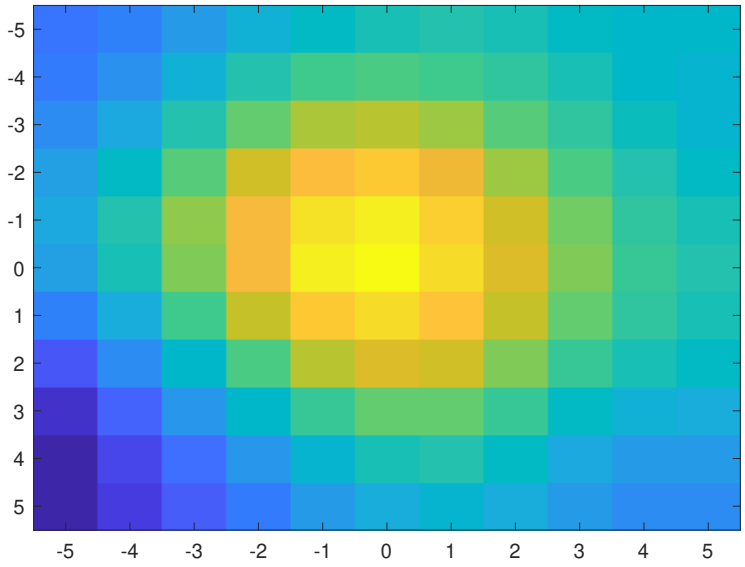

Figure 2.2: Example ZNCC correlation map on speckle image

which is the correlation peak, and the neighbouring values, $C_{+1} = C(i_0 + 1, j_0)$, and $C_{-1} = C(i_0 - 1, j_0)$. To then estimate the sub-pixel offset, the peak of the parabola is needed. This is achieved by finding the derivative of the parabola giving the offset in the x direction as [23]

$$d_x = i_0 + \frac{C_{+1} - C_{-1}}{4C_0 - 2(C_{+1} + C_{-1})} \tag{2.3}$$

and the same can be performed to solve for the y offset.

However, using a 3-point parabola is not an ideal case as the parabola will always fit the values, see figure 2.3. This implies that if there is any noise in the correlation map,

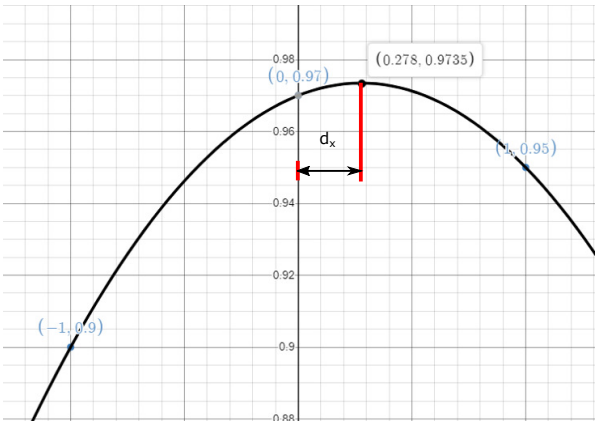

Figure 2.3: Example of a parabola fitting to 3 points

then it will fit the parabola to that noise as well. To mitigate this, the parabola can be extended to be fitted over 5 points rather than 3. The idea behind this is that more information from more data points will reduce the effect of noise. However, a parabola is far less likely to fit perfectly across all 5 points, so it becomes a line of best fit solved by least squares. Solving for the peak of this parabola gives the offset as

$$d_x = i_0 + \frac{7(-2C_{-2} - C_{-1} + C_1 + 2C_2)}{5(2C_{-2} - C_{-1} - 2C_0 - C_1 + 2C_2)} \tag{2.4}$$

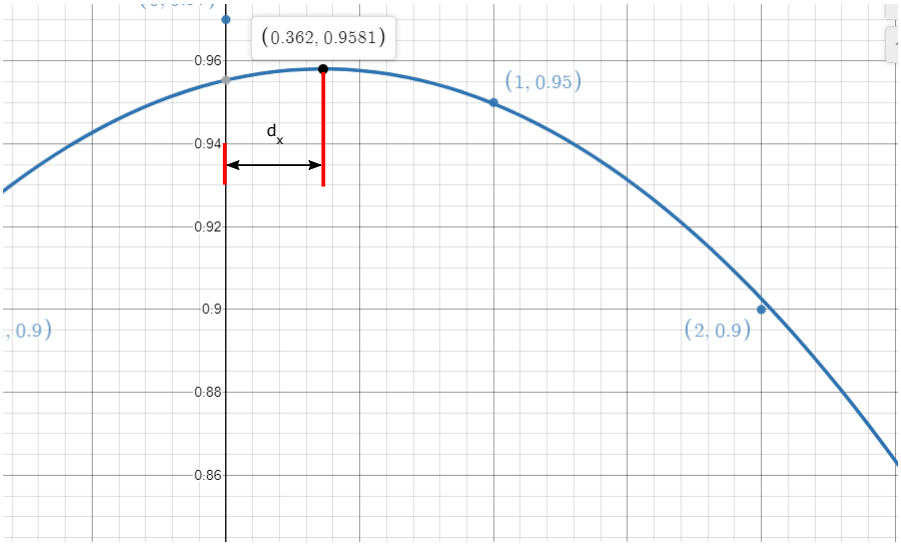

Figure 2.4: Example of a parabola fit by least squares to 5 points

Fitting curves in one dimension is based on the assumption that the shifts are completely independent of each other. That is a shift in y has no effect on the values in x, which is often not the case for real examples.

2.2.2 2-Dimensional surface fitting

To improve the accuracy of the results from one-dimensional curve fitting, a 2-dimensional parabolic model can be used.

$$C(i,j) = a_0 i^2 + a_1 j^2 + a_2 ij + a_3 i + a_4 j + a_5 \tag{2.5}$$

For a 2D parabolic model, a minimum of 6 points are required to calculate the coefficients. The peak, and its immediate neighbours in x and y provide 5 of the 6 points needed. This final required point is a corner (e.g. $C(i_0 + 1, j_0 - 1)$). Only using one corner skews the results in that direction. The coefficients can be calculated for all 4 corners, and then by taking the average estimation of the sub-pixel peak removes the

skew [24]. If all 9 points within a 3×3 window are used, the coefficients can be solved by least squares [25], although this has been proven to be less accurate[24].

With the 6 coefficients calculated, calculating the peak of the 2-dimensional parabolic model can be found by differentiating solving the resulting system of linear equations.

$$d_x = i_0 + \frac{2a_1a_5 - a_2a_4}{a_4^2 - 4a_3a_5} \qquad (2.6) \qquad d_y = j_0 + \frac{2a_2a_3 - a_1a_4}{a_4^2 - 4a_3a_5} \qquad (2.7)$$

This makes the calculations still very fast with the limitation being on calculating the correlation map.

2.3 Phase Correlation Methods

To perform phase registration there are two techniques to calculate the offsets: working directly in the Fourier domain, or taking the inverse Fourier transform an performing the sub-pixel estimation on the phase only correlation (POC). These methods are based on the shifting theorem of the Fourier transform which states *"a shift in the coordinate frames of two functions is transformed in the Fourier domain as linear phase differences"* [26].

As these phase registration techniques require working in the Fourier domain, it is beneficial to utilise a windowing function. This windowing function reduces the amount of spectral leakage that is introduced to the frequency response of the signal. The Fourier transform assumes the signals are periodic and as the images are different at the edges, there is a discontinuity that introduces spectral leakage. Applying a window function to the image before the Fourier transform, attenuates the signal at the edges, making them more similar, hence reducing spectral leakage. For this implementation a Hanning window was utilised. This was due to the Hanning window starts and ends at 0, nullifying the edge discontinuity. On top of this, it has very low side lobes, this results in less chance of any peaks being masked by the spectral leakage.

The direct Fourier domain method requires both images undergo the discrete Fourier transform so a plane can be fit to the magnitude of the phase differences such that:

$$\mathfrak{F}(R(x,y) = I(x + d_x, y + d_y)) \rightarrow R(u,v) = I(u,v)e^{-i2\pi(u+d_x, v+d_y)} \qquad (2.8)$$

$$\angle R(u,v) - \angle I(u,v) = P(u,v) \qquad (2.9)$$

where d_x and d_y are the image shifts, u and v are spatial frequency co-ordinates and P is the phase difference. By fitting a plane to the phase difference and then finding the gradient of the plane in both x and y directions, results in the spatial domain shift in pixels.

The POC method has the same underlying method as the correlation method, except that the correlation is performed in the Fourier domain rather than the spatial domain. The equivalent of the normalised cross correlation in the Fourier domain is the normalised cross spectrum, which instead of comparing the pixel intensity values, compares the strength of the frequencies in both the reference and objective windows.

$$G(u,v) = \frac{\angle R(u,v) \times \angle I^*(u,v)}{|R(u,v) \times I^*(u,v)|} \tag{2.10}$$

Where I^* is the complex conjugate of I. The POC however, will only give results to a pixel level shift in the spatial domain. This can be extended to give the registration to a sub-pixel accuracy, by fitting a 2-dimensional surface to the peak [24][4].

$$\mathfrak{F}^{-1}(G(u,v)) \rightarrow POC \tag{2.11}$$

A 2-dimensional sinc model performs the best [24] due to a pixel level shift generating an impulse in the Fourier domain, which when inverted to the spatial domain becomes a sinc function. The response will not be a perfect impulse for a sub-pixel shift, but it still the most similar response.

When comparing spatial domain correlation and POC, the difference is quite significant. For the spatial domain, the correlation peak is quite broad as the images at integer displacements are still quite similar, whereas the phase correlation has a much sharper peak due to the inverse Fourier transform of a shift being a delta function. This can be seen in figure 2.5.

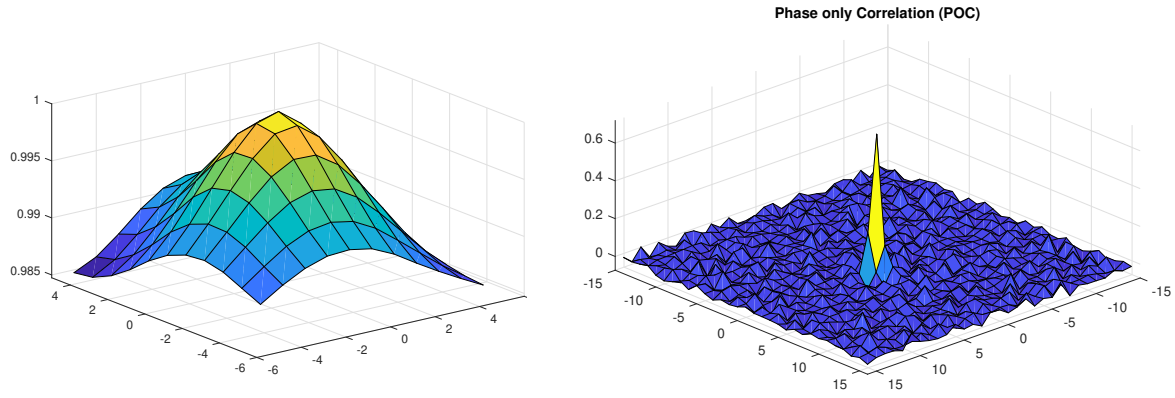

Figure 2.5: Spatial domain cross correlation (Left) Phase only correlation (Right)

2.4 Root Finding Methods

The correlation map can be found to a sub-pixel accuracy by interpolating the correlation map as seen in the curve and surface fitting techniques, or by interpolating the objective image, which can then be used to calculate the correlation map at non-integer locations. An iterative search is then used to find the best offset. An example of this is the Newton-Raphson method for image registration [27].

The Newton-Raphson method [28] is a well-known algorithm used for finding the zero crossings of a function. This is an iterative method that utilises the current value and the gradient of the function at that point to calculate the next, closer estimation. This is depicted in figure 2.6

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)} \tag{2.12}$$

This incremental estimation will only find the nearest zero crossing so if there are multiple zero crossings a close initial guess is required.

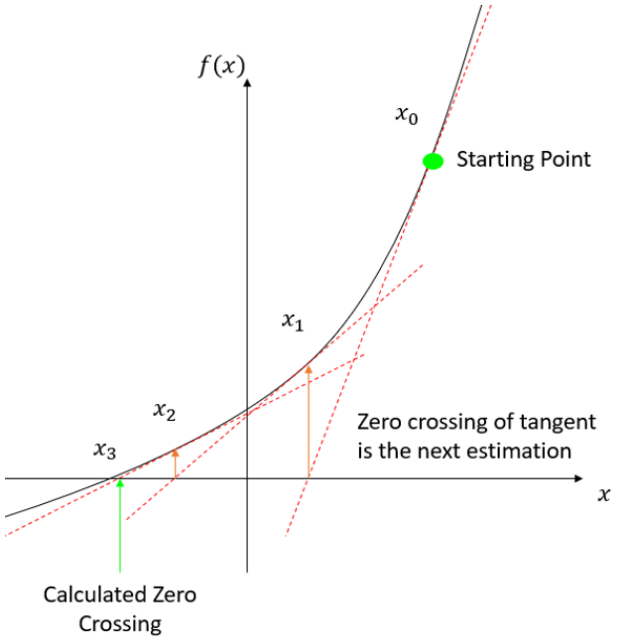

Figure 2.6: Example of Newton-Raphson root finding algorithm on a single variable function

In terms of image registration however, the goal is to find the peak of a correlation map (2.1)[12], rather than the zero crossings. To adjust for this, the derivatives of the correlation map are used. The zero crossings of the derivative correlation map is the

peak of the derivative. Applying this change gives the equation

$$(d_{x(n+1)}, d_{y(n+1)}) = (d_{x(n)}, d_{y(n)}) - \frac{C'(d_{x(n)}, d_{y(n)})}{C''(d_{x(n)}, d_{y(n)})} \tag{2.13}$$

which in 2-dimensions can be seen as

$$\begin{bmatrix} d_x(n+1) \\ d_y(n+1) \end{bmatrix} = \begin{bmatrix} d_x(n) \\ d_y(n) \end{bmatrix} - \mathbf{H}^{-1}(d_x(n), d_y(n))J(d_x(n), d_y(n)) \tag{2.14}$$

where J and $\mathbf{H}$ are the Jacobian and Hessian matrices containing the first and second partial derivatives

$$\mathrm{J} = \begin{bmatrix} \frac{dC}{di} \\ \frac{dC}{dj} \end{bmatrix} \qquad \mathbf{H} = \begin{bmatrix} \frac{d^2C}{di^2} & \frac{d^2C}{djdj} \\ \frac{d^2C}{didj} & \frac{d^2C}{dj^2} \end{bmatrix} \tag{2.15}$$

After estimating the shift values, the objective image is interpolated using a bicubic interpolation kernel, which is performed after every estimation.

2.4.1 Forward Additive - Newton-Raphson Method

The forward additive - Newton-Raphson (FA-NR) method was the first sub-pixel registration algorithm based on the Newton-Raphson root finding algorithm[27]. The integer-level correlation peak position is used as the initial approximation. Each iteration requires the first and second derivatives of the resulting correlation map[27] [18]. These are estimated by interpolating the objective image with a very small shift of $\pm\delta$

$$\frac{dC}{di}(i,j) = \frac{C_{(i+\delta,j)} - C_{(i,j)}}{\delta} \tag{2.16}$$

$$\frac{dC}{dj}(i,j) = \frac{C_{(i,j+\delta)} - C_{(i,j)}}{\delta} \tag{2.17}$$

$$\frac{d^2C}{di^2}(i,j) = \frac{(\frac{C_{(i+\delta,j)} - C_{(i,j)}}{\delta}) - (\frac{C_{(i,j)} - C_{(i-\delta,j)}}{\delta})}{\delta} \tag{2.18}$$

$$\frac{d^2C}{dj^2}(i,j) = \frac{(\frac{C_{(i,j+\delta)} - C_{(i,j)}}{\delta}) - (\frac{C_{(i,j)} - C_{(i,j-\delta)}}{\delta})}{\delta} \tag{2.19}$$

$$\frac{d^2C}{didj}(i,j) = \frac{C_{(i+\delta,j+\delta)} + C_{(0,0)} - C_{(i,j-\delta)} - C_{(i-\delta,j)}}{\delta^2} \tag{2.20}$$

These can then be stored in a Jacobian vector and Hessian matrix respectively:

$$J = \begin{bmatrix} \frac{dC}{di} \\ \frac{dC}{dj} \end{bmatrix} \tag{2.21}$$

$$\mathbf{H} = \begin{bmatrix} \frac{d^2C}{di^2} & \frac{d^2C}{didj} \\ \frac{d^2C}{didj} & \frac{d^2C}{dj^2} \end{bmatrix} \tag{2.22}$$

These are then substituted into 2.13 to give

$$\begin{bmatrix} d_{x(n+1)} \\ d_{y(n+1)} \end{bmatrix} = \begin{bmatrix} d_{x(n)} \\ d_{y(n)} \end{bmatrix} - \mathbf{H}^{-1}J \tag{2.23}$$

The shift values of the current estimate are applied to the objective image, bringing it closer to the reference image. This iteration is then terminated when the next increment is too small to make a difference ($|\mathbf{H}^{-1}J| \leq \delta$), or after 15 iterations. This hard iteration limit is to stop any infinite loops caused by diverging estimations[27].

2.5 Interpolation based matching

Another method to achieve a sub-pixel accuracy is to use some form of interpolation to the pixel level correlation (2.2). This is achieved by interpolating the objective image I using any interpolation method, such as bilinear, bicubic, splines etc. This interpolation should increase the correlation values which will be maximised when:

$$R(x,y) = I(x + d_x, y + d_y) \tag{2.24}$$

This usually requires an iterative search, such as the afore mentioned Newton-Raphson based methods, to find the sub-pixel offset [29][30]. There are two main issues with this approach though, first is choosing the best choice of interpolation kernel for the given situation. Second is that the filter coefficients are generally non-linear functions of the offset. The error function is therefore non-convex and may contain many local minima/maxima, although an iterative search will usually converge if the starting point is close enough to the peak.

2.5.1 Optimal Interpolation Filtering

One solution to avoid an iterative search and selection of an arbitrary interpolation function is to implement the "optimal" interpolation filter. As outlined in [31][32] the optimal interpolation filter *"attempts to predict the pixel outputs for the second image as a function of the pixel values in the reference image"* To determine the filter coefficients

h, least squares fitting is used.

$$h = \underset{h}{\operatorname{argmin}} \sum_{m=\frac{-w}{2}}^{\frac{w}{2}} \sum_{n=\frac{-w}{2}}^{\frac{w}{2}} (I * h - R)^2 \tag{2.25}$$

For a linear interpolation filter (i.e. one that can be represented as a convolution), this is convex with a single global minimum, which can be solved analytically for each patch. Representing this filter in matrix form,

$$I_h \cdot h = R_h \tag{2.26}$$

where I_h is a matrix of the normalised objective patch values where each row contains the input pixel for calculating a single output pixel, h is the vector of filter coefficients, and R_h is a vector of the reference image patch pixels. Equation (2.26) can be solved in a least squares sense to give the optimal filter values:

$$h = (I_h^T \cdot I_h)^{-1} \cdot I_h^T \cdot R_h) \tag{2.27}$$

The sub-pixel shift can then be calculated from the filter coefficients by super-position [32]:

$$dx = \sum i \times h(i,j) \tag{2.28}$$

$$dy = \sum j \times h(i,j) \tag{2.29}$$

Filters with an even number of coefficients introduce less bias into the results as demonstrated in [31] which is why the 2×2, 4×4 and 6×6 optimal filters are considered. Previously the optimal interpolation filter method was used to register whole images in the context of super-resolution [32]. This project has explored its applicability for local image registration using small patch sizes.

2.6 Testing Methodology

To test the feasibility of these sub-pixel registration techniques, it is essential to have an accurate, known sub-pixel shift. This is very hard to achieve mechanically when the sub-pixel step is $60\mu m$. Synthetic images were a solution to this problem. By making the image 20 times larger and then taking every 20^{th} sample, we can get accurate shifts of 0.05 pixels.

To initially test these registration methods, we used a minimum sub-pixel accuracy of 0.05 pixels. These registration methods were initially tested with a uniform shift, that is: everything in the image was shifted by multiples of 0.05 pixels, ranging from

0 to 0.5 pixels. This range was chosen as positive and negative shifts are symmetrical, and greater than 0.5px will change the integer level registration from the origin to the neighbouring location (e.g. a shift of 0.8 pixels from the origin becomes -0.2 pixels with regards to the integer registration). These sample images were generated with a grain size of 10, and a window size of 31, to allow for multiple grains to fit within the window, increasing the accuracy. For a 500×500 window 6400 grains were used, this allowed for the grains to overlap, creating shapes other than perfect spheres and hence making a more realistic synthetic image. To compare these methods, the mean and standard deviation of the x differential phase was calculated for each of the registration methods. This shows the registration bias (mean), and random error (standard deviation), which can then be used to compare the feasibility of these methods.

It should be noted at this point that there are plenty of papers outlining the accuracy of these methods [18][33][31]. However, these methods have not been explored or compared for X-ray phase contrast imaging using shadow-based intensity modulation. There are a few registration methods, notably the optimal interpolation filtering and phase correlation, which have not been widely explored with speckle-like images. While the Newton-Raphson based methods are predominantly used with speckle based images in digital image correlation [13][33][27].

2.7 Results

The preliminary results give an interesting insight into which of the registration methods are best suited for X-ray phase contrast imaging. As expected, the simplest methods, curve fitting, and surface fitting perform the worst, when compared to the iterative Newton-Raphson based methods and the "optimal" interpolation filter. This also highlighted a method that is not suitable for X-ray phase contrast imaging which is the phase registration method. The results seen here are for only shifts in one direction, by simply shifting x direction, and keeping the y shift at 0 (or vice versa), the detected shift is most accurate as it is not affected by the opposite direction. A 2-dimensional map of registration offsets can be seen in appendix A.1 for all the different registration techniques.

2.7.1 Curve Fitting

As seen in figure 2.7 the systematic bias (mean error) is far larger for the curve fitting methods with both 3 and 5 values from the correlation map. These errors are consistently larger than the other methods, and the values are different to what was expected. As the 5 point curve has more data points, it was expected that the shift values would be more accurate, and this is the case for shifts up to 0.3, however for

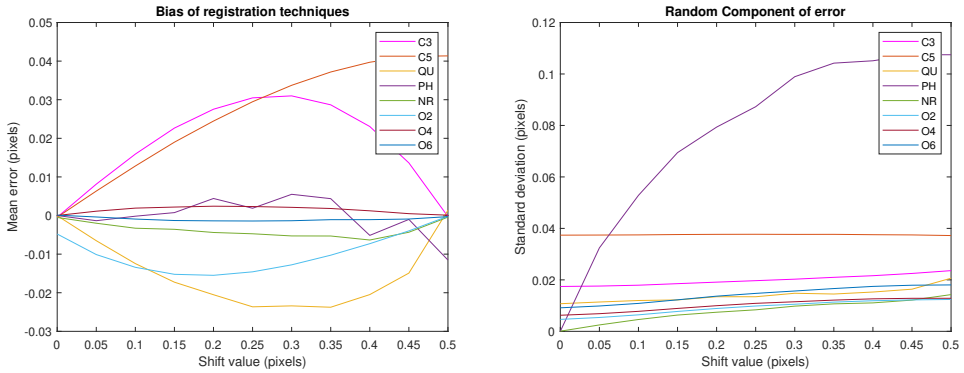

Figure 2.7: Results of registration error for each of the different sub pixel registration techniques. C3 = 3 point curve fitting, C5 = 5 point curve fitting, QU = quadratic surface fitting, PH = phase only correlation, NR = Newton-Raphson, O2 = 2×2 optimal interpolation filter, O4 = 4×4 optimal interpolation filter, O6 = 6×6 optimal interpolation filter

shifts larger than this the 3-point curve tends to increase in accuracy, while the 5 point curve fitting continues to get worse. Our tests were performed in steps of 0.05, as this was discussed by our team to be a good base for a detectable shift difference. To consider a shift of 0.05 detectable, the mean error needs to be less than half of the shift, in this case 0.025. Both of the curve fitting methods go higher than this after a shift of 0.2. This implies these methods are not capable of detecting sub-pixel shifts with an accuracy of 0.05 pixels. This is solidified by the standard deviation of values within the image, which for the 5-point curve fitting is quite consistent around 0.03, while the 3-point curve fitting has a lower standard deviation of about 0.01. For noise free images this makes sense, as with 3 data points a parabola will fit perfectly, and for 5 data points, the parabola becomes a curve of best fit so has a higher variance due to the smaller imbalances in the correlation map. It is obvious from this that they perform worse than the more complex and accurate methods such as the Newton-Raphson and optimal filtering techniques.

Due to the 2-dimensional nature of the surface fitting method, and the least squares minimisation of the surface, it performs much better than both of the curve fitting techniques. This can be seen by the mean error being less than 0.025, and the standard deviation of 0.02 which implies better performance than the 1 dimensional curve fitting counterparts.

2.7.2 Phase Registration

The mean error, or bias, of the phase registration is significantly smaller than the curve fitting and surface fitting techniques, even though they also fit to a correlation map

with no iterations. The bias is more comparable with the optimal filters 4×4, and 6×6. However, unlike the other sub-pixel registration techniques, there is not a smooth curve, it instead centred around 0, but has a larger error at higher shifts due to larger standard deviation.

However, where the phase registration performs poorly is the random component of the error, shown by the large standard deviation. The phase registration very quickly has the worst standard deviation ≈ 0.11. This high random component of the error is due to the POC producing a sinc function where the peak is clearly defined as seen in figure 2.5, but due to this, small deviations of the neighbouring pixels to the peak of the POC have a large effect on the sub-pixel registration, resulting in a higher standard deviation.

This high random component of the error is detrimental to the quality of phase contrast images. The low bias implies it is capable of measuring a differential phase shift, however the high random component implies it has a low precision with large deviations around the shift, which when integrated will produce a poor phase contrast image. This results in a limited potential for X-ray phase contrast imaging using shadow-based intensity modulation.

2.7.3 Newton-Raphson

The Newton-Raphson method performs much better than curve fitting techniques, with the maximum error is less than 0.01, this implies it can accurately detect different shifts of 0.05 with no overlap. This high accuracy is to be expected as the Newton-Raphson method is an iterative solution, so continues updating the shift estimate until the best solution is found. These error values can be improved by increasing the complexity of the interpolation method. Bicubic interpolation is used in this method as a good trade-off between computation time and accuracy, a more complex interpolation method, such as biquintic interpolation or spline interpolation, should introduce less bias and hence result in a lower registration error.

When looking at the standard deviation, it can be seen that the Newton-Raphson method performs better with smaller shift values, as the standard deviation begins to increase when the shift is approaching 0.5.

2.7.4 Optimal Filters

As to be expected from their name, the “optimal” interpolation filters perform as well as, or better than the Newton-Raphson based methods which are considered to be state of the art for many digital image correlation problems. The 4×4 and 6×6 optimal filters have the lowest registration error, staying under 0.005 pixels while the 2×2 optimal

filter performs more similarly to the Newton-Raphson registration method. The largest optimal filter has the lowest registration error as it contains more information in the filter values. The 4×4 optimal interpolation filter has the same number of filter coefficients as the Newton-Raphson method, but as the optimal interpolation filter is not limited to a choice of interpolation method, the registration bias is significantly lower, and they share similar standard deviations, until the higher shift values, where the Newton-Raphson method performs worse.

When looking at the standard deviation of the optimal interpolation method, the 6×6 filter performed the worst of the optimal filters, while the 2×2 had the lowest standard deviation of the optimal filters. This is also correlated to the filter size, as the values of the filter co-efficients are weighted, small changes at the edges of the filter will be magnified more than small changes near the centre, resulting in a larger standard deviation. As the mean is significantly smaller than the standard deviation, the performance of the optimal interpolation filters is dominated by the random error component.

2.7.5 Conclusions

This uniform shift test has demonstrated that a range of different registration techniques are feasible for X-ray phase contrast imaging using shadow-based intensity modulation, while there are some registration techniques that are less reliable. These conclusions are based on the random component of error as it tends to be larger than the registration bias.

The phase registration performed the worst out of all the methods explored. This is due to the surface being most similar to a 2-dimensional sinc function that results in a higher random component of error, and hence a poor differential phase image. The Curve fitting methods, with 3 and 5 data points, are borderline feasible. They have a larger random component than most of the other techniques and hence there are many other techniques that would perform better.

Of the correlation methods, the surface fitting method has the most potential for phase contrast imaging using shadow-based intensity modulation as shown by the consistently lower random component of error when compared to the 1-dimensional curve fitting counterparts. It performs similarly to the 6×6 optimal interpolation filter. However, there is a large difference in bias between the two methods which favours the optimal interpolation filters, implying the surface fitting with a 2-dimensional parabola performs worse.

The methods that show the most promise for X-ray phase contrast imaging using

shadow-based intensity modulation are the Newton-Raphson method, and the optimal interpolation filters. These outperformed the other methods in terms of random component which has the larger effect on the accuracy of the registration. The Newton-Raphson based methods are often considered state of the art for digital image correlation problems due to their iterative nature, resulting in high accuracy. The optimal interpolation filter, has only been explored for image registration [32] previously, but shows promise to be as capable as the Newton-Raphson based methods for X-ray phase contrast imaging. A benefit of the optimal interpolation filters is that it is a direct method. This means it does not require any iterations that is apparent in the Newton-Raphson based methods, which reduces the computational complexity of the optimal interpolation filter.

Chapter 3

Synthetic Images

Generating synthetic images was imperative for testing. This was due to their ability to give images with an accurate known shift which allows for quantitative testing and comparative assessment of different sub-pixel registration methods. Synthetic images also give complete control over the important features of the sandpaper. This includes:

- grain attenuation [34]
- average grain size [35]
- grain density [36]

Having control over these features gives the opportunity to explore how these affect the quality of the image registration. This will give insight into the selection of the best features of the shadow reference for optimising the quality of the differential phase images and in turn the phase contrast image.

3.1 Accurate sub-pixel shifts

To achieve this precise testing, pairs of images with an accurate, known, localised sub-pixel shift are needed. This known and accurate shift is hard to achieve with true images. This is due to the X-ray sensors having a pixel size of $\approx 100\ \mu m$ resulting in the sub-pixel offset requiring movement of the sandpaper by a few μm. This can be achieved with a piezoelectric actuator, which can achieve mechanical translation in the nanometer range [37]. However, the drawback of a piezoelectric actuator is that it only applies a uniform shift, rather than having the ability to apply small localised shifts. Therefor the synthetic images were the only option for accurate testing.

Synthetic images solve this by generating the images to be $n\times$ times larger, filtering with a $n \times n$ box filter to simulate area sampling, and then taking every n^{th} sample to

generate the correct sized images. By starting at an offset of d pixels, the downsampled images then have a shift of $\frac{d}{n}$.

A further benefit to the synthetic images is that it allows for sub-pixel shifts to be localised rather than shifting the entire image (piece of sandpaper) like a piezoelectric actuator does. This is beneficial as the differential phase will change throughout the image. The differential phase shifts induced by the object are going to be localised to the different thicknesses/materials/tissues etc. This is achieved by only applying the offset in the required region. This also furthers the testing of registration methods, as small windows are used for local registration. As the window moves over the edge of the shifted region, there will be some blurring effect at the edges of the object. The blurring is potentially different for each of the different registration methods, and synthetic images make it possible to quantify and assess these effects.

With the local shift value known, it is then relatively easy to quantify the error using any of the many pre-established error metrics such as mean square error (MSE), mean absolute difference (MAD), sum squared differences (SSD) etc.

3.2 Control of Sandpaper Features

The quality of speckle, which the shadow images are mimicking, used in image registration and digital image correlation has a significant effect on the accuracy of the results [4][33]. These synthetic images provide the accessibility to adjust the features and explore how changing these features (most notably grain size and grain density) affects the results of the registration.

3.2.1 Grain Material

Commercial sandpaper is usually made with silicon carbide, which has a relatively low X-ray absorption coefficient. This results in a low contrast between the grains and background within the shadow image, which is not ideal [33][34]. The registration error is much larger with a lower contrast, this implies that a higher contrast is beneficial. At 100KeV tungsten has a X-ray absorption coefficient(μ) of $4.483\frac{cm^2}{g}$, compared to silicon which has $0.1835\frac{cm^2}{g}$ [38].

Due to this, our research group decided to explore the use of tungsten powder for making our own shadow images that provide a stronger contrast than the sandpaper, hence providing us with higher quality phase contrast imager. The Tungsten particles are 10 μm in diameter which are spread onto a sheet of polymethyl-methacrylate (PMMA) substrate with epoxy.

3.2.2 Average Grain Size

When generating speckle/shadow images, the grain size should be kept as small as possible. This is due to the speckle size defining the minimum size of the window used for local registration, which in turn limits the spatial resolution [35]. The window must be larger than the largest grain size, as it is the grain edges that contain the information the provides the features for regsitration. If the window is smaller than a grain, the registration becomes much less accurate.

As with all digital signal processing, we need to avoid aliasing, the absolute minimum grain size is two pixels [39]. This comes from the Nyquist sampling theorem which states:

"if a signal $x(t)$ has a maximum frequency of B hertz, it is completely determined by giving its ordinates at a series of points spaced $\frac{1}{2B}$ seconds apart"

$$f_s \geq 2B \tag{3.1}$$

This was stated for a 1D signal, but can be easily extended to 2D. This extension to imaging, results in the sampling points being pixels. To achieve this Nyquist sampling rate an absolute minimum of two pixels per grain is required for perfect reconstruction. Anything larger than this is also sufficient as there is more than two samples per speckle.

3.2.3 Grain Density

The grain density is important when generating the synthetic images [36]. This is due to the edges of the grains containing the most useful information for the registration. If there are windows with no grains or edges, the amount of information is less, which reduces the reliability of the registration.

This implies that for a low grain density a larger local window will be required. This larger window results in a lower resolution of differential phase images, this is due to the percentage of shifted pixels within the window being much lower, and hence the average shift in the window is much lower and smaller shifts get lost more easily. On the contrary a high grain density is desirable up to a point. This will result in a smaller window, increasing the resolution of the differential phase and by extension phase contrast images.

3.3 Assumptions

These synthetic images, while essential for preliminary testing, are simplified copies of shadow images. There were a few obvious assumptions, that are different to the

true shadow X-ray mask image. However, they are not different enough to be a major concern. These are

1. Substrate and glue has negligible, and spatially uniform attenuation
2. Grains are perfectly spherical
3. The X-ray beam is parallel

3.3.1 Negligible Attenuation

As one of the requirements for a good speckle/shadow image is high contrast [33], it was decided earlier in the project to use a tungsten powder rather than store bought sandpaper. This has a much higher X-ray absorption and hence has a much higher contrast than any variations in the substrate and epoxy.

3.3.2 Spherical Grains

The grains were made to be spherical for simplicity. This is not the case for the Tungsten shadow images as the tungsten particles are all different, imperfect shapes to increase abrasiveness. This can be countered in our synthetic images by allowing the grains to overlap, which they do not do in sandpaper. This overlap makes the combined grains less spherical as they then become a random shape with differing regions of thickness. While there is no explicit dents or cavities, the differing thickness mimics this suitably.

3.3.3 Parallel X-ray Beams

When making these synthetic images, it was decided to approximate the X-ray beam as parallel rather than a fan beam. The difference between a fan beam and parallel beam can be seen in figure 3.1. A fan beam produces X-rays from a single point and the beam expands towards the detector. The tungsten (or other shadow mask) is early in the imaging set up and is a single absorbing plane. This results in the the distance to the detector, and the in-plane distance from the point source, having a blurring effect on the due to geometrical magnification at the detector The main practical difference is that the fan beam will introduce some slight blurring at the edges of the sensor. This can be seen in figure 3.1 , as the tungsten particles in the shadow mask get further away from the X-ray point source, in the same plane, the distance to the sensor is much larger resulting in a blur at the sensor.

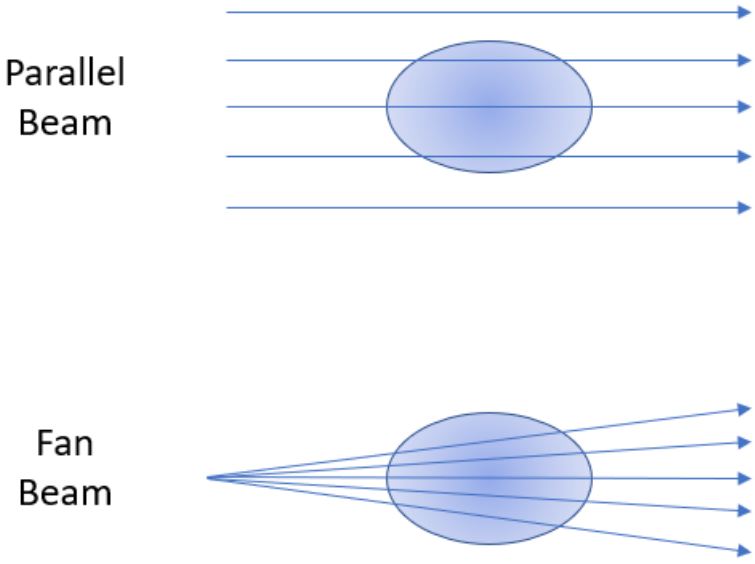

Figure 3.1: Fan beam vs parallel beam

3.4 Generating Synthetic Images

To generate these synthetic images, the first step is to decide what the desired sub-pixel shift P is for testing. A mask of a circle is generated. These are not wanting to be all the same size as there needs to be some distribution of sizes. To achieve this the grain size is randomly generated following a normal distribution, with a mean of pixels ($\frac{M}{P}$ pixels after downsampling) and a standard deviation of S. The tungsten powder that was used for generating the shadow images has a size of 10 μm.

The thickness t of the sphere can be calculated based on the radius r and distance from the centre in both x ($|x_0 - x_i|$) and y ($|y_0 - y_i|$) directions such that:

$$t(x_i, y_i) = \begin{cases} 2\sqrt{r^2 - (|x_0 - x_i|^2 + |y_0 - y_i|^2)}, & \text{for } |x_0 - x_i|^2 + |y_0 - y_i|^2 < r^2 \\ 0, & \text{elsewhere} \end{cases} \tag{3.2}$$

As the assumption has been made that the grains are perfect spheres, the thickness is a combination of the radius r and the distance from the centre of the sphere. Pythagoras theorem is applied twice to calculate the thickness. First to calculate the shortest distance from the centre, and then the root of the difference between the radius squared and distance squared gives the thickness. This then needs to be multiplied by 2 as that

calculates the thickness of a semicircle. an example of a spherical grain can be seen in figure 3.2.

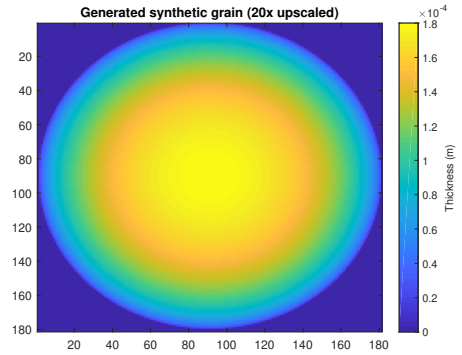

Figure 3.2: Final grain estimation, pixel value is proportional to the thickness of the grain

Once this sphere has been generated, it is inserted into another image that is $M\times$ larger than the desired output image. This is achieved by generating two random numbers that become the centre values for the x and y position, the sphere is then added to the output image. Simple addition can be used as the borders are 0 so adding them on top of grains has no affect, while adding grains to each other adds to the thickness. These three steps are then repeated until the desired grain density has been achieved, and the final thickness image has been generated, as in figure 3.3.

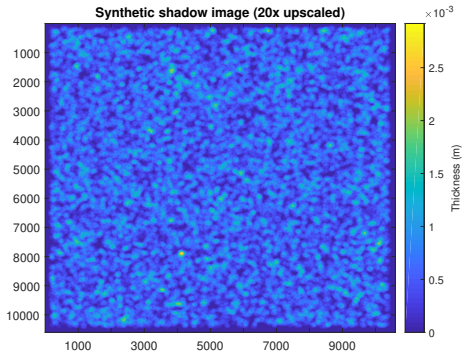

Figure 3.3: Final image showing the thickness of the sandpaper

This thickness image then needs to have the borders clipped so there is no blank regions around the edge of the paper, and it can then be transformed to an X-ray absorption image. The attenuation of each pixel can be calculated using the X-ray absorbance coefficient which can be found for many materials from Argonne National Laboratory [40]. This can then be used to transform the thickness image t into a

transmission image T:

$$T(x, y) = e^{-2\mu t(x,y)} \tag{3.3}$$

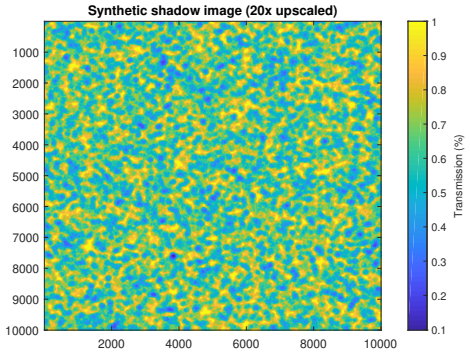

Figure 3.4: Synthetic image, mimicking transmission of 100keV X-rays through 10 μm tungsten

Once the whole image has been transformed to synthesize the transmission of X-rays through the 10 μm tungsten powder, the whole image is box filtered to simulate image capture on a digital sensor resulting in a high resolution (HR) image. Then the image is downsampled by taking every M^{th} pixel, and the objective image is downsampled the same way, taking every M^{th} sample but instead of starting at the same position, it starts at an offset of d. This gives a $\frac{d}{M}$ offset to get a uniform shifted image.

3.4.1 Phantom Images

The main downside to the piezoelectric actuator is that it shifted the entire image. One of the main benefits to using synthetic images is that it makes small localised shifts available for testing. These are known as phantom images. A phantom image has one or more regions within an image that has an offset. These regions can be any shape and have different shift values allowing for more complex testing.

To generate a phantom image, the synthetic images are generated identically to those above until the high-resolution image stage. A shift map of localised shifts, for both x and y differential phase images, needs to be generated before downsampling. These maps contain the shift S in terms of HR pixels. To achieve the shift in the downsampled images, every $M^{th} + S$ sample is taken as the downsampled value. As there are two shift maps, they are applied separately (x then y or vice versa).

Chapter 4

Effects of Noise

Noise is an unwanted signal that distorts the original signal, making it harder to identify and distinguish the true information [41][42]. For imaging, noise results in sharp changes in luminance. Noise can occur from many different sources within an image, most dominantly:

- Thermal noise and dark current noise are inherent within the sensing devices and circuitry.
- Shot noise which is related to the number of photons received during the exposure time.
- Quantisation noise occurs due to the analog to digital conversion.

These sources of noise will reduce the quality of the images and hence will reduce the accuracy of the sub-pixel registration [33][12]. This is problematic as the sub-pixel registration is related to the differential phase. Therefore, with less accurate registration, the differential phase images and by extension the phase contrast images become less accurate. For this reason, it was important to explore the effects of noise and denoising methods on these sub-pixel registration methods, with a primary focus on Poisson noise, which is the dominant type of noise in a low dosage X-ray shadow image.

4.1 Poisson Noise

As the project is based on low dosage X-ray phase contrast imaging using shadow-based intensity modulation, shot noise is going to be the most dominant noise in the images [41][43]. Shot noise follows the Poisson distribution. This distribution *"models the number of times an event occurs in an interval of time or space"*. This distribution

can be described as:

$$P(k) = \frac{\lambda^k e^{-k}}{k!} \tag{4.1}$$

Where λ is both the mean and variance and k is the sample value [44]. In the case of imaging, the Poisson or shot noise occurs due to the discrete differences in photon counts at each pixel in the sensor [41]. To meet the requirements of the Poisson distribution, these criteria must be met:

- "k" is an integer number of occurrences – an integer number of photons will occur at each pixel, no such thing as a half or a third etc. of a photon.
- The occurrences are independent - The photons sensed by each pixel are independent, one photon arriving does not change the probability of a second occurring.
- The average sample rate is constant - the X-ray source produces photons at a constant rate for the entire exposure.
- Two events cannot occur at the same time – two photons cannot arrive at exactly the same time.

Each of the pixels will have a similar but different value depending on the number of photons due to the discrete generation of photons from the X-ray source. As shot noise follows a Poisson distribution, the signal to noise ratio (SNR) is described as:

$$SNR = \frac{\mu}{\sigma} = \frac{\lambda}{\sqrt{\lambda}} = \sqrt{\lambda} \tag{4.2}$$

As the photon count decreases, then so does the SNR (e.g. $\lambda = 121 \rightarrow SNR = 11$ $\lambda = 49 \rightarrow SNR = 7$). Since low dosage X-ray images have fewer photons detected and hence the shot noise will result in the image having a low SNR.

4.2 Applying Noise to Synthetic Images

As the initial testing was performed on synthetic images it made sense to use the same synthetic images to test the susceptibility of different sub-pixel registration methods to noise. As the synthetic images were previously generated as a transmission ratio, they need to be scaled by the X-ray beam intensity and exposure time to give an expected photon count detected by each pixel. the expected photon count gives λ for the distribution 4.1 from which the noisy pixel value can be drawn.

The range of maximum photon count used was 10 to 50,000 where 10 photons is dominated by noise, and 50,000 photons is visually noise free as seen in figure 4.2. These images match what is expected from shot noise, a low SNR for low photon counts and at 50000 photons the images are visually noise free.

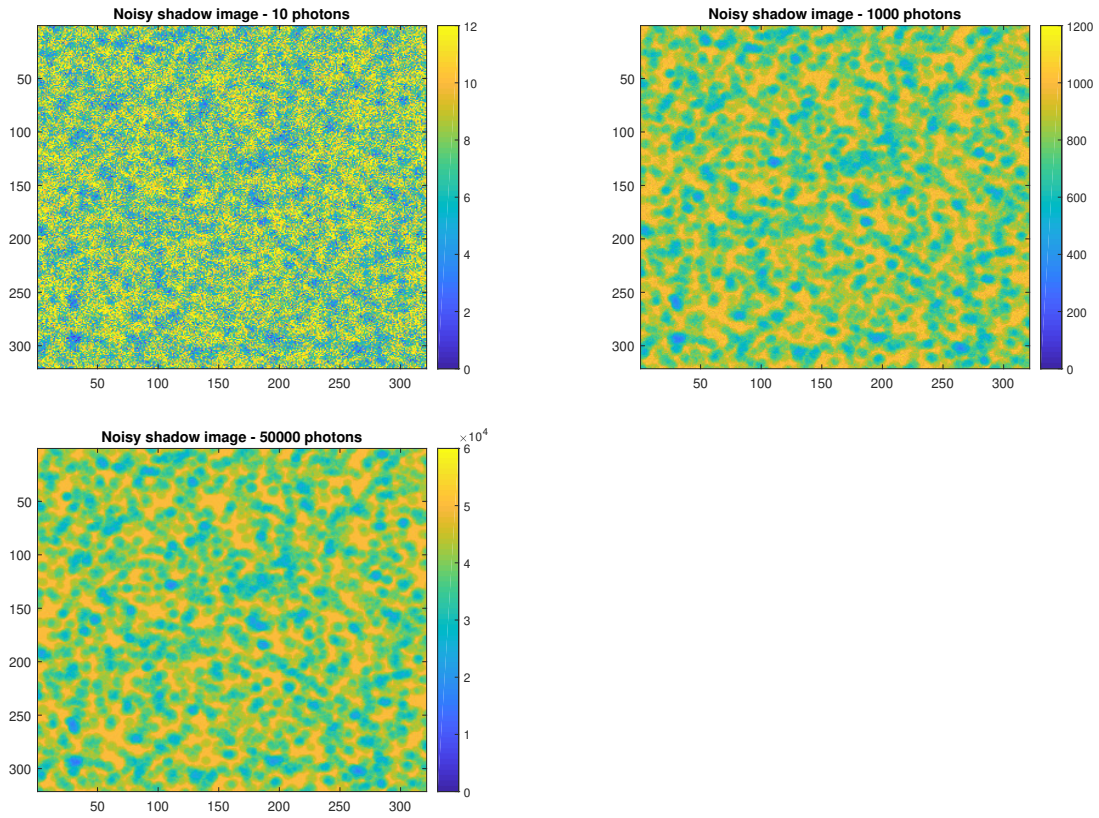

Figure 4.1: Example of synthetic images with a noise free maximum of 10,1000,50000 photons, corrupted by shot noise

4.3 Noise Mitigation

As noise is an unwanted signal [41][42], mitigation of this noise will not only improve the visual quality of the images, but also should result in an improvement of the sub-pixel registration, and by extension the phase contrast image. There has been a lot of research into denoising techniques with a wide range of general and image specific denoising techniques. The aim of denoising is to return the image to its noise-free state, or as close as possible. These denoising techniques can be assessed by looking at the following criteria:

- Flat regions in an image should be smooth.
- Edges should be sharp.
- Textures should be preserved.
- Not introduce new artefacts.

The majority of denoising techniques do not meet all of these criteria but instead

choose what is most important for a specific function. The "traditional" denoising method is achieved by low pass spatial domain filtering which blurs edges. There are more complex denoising techniques, such as wavelet transform denoising which relies on the sparsity of signals in the wavelet domain. Poisson and Gaussian noise affect the entire frequency range equally because the noise is independent so noise in one pixel does not affect the adjacent pixels. More recently there has been an increase of research into neural network denoising for image specific data, which given good training data can achieve good denoising while meeting the required criteria [45].

The majority of denoising research has been applied to additive white Gaussian noise (AWGN). As shot noise is going to be more dominant in low dosage X-rays, the standard denoising techniques are not optimised for Poisson noise. Due to this, it may be beneficial to transform the shot noise into Gaussian noise by applying a variance stabilising transform such as the Anscombe transform [46]. This Anscombe transform can be described as:

$$A : x \to 2\sqrt{x + \frac{3}{8}} \tag{4.3}$$

For Poisson noise, the mean and variance are not independent, while for Gaussian noise they are. The Anscombe transform is a variance stabilising transform that transforms Poisson noise into Gaussian noise by making the standard deviation of the noise essentially constant rather than changing with the pixel value as it does with the Poisson distribution.

This project explored a wide range of denoising techniques including:

- box filtering
- weighted low pass filter
- Wiener filter
- wavelet denoising
- BM3D denoising

As mentioned above, some of these are specialised to work on AWGN rather than shot noise, so apart from the wavelet denoising and neural network denoising, the images undergo the Anscombe transform before denoising. The main goal for many of these methods is to limit the high frequencies as noise affects the entire range of frequencies, images tend to have most of their information maintained in the low frequencies, so by attenuating the high frequencies it is primarily removing the noise.

4.3.1 Box Filtering

Box, or average filtering [47] is the simplest of the spatial domain filtering techniques. It simply takes the average of all the values within a window such that the filter coefficients are equal and sum to 1, for example a 3×3 box filter would be:

$$\begin{bmatrix} \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \end{bmatrix}$$

By convolving this filter with the noisy image, the sharp intensity fluctuations within the image caused by the noise, are minimised as the noise values are independent, the values higher and lower than the average should cancel out within the window.

4.3.2 Weighted Low Pass Filter

Much like the box filter, the weighted low pass filter is a combination of the values in a local window. The main difference is the weights are not the same for all values, the centre pixel is considered to be most similar to the actual value and has the highest weighting, the direct neighbours (vertical and horizontal) are closest spatially to the centre pixel and should have similar values, and hence have a higher weight than the corners. The filter can be seen as:

$$\begin{bmatrix} \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \\ \frac{2}{16} & \frac{4}{16} & \frac{2}{16} \\ \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \end{bmatrix}$$

This achieves marginally less noise mitigation, depending on the filter weights, but also blurs the image less which depending on the application and amount of noise, may perform better than the box filter.

4.3.3 Median Filter

The median filter is a non-linear filter. It is not a combination of all values within the filter, instead, it simply takes the middle value within the sorted list of pixel values within the filter [47] [48]. By taking the median, the extreme high and low values caused by the noise are completely ignored. This can be seen in figure 4.2. Depending on noise level, the median filter will produce less blurring than the low pass and box filters, but will also be less effective at removing Gaussian noise than the box filter for

a given window size. This is the same for low light images and Poisson noise as well. At low photon counts, the Poisson distribution will have a strong skew towards the lower photon counts, introducing a bias. For the higher photon counts, the Poisson noise becomes more Gaussian,

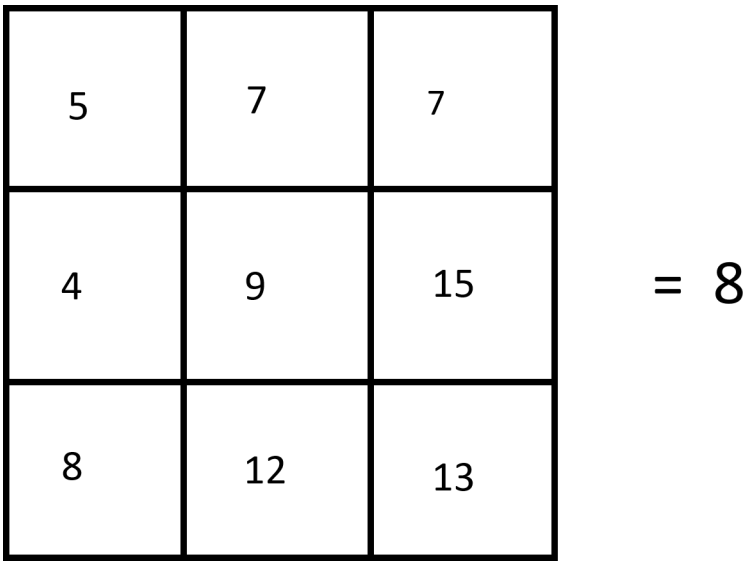

Figure 4.2: Example of a median filter

4.3.4 Wiener Filter

The Wiener filter is proven to be an effective filter for removing Gaussian noise in images [49][47][50]. The method is simple, but acquires very good results, and performs much better than direct low pass filtering, and median filtering, though the Wiener filter has the potential to blur edges if the noise is overestimated. The Wiener filter is based on the minimising the mean square error such that

$$e^2 = E[(f(x,y) - \hat{f}(x,y))^2] \tag{4.4}$$

Where E is the expected value, $f(x,y)$ is the noise free signal, and $\hat{f}(x,y)$ is the reconstructed signal. To solve for the reconstruction filter, the Fourier transform is used [50] such that

$$\hat{F}(u,v) = \frac{H^*(u,v)S_f(u,v)}{S_f(u,v)|H(u,v)|^2 + S_n(u,v)}G(u,v) \tag{4.5}$$

In this $G(u,v)$ is the noisy image, S_n is the noise power spectrum and S_f is the image power spectrum, which is the Fourier transform of the autocorrelation of the input signal. Taking the inverse Fourier transform of $\hat{F}(u,v)$ gives the denoised image.

4.3.5 Wavelet Denoising

One of the most investigated transform domain techniques for image denoising is the wavelet transform. The wavelet transform decomposes an image by passing it through layers of high and low pass filters. There are two main techniques [51] to denoise an image in the wavelet domain, these include: threshold denoising, and correlation denoising.

Threshold Denoising

Thresholding is the most simple of these denoising techniques. The wavelet transforms decomposes the image into a sparse representation of coefficients [51]. Most of the information contained within the image can be represented by a few wavelet coefficients, while the noise will affect all the coefficients equally. By thresholding the coefficients and setting the smaller values to 0, the noise is limited while preserving the image detail[52][53]. This method can be split into two different types of thresholding, hard thresholding (D^H) and soft thresholding (D^S). Hard thresholding sets any coefficients (d) less than the threshold value (λ) to 0, eliminating them as they are more likely to be noise. this can be written as:

$$D^H(d|\lambda) = \begin{cases} 0, & \text{for } |d| \leq \lambda \\ d, & \text{for } |d| > \lambda \end{cases} \tag{4.6}$$

In the case of soft thresholding, the values less than the threshold are set to 0, while the higher coefficient values are also shrunk by the amount of the threshold. Which results in the equations

$$D^S(d|\lambda) = \begin{cases} 0, & \text{for } |d| \leq \lambda \\ d - \lambda, & \text{for } d > \lambda \\ d + \lambda, & \text{for } d < -\lambda \end{cases} \tag{4.7}$$

Correlation Denoising

Wavelet decomposition transforms the original image into coefficient images for multiple different levels. At each different level, the coefficient images are scaled down by a factor of two. As the noise is independent per pixel, it will not shrink like the desired information will, so will have different wavelet coefficients for the same spatial positions. By looking at the correlation of these corresponding wavelet coefficients between levels, the noise can be separated as the correlation values will be much lower than that of the true signal. By then removing the low correlated wavelet coefficients most of the noise

is removed.

4.3.6 BM3D Denoising

Block matching and 3D filtering (BM3D) [5][54][49] groups similar blocks from within the image that are similar and then performs the filtering on all the blocks at once. This has been shown to perform better than spatial filtering as the group denoising performs a 3D transform and applies threshold denoising similar to wavelet denoising. This group denoising has been shown to be improved by applying collaborative Wiener filter instead of hard thresholding. BM3D is a two-pass method, the first pass creates a basic estimate of the signal, while the second pass creates the final estimate based on the basic estimate. BM3D and its derivatives are often considered state of the art for "traditional" denoising techniques, only potentially surpassed by neural networks.

Basic Estimate

To perform the first pass, and calculate the basic estimate, the images are first transformed into blocks. These similar blocks are then stacked into a 3-dimensional array referred to as a group. 3-dimensional denoising has proven to be effective by utilising a sparse representation in many video denoising algorithms, as well as compression such as MPEG. Block matching is achieved by grouping together blocks with similar information that are not necessarily spatially close within the image. This can be seen in figure 4.3 Due to these blocks being similar, the groups have a high correlation in the direction of the stacking. This can then be exploited by applying a 3D decorrelating unitary transform to produce a sparse representation of the information within the blocks.

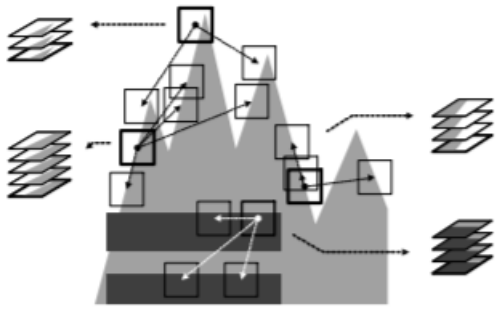

Figure 4.3: Example of block matching technique from [5]

The noise within these blocks will be independent of each other, so when the group is transformed to create the representation coefficients, it will have a high sparsity. That is, there will be a few coefficients that have a high strength and many coefficients

that have low strength as they are mostly the random noise between blocks. This noise can be minimised by applying a hard thresholding technique to the coefficients of the sparse representation. Hard thresholding is achieved by applying a value of which anything under this is considered to be noise and non-essential to reconstructing the main information in the image. By minimising the amount of less common coefficients the noise is reduced drastically.

The values that are considered to be signal, can be used to reconstruct the noise reduced signal by performing the inverse 3D transform. These reconstructed blocks can be replaced into their respective spatial positions. As there is a fair amount of overlap from similar blocks, a weighted average is taken to find the overlapping pixel values, and once that is complete, the output is the basic estimate. A pictorial description can be seen in fig 4.4

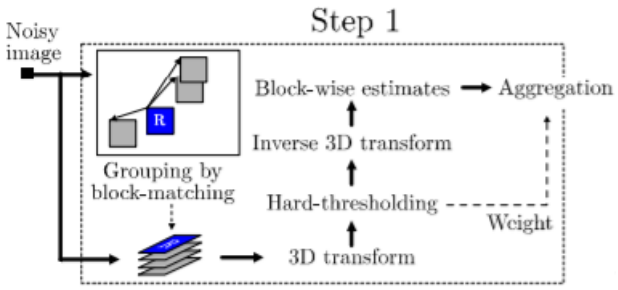

Figure 4.4: pictorial description of BM3D basic estimate, from [5]

Final Estimate

Calculating the final estimate has a lot of the same steps as the basic estimate, except it uses the output from the basic estimation and the noisy image to complete the denoising. The basic estimate is a relatively noise-free image, and so that is used to achieve higher accuracy in the block matching and grouping step. These more accurate groups produce a better, more sparse representation. By performing the block matching on the basic estimate first, the noisy image can have the same spatial positions turned into the same groups [5]. However, once the 3D transform of each group has been performed, instead of denoising via hard thresholding, Wiener filtering is performed to achieve better denoising results. This requires the basic estimate groups to be used as the estimation of the signal, and the noisy groups as the signal to be denoised.

4.4 Results

The effects of noise and denoising techniques were tested on all the registration methods. The standard deviation will show the effects of noise more reliably rather than the mean error values (bias) for the noisy results. This is due to the standard deviation being more stable due to it being much less affected by extreme outliers as seen in fig 4.5.

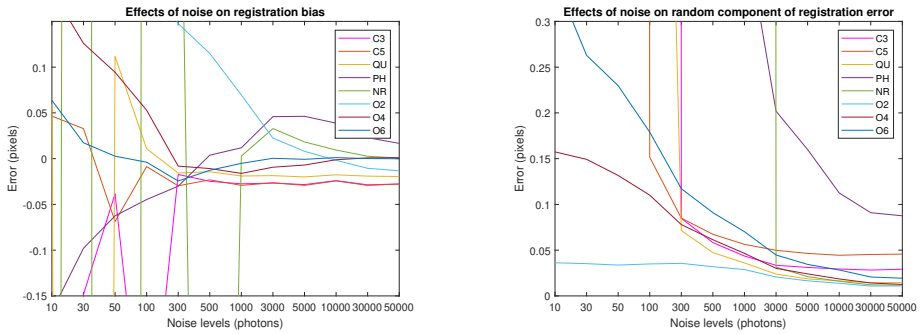

Figure 4.5: Results of registration methods, with shadow images corrupted by Poisson noise. C3 = 3 point curve fitting, C5 = 5 point curve fitting, QU = quadratic surface fitting, PH = Phase registration, NR = Newton-Raphson, O2 = 2×2 optimal interpolation filter, O4 = 4×4 optimal interpolation filter, O6 = 6×6 optimal interpolation filter

The effects of noise on the different registration techniques can be seen in figure 4.5. This shows that the error is much larger for the correlation based methods. This is due to the noise corrupting the correlation peak that is used to estimate the sub-pixel shift, and hence the differential phase images. It gives the correlation map a larger number of potential peaks as seen in figure 4.6 where there are more than 3 potential peaks. This can be seen very clearly with the Newton-Raphson (NR) method, where the error is amplified by the iterative process, and hence performs worse than all the other techniques in the presence of noise, not stabilising for photon counts less than 3000 photons, whereafter the results are very reliable.

The acceptable standard deviation for comparing these methods, and assessing their suitability was 0.1px. This is worse than the 0.05 sub-pixel shift we were testing for with the noise-free data, but the results are also much worse than the noise-free counterparts. A standard deviation of 0.1px ensures that 95% of the data is less than 0.2px from the measured shift. The cut off for mean error has been considered to be 0.05px as that is the shift we were wanting to detect. This is not ideal, but it is unrealistic to expect the results to be as good as noise-free.

The curve fitting (C3,C5) and quadratic surface fitting (QU) methods also use the correlation map to estimate the sub-pixel offset, shown by poor results at low photon

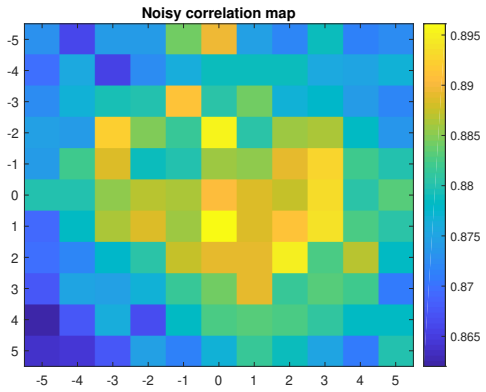

Figure 4.6: Example of noisy correlation map with multiple peaks

counts. As they are not iterative methods, they are not as affected as the Newton-Raphson method, though the results under 300 photons are unreliable due to a large mean error and standard deviation.

As expected the 5-point curve fitting method performs better than the 3-point curve fitting. This is due to the curve always fitting the 3 points perfectly, whereas the 5 point curve becomes a curve of best fit so is less affected by the noise. This is seen in both the mean error and also standard deviation performing better in low photon count images.

When corrupted by noise, the POC phase registration method does not perform as well as expected. The benefits of performing the correlation in the Fourier domain is not only producing an impulse or sinc functions with a much sharper peak for calculating the sub-pixel shift, but it is also resilient to noise. This is due to noise having a higher frequency than the information within the image and can therefore be limited by taking a smaller window within the POC that removes some of the high frequencies which is predominantly noise. It appears that even at high photon counts the error is much higher than the spatial domain methods. However, the results from the ideal case also had a higher random error at approximately 0.1 pixels so it is to be expected that the noisy results are above this. Even allowing for the higher random error, the POC phase registration method still performs poorly with the random error increasing much more rapidly and evenly than the other techniques. This results in the POC phase registration performing the worst except for photon counts less than 3000 when it is better than the Newton-Raphson method, but the error is still too large to be viable.

The optimal interpolation filtering methods with a 2×2 (O2),4×4(O4), and 6×6(O6), all perform much better than the correlation based methods, and the Newton-Raphson method especially for the exceptionally low photon counts (less than 100 photons).

This, however, does not make a difference as the results are still too poor to be utilised. When looking at the standard deviation, the 2×2 optimal filter appears to be amazing, as the values are all well under the 0.1px threshold, however, for many of the lower pixel values, the 2×2 filter has a bias that is far too large, greater than 0.05 until ≈ 3000 photons. The 4×4 optimal filter has a standard deviation less than 0.1, and mean error less than 0.05 at ≈ 200 photons, and ≈ 400 for the 6×6 optimal filter.

4.4.1 Directly Denoising Poisson Noise

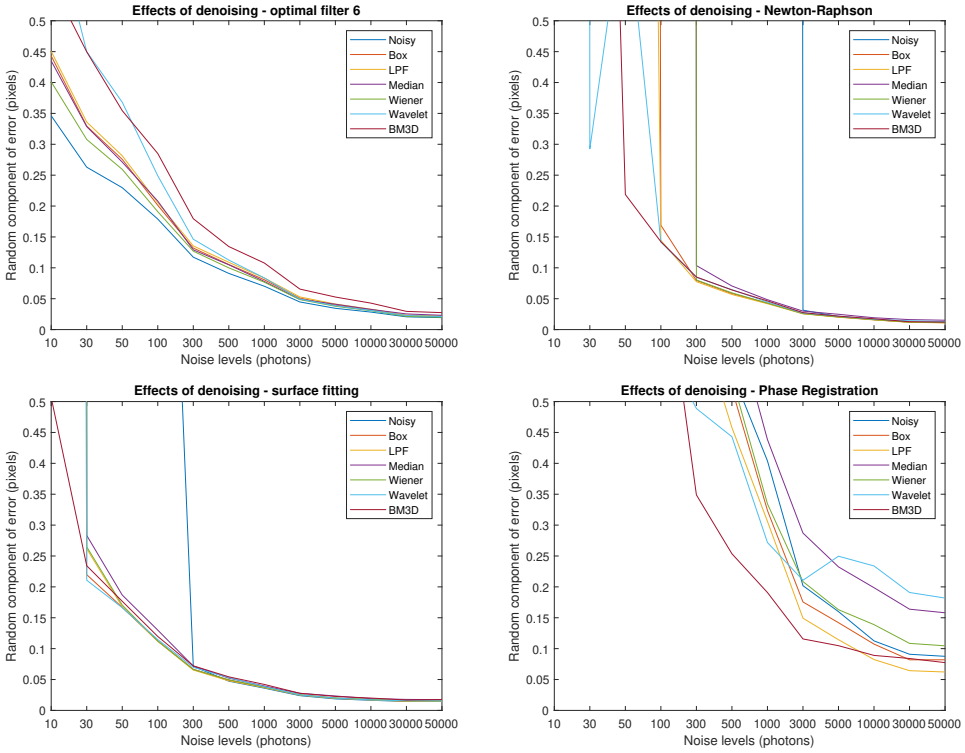

Figure 4.7: Effects of denoising techniques on Poisson noise

The images with photon counts less than 100 photons per pixel are essentially useless as there is not enough to distinguish between the grains and empty space due to the low SNR. This can be seen by the large mean error being often 5-10$\times$ larger than the noise-free errors, and this is further confirmed by the standard deviation being exceptionally large. This is to be expected as the images at that photon count are dominated by noise. This larger standard deviation error is much more problematic than the larger error, this is due to the increased uncertainty for what should be a uniform shift. This can be seen in figure 4.7. Any standard deviation above 0.3, is no longer considered sub-pixel registration as it has a range close to 1 pixel implying it is more of a guess than a calculation. To be considered feasible the cut off for standard deviation was chosen to be 0.1. The registration techniques that are not shown here, can be found in

appendix A.2 and A.3.

One trend that is common amongst all the spatial domain registration methods is that the more complex denoising methods, achieve only minimal improvements over the "traditional filtering" methods. The most complex of these methods tested was the BM3D filtering, which is also considered to be state of the art. In the case of X-ray phase contrast imaging using shadow images, it only has a significant effect on the Newton-Raphson method, it increases the standard deviation of the optimal interpolation filters while having minimal effect on the mean error, and for the surface fitting method, it performs similarly to all the other denoising techniques.

With the synthetic images being made of perfect spheres, the block matching will have the best opportunity to perform as the edges will all be the same for spheres of the same size. Where there is overlap, the edges will differ, but there are multiple edges that will be identical. As this is the basis for the BM3D method, the results are likely to get worse for real data. The BM3D method is computationally expensive when compared to the other filtering methods due to it taking two passes to achieve the final estimate. As the results are not significantly better, it is therefore not worth the extra computation in this application.

As for the other filtering denoising methods (box filtering, low pass filtering, Wiener filtering), they all have a positive effect on the results, but there is no one method that performs significantly better than the others.

A further observation is that denoising on the optimal interpolation filter appears to make the standard deviation worse, and the more complex the denoising technique is, the worse it becomes. This is due to the fact that the registration is already optimised in a least squares sense, so any denoising introduces a blur which has a more negative effect than the noise itself.

The Phase only correlation method is the obvious outlier when looking at the effects of denoising on registration accuracy. The values of the random component of the error are not as low as the spatial domain techniques similar to the results explained in the registration chapter. It is the only registration technique which the results of registration after denoising are not tightly grouped together.

This implies that the choice of denoising techniques has a significant effect on the quality of the registration and hence phase contrast images. This improvement is due to the noise resulting in a lot more high frequencies than the noise-free equivalent, by removing these high frequencies the impulse peak of the cross-spectrum is more significant which results in the sinc function fitting the POC more accurately. The BM3D denoising has the largest effect on the registration error, except for at high

photon counts where the low pass filtering performs slightly better. this makes sense as the BM3D denoising is more efficient at removing noise than many spatial domain techniques [5] and hence performs better for the lower photon counts. The low pass filtering surpasses this at higher photon counts as there is less noise, and it introduces less blur than the box median and Wiener filters.

4.4.2 Denoising Anscombe Transformed Poisson Noise

The Anscombe transform is a variance stabilising transform which transforms the Poisson noise to Gaussian noise. As a lot of the denoising techniques are designed and optimised for Gaussian noise, it is expected that this transformation from Poisson noise to Gaussian noise will improve the effects of most of the denoising methods, specifically the Wiener filter, and the BM3D denoising which uses the Wiener filter. The box filter and wavelet shrinkage will not be as affected. This is due to the box filter just taking the average of the values within the window, as Poisson noise is random around the expected value (evenly higher and lower), it should cancel out as well as Gaussian noise. Wavelet shrinkage is independent of the type of noise introduced, just the combination of wavelets which can accurately reconstruct the image.

To perform this test, the first step was to apply the Anscombe variance stabilising transform, which transforms the noise from Gaussian into Poisson noise. Then the different denoising techniques were applied to copies of the same noise corrupted image so they had the same grain position and noise applied to it. Then to maintain consistency for all the different registration techniques, a window size of 31×31 was used. The grain size follows a normal distribution centered around 10 pixels with a standard deviation of 1 providing a range of approximately 7 to 13. The image size that was used for testing the registration was 500×500 pixels with a uniform shift of 0.25 pixels.

The results in figure 4.8 show minimal to no improvement for any of the standard deviation or means of the different registration techniques. This is more likely due to the limit of Poisson distribution becoming Gaussian, meaning the only noticeable positive effect is in the lower, less usable photon counts (< 300photons). This is most noticeable in the Newton-Raphson method, where the standard deviation values of the low pass and Wiener filters have improved from the last visible value being 100 and 300 photons respectively with no Anscombe transform, to 50 photons and 100 photons with the Anscombe transform. The registration techniques that are not shown here, can be found in appendix A.4 and A.5.

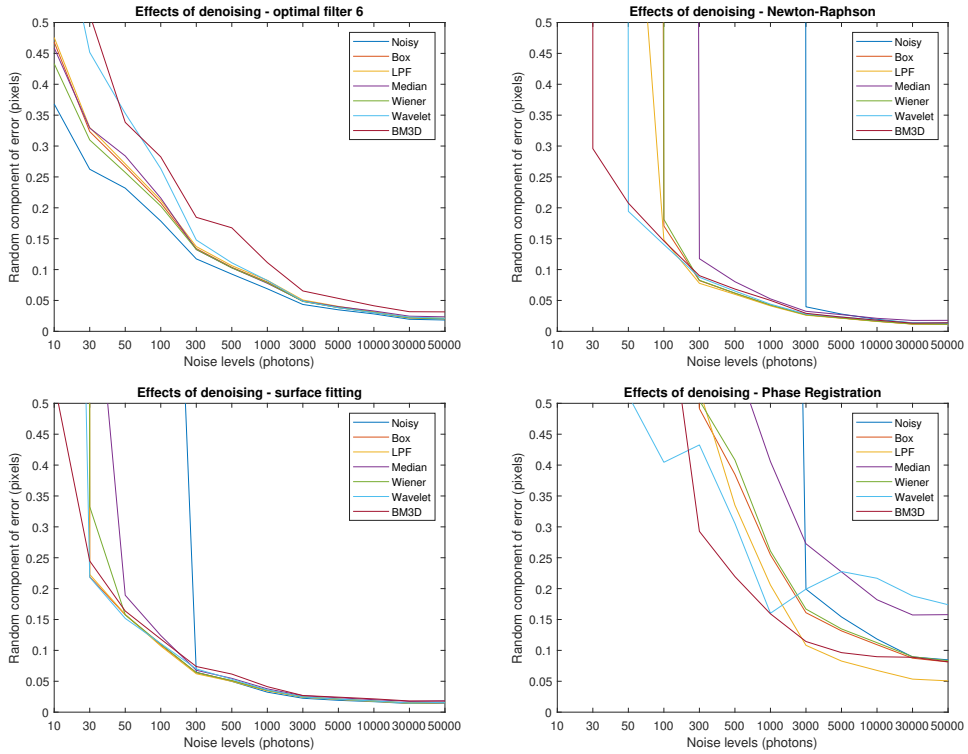

Figure 4.8: Effects of denoising techniques on Anscombe transformed Poisson noise

4.5 Phantom images

A phantom image was generated with a phase shift in the shape of a square based pyramid, as can be seen in figure 4.9. As the pyramid has evenly increasing sides, the differential phase image of the sides will be a constant value. These were used in the noise testing as it is a provides a quantifiable comparison across all noise levels for a "real" object. Each side of the pyramid has a shift of 1.3 pixels in its respective direction. The ideal case is that the x direction detects none of the shift in the y direction as seen below.

To compare the overall quality of this the reconstruction, structural similarity image quality metric [55] was utilised. The structural similarity metric was chosen over the more simple methods such as mean square error (MSE), peak signal to noise ratio (PSNR), as SSIM is a good full-reference image quality metric that better correlates to human perception. This is due to that considers all 3 of the mean or luminance (l), variance or contrast (c), and cross-correlation or structure (s) of the images in calculating the similarity such that :

$$SSIM(x,y) = [l(x,y)]^{\alpha} \cdot [c(x,y)]^{\beta} \cdot [s(x,y)]^{\gamma}. \tag{4.8}$$

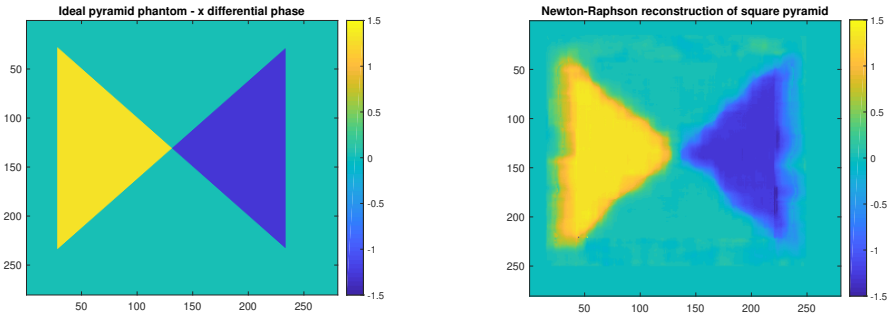

Figure 4.9: Phantom of square based pyramid(left), Noise free reconstruction from optimal filter(right)

Registration method	SSIM
Optimal Filter 6×6	0.4919
Optimal Filter 4×4	0.4865
Optimal Filter 2×2	0.4296
Newton-Raphson	0.5365
Phase only correlation	0.2568
Surface fitting	0.3695
Curve fitting 3	0.2391
Curve fitting 5	0.1725

Table 4.1: SSIM results of noise free registration when compared to ideal phantom

α, β, γ are weights, which can be adjusted depending on what is more important for different functions, for this test the values are all set to 1 so they all have an equal effect on the SSIM value. SSIM was chosen over more complex methods such as the natural image quality valuator (NIQE) [56], and blind/referenceless image spatial quality evaluator (BRISQUE)[57], which are no-reference image quality metrics, that grade quality based on natural statistics of the image.

The SSIM metric gives a value of 1 when the images are identical, and 0 when there is no similarity, and -1 when comparing to the negative of the image. When comparing all the registration methods with noise-free images (table 4.1) to the ideal phantom, we can see that the Newton-Raphson method performs best, followed by the optimal filters 6×6 and 4×4.

4.5.1 Testing Methodology

A single phantom image was used for all the registration techniques so that the slight fluctuations due to the shadow patterns were consistent across the entire range of photons. The same range of photon counts was used as the uniform tests($10 \rightarrow 50,000$), and all the registration methods used the same window size of 31×31 for consistency.

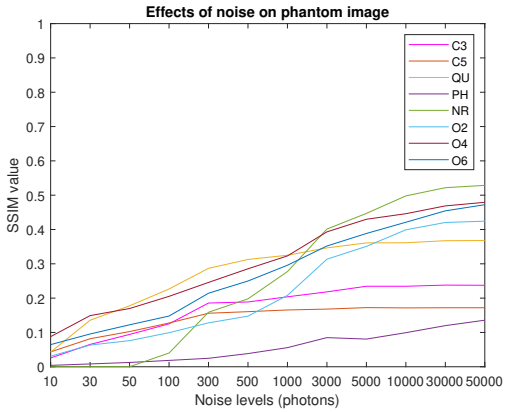

Figure 4.10: Results of SSIM of noisy differential phase images using ideal phantom as a reference.

These registration techniques were compared to the ideal x differential phase.

4.5.2 Results

As we can see in figure 4.10, the SSIM results are rarely over 0.5, and the Newton-Raphson method is the only method to achieve this while the optimal filters 4×4 and 6×6 filters get close at 0.48 and 0.47 respectively. While the surface fitting method (QU) is just above 0.38, and the curve fitting methods are approximately 0.2. The phase only correlation is constantly the lowest due to the significantly higher random component of error in the presence of noise resulting in the differential phase image often having a large error. These values are low, due to the small fluctuations in calculating the differential phase shift for each pixel, the blurring around the edges due to the window size, and detecting small changes due to the Y differential phase shifts as there is a noticeable difference from the 0 shift regions in the background.

For the photon counts in the range of 50 through to 1000, the surface fitting method is most achieves the highest SSIM, after which the Newton-Raphson achieves the highest SSIM. The optimal filter 4 is the most consistent over the full range of photons, due to the Newton-Raphson performing poorly at low photon counts, and the surface fitting method, not achieving a high SSIM value. The phase only correlation and the curve fitting techniques are consistently the worst 3 methods, which are less than 0.3.

This highlights the significant improvement of the more complex methods (optimal interpolation filters and the Newton-Raphson method) and shows how the iterations of the Newton-Raphson method improve upon the correlation map surface fitting for medium to high photon count images. Of the optimal interpolation filters, the 4×4 filter is consistently performing better than the 6×6 and 2×2 counterparts, which has

the advantage of a low bias and random component of error. Compared to the 6×6 has a lower bias at the cost of a higher random component and the 2×2 filter has a lower random component of error, but much larger bias, as seen in figure 2.7.

4.5.3 Conclusions

From looking at the accuracy of the denoised registration methods, it can be concluded that the Anscombe transform is not required as it makes little to no difference to the mean error or standard deviation at the "usable" photon counts. This is due to the fact that at higher photon counts Poisson noise becomes more similar to Gaussian noise. The denoised registration tends not to reach a similar accuracy as the noise-free results until approx 30,000 photons. If we look at where the standard deviation crosses 0.05, the minimum photon count is 500 for the Newton-Raphson method, surface fitting, and curve fitting methods, whereas the 6×6 optimal filter it is 3000 photons, 4×4 filter is 1000 photons, and the 2×2 optimal filter has a standard deviation always under 0.05 for noisy images. This implies the optimal interpolation filters, while performing better for very low photon counts, the other methods are more reliable in the medium, more feasible photon counts.

When looking at both the standard deviation and mean of the optimal interpolation filters, the 4×4 filter shows the most promise as it has both a low mean error and standard deviation, whereas the 6×6 optimal filter has too high of a standard deviation and the mean error of the 2×2 optimal filter is too large making them less than ideal. For all the different filter sizes, the denoising techniques provide minimal to no improvement on the mean, and makes the standard deviations worse than the noisy results. It is therefore detrimental to perform any denoising if the optimal filtering techniques are to be used.

If the Newton-Raphson method is to be used for low power X-ray phase contrast imaging, then it needs to be paired with a strong denoising algorithm, such as BM3D or wavelet denoising. These increase the computation time and complexity. There have been improvements to the Newton-Raphson method, most notably the Inverse Compositional - Gauss Newton (IC-GN) method, which have a higher resilience to noise by a factor of approximately 3 [33], but as it is still iterative and based on a correlation map, the low photon counts will still be very poor. It is assumed the IC-GN method will perform similarly to the Newton-Raphson method, after the denoising has been applied.

The results of the phase only correlation in the presence of noise and Anscombe transformed Gaussian noise perform significantly worse than all the other explored registration techniques. This can be seen by the larger standard deviation. Of the

denoising techniques, the BM3D performs the best for reducing the random component of the error, however it is still not feasible with photon counts under 5000 as the random component of the error is larger than the specified 0.1px. As the POC technique performs poorly both with and without noise, it is becoming less feasible to suggest as a viable sub-pixel registration technique for X-ray phase contrast imaging using shadow based intensity modulation.

This phantom test was performed to explore the effects of noise on a phantom image rather than a uniform shift as it not only explores multiple shifts in the same image, but also the effects of noise on edges which will be within a real phase contrast image. These results show that the optimal filter 4×4 performs the most consistently across the entire range of photons. This is to be expected as it had a low bias and random component of error in the noise free test in chapter 2, however, it performs better than the 6×6 optimal filter in the presence of noise due to the reconstruction filter being smaller and hence less room for noise. At higher photon counts, the Newton-Raphson method performs better due to the better random component of error in the ideal case.

Chapter 5

Effects of Grain and Window Size

As stated in chapter 2, one of the benefits of utilising synthetic images is that we have the ability to explore the effects of window and grain size on the registration quality, and by extension the differential phase and phase contrast imaging. To achieve a good registration quality, the speckle/shadow size wants to be as small as possible [35] while not being so small that it still meets the Nyquist theorem to avoid aliasing [39]. Finding the optimal relationship between the grain and window size, for absorption shadow images, will improve the registration quality, and hence the quality of phase contrast images.

5.1 Relationship Between Grain Size and Window Size

To explore the relationship between the grain area and window size, they were investigated together with a range of 5 through to 35 for the grain size, and the window size used is 11 through to 71 in steps of 2. This was to maintain the current pixel in the centre of the window. The effects of these are explored with two different synthetic image tests, the first test is performed with a uniform shift applied to the entire image, and the second test was a step edge within an objective image with constant phase shift on each side, such that there is a just a line of shifted pixels in the differential phase image. This step edge test is mimicking the regions within the objective image with a constant phase shift (ie: regions within the image of same material and thickness). This results in the differential phase occurring only at the edges. As this will only occupy a small proportion of the window, this will induce some blurring across the windows and hence the larger windows will incur more blurring.

The optimal interpolation filtering registration methods are limited by the size of their filter. The 6×6 filter requires 36 filter coefficients, therefor each line in R_h

requires a 6×6 window. To solve for the filter coefficients, R_h must have at least 36 rows which can be achieved by moving the window in a 6×6 grid, giving a minimum patch size of 11×11 pixels.

5.1.1 Uniform Shift

It is stated in many papers that the grain size needs to be less than that of a window [4][33][12]. However, they never give a distinct ratio or relationship between them. For the uniform shift, it would be expected that a smaller grain results in a higher registration accuracy as it is the grain edges that contains the useful information, therefore more edges should provide more information. As for the window size, in contrast to grain size, a larger window size should be more beneficial for much the same reason, a larger window allows for more grain edges and more information.

To explore the effects of window and grain size, a shift of 0.25 pixels was applied to the entire image (256×256 pixels). This value was chosen as for many techniques at this point the registration methods are performing the worst in terms of bias while for many techniques the random component of error is fairly constant for the entire range, as seen in figure 2.7 and for the sub-pixel registration techniques where it is not the worst, the error is still large enough that there is significant improvements to be had in terms registration bias. As the shift is uniform across the entire image, taking the 2-dimensional mean and standard deviation of the differential phase images will provide an insight into how the number of grain edges affects the quality of registration as a whole. This test was repeated $5\times$ with the average value being the final value. This was then performed for the afore-mentioned range of grain and window size, and by placing the values into a matrix they can be directly compared and identify any obvious trends.

The results shown in figure 5.1 indicate that the error introduced from registration bias is negligible for the Newton-Raphson and the optimal filter methods and the effects of changing the window and or grain size has no effect on the results. However, for the surface fitting, the results are clearly worse when there is a large grain and small window size, shown by the values being $2 - 3\times$ worse than the results when the grain is bigger than the window size. The quality of the registration is therefore given by the precision (standard deviation) which can be seen in figure 5.2. As expected, the random component of the error is lower when there is a large window and smaller grains, as seen by the darker regions in the bottom left corner of the plots. The values of the random component of the phase only correlation is significantly higher than the other techniques.

These results highlight that the window size has a larger effect on reducing the error

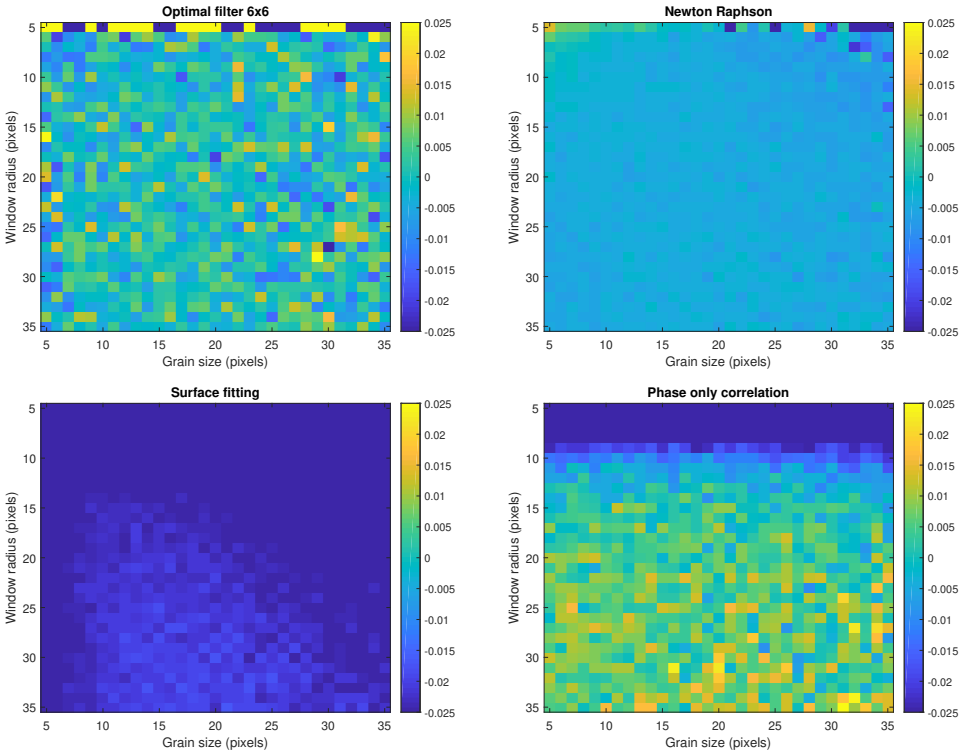

Figure 5.1: Effects of grain and window size (2×window radius +1) on registration accuracy (mean error) for a uniform shift

of the registration than the grain size, as the values change much more rapidly with window size than the grain size. This is to be expected, as the best results contain the most information, both from a large physical window size, and an increase in grain edges.

The plots displaying the standard deviations are clipped marginally higher than the standard deviation of the results in the registration chapter with the grain size of 10, window size of 31, and a 0.25 pixel shift. The results are clipped at 0.1 to allow for the observed higher standard deviation of the phase only correlation. There are values higher than the clipped values shown by the large bright yellow regions. When the window size is small and the grains are large, the registration performance is poor for all the registration techniques, this can be seen as the top right corner of all these images are clipped. This is due to the windows being entirely within the grains, and hence often have minimal or no grain edges which contain the most information, reducing the registration accuracy. When this is the case, the Newton-Raphson method has a random component of error greater than 0.5 pixels. This is due to the iterations magnifying the uncertainties, due to the Newton-Raphson method diverging rather than converging on the estimated shift values. The surface fitting and optimal filter

methods do not perform as poorly as the Newton Raphson method in this case, but they are poor as shown by the standard deviations being significantly larger than the results discussed in the chapter 2 (figure 2.7).

The 6×6 optimal interpolation filter was used here even though the 4×4 appears to have better performance in noise. For the case where the window size is 11,the values are meaningless as they are significantly larger than the rest of the results. This is due to the 6×6 filter failing as there is not enough information to reconstruct the filter, which results in the matrix often becoming singular. The rest of the results in this chapter will use the 4×4 optimal filter as it has a better noise performance and differential phase reconstruction as seen in table 4.1.

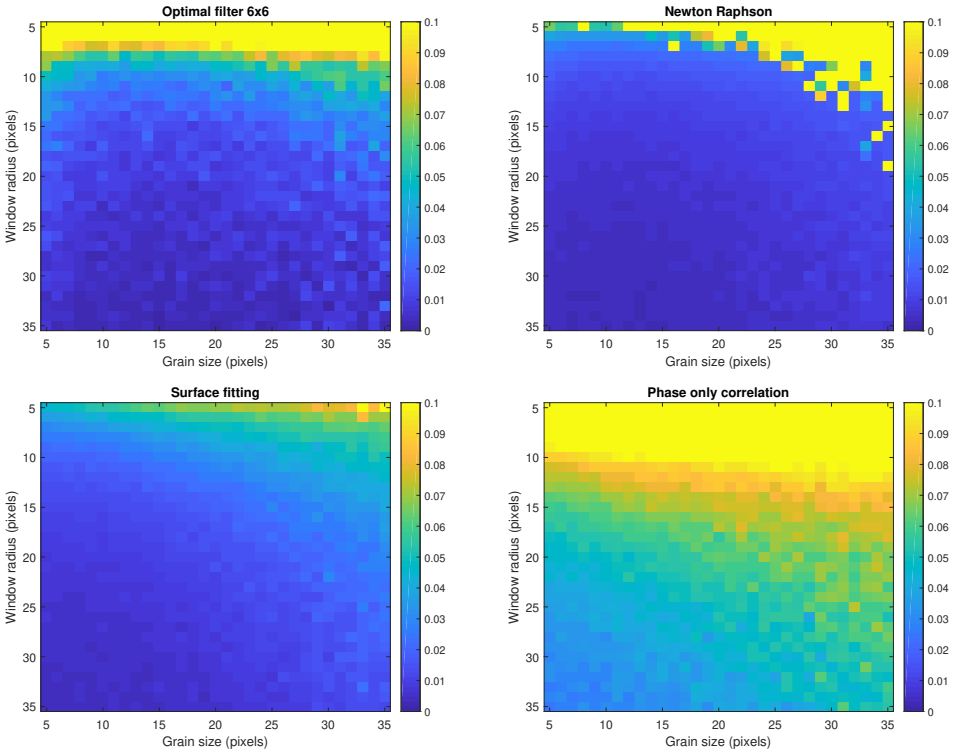

Figure 5.2: Effects of grain and window size (2×window radius +1) on registration precision (standard deviation) for a uniform shift

This test has confirmed the expectations that a larger window size and a small grain size provide the best results in terms of registration error for a uniform shift, or a large similarly shifted region within the objective image, seen by the reduction of the random component of the error. This is due to both of the variables resulting in more information to calculate the sub-pixel displacement more accurately.

This uniform test is not perfect however, as a realistic X-ray image will not be a large uniformly shifted image. Within an X-ray image that is going to be transformed

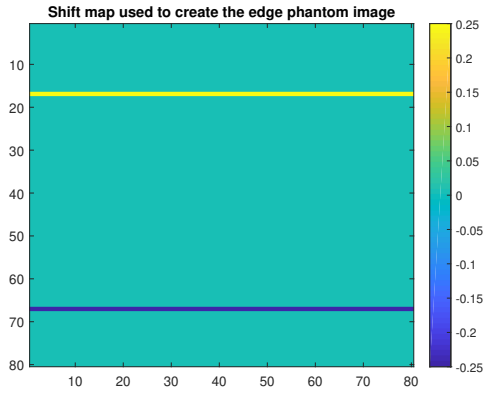

Figure 5.3: Example of ideal y differential phase given the phantom image

to produce a phase contrast image, there will be non-uniform shifts due to different material composition or thickness within the object of interest. These changes result in phase shifts to the X-ray wavefront resulting in a differential phase shift which will be detected as an edge. As many registration techniques utilise localised windows for image registration, and the shifted regions are a small proportion of the window, a larger window will introduce more blurring degrading the differential phase and phase contrast image.

5.1.2 Edge Testing

The goal of this investigation was to explore the effects of the grain and window size for these shadow images containing edges. The setup was relatively simple, instead of shifting the entire image or a large region of pixels within the image, only a single line of pixels are shifted to create a sharp step edge. The same shift value of 0.25 pixels was used for this test as previously mentioned this is where the bias can benefit the most from improvement for many of the registration techniques. It also maintains consistency between tests as the uniform shift test also used a 0.25 pixel shift. This test is important as it shows how the edges will be blurred due to the window size, which may change for the different registration techniques.

The ideal y differential phase can be seen in figure 5.3, it has a background of 0 and the applied shift of 0.25 pixels for a single line of pixels. When the registration techniques are tested with this phantom, the phase only correlation, Newton-Raphson, and all the optimal filters are able to detect a shift even with worst-case where the grains are larger than the windows. This detected shift can be seen in figure 5.4 as the blurred region centered around the correct position. However, when it comes to the surface and curve fitting, it appears, visually, the shift has not been detected. The

images displayed here are using a window size of 15 and a grain size of 5 though the trend is consistent across the full range of window and grain sizes.

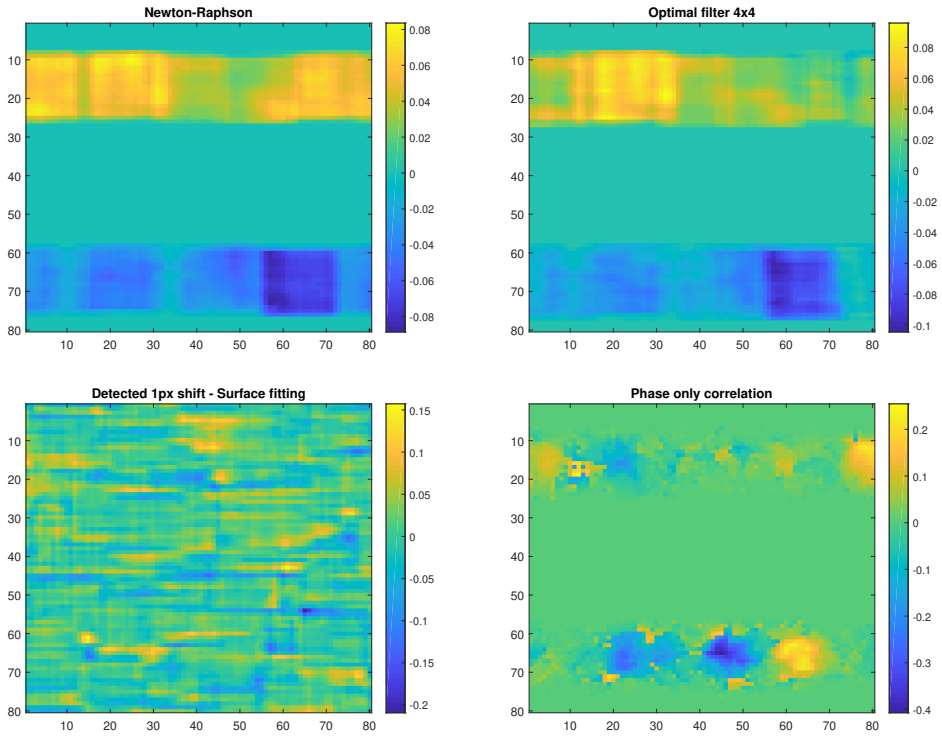

Figure 5.4: Y differential phase images of registration techniques detecting a sharp edge. The examples shown here have a grain size of 5 and window size of 15.

For the registration techniques that can detect a shift, the detected shift values will not be accurate to the same value, but instead a smaller percentage of the shift value. This is due to the shifted values being only a small portion of the window. The fluctuations in the detected shifted region are due to the amount of grain edges within the window. When looking at the darker regions of the differential phase, there is a much lower grain density, and hence fewer grain edges for accurate registration, whereas the brighter region has a higher grain density, and a more accurate registration. The range of shifted values for the surface fitting and the phase only correlation also highlights the detriment of the higher standard deviation, especially compared to the Newton-Raphson method for the window size of 15 in figure 5.4.

To explore this quantitatively, the y differential phase is compared with the ideal phantom using the structural similarity (SSIM) image quality metric rather than mean square error. As mentioned in chapter 4, the SSIM quality metric was utilised due to it being a high fidelity image quality metric that considers not only luminance but also the contrast and cross-correlation to better match human perception[55].

This edge test was completed primarily to compare the effects of changing both the window and grain size on the quality of reconstruction rather than just registration error, but also has the benefit of highlighting techniques that fail to detect a small shift rather than an entire image. The grain sizes explored was 5 through to 25, and the window sizes ranged from 11 through to 51 in steps of 2.

This is a smaller range than the uniform test due to the large grains being detrimental so there was no benefit to exploring the grain sizes much larger. While the window size shows some incremental improvement the larger it gets, a window size of 70 is a large proportion of the full X-ray image, and if it is that large, it will result in fine details being lost in the window. It can be seen in figure 5.5, that it is again the window size that has the main effect on the results as the SSIM is highest when the window is the smallest while being fairly consistent as the grain size increases. The SSIM values of the surface fitting are, not only significantly worse than the Newton-Raphson and the optimal filtering methods due to the larger standard deviation for the background or consistent regions but also shows the opposite trend.

The reducing SSIM values, of the Newton-Raphson and optimal filters, as the window size increases is due to the blurring of the detected shift, and the shift value being less similar to the single row shift of 0.25 pixels. This is due to the percentage of shifted values within the window being smaller, resulting in the detected region being larger, and the calculated shift values being lower than the true values as the shifted region within the local window is smaller. With the surface fitting method the SSIM values increase with window size as it is not detecting the shift, because of this, the SSIM values get higher when the window size is larger as the registration gets more accurate and hence the background values become more accurate, increasing the SSIM values, much like the results of the uniform test above. The curve and surface fitting techniques do not work for detecting edges, this is seen in figure 5.4, and confirmed by the poor SSIM values, and the effects of window and grain size not following the same, expected, trend as the registration techniques that successfully detect edges.

The structural similarity metric turned out to not be ideal for this test, as it can be skewed by the amount of background information. This is due to the background being consistent and centered around 0, while only a single line of pixels is shifted. As the luminance contrast and structure all had equal weights, the larger and more constant background (luminance) region increases the SSIM value much more than the blur (structure) decreases the SSIM value. This is why the SSIM values are mostly between 0.85 and 0.7, a small grouping considering the difference in blur sizes. This does not mean the results are useless as they are all collected with the same phantom images, the results can still be directly compared to each other. However, the SSIM

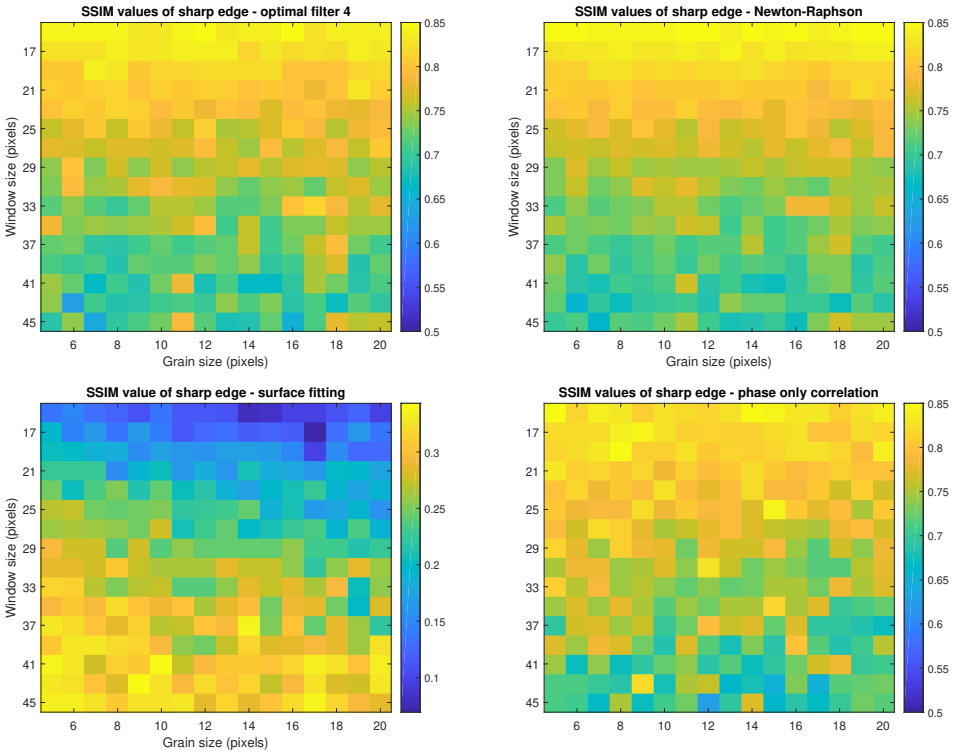

Figure 5.5: Structural similarity metric of registering detecting a sharp edge, with changing window and grain size

image quality metric may be better suited for more complex differential phase images with many edges, or with the phase contrast images where the background is a smaller proportion of the image.

Results of phase contrast image

Detecting the differential phase is only half of the phase contrast imaging modality. The phase contrast image is the integration of both x and y differential phase, so looking at only one of the differential phase images is incomplete for phase contrast imaging. The example images (fig 5.6) had a shift of 0.25 and -0.25 pixels for one line each at 31 and 81 pixels, the grain size was 5 and a window size was 21.

This gives an interesting insight as the surface fitting method, even though figure 5.4 shows the sharp edge is not visually detected, there is a phase shift that has been detected. This can be seen in figure 5.6 as the phase contrast image has a brighter middle region than the background regions, which is where the phase shift would be. It is not as clear or consistent as the other registration techniques, it appears a lot more broken due to the higher standard deviation of the surface fitting methods. This is the same for the phase only correlation as well, the phase contrast image values

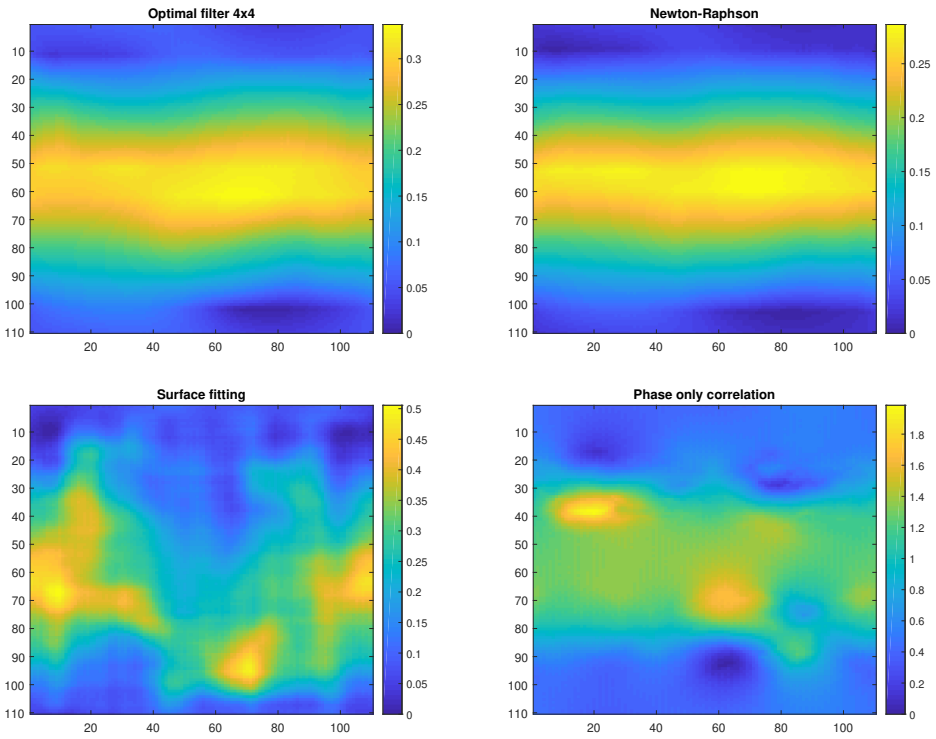

Figure 5.6: Examples of phase contrast images with a sharp edge. The displayed images have a grain size of 5 and window size of 21

are significantly larger due to the larger standard deviation of the differential phase images. Though the differential phase images have shift values that are a proportion of the window size, the phase contrast image values are close to the expected shift due to the integration.

It can be observed that window size affects physical size of the final detected phase shift in the phase contrast image as seen in figure 5.6 as the values between rows 31 and 81 should all have a shift of 0.25. However, the methods that are closest to the real value plateau are between $\approx$ 45 and 70. The larger window results in a significantly smaller phase shifted region as the edges of the blurred regions are closer together, and the blurred regions in the differential phase image creates a ramp up to the detected shift. However, it also results in the shift values having a lower random component of error due to more information within the window, which results in the shifted region being calculated more accurately, increasing the SSIM values. This can be seen in figure 5.8.

The SSIM values of the sharp edge phase contrast image are shown in figure 5.8. In contrast to the differential phase images the results show that a larger window size is more beneficial for the final phase contrast images. This is due to the calculated

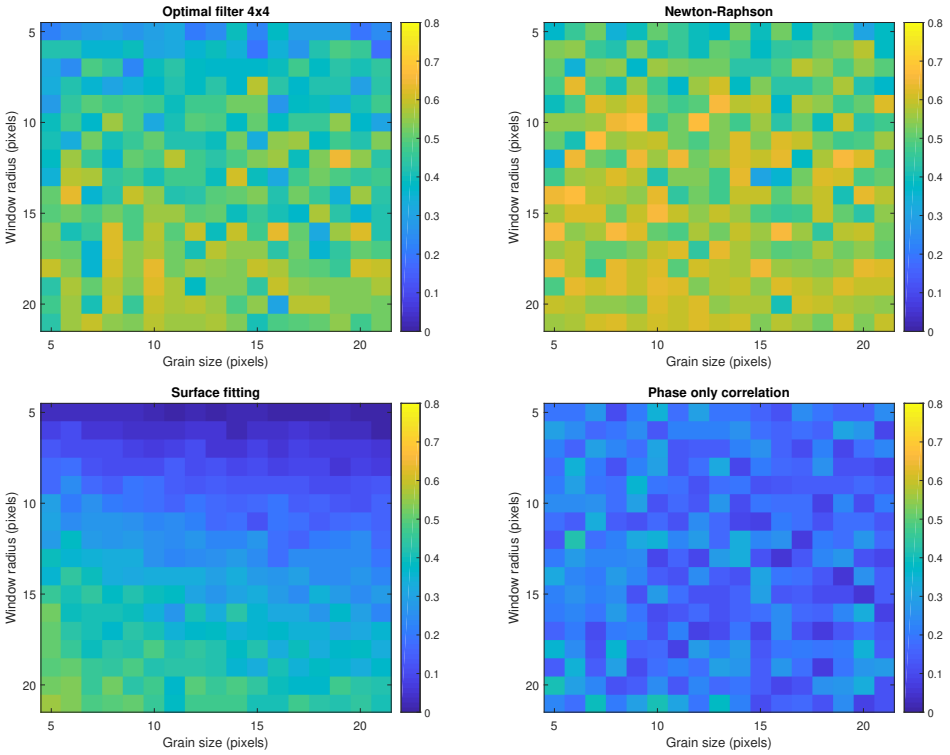

Figure 5.7: SSIM values of phase contrast image containing a sharp edge, with changing window and grain size

phase shift values being more similar to the generated phase shift. This implies that the blurring introduced from the window size is not the most important in terms of phase contrast imaging, as long as the shift is detected.

This test was repeated with a softer edge of 3 lines of shifted pixels instead of the sharp edge with 1 line of shifted pixels. This is a possibility due to the magnification of the image, and the small pixel size on the sensor. This increased the SSIM values imply an increase in quality of the phase contrast image, this is due to the larger proportion of the window containing the shifted pixels, which increases the quality of the registration and hence the differential phase and phase contrast images.

The phase registration method had the highest standard deviation in the registration test, and it is apparent here as well. This results in the differential phase values often being much higher than the shift, producing phase contrast images with significantly larger values than the expected shift. This decreases the quality of the phase contrast image in terms of structural similarity as there appears to be multiple different phase shifts where there is only one. It is worth noting that the surface fitting method detects the shift much more clearly with a softer edge, producing better results. However, the larger standard deviation (when comparing to the better performing Newton-Raphson

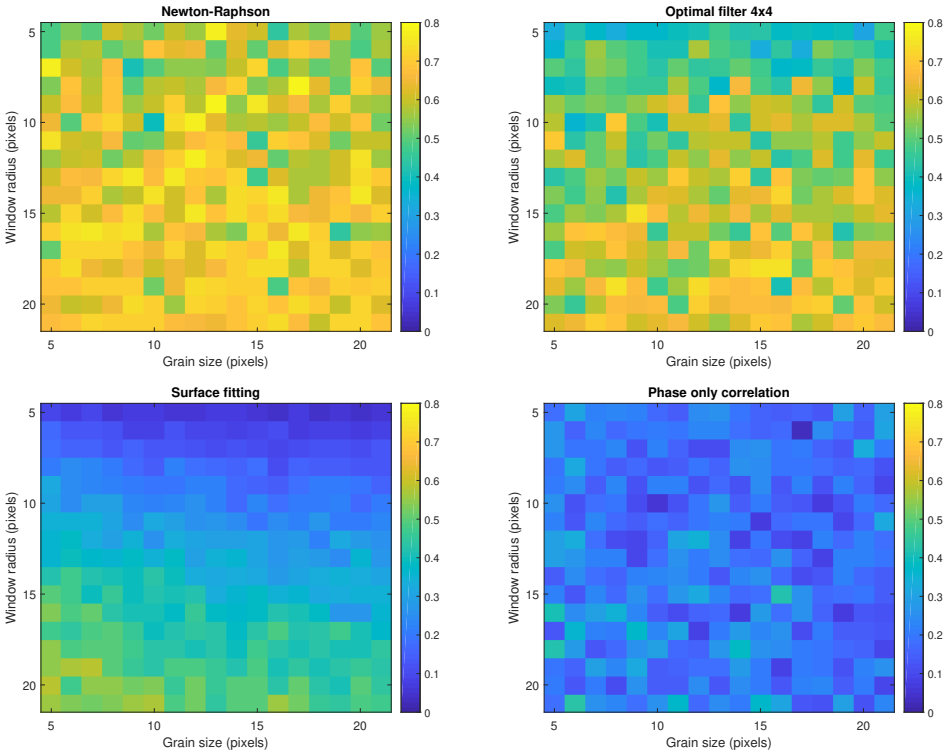

Figure 5.8: SSIM results of blurred edge phase contrast image, with changing window and grain size

and optimal filtering methods) results in the phase contrast image being less smooth than the better performing techniques, hence the SSIM values are lower.

5.2 Conclusions

These tests have given us a good insight into how the size of the window and grains within the shadow images affect the quality of the registration techniques which correlate to the differential phase which affect the quality of the final phase contrast images.

For a uniform shift, the registration techniques all perform the best when the grain size is small and the window size is largest. This was to be expected as smaller grains result in more grain edges within a window, and a large window increases the number of grain edges even more. This results in there being the most information to accurately calculate the offset.

The edge test is more realistic than a uniform shift for a phase contrast image as the object of interest will introduce shifts based on its composition. This involves using the structural similarity image quality metric to compare a differential phase image that is only a single line with the shift map.

This test showed that the surface fitting technique is not capable of clearly detecting a shift with a sharp edge. This implies the surface and curve fitting techniques will fail to detect some, if not all of the fine details within a phase contrast image, limiting their feasibility. However, the optimal interpolation filters, Newton-Raphson, and phase only correlation techniques detect the shift, however the values are not accurate to the expected shift and are instead a proportion of the shift based on the window size. For this edge test, it appears that the window size wants to be as small as possible to achieve the best SSIM results for the differential phase image. This is contrasting to the uniform shift, but makes sense as the blurring of a shift across the window is going to make the differential phase image less similar.

When this is extended to explore the effects of window and grain on the phase contrast image, the results again indicate that a larger window provides better results than a smaller window. This is primarily due to the decrease in random component which results in the detected phase shift being much more consistent, resulting in a more smooth, more accurate phase contrast image after integration. It also shows that softer edges (larger than 1 pixel) provide much better results due to the shift being a larger proportion of the window. Due to the magnification of the object of interest in the phase contrast imagine set-up, and the small pixel sizes ($\approx 100\mu m$) the edges are likely to span multiple pixels similar to the soft edge test.

This test could be continued with multiple, potentially overlapping shifts in the same phase contrast image to better recreate a real phase contrast image which might perform differently. This test also excluded the possibility of the window size being large enough to contain both the positive and negative shift, which may occur in fine details within a realistic phase contrast image.

Chapter 6

Summary and Conclusions

6.1 Summary

This thesis has explored the feasibility of several different image registration techniques for X-ray phase contrast imaging using shadow-based intensity modulation.

X-ray phase contrast imaging is an imaging modality that increases the contrast between materials with similar X-ray attenuation by looking at their phase shift instead. In a medical sense, this can be used to inspect soft tissues, like a CT scan, rather than just seeing the bones like a regular X-ray image. This phase shift cannot be measured directly as the sensors are only able to measure the intensity of the X-rays. To counter this, a pattern can be used, and the translational shift is related to the refractive index of the different materials and hence the differential phase.

Currently the two techniques for generating phase contrast images include monochromatic and coherent light, generated from a synchrotron, and utilising a gold diffraction grating. Both of these techniques have their own flaws, there are only ≈60 synchrotrons worldwide, making them expensive and not feasible for every day medical applications, and the gold grating is expensive due to the precision engineering required the generate the 2 micrometer grating period.

X-ray phase contrast imaging using shadow-based intensity modulation is a newly researched modality that aims to make X-ray phase contrast imaging more widely available. It replaces the expensive gold grating with a sheet of sandpaper or something that generates a random speckle-like pattern due to absorption. As the X-rays travel through the object of interest and the X-ray wave front is refracted, and this results in the attenuation pattern being distorted in the form of a translational shift. This pattern shift can then be tracked via image registration to estimate the differential phase image.

The registration techniques explored in this thesis range from curve fitting to the correlation map, Fourier based image registration, Newton-Raphson based methods, and finally the optimal interpolation filter. These registration techniques are tested with synthetic images modelling a tungsten powder random intensity pattern. These synthetic images have the benefit of being able to apply a known shift which is very hard achieve with real images. This allows for quantitative assessment, allowing them to be compared directly with each other using a pair of shadow based intensity modulation images.

Due to this project being based around low-dosage X-ray phase contrast imaging, these registration techniques were then compared in the presence of Poisson noise, as it will dominate the low light images. To try and improve this, several denoising techniques were applied and an investigation of the benefits of a variance stabilising transform was explored. The denoising techniques applied include a simple box filter, a weighted low pass filter, a median filter, the Wiener filter, wavelet denoising, and block matching and 3-dimensional filtering.

To improve the quality of the differential phase and hence the phase contrast image, it is important to utilise an optimal ratio between grain and window size. These both have an effect on the quality of the registration, as they limit the amount of useful information (grain edges) are used to calculate the phase shift. These features were explored together as they have similar effects, affecting the registration based on the amount of information within a window.

6.2 Conclusions

6.2.1 Registration accuracy

The initial registration test was performed on a uniform shift (the entire image shifted by the same amount) to ensure the registration techniques were all suitable with shadow-based intensity modulation images, and within windows rather than the entire image. The results showed that the random component of error is often higher than the registration bias and is therefore what was primarily used to compare the registration techniques and their feasibility for X-ray phase contrast imaging using shadow based intensity modulation. The initial uniform test showed that the most viable sub-pixel registration techniques are the 4×4 and 6×6 optimal interpolation filters as well as the Newton-Raphson method. These all have a low bias and random component of error implying that they measure the shift accurately, and they are consistent across the entire range of shifts.

As for the other techniques, the surface fitting technique performs the best out

of the curve fitting methods, lower in both the bias and random component, but is noticeably worse than the optimal filters and Newton-Raphson method. The curve fitting techniques with 3 and 5 points perform poorly due to the high bias and random component of error, with no benefit to using them when surface fitting method has the same underlying method but performs better. The phase registration using the phase only correlation performs the worst of the registration techniques as the peak is a sinc function which results in a low bias, but it has a high random component of error due to small deviations of the neighbouring pixels having a larger effect on the estimated shift.

Though there are some techniques that are better than others, all these registration techniques are capable of detecting a shift using shadow-based intensity modulation images.

6.2.2 Effects of noise

The Newton-Raphson, curve and surface fitting methods become unreliable in the presence of noise. This is due to the noise corrupting the correlation map, which often results in a correlation map with multiple peaks. When a noise peak has a higher correlation than the actual peak, the sub-pixel shift error is significantly larger. The Newton-Raphson is the most affected by this as the larger error is often then amplified by the iterative process, as the results do not converge. This results in an exceptionally large bias and random component of error. However, the results improve for all these techniques when the images undergo denoising. Of the denoising techniques, it can be seen that the more complex denoising techniques such as BM3D and wavelet threshold denoising are required for exceptionally low photon counts. However, for reliable differential phase images the standard deviation of the results should be less than 0.1 pixels which is approximately 300 photons. At this point all of the denoising techniques are essentially performing the same. This implies there is no reason to be implementing the more complex filters when a simple box filter achieves the same results.

The optimal interpolation filtering methods do not benefit from the images being denoised, in fact it reduces the quality of the registration as the filter is already optimised in a least squares sense. This results in the noise fluctuations having less of an effect on the registration results than the blurring introduced during the denoising process. This can be seen by the increase in random component of error in section 4.

The phase only correlation technique already had the worst accuracy of registration shown by the large random component of the error in chapter 2, and this only gets worse in the presence of Poisson noise. The random component of the POC registration is twice as bad as any of the other techniques, it appears to less feasible. It is the only

registration technique that benefits from a specific denoising technique, which is the BM3D denoising, but even that does not increase the quality enough to improve its feasibility.

These registration techniques were also tested on a phantom image rather than a uniform shift to ensure the reconstruction is accurate. To compare this, the structural similarity image quality metric was used, this compared the calculated differential phase images with the ideal differential phase phantom. This showed that of the optimal filters, the 4×4 optimal filter is the most consistent across the entire range of photon counts, with the only 6×6 optimal filter and Newton-Raphson registration techniques performing better at high photon counts. This also highlights the failure of the curve fitting techniques as the SSIM values are so low. It also shows that the even though the 6×6 optimal filter performs better than the 4×4 optimal filter when noise free, when there is noise introduced into the images, the 4×4 optimal filter performs better due to the smaller number of filter coefficients so it is less affected by noise.

6.2.3 Window and Grain

An experiment was undertaken to explore the effect of window and grain size on the accuracy of these sub-pixel registration techniques. These were explored together as they both affect the amount of information, in the form of grain edges, within the windows. It is obvious from both the uniform test and the edge test that the window size has the largest affect on the registration accuracy. This can be seen when looking at the random component of the error, the error is fairly constant within the same window size while changing the grain size, but maintaining the same grain size and increasing window size clearly reduces the random component of error.

For the uniform shift, the optimal ratio between grain and window size for the Newton-Raphson and surface fitting methods is to have the grain as small as possible with a large window. This is to be expected as it allows or the most information within a window which increases the accuracy of the registration. For the optimal interpolation filters, the registration accuracy is not highest when the grains are smallest, but instead at about half of the window size.

When looking at more realistic image however, there is going to be a number of edges within the images, and it is the detection of these edges that generates the differential phase images and hence the quality of the detection of these edges affects the quality of the phase contrast images. Instead of measuring the registration accuracy directly, the SSIM image quality metric is used. This is due to there being different shifts in the same image and therefore the mean and standard deviation are not fair tests.

The Newton-Raphson, optimal filters, and phase registration techniques perform

differently to the results of the uniform shift when there is an edge involved. The window size is still the main contributor to affecting the quality of the results, however, it is the opposite of the uniform shift as the smaller window size performs better than the larger windows as it produces less blur, and a more accurate shift as the percentage of the shifted values in the windows is higher. This smaller blur and more accurate estimated shift values results in a higher SSIM value.

This edge test highlights the failure of the curve fitting and surface fitting techniques as they are unable to visually detect a shift. This can be seen by the exceptionally low SSIM values, and that they share the same trend as the uniform shift of the larger windows perform better as the registration is more accurate for the background, completely ignoring the edge.

The phase contrast images show that the effects of blurring introduced by the larger window size is not as important as initially thought. This is shown by the results increasing with window size. Though this test should be improved with multiple, potentially overlapping shifts in the same phase contrast image to better recreate a real phase contrast image which might perform differently.

6.2.4 Final Conclusions

Throughout this thesis, there are 2 methods that constantly perform better than the others, and are hence the best choice for X-ray phase contrast imaging using shadow based intensity modulation. These are the 4×4 optimal interpolation filter and the Newton-Raphson method. The phase registration has a low bias but the standard deviation is the highest of all the explored techniques, and it performs poorly in the presence of Poisson noise. There is no benefit to performing 1 dimensional curve fitting over a 2 dimensional surface fit as they have the same underlying algorithm, but the surface fitting performs better. This however is still not the most feasible as it struggles to detect sharp edges in an image and so will lose a lot of fine detail in the phase contrast images.

The Newton-Raphson method has a low bias and high registration accuracy, however due to its iterative nature, if it is going to be used for low photon count imaging then the images need to be denoised or it will not converge. It can be seen in chapter 4 that this can be achieved with any denoising method to achieve similar results. Even with denoising, some outlier minimisation needs to be implemented as well. This is due to the edges of an object, even without noise, often resulting in diverging results with the values being outliers. When integrated the values of the phase contrast image are significantly affected, resulting in the shape being similar to what is expected, but the values are significantly higher, even from just a few outlying pixels.

Of the optimal filtering techniques, the 4×4 appears to be the most consistent, with a filter size large enough to have a reasonably low bias and high registration accuracy, not too strongly affected by noise, and requires a smaller window size than the 6×6 optimal filter. Due to the performance of the 4×4 optimal filter in the presence of noise, and comparable bias and registration accuracy, it is the most consistent of all the explored techniques and hence the best choice, of the explored registration techniques, for X-ray phase contrast imaging using shadow based intensity modulation.

6.3 Future work

This work has the potential to be furthered by exploring the feasibility of neural networks. Neural networks have the potential to achieve state of the art performance for class specific data. This implies that with enough training data, and the right weights, a convolutional neural network may perform better in both the registration[58] and denoising[45] steps for this specific data and application.

These registration techniques need to be tested with some shadow based intensity modulation X-Ray images to ensure that the results are consistent with the synthetic images. This can also be extended to testing with a simple phase contrast image in which the results are known so a comparison can be drawn.

There has been improvements to phase contrast imaging using the gold gratings by applying a phase stepping technique[4]. This involves using multiple images that are shifted slightly to measure the phase differences. This technique may provide some improvement to the phase contrast image rather than relying on the single shot technique.

Appendix A

Supporting Images

A.1 Noise free shift accuracy

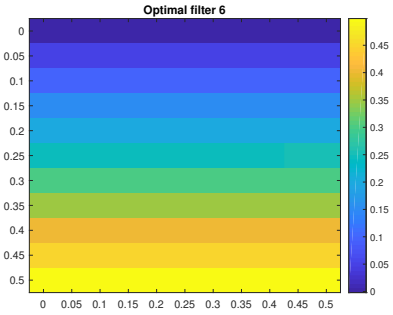

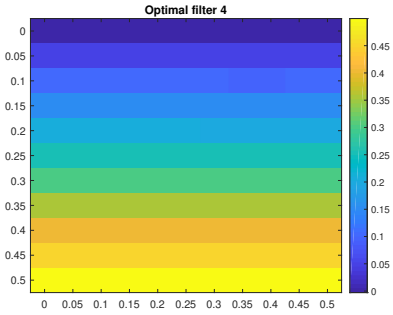

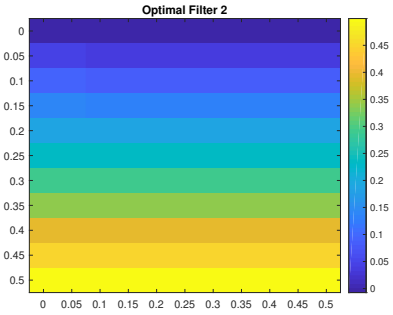

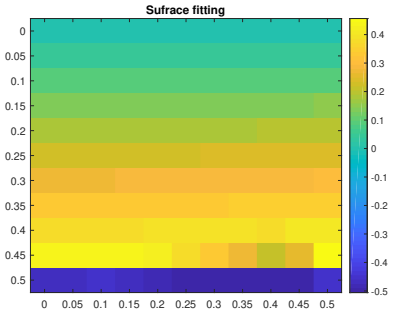

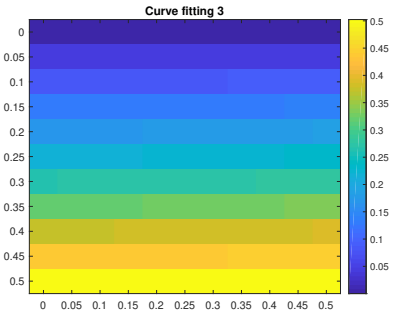

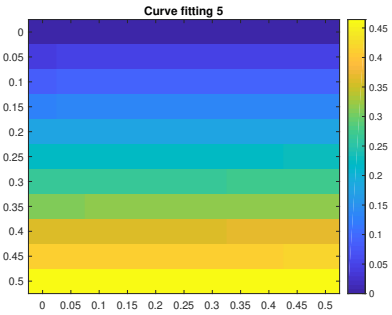

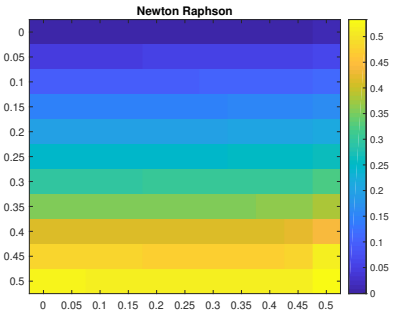

A.2 Denoised - Poisson optimal filter

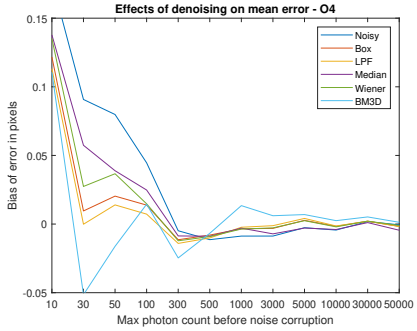

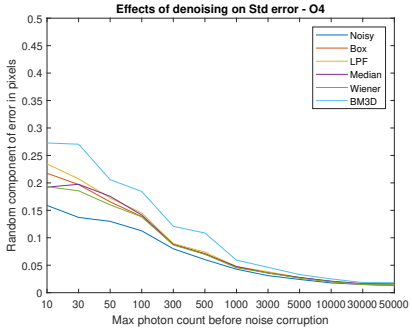

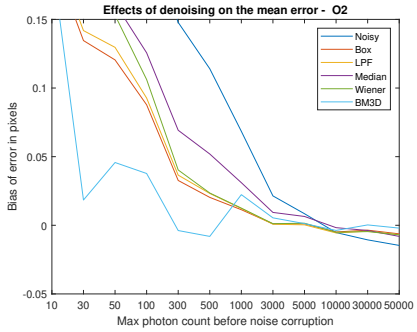

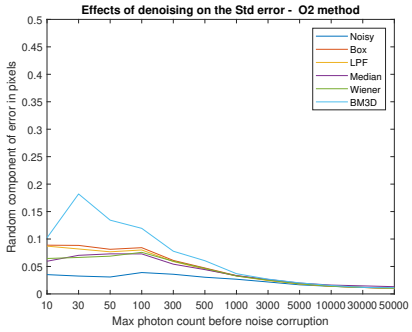

A.3 Denoised - Poisson Curve fitting

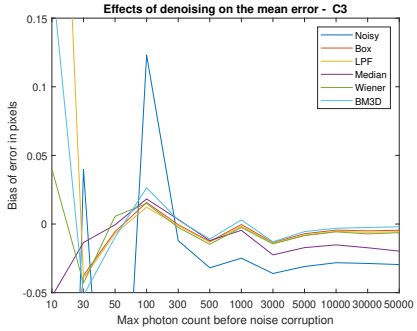

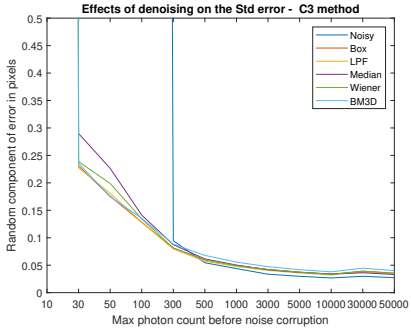

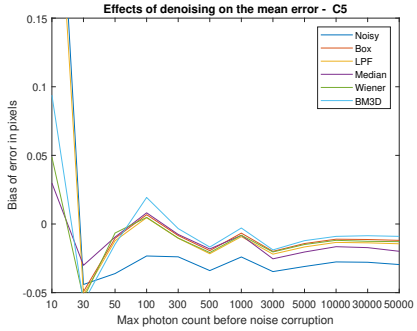

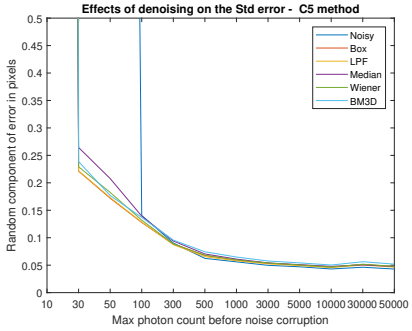

A.4 Denoised - Anscombe Optimal filter

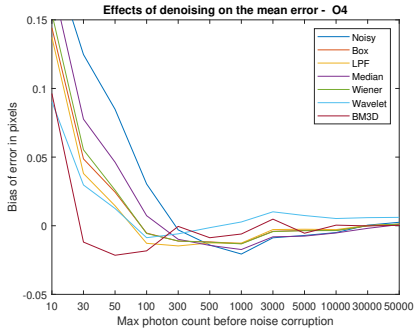

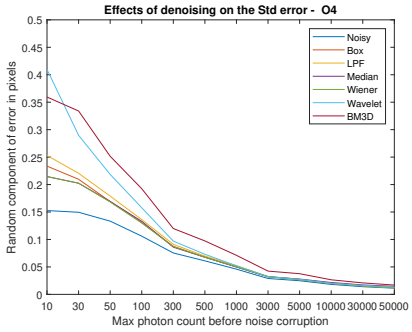

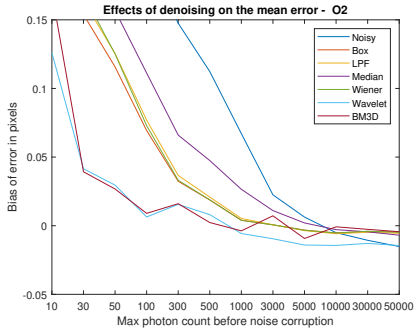

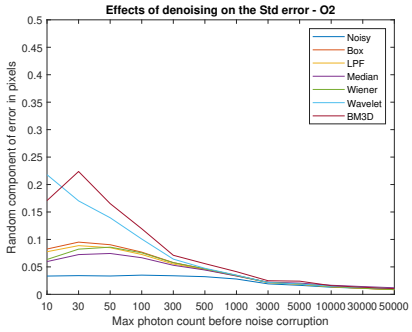

A.5 Denoised - Anscombe Curve fitting

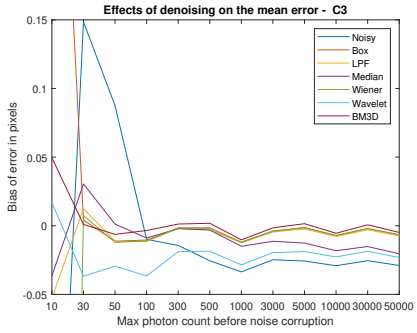

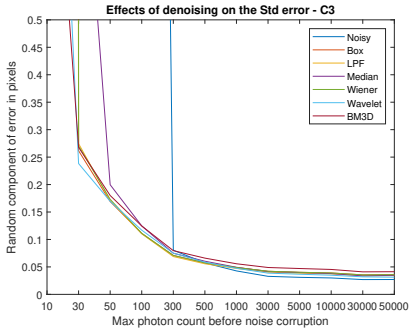

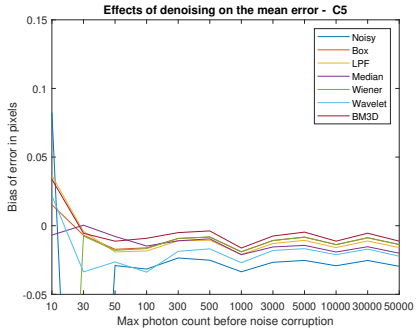

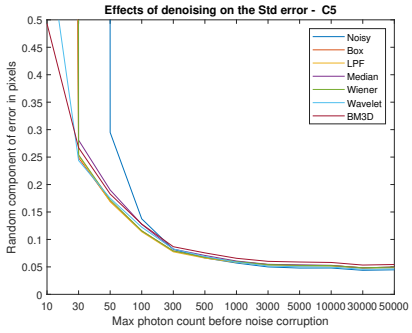

References

[1] C. David, T. Weitkamp, F. Pfeiffer, A. Diaz, J. Bruder, T. Rohbeck, A. Groso, O. Bunk, M. Stampanoni, and P. Cloetens, "Hard x-ray phase imaging and tomography using a grating interferometer," *Spectrochimica Acta Part B: Atomic Spectroscopy*, vol. 62, no. 6-7, pp. 626–630, 2007.

[2] Wikipedia, "Talbot effect.", *Retrieved from:* https://en.wikipedia.org/wiki/Talbot_effect.

[3] Wikipedia, "Moiré pattern.", *Retrieved from:* https://en.wikipedia.org/wiki/Moir%C3%A9_pattern.

[4] M.-C. Zdora, "State of the art of x-ray speckle-based phase-contrast and dark-field imaging," *Journal of Imaging*, vol. 4, no. 5, p. 60, 2018.

[5] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, 2007.

[6] W. Rontgen, "On a new type of rays," *Science*, vol. 3, p. 6, 1896.

[7] P. Zhu, K. Zhang, Z. Wang, Y. Liu, X. Liu, Z. Wu, S. A. McDonald, F. Marone, and M. Stampanoni, "Low-dose, simple, and fast grating-based x-ray phase-contrast imaging," *Proc Natl Acad Sci USA*, vol. 107, no. 31, pp. 13576–81, 2010.

[8] A. Momose, S. Kawamoto, I. Koyama, Y. Hamaishi, K. Takai, and Y. Suzuki, "Demonstration of x-ray Talbot interferometry," *Japanese Journal of Applied Physics*, vol. 42, no. Part 2, No. 7B, pp. L866–L868, 2003.

[9] T. Weitkamp, A. Diaz, C. David, F. Pfeiffer, M. Stampanoni, P. Cloetens, and E. Ziegler, "X-ray phase imaging with a grating interferometer," *Optics Express*, vol. 13, no. 16, pp. 6296–6304, 2005.

[10] H. Itoh, K. Nagai, G. Sato, K. Yamaguchi, T. Nakamura, T. Kondoh, C. Ouchi, T. Teshima, Y. Setomoto, and T. Den, "Two-dimensional grating-based x-ray

phase-contrast imaging using Fourier transform phase retrieval," *Optics Express*, vol. 19, no. 4, pp. 3339–3346, 2011.

[11] A. Momose, W. Yashiro, H. Maikusa, and Y. Takeda, "High-speed x-ray phase imaging and x-ray phase tomography with Talbot interferometer and white synchrotron radiation," *Opt. Express*, vol. 17, pp. 12540–12545, Jul 2009.

[12] B. Zitová and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, 2003.

[13] Z. Wang, H. Kieu, H. Nguyen, and M. Le, "Digital image correlation in experimental mechanics and image registration in computer vision: Similarities, differences and complements," *Optics and Lasers in Engineering*, vol. 65, pp. 18–27, 2015.

[14] S. Dawn, V. Saxena, and B. Sharma, "Remote sensing image registration techniques: A survey," in *International Conference on Image and Signal Processing*, pp. 103–112, Springer, 2010.

[15] S. Wang, D. Quan, X. Liang, M. Ning, Y. Guo, and L. Jiao, "A deep learning framework for remote sensing image registration," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 145, pp. 148–164, 2018.

[16] M. Irani and S. Peleg, "Improving resolution by image registration," *Graphical Models and Image Processing*, vol. 53, pp. 231–239, 1990.

[17] Y. Tian and K.-H. Yap, "Joint image registration and super-resolution from low-resolution images with zooming motion," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 23, no. 7, pp. 1224–1234, 2013.

[18] P. Bing, X. Hui-min, X. Bo-qin, and D. Fu-long, "Performance of sub-pixel registration algorithms in digital image correlation," *Measurement Science and Technology*, vol. 17, no. 6, pp. 1615–1621, 2006.

[19] B. Pan, K. Li, and W. Tong, "Fast, robust and accurate digital image correlation calculation without redundant computations," *Experimental Mechanics*, vol. 53, no. 7, pp. 1277–1289, 2013.

[20] J. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, no. 1, pp. 1 – 36, 1998.

[21] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: a survey," *IEEE Trans Med Imaging*, vol. 32, no. 7, pp. 1153–90, 2013.

[22] H. Wang, Y. Kashyap, and K. Sawhney, "From synchrotron radiation to lab source: Advanced speckle-based x-ray imaging using abrasive paper," *Scientific Reports*, vol. 6, p. 20476, 2016.

[23] D. G. Bailey, "Sub-pixel estimation of local extrema," in *Image and Vision Computing New Zealand*, 2003, pp. 414–419.

[24] A. Alba, J. F. Vigueras-Gomez, E. R. Arce-Santana, and R. M. Aguilar-Ponce, "Phase correlation with sub-pixel accuracy: A comparative study in 1D and 2D," *Computer Vision and Image Understanding*, vol. 137, pp. 76–87, 2015.

[25] E. Baugh and E. F. Talke, "Biquadratic surface fits for two-dimensional interferometric head/tape spacing," *IEEE Transactions on Magnetics*, vol. 30, pp. 4188–4190, 1994.

[26] H. Foroosh, J. B. Zerubia, and M. Berthod, "Extension of phase correlation to subpixel registration," *IEEE Transactions on Image Processing*, vol. 11, pp. 188–200, 2001.

[27] H. Bruck, S. McNeill, M. Sutton, and W. Peters, "Digital image correlation using Newton-Raphson method of partial differential correction," in *Experimental Mechanics*, pp. 261–267, 1989.

[28] N. Kollerstrom, "Thomas Simpson and 'Newton's method of approximation': An enduring myth," *British Journal for the History of Science*, vol. 25, pp. 347–354, 1992.

[29] P. Thevenaz, U. E. Ruttimann, and M. Unser, "A pyramid approach to subpixel registration based on intensity," *IEEE Transactions on Image Processing*, vol. 7, no. 1, pp. 27–41, 1998.

[30] D. G. Bailey, A. Gilman, and R. Browne, "Bias characteristics of bilinear interpolation based registration," in *TENCON 2005 - 2005 IEEE Region 10 Conference*, pp. 1–6, 2005.

[31] D. G. Bailey and A. Gilman, "Bias of higher order predictive interpolation for sub-pixel registration," in *2007 6th International Conference on Information, Communications Signal Processing*, pp. 1–5, 2007.

[32] A. Gilman, *Least-squares Optimal Interpolation for Direct Image Super-resolution.* Phd thesis, School of Engineeing and Advanced Technology, Massey University, 2009.

[33] B. Pan, "Digital image correlation for surface deformation measurement: historical developments, recent advances and future goals," *Measurement Science and Technology*, vol. 29, no. 8, 2018.

[34] P. Reu, "All about speckles: Contrast," *Experimental Techniques*, vol. 39, no. 1, pp. 1–2, 2015.

[35] P. Reu, "All about speckles: Speckle size measurement," *Experimental Techniques*, vol. 38, no. 6, pp. 1 – 2, 2014.

[36] P. Reu, "All about speckles: Speckle density," *Experimental Techniques*, vol. 39, no. 3, pp. 1–2, 2015.

[37] S. C. Burhanettin Koc and K. Uchino, "A piezoelectric motor using two orthogonal bending modes of a hollow cylinder," *IEEE transactions on ultrasonics, ferroelectrics and frequency control*, vol. 49, p. 6, 2002.

[38] National Institute of Standards and Technology, "X-ray mass attenuation coefficients," 2019, *Retrieved from:* https://11bm.xray.aps.anl.gov/absorb/absorb.php.

[39] P. Reu, "All about speckles: Aliasing," *Experimental Techniques*, vol. 38, no. 5, pp. 1 – 3, 2014.

[40] Argonne National Laboratory, "Advanced photon source: Compute x-ray absorption," 2013, *Retrieved from:* https://11bm.xray.aps.anl.gov/absorb/absorb.php.

[41] A. K. Boyat and B. K. Joshi, "A review paper: noise models in digital image processing," *arXiv preprint arXiv:1505.03489*, 2015.

[42] A. C. Bovik, *The Essential Guide to Image Processing.*, vol. 2nd ed. Academic Press, 2009.

[43] S. Kaur, "Noise types and various removal techniques," *International Journal of Advanced Research in Electronics and Communication Engineering (IJARECE)*, vol. 4, no. 2, pp. 226–230, 2015.

[44] W. J. Thompson, "Poisson distributions," *Computing in Science Engineering*, vol. 3, no. 3, pp. 78–82, 2001.

[45] T. Remez, O. Litany, R. Giryes, and A. M. Bronstein, "Deep convolutional denoising of low-light images," *arXiv preprint*, no. 1701.01687, p. 11 pp, 2017.

[46] F. J. Anscombe, "The transformation of Poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3/4, pp. 246–254, 1948.

[47] B. Goyal, A. Dogra, S. Agrawal, B. Sohi, and A. Sharma, "Image denoising review: From classical to state-of-the-art approaches," *Information Fusion*, vol. 55, pp. 220–244, 2020.

[48] N. Ishak, M. J. Gangeh, and R. Logeswaran, "Comparison of denoising techniques applied on low-field MR brain images," in *2008 Fifth International Conference on Computer Graphics, Imaging and Visualisation*, pp. 345–349, IEEE, 2008.

[49] L. Fan, F. Zhang, H. Fan, and C. Zhang, "Brief review of image denoising techniques," *Visual Computing for Industry, Biomedicine, and Art*, vol. 2, no. 1, p. 7, 2019.

[50] J. Mohan, V. Krishnaveni, and Y. Guo, "MRI denoising using nonlocal neutrosophic set approach of Wiener filtering," *Biomedical Signal Processing and Control*, vol. 8, no. 6, pp. 779–791, 2013.

[51] J. L. Song, M. J. Chen, C. Jiang, Y. X. Huang, Q. Liu, Y. Meng, W. Q. Mo, K. F. Dong, and F. Jin, "Research on image denoising method based on wavelet transform," in *2018 37th Chinese Control Conference (CCC)*, pp. 7354–7358, 2018.

[52] X. Qin, Y. Yue, X. Dong, X. Wang, and Z. Tao, "An improved method of image denoising based on wavelet transform," in *2010 International Conference on Computer, Mechatronics, Control and Electronic Engineering*, vol. 5, pp. 167–170, IEEE, 2010.

[53] W.-S. Lu, "Wavelet approaches to still image denoising," in *Conference Record of the Thirty-First Asilomar Conference on Signals, Systems and Computers (Cat. No. 97CB36136)*, vol. 2, pp. 1705–1709, IEEE, 1997.

[54] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "BM3D Image Denoising with Shape-Adaptive Principal Component Analysis," in *SPARS'09 - Signal Processing with Adaptive Sparse Structured Representations* (R. Gribonval, ed.), (Saint Malo, France), Inria Rennes - Bretagne Atlantique, Apr. 2009.

[55] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[56] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a "completely blind" image quality analyzer," *IEEE Signal processing letters*, vol. 20, no. 3, pp. 209–212, 2012.

[57] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Transactions on image processing*, vol. 21, no. 12, pp. 4695–4708, 2012.

[58] S. Guan, C. Meng, Y. Xie, Q. Wang, K. Sun, and T. Wang, "Deformable cardiovascular image registration via multi-channel convolutional neural network," *IEEE Access*, vol. 7, pp. 17524–17534, 2019.