

Modelling the Penumbra in Computed Tomography ^{*}

Audrey Kueh^{†1}, Wilfrid S. Kendall^{‡1}, and Thomas E. Nichols^{§1,2}

¹Statistics Department, University of Warwick

²Warwick Manufacturing Group, University of Warwick

November 6, 2014

1 Introduction

The aim of this paper is to quantify the blurring in the image which is caused by the X-ray source not being a point source. We do this by studying the *penumbra*, which is the blurring specifically at the edges of the object image. If the source were a point, there would be sharp discontinuity at the boundary of the image. However, the data shows that this boundary instead exhibits almost linear decay, indicating that the source is not sharp. Indeed, the steepness of the slope should give an insight into the spot geometry, which is the area from which the X-rays originate.

Our aims are two-fold. Firstly, we aim to quantify blurring if spot geometry is known. This is crucial information when calculating the size of confidence intervals for any measurements. Secondly, we aim to confirm spot geometry from penumbra data. The spot geometry is affected both by measurable factors, such as the electron beam fired at the source, as well as by unmeasurable sources of variation, such as ambient temperature and the pitting of the source. Consequently, the spot geometry may fluctuate from scan to scan. More reliable methods to measure the spot geometry such as the coded aperture mask detailed by Russo and Mettivier[4] can only serve as a guide to the conditions in any subsequent experiments. We seek to be able to confirm spot geometry from penumbra data, and thus to learn the spot geometry for the current experiment.

We therefore examine penumbra from different power settings of the electron beam. Single images of a rotating 2mm cylinder at 16x magnification are taken at each of following 5 different time/power settings: 0.25s/40.6W, 0.5s/22.8W, 0.7s/16.8W, 1.4s/10.4W, 2s/9.2W. Time/power settings are chosen so that the energy outputs from the source are approximately equal. Multiple images are taken to capture variance from arising from unmeasured sources of variation.

Since the size of penumbra is proportional to the spot size, one way of deducing the spot geometry is by measuring penumbra size. However, this is fraught with difficulty because boundary

^{*}This work was funded by EPSRC grant number EP/K031066/1. We would also like to thank Dr John Thornby of WMG for providing us with the data.

[†]a.kueh@warwick.ac.uk

[‡]w.s.kendall@warwick.ac.uk

[§]t.e.nichols@warwick.ac.uk

detection is difficult as exemplified by Figure 1. This is further complicated because the data itself also contains noise from the environment and scatter, as well as measurement errors from the detector. There is no known robust way for defining the endpoints of the penumbra and measuring it. Hence, we have resorted to using a more complicated model, which is discussed in the next section.

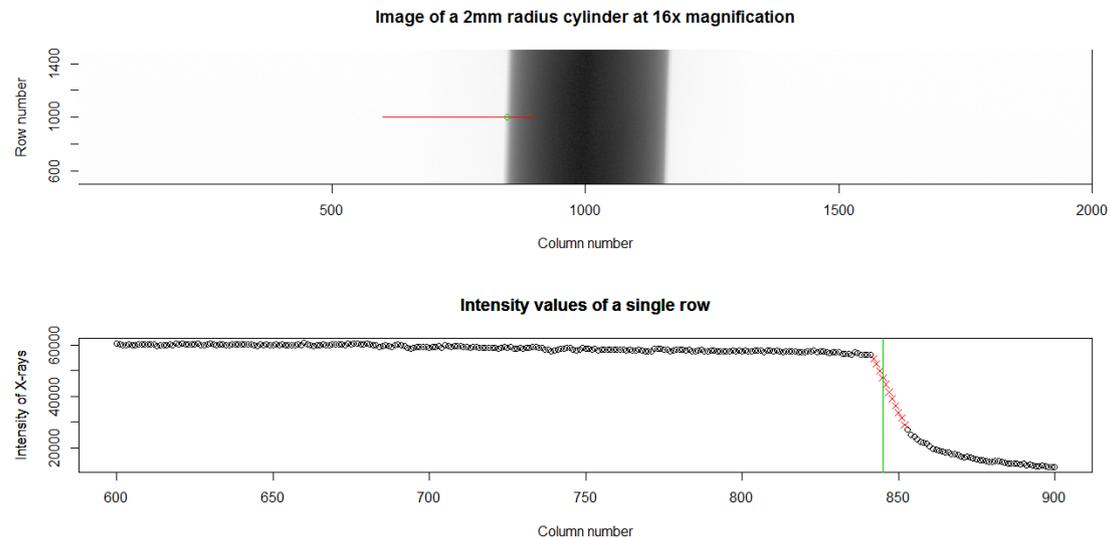


Figure 1: The upper subfigure displays the CT image of a cylinder, where darkness indicates a lower intensity. The horizontal line indicates a row of pixels, for which the image intensity is plotted in the lower subfigure. Later in this paper, we will show how to locate the approximate object boundary and penumbra, indicated by a line and crosses respectively.

2 The Model

Several simplifying assumptions are imposed. It is assumed that the X-ray spectrum is monochromatic, and that there is no absorption in air. Variability due to the detector or scatter is also ignored. Finally, it is assumed that the detector is unsaturated, that is, signal strength of the detector increases linearly with X-ray intensity. This implies that the proportion of X-rays which pass through the object to register on the detector is $\exp(-\mu L)$, where μ is the attenuation coefficient and L is the length of the intercept. Later on, these assumptions will be tested by examining the residuals.

The model is described in two stages. In the first case, we assume that the source is a point. Then, we relax this assumption to allow the source to be a disk, thus providing a closer approximation to reality.

2.1 Single source model

We first assume that the source is located at a point c units away from O , the centre of the cylinder. The radius of the cylinder is r . The origin is taken to be the point at which the ray through the source S and the centre of cylinder meets the detector. The distance between the source and the origin is mc , where m is the magnification level, and the angle between this segment and the detector is θ . In view of the assumptions, the proportion of X-rays from this point source to hit the origin is $\exp(-2\mu r)$. We need to calculate the proportion of X-rays to hit the detector a units from the origin at point K , where a is directed positive in the downward direction (rightward in the real world). In the case of Figure 2, a would be negative.

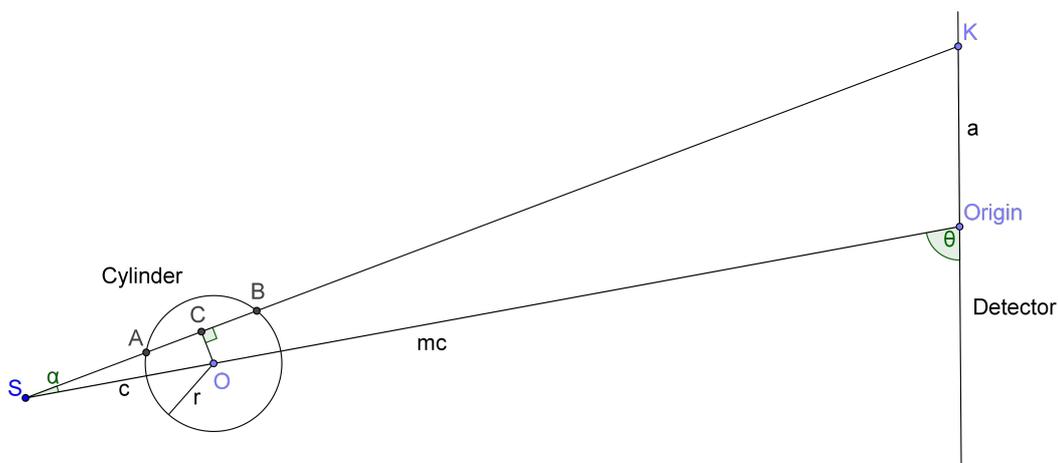


Figure 2: Cross-section of the point spot setup.

To do this, the length of the intercept between K and the source, which in the diagram is the distance \overline{SK} , must first be found. Now, by simple trigonometry,

$$\overline{OC} = c \sin \alpha.$$

By Pythagoras theorem, for $|c \sin \alpha| < r$,

$$\overline{AB} = 2\overline{BC} = 2\sqrt{r^2 - c^2 \sin^2 \alpha}$$

By the cosine rule, bearing in mind the sign convention for the distance a ,

$$\overline{SK}^2 = m^2 c^2 + a^2 - 2amc \cos \theta.$$

Hence, by the sine rule,

$$\sin^2 \alpha = \frac{a^2 \sin^2 \theta}{m^2 c^2 + a^2 - 2amc \cos \theta}.$$

Thus, for $|c \sin \alpha| < r$, the proportion of X-rays to hit the detector at point K is:

$$\exp(-2\mu\overline{AB}) = \exp\left(-2\mu\sqrt{r^2 - \frac{a^2 c^2 \sin^2 \theta}{m^2 c^2 + a^2 - 2amc \cos \theta}}\right). \quad (1)$$

The following plot shows the varying intensity of a row in our spot model when the baseline intensity is fixed at 60000. Notice the two gradient discontinuities and associated penumbra occurring at 160-170 pixels away from the centre of the image.

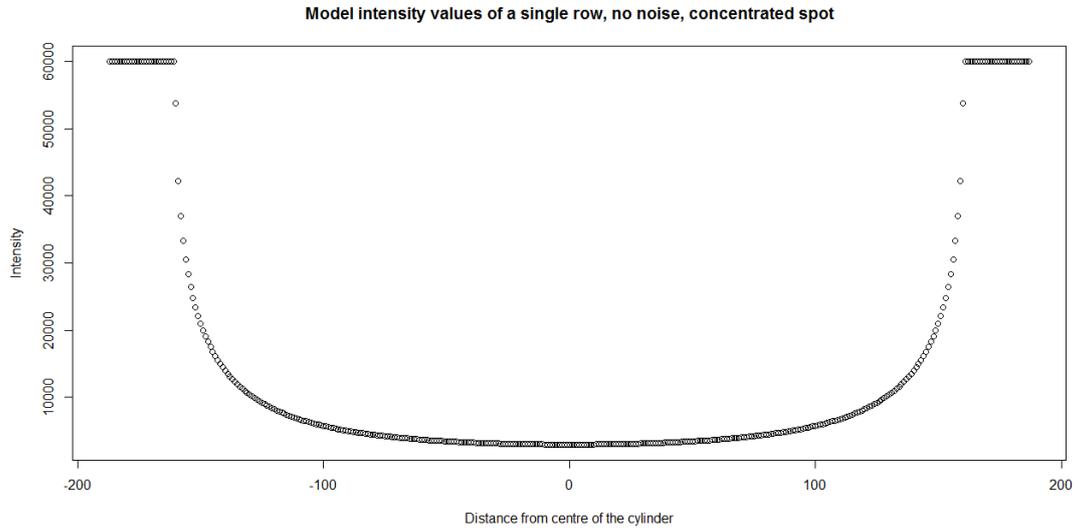


Figure 3: Intensity of a single row of pixels, concentrated spot.

2.2 Disk source model

Now, suppose X-rays originate from a line segment as opposed to a point source. Fix points O_1 and O_2 to be the feet of the perpendiculars from the centre of the cylinder onto the source line and detector respectively, with $\overline{O_1 O_2} = l$. We want to calculate the intensity of a row from a point source at an arbitrary b from O_1 to a point k units away from O_2 , as shown in Figure 4; this will then be integrated over the line segment corresponding to the source.

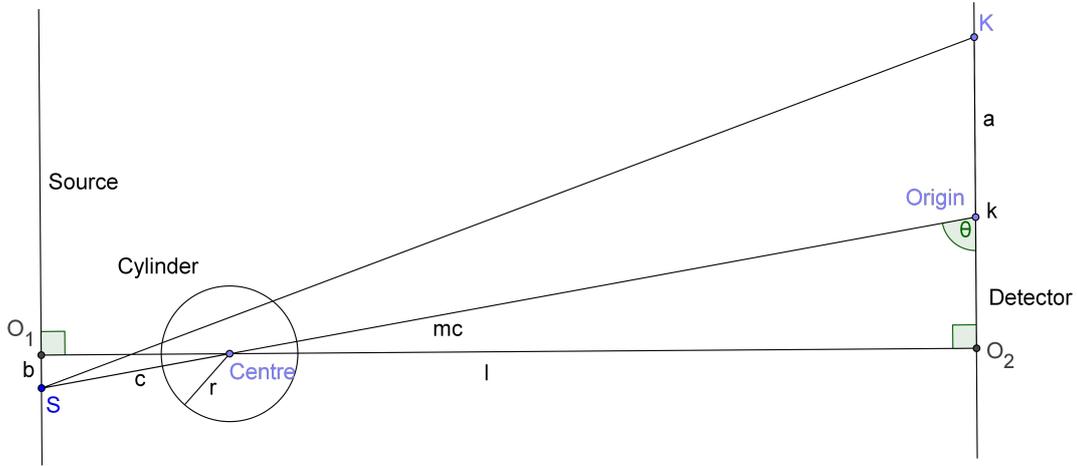


Figure 4: Cross-section of the experimental setup with source represented by a line segment.

By simple triangle geometry,

$$mc = \sqrt{l^2 + (bm)^2}, \cos \theta = \frac{bm}{\sqrt{l^2 + (bm)^2}}, \sin \theta = \frac{l}{\sqrt{l^2 + (bm)^2}}, a = k + b(m - 1).$$

Hence, by substituting the above values into Equation 1, the proportion of X-rays to hit the detector at a point K (located k units from O_2) is:

$$\exp \left(-2\mu \sqrt{r^2 - \frac{(k + b(m - 1))^2 \frac{l^2 + (bm)^2}{m^2} \frac{l^2}{l^2 + (bm)^2}}{l^2 + (bm)^2 + (k + b(m - 1))^2 - 2(k + b(m - 1))\sqrt{l^2 + (bm)^2} \frac{bm}{\sqrt{l^2 + (bm)^2}}}} \right).$$

This simplifies to

$$\exp \left(-2\mu \sqrt{r^2 - \left(\frac{k + b(m - 1)}{m} \right)^2 \frac{l^2}{l^2 + (k - b)^2}} \right).$$

Hence, if the source is a uniform spot of power N with endpoints $x < y$, then the intensity k units from O_2 will be,

$$\frac{N}{y - x} \int_x^y \exp \left(-2\mu \sqrt{r^2 - \left(\frac{k + b(m - 1)}{m} \right)^2 \frac{l^2}{l^2 + (k - b)^2}} \right) db.$$

However, the spot is better modelled by a disk rather than a line. Thus, the disk is projected onto a line segment, where each point of this line segment is weighted by the length of the projected chord. This gives the following expression,

$$\frac{N}{y-x} \int_x^y \frac{4}{\pi} \sqrt{1 - \left(\frac{2b - (x+y)}{y-x}\right)^2} \exp\left(-2\mu \sqrt{r^2 - \left(\frac{k + b(m-1)}{m}\right)^2 \frac{l^2}{l^2 + (k-b)^2}}\right) db. \quad (2)$$

The intensity of a row in our disk model where the baseline intensity is 60000 is plotted in Figure 5 on the next page. Notice the near linearity in the penumbra 150-170 pixels from the centre of the row, which is different from the gradient discontinuity penumbra in the single spot case in Figure 3.

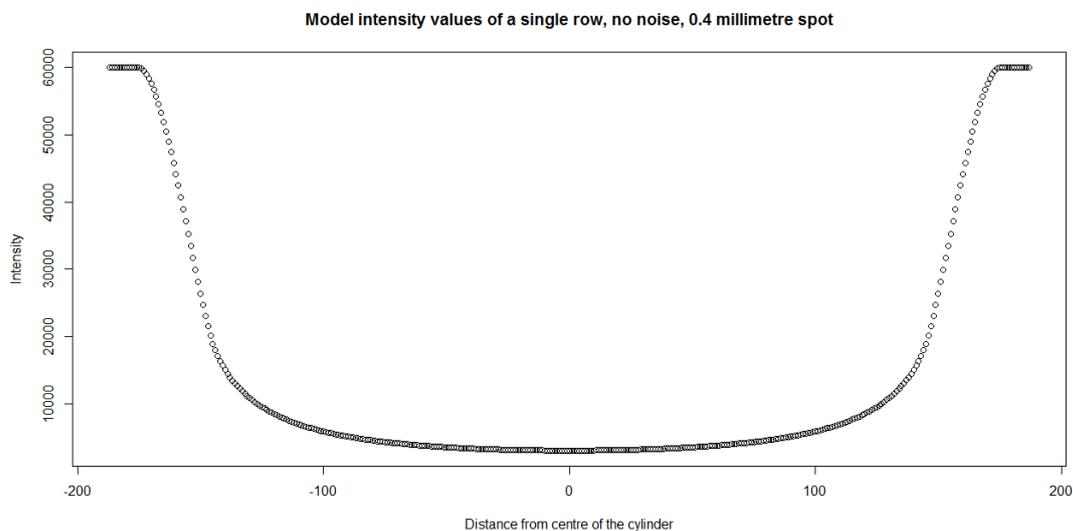


Figure 5: Intensity of a single row of pixels with source represented by an 0.4 millimetre spot disk.

We assume that the object image is centred at O_2 . The experiment is set up to achieve this, and at this preliminary stage we assume this to be adhered within negligible error. We will need to estimate this centre by a median calculation. Then, the unknowns in this model are l , x , y , μ , N . The distance between source and detector l can be measured from the experimental setup. There is information on the baseline intensity N from the intensities at the borders of the images where there is no object between the detector and the source. However, the endpoints of the source x , y , and the attenuation coefficient μ will have to be estimated from the data. This will be the subject of the next section.

3 Model Fitting

We fit the model in three steps: first the baseline intensity N is estimated using information from the borders of the images, then the projections of the centres of the cylinder onto the detector are found by a median calculation, and finally non-linear regression is used to find the endpoints of the source x , y , and the attenuation coefficient μ .

3.1 Finding the baseline ‘empty air’ intensity

There are large variations in the baseline ‘empty air’ intensity. The image is often brighter in the middle than at the sides. This is because X-rays hit the sides at an angle, and so the flux at the sides is lower. The detector may also have different sensitivities, causing different intensity readings from the same X-rays. This problem is largely dealt with through the use of a shading correction. Renormalisation factors for each pixel are found just prior to the experiment by measuring the intensity of X-rays through ‘empty air’. These factors are then applied such that the baseline level is 60000.

However, there may have been a change in the baseline levels in between the experiments. We therefore recommend renormalisation of the baseline level for each particular experiment.

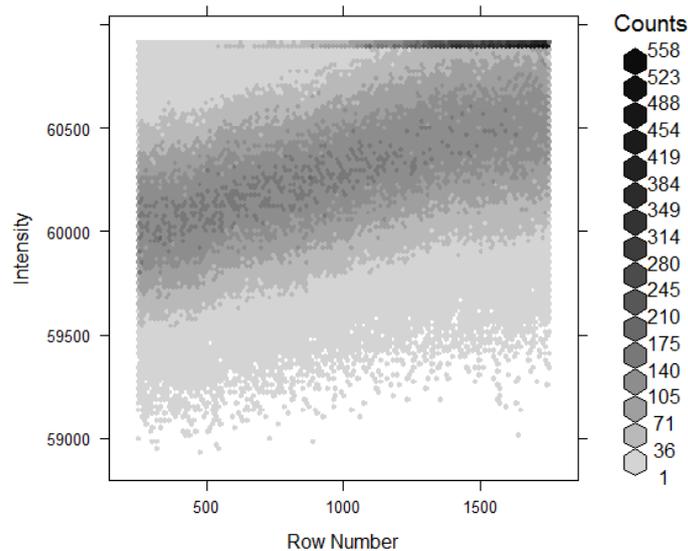


Figure 6: Gray-level map of number of pixels in various rows at various intensities. Note that the median intensity increases as the row number increases.

We can approximately determine the baseline intensity of a row by taking the median of the intensities at the border of the image. However, the baseline level varies with the height of

the row. Figure 6 provides a gray-level map for the 400 values at the corners of each row; the intensities are clearly higher for the larger row numbers.

Thus as a compromise, a linear variation between the rows is allowed by doing a linear fit using 400 values from each row. Doing this uses all the information in the image matrix, and yield fitted values for each row which are used to normalise the intensities of X-rays in each to 60000. This eases the calculations without increasing error (since the normalisations in the detector are already arbitrary).

N.B: The median baseline intensities in each row are plotted to check the above fit. In fact, the regressions show reasonable agreement with the row data, as seen in Figure 7, so the fit is adequate.

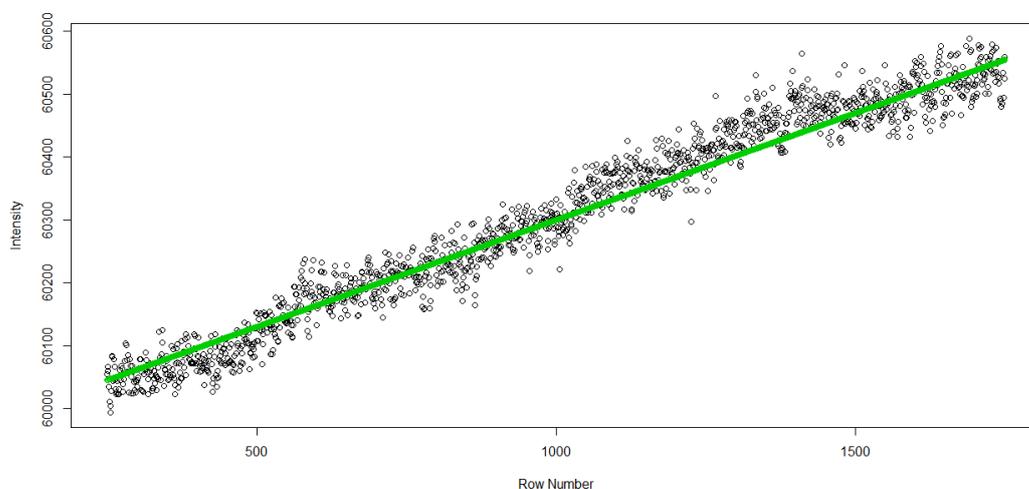


Figure 7: Scatterplot shows the individual row medians; the fitted values are plotted on a line.

3.2 Registration

We noted, at the end of Section 2, that the object image is assumed to be centralized at O_2 . Inspection of the image makes it clear that this is not completely correct, as typically the cylinder image is slightly slanted. In this preliminary analysis, we deal with this in an *ad hoc* way by estimating the central axis of the rod using a simple least-squares technique, and adjust the image accordingly. For the present, we assume this *ad hoc* approach deals with all but a negligible part of the resulting discrepancies. A later analysis will deal with this more coherently.

Let f_{ij} be the response at row i , column j , after the baseline intensity has been normalised to 60000. This means that the object mass intercepted by the X-rays hitting row i , column j on the detector is proportional to $60000 - f_{ij}$. Furthermore, we want to ignore low intensities which are due to noise or penumbra. Reasonable changes of threshold amount to adding or removing points of low intensity, and therefore do not greatly affect the calculation of the median. Here we require $60000 - f_{ij} > 18000$. Thus, the basic technique to find the centre for a row of values

$\{f_{i1}, f_{i2}, \dots, f_{in}\}$ is to find

$$\arg \min_m \sum_j I(f_{ij} < 42000)(60000 - f_{ij})(j - m)^2.$$

For a single row, a closed form solution for this minimisation is

$$m = \frac{\sum_k k I(f_{ij} < 42000)(60000 - f_k)}{\sum_k I(f_{ij} < 42000)(60000 - f_k)}.$$

However, following similar considerations to those applied when finding the intensity of air, we prefer to use all available information at once, especially since the centres must form a line due to the geometry of a cylinder. Hence, we instead solve the following minimisation problem

$$\arg \min_{a,b} \sum_{i,j} I(f_{ij} < 42000)(60000 - f_{ij})(j - a - ib)^2.$$

The centres of each row are plotted to check the fit in Figure 8; clearly, the regressions agree well. Note that the regression removes the pixelation effects introduced from finding the centre values individually.

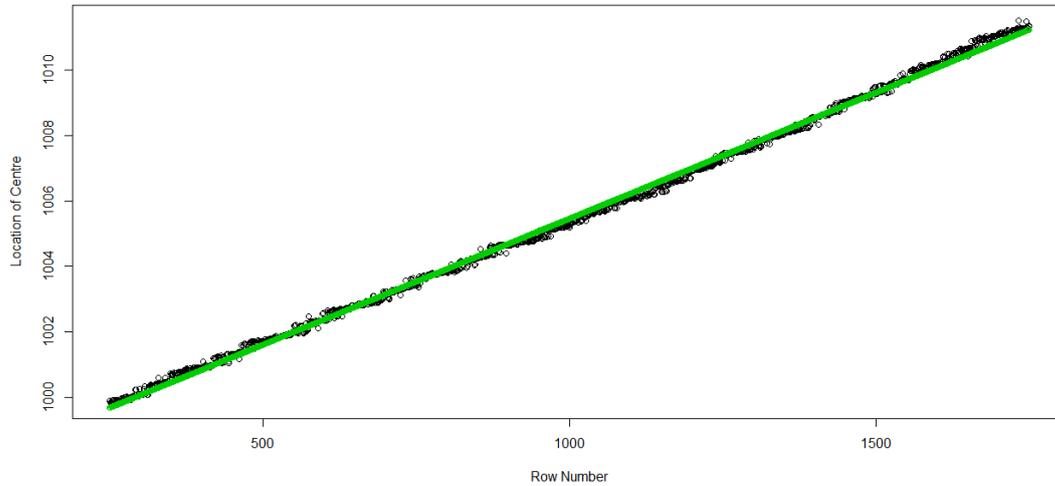


Figure 8: Scatterplot shows the individual row centres; the fitted values are plotted as a line.

3.3 Fitting x, y, μ

We have thus found the centres of each row c_i . The next task is to fit the data $\{f_{ij}\}_{(i,j) \in A}$ for some set A to a model of the form

$$f_{ij} \sim M(j - c_i, x, y, \mu) + \epsilon_{ij}, \text{ for } (i, j) \in A.$$

where $\{\epsilon_{ij}\}_{(i,j) \in A}$ are independent and identical normals. This is a non-linear regression problem, and thus can be solved by the Gauss-Newton method, an iterative process to improve on an initial guess.

We first start the Gauss-Newton method with an initial guess $\{x_0, y_0, \mu_0\}$. By Taylor's theorem,

$$\begin{aligned} M(j - c_i, x, y, \mu) - M(j - c_i, x_0, y_0, \mu_0) &\approx (x - x_0) \left. \frac{\partial}{\partial x'} \right|_{x_0} M(j - c_i, x', y, \mu) \\ &+ (y - y_0) \left. \frac{\partial}{\partial y'} \right|_{y_0} M(j - c_i, x, y', \mu) \\ &+ (\mu - \mu_0) \left. \frac{\partial}{\partial \mu'} \right|_{\mu_0} M(j - c_i, x, y, \mu'). \end{aligned}$$

Thus, a related linear fit problem is as follows,

$$\begin{aligned} f_{ij} - M(j - c_i, x_0, y_0, \mu_0) &\sim \delta x \left. \frac{\partial}{\partial x'} \right|_{x_0} M(j - c_i, x', y, \mu) \\ &+ \delta y \left. \frac{\partial}{\partial y'} \right|_{y_0} M(j - c_i, x, y', \mu) \\ &+ \delta \mu \left. \frac{\partial}{\partial \mu'} \right|_{\mu_0} M(j - c_i, x, y, \mu'). \end{aligned}$$

Assuming the derivatives can be computed, standard methods can be used to find the best values for δx , δy and $\delta \mu$ which have the smallest least-squares error. The next estimate would be

$$\begin{pmatrix} x_1 \\ y_1 \\ \mu_1 \end{pmatrix} = \begin{pmatrix} x_0 \\ y_0 \\ \mu_0 \end{pmatrix} + \lambda \begin{pmatrix} \delta x \\ \delta y \\ \delta \mu \end{pmatrix}.$$

Ordinarily $\lambda = 1$, but this choice of λ may increase the least-squares error in the original problem if the surface varied from the tangent space enough. If this were so, λ would be successively halved until a choice of λ which decreases the least-squares error is found, yielding a new guess for the parameters. Then, a convergence check is done. If the residuals are nearly orthogonal to the tangent space, this indicates that a critical point has been reached and thus the procedure is terminated. This is done by checking if the ratio between the scaled length of the tangent plane component and the scaled length of the orthogonal component is less than a fixed convergence tolerance. The statistical package `nls` in R sets this to a default value of 10^{-6} .

The Gauss-Newton method can fail in two ways. Firstly, the linear problem above may not have full column rank, and thus standard methods will not work. In R, this raises a ‘singular gradient matrix at initial parameter estimates’ error. This is avoided by choice of better starting parameters. Secondly, the algorithm may experience extreme slowdown as step size becomes very small. In R, one of the controls of the `nls` function is the size of the minimum step, and if the step size is too small, the algorithm halts and raises an error ‘step factor reduced below minFactor’. This problem is linked to the default choice of the convergence tolerance of 10^{-6} which is too demanding. It is possible that the final estimate is within the minimum step size from the optimal solution, and so is entirely adequate. Bates and Watts[1] argue that a convergence tolerance of 0.001, or 10^{-3} is sufficient as ‘any inferences will not be affected materially by the fact that the current parameter vector is less than 0.1% of the radius of the confidence region from the least squares point’. Indeed, the achievable convergence tolerance is dependent on the problem. For example, we simulated toy data from simple exponential models with normal noise. It is found that the algorithm only attained convergence tolerance of about 10^{-3} when applied to the toy data. The context must thus be considered when choosing the convergence tolerance threshold. We propose the following strategy. First, good initial conditions are chosen to ensure that the gradient matrix is invertible. Then, the algorithm is run on different datasets until the algorithm halts due to small step size. Next, the residuals are examined to check the adequacy of fit. Finally, the achieved convergence tolerance of the adequate fits are found. The convergence tolerance threshold is then be chosen as the maximum of these achieved convergence tolerances.

3.4 Implementing the Gauss-Newton method

We have seen how we can approximate the centre of a row, which for purposes of illustration we take to be at column 1005. Since the radius of the cylinder is 160 pixels, the approximate boundary of the cylinder image is at column 845, as indicated in Figure 9.

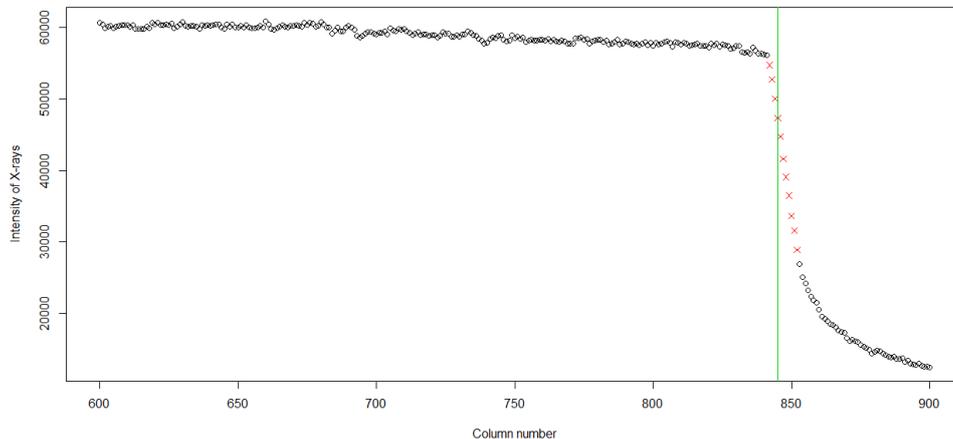


Figure 9: This displays an estimate of the approximate object boundary (indicated here by a vertical line) and of the location of the penumbra (plotted using crosses).

We expect the curve formed by the data points to exhibit nearly linear behaviour over a range of several columns around the cylinder boundary, indicating a penumbra effect. The left edge of this penumbra can be approximated from the sharp gradient change, but the right edge of this penumbra is more difficult to detect. A possible estimate for the penumbra is marked as crosses in Figure 9.

However, there is a non-zero slope to the left of the penumbra, as is apparent in Figure 9. By simple geometry, if the spot size were smaller than 1mm, then the intensity curve this far from the object should be more or less level. Thus, the model cannot be a good fit to the data as a small but significant amount of X-rays is coming from a more diffuse spot. This suggests a mixture model as noted by Dong et al.[2]: most of the X-rays come from a more concentrated spot but a small amount come from a larger region. In this new model, the attenuation constant μ is still included in our model. However, instead of estimating x, y , we instead estimate the limits of the outer spot x_1, y_1 , the limits of the inner spot x_2, y_2 and the proportion of X-rays coming from the outer spot θ .

Initial explorations indicate that this is the case, as we can find parameters for this that give a good visual fit for the data. We thus are able to use the parameters for the first fit as our initial conditions for our fit.

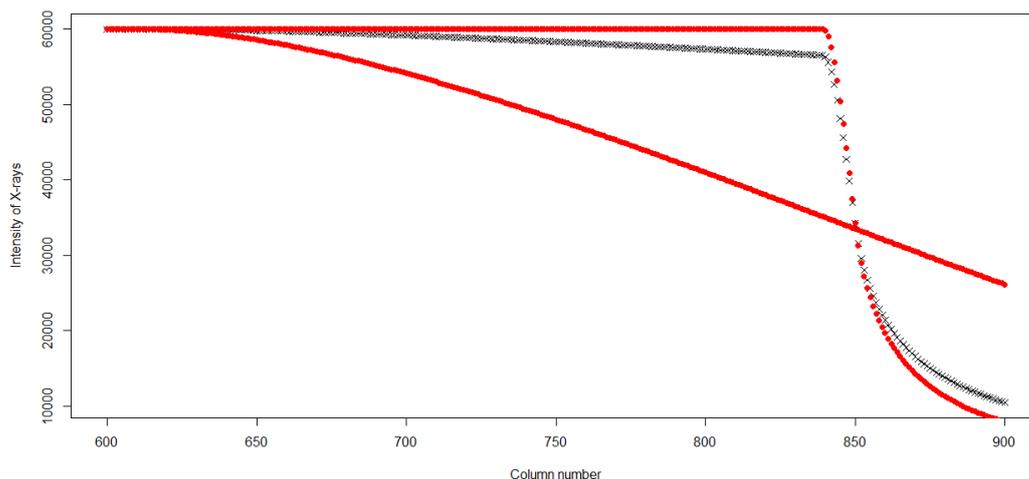


Figure 10: The two models with different spot sizes are plotted using solid circles. Their mixture is plotted using crosses and produces a promising first guess for the data from Figure 9.

Now, a matrix contains 2000 rows of data. However, using all of this data slows down the Gauss-Newton algorithm. Hence for the fit, the outer 249 rows on each side are discarded since detector performance can tail off to the edges. Furthermore, the analysis will only use every k th row, where $k \geq 5$, to try to mitigate the correlation between intensity values. We will vary k to explore the effect of sampling more rows on the computation time as well as the convergence tolerance.

4 Results

Computing was carried out on the departmental computing cluster Buster, which consists of 5 Dell PowerEdge R420 servers with 12x2.40GHz processors (E5-2440) and 6 Dell PowerEdge R410 servers with 12x2.80GHz processors (X5660).

First we discuss the validity of the model by looking at the residuals. Then, we look at the effect of the number of rows on computation time and convergence tolerance. Finally, we discuss the estimates found.

4.1 Residuals

Most residual plots raise no apparent issues. A characteristic fit is shown in Figure 11. Calculation of the X-ray intensities in the absence of an object yield a standard deviation estimate for baseline values of about 250. Hence, errors of order 1000 are expected given the number of data points. The residuals display a kink about 330 pixels from the centre, but this is expected characteristic behaviour of non-linear regression since the model is not smooth at this point. Since the model is concave at this point, we expect our residuals to be positive here, and this is indeed the case. Finally, there is a large spike in residuals within the object between 150 and 170 pixels from the centre. This is to be expected since the errors in the registration of centres will be magnified by the large slope in our primary penumbra. Bearing these points in mind, the reported standard errors of order 2000 are satisfactory.

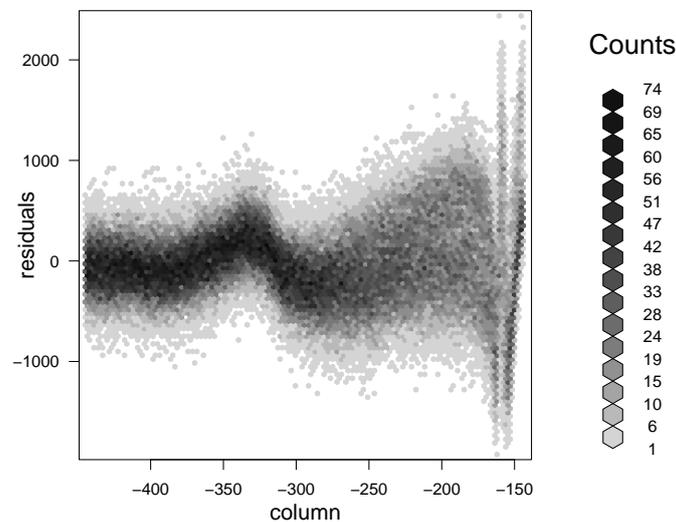


Figure 11: Fitted residuals when taking every 5 rows.

However, two of the residual plots exhibit anomalies and we have had to exclude some of the results from these two data sets. In these two data sets, the data is remarkably constant at intensity level 60000, at locations left of approximately 300 units from the centre. This signals saturation: the true intensity is larger than 60000, but the detector can only measure values up to 60000. Thus, the true range of the outer penumbra is concealed, and so these data sets cannot yield good values for the size of the outer spot and the proportion of power. We therefore discard these values.

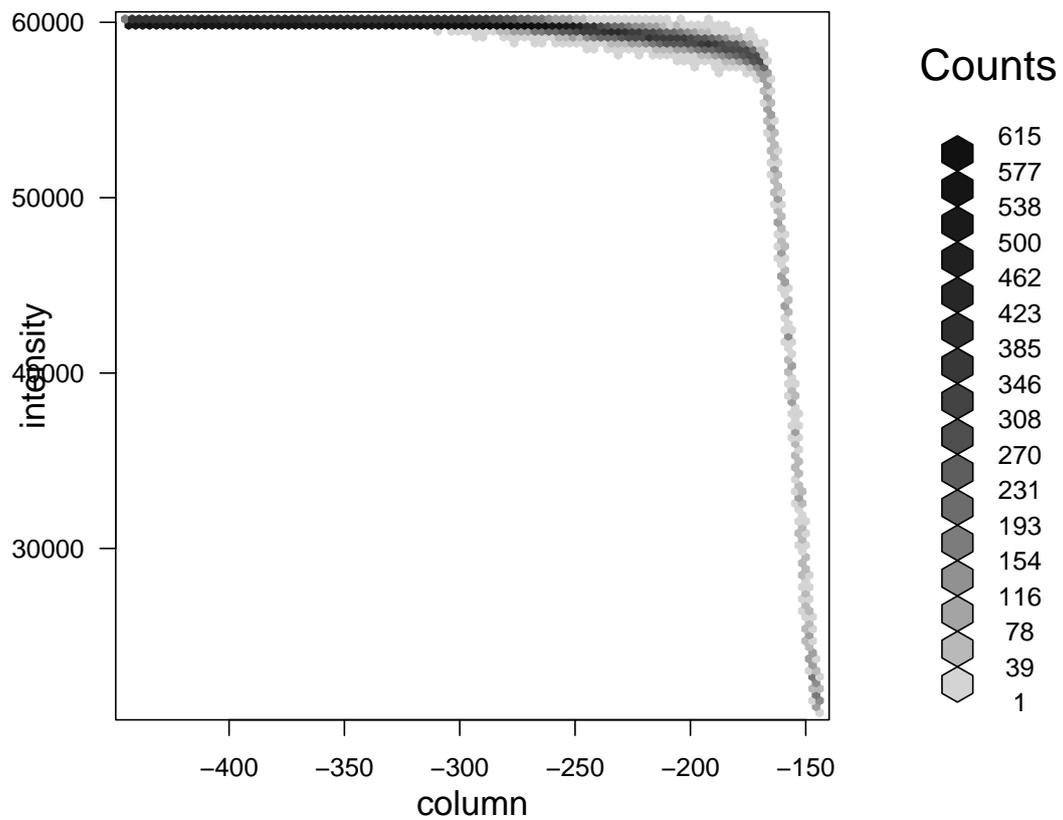


Figure 12: Data from a plot suffering from saturation. Note the constant intensity of 60000 from column -450 to column -300.

4.2 Choice of number of rows

We have repeated our nonlinear regression method for an image matrix corresponding to each experimental set of parameters, using successively every 50th, then every 20th, then every 10th, then every 5th row from that matrix. Generally, the achieved convergence tolerance decreases as more rows are sampled. However, this is rather erratic, and there appears to be no general rules about the right number of rows to sample. Furthermore, no matter how many rows are taken, the convergence tolerance does not reduce below 0.003. This suggests that a suitable choice of convergence tolerance threshold is of the order of 0.005.

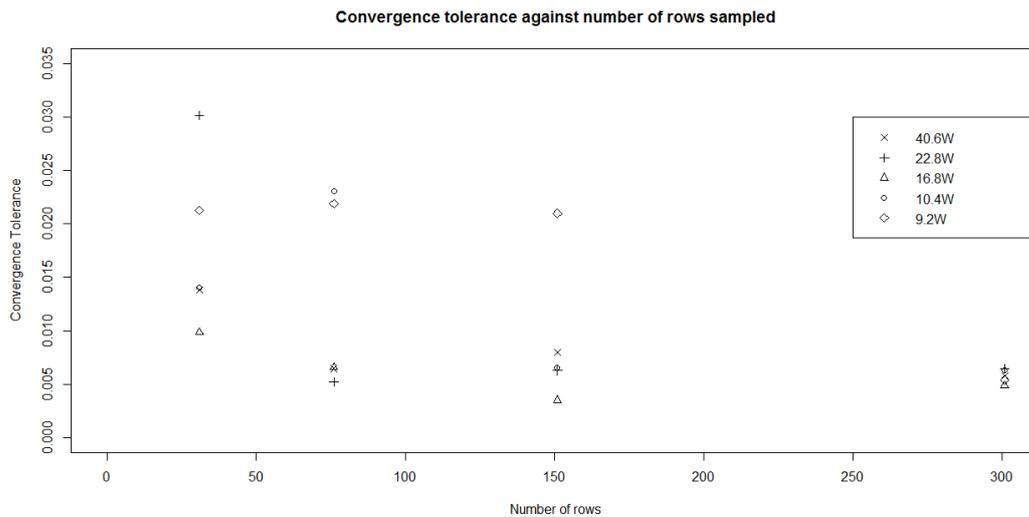


Figure 13: Convergence tolerance achieved by the algorithm for differing matrices and number of sampled rows.

However, the time taken to run the method appears to increase linearly with the number of rows. We plotted the regression line and it indicates that analysis for a row takes approximately 15 minutes.

The aim is to achieve a satisfactory convergence tolerance without analysing unnecessarily many rows when a few would suffice. We propose to first run the method on every 50th row and see if we achieve a convergence tolerance of less than 0.005. If this is not sufficient, we repeat the process on every 20th, then every 10th, then every 5th row. We cap this process at every 5th row because the rows are correlated.

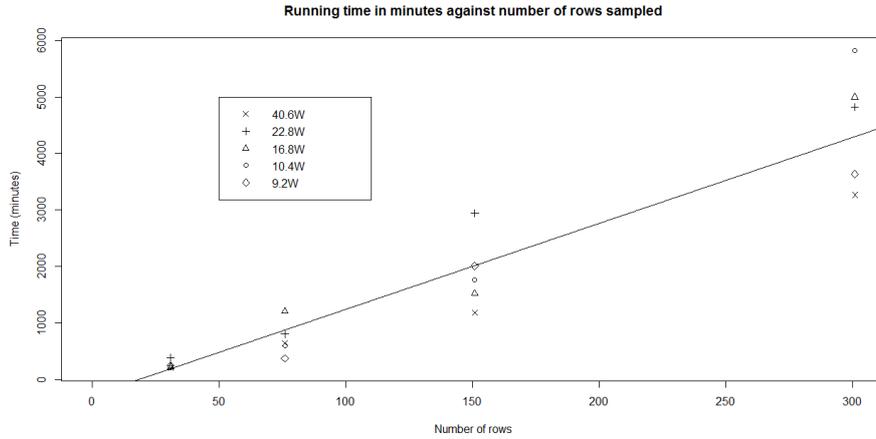


Figure 14: Time taken by the algorithm for differing matrices and number of sampled rows.

4.3 Parameter estimates

We run the algorithm on 15 different matrices using the maximum 301 rows and consider at the parameter estimates for each time/power settings: 0.25s/40.6W, 0.5s/22.8W, 0.7s/16.8W, 1.4s/10.4W, 2s/9.2W. In many of these cases, we have not reached a convergence tolerance of less than 0.005, in some we have not even reached a convergence tolerance of 0.01. Hence, we will not be able to make any strong statements about the parameters and their standard errors, but we can show some general trends, especially when the evidence is striking. In the following graphs, we will indicate the convergence tolerance of each point.

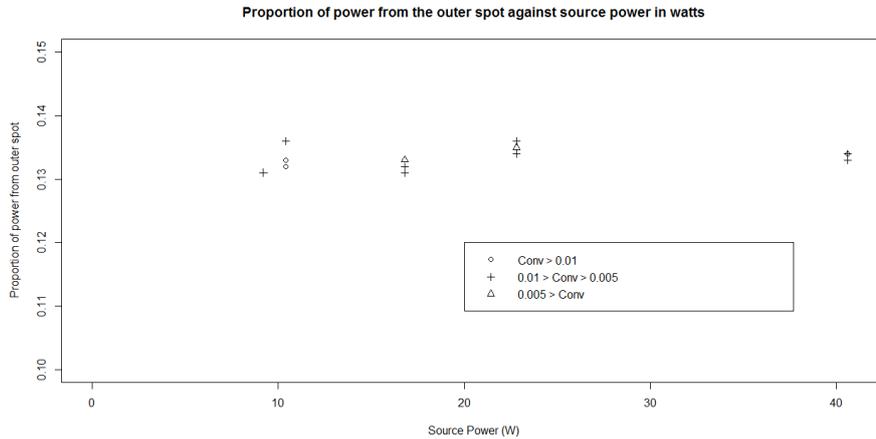


Figure 15: Proportion of power in outer spot for experiments of different power, discarding two data points due to saturation. Standard error is less than 0.002.

Recall that for analysis of the proportion of power and the size of outer spot, we have discarded two of the experiments due to failures in the detector. The first thing we notice is that the proportion of power in the outer spot is similar for the experiments recorded in Figure 15. Furthermore, the size of the outer spot is consistently estimated at 6.20mm with standard errors all less than 0.002mm. This indicates that our mixture model is working, and furthermore, that the power and size of the outer spot is unaffected by the power of the source.

Now, the size of the inner spot decreases with the power of the source. This contradicts the common assumption that the size of the inner spot increases with the power of the source (see for example Müller et al.[3]). This requires further investigation.

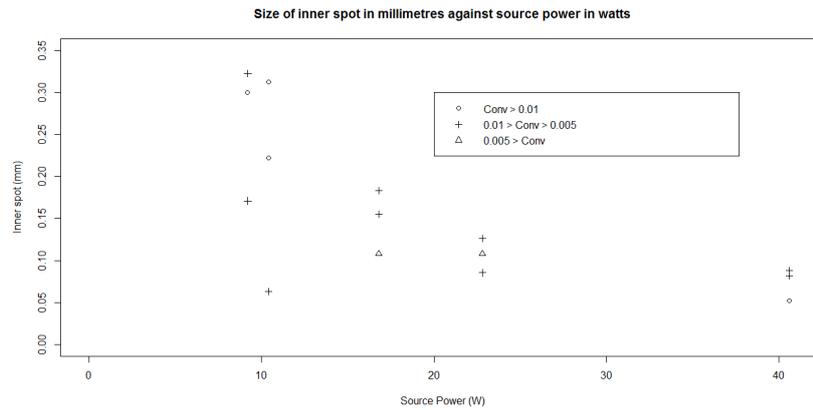


Figure 16: Size of the inner spot for experiments of different power. Standard error is less than 0.001mm.

In fact, there is an acceptable fit between the inverse of the size of the inner spot and power as we can see from Figure 17.

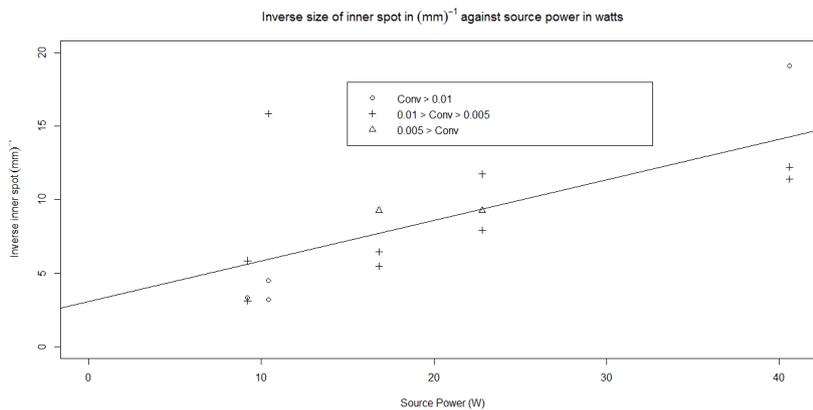


Figure 17: Inverse size of the inner spot for experiments of different power.

Finally, we consider estimation of the attenuation coefficient. This is the parameter that will change the most from experiment to experiment since it depends on the spectrum of X-rays emitted from the source. As the power goes up, the spectrum of X-rays will generally move towards more energetic ones, and we would then expect fewer of these X-rays will be blocked, thus we will expect the attenuation coefficient to generally decrease. However, this is not the case, we instead see that the attenuation coefficient in fact tends to increase as the power goes up. This means that a higher proportion of the X-rays are being blocked for higher powers at the source. This is curious and requires further investigation.

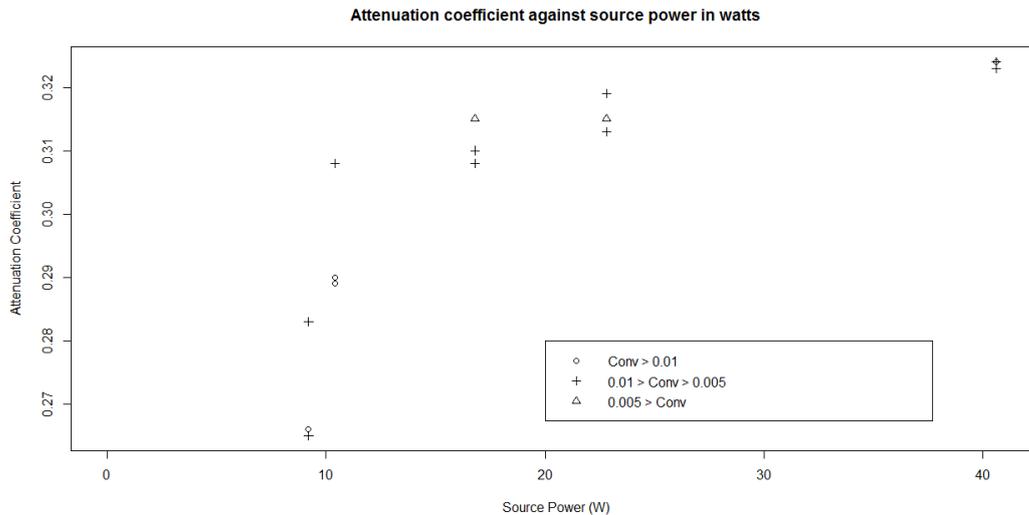


Figure 18: Attenuation coefficient for experiments of different power. Standard error is less than 0.0002.

5 Conclusions

From the residuals, the model appears to be generally adequate. However we plan to develop this analysis to remove the assumption mentioned at the end of section 2, that the object image is assumed to be centralised at O_2 .

There is a 6.2mm outer spot with about 13% of the power attributed to it, and these values are unaffected by changing the power of the source. Furthermore, this size is close to the size of window in front of the source from which the X-rays emerge. This suggests that there is scatter within the source which then permeates out of the source window uniformly. Now, these experiments were done without a prefilter (material placed at the source window to help cut out the scatter). One next step will be to repeat these experiments with a prefilter. If we see that the outer spot disappears, this would provide strong evidence that the outer spot is due to scatter within the source.

There is also a problem with the higher power at the source yielding a smaller spot size as well as a higher attenuation coefficient. This is very contrary to expectations as it is generally thought

that higher power at the source will result in a more diffuse but more powerful spectra, leading to a lower attenuation coefficient. It has been suggested that it is possible that the machine's shading corrections were not set up correctly, or that there are problems with the focus matrix.

It would be useful to test both these issues on other machines. In general, the method explained here will work on any cylinder of constant radius. The first step is to choose a magnification level suitable for the detector and object. Then, at different power levels, many independent images would be found. One flaw in the above experimental design was that all images of the lowest power setting were found before moving on to the next higher power setting and so on. This means that if there is any time-sensitive variation in the conditions, for example temperature change, would be mixed into the experimental results. An ideal experimental setup is for an image of each power setting would be found in turn before repeating the experiment. This would mean that the experiment would take a lot longer because we will need to let the machine 'warm up' to the appropriate power setting. Instead, we propose a deliberate nonmonotonic ordering; a possible order could be 2s/9.2W, then 0.25s/40.6W, then 1.4s/10.4W, then 0.5s/22.8W, and finally 0.7s/16.8W. This would make any variation seem like random noise for any regression, but it also allows us to extract this variation through sorting the residuals in time order. After collecting the results, we would then run it through the method detailed in this paper, allowing us to generalise the results given above. This would give a protocol for confirming spot geometry from penumbra data in each experiment, in view of constructing a data-driven estimate for the blurring.

References

- [1] D. M. BATES, D. G. WATTS. *Nonlinear Regression Analysis and Its Applications*. John Wiley & Sons 1988.
- [2] X. DONG, T. NIU, X. JIA, L. ZHU. Relationship between X-ray illumination field size and flat field intensity and its impacts on X-ray imaging. *Medical Physics* 39(10) 2012, 5901-5909.
- [3] P. MÜLLER, J. HILLER, A. CANTATOR, M. BARTSCHER, L. DE CHIFFRE. Investigation on the influence of image quality in X-ray CT metrology. *Conference on Industrial Computed Tomography, Wels, Austria* 2012, 229-238.
- [4] P. RUSSO, G. METTIVIER. Method for measuring the focal spot size of an X-ray tube using a coded aperture mask and a digital detector. *Medical Physics* 38(4) 2011, 2099-2115.