Correction of Beam Hardening Artefacts in Microtomography for Samples Imaged in Containers

Jeremy Holt, Mahsa Paziresh, Andrew Kingston, and Adrian Sheppard

Dept. Applied Mathematics, Research School of Physics and Engineering, The Australian National University, Canberra 0200, Australia

ABSTRACT

We explore the use of referenceless multi-material beam hardening correction methods, with an emphasis on maintaining data quality for real-world imaging of geologic materials with a view towards automation. In particular, we consider cases where the sample of interest is surrounded by a container of uniform material and propose a novel container-only pre-correction technique to allow automation of the segmentation process required for such correction methods. The effectiveness of the new technique is demonstrated using both simulated and experimental data.

Keywords: x-ray tomography, beam hardening, reprojection

1. INTRODUCTION

Computed Tomography (CT) has been applied to the study of geological samples such as fossils, reservoir rocks and soils as a non-destructive image technique for the past three decades.[1] The number of possible applications has grown with increasing resolution, with computed microtomography (μCT) imaging techniques now able reach sub-micron resolutions, *e.g.*, Arns *et al.*[2] One technique seeing increased application is multiple state imaging, in which the same sample is imaged under different conditions such as varying pressure,[3,4] saturation states[5] or weathering.[6] When performing such analyses, a cylindrical container is required to maintain pressure or hold saturating fluid. For high-magnification, high cone-angle imaging, such as that which we perform at the ANU μCT facility, the sample is extremely close to the X-ray source. Containers are used in this case to dissipate the heat radiating from the source which to prevent sample movements.[7] Containers are thus seeing increased use in μCT imaging; as we discuss, they introduce new challenges and opportunities, specifically with regard to the correction of beam hardening artefacts.

The standard linearisation pre-processing step that is performed prior to tomographic reconstruction is based on the Beer-Lambert law which assumes an exponential relationship between beam intensity and attenuation path length:

$$I(t) = I_0 e^{-\int_0^t \mu(x)dx}. \tag{1}$$

where I_0 is the initial beam intensity, x is the depth into the material, $\mu(x)$ is the attenuation coefficient at a depth x and $I(t)$ is the beam intensity after passing through a material thickess t. This assumes that the only effect of the material on the beam is to reduce its intensity; *i.e.*. that the attenuation within each volume element is independent of how much material it has already passed through. However, the attenuation coefficient μ is also a function of x-ray beam energy, with attenuation generally decreasing with energy. Because of this, the lower energy components of a polychromatic beam (such as produced by lab-based microfocus x-ray sources), are preferentially absorbed, causing the beam to increase in average energy (become *harder*) as it passes through a sample. This effect is known as *beam hardening* and can cause artefacts such as *cupping*, (the apparent higher density near the edge of a sample), and *streaking*, (dark shadows between areas of relatively high density).[1,8–10] Such artefacts reduce image fidelity and can have a dramatic effect on downstream quantitative analysis, especially the process of intensity-based *segmentation* in which regions of the tomogram are labelled to represent different materials.

Further author information:
A.P.S.: E-mail: adrian.sheppard@anu.edu.au

Developments in X-Ray Tomography IX, edited by Stuart R. Stock, Proc. of SPIE Vol. 9212, 92120A · © 2014 SPIE · CCC code: 0277-786X/14/$18 · doi: 10.1117/12.2062393

There is, in general, insufficient information to perform tomographic reconstruction from polychromatic attenuation data, even when the beam spectrum and detector response are known exactly. Therefore, beam hardening is an insoluble problem unless assumptions can be made about the sample. While numerous methods have been proposed that target specific cases (such as imaging of bone and tissue), no method has seen widespread adoption for μCT imaging. Therefore, as other artefacts of x-ray imaging have been dramatically reduced in recent years, beam hardening is now the major source of imaging artifacts in many situations and it remains one of the key challenges for x-ray tomography.

In this paper we work towards an automated method for beam hardening artefact correction that works in situations that are of widespread interest for μCT imaging. The specific technique investigated and developed here is that of referenceless post-reconstruction correction by Krumm, Kasperl, and Franz[11] (KKF method). We first examine the validity of the KKF method when a multi-material sample is treated as a single material. We then show that when a container is introduced, the resulting beam hardening artefacts can make the segmentation of materials inside the container difficult. To resolve this issue we present a modification of the KKF method that allows for accurate automatic segmentation of a sample within a container. We also include a discussion on the implications of making corrections in intensity space or attenuation space.

2. EXISTING BEAM HARDENING CORRECTION TECHNIQUES

Various methods of correcting beam hardening artefacts in μCT images have been proposed (see[9] for a review). A common technique to reduce the artifact is to filter the beam,[1,12] where a sheet of material is placed before or after the sample to remove those low energy x-rays that would suffer rapid absorption within the sample. This method has a couple of shortcomings; it reduces severity of the artefacts but does not eliminate beam hardening; it causes a reduction in overal flux reaching the detector, reducing signal to noise ratio and increasing required acquisition time for the same results. When a container is present, especially if it is highly attenuating as is required for high pressure imaging, the flux is reduced further, causing additional loss of contrast. A dual energy approach can also be taken,[13,14] in which the same sample is imaged at two different energies, and the results combined under the assumption of a linear decomposition into basis functions. However this requires an additional μCT scan, which can be time consuming and expensive, especially if two scans are already required for dual-state imaging.

Other correction techniques occur after the acquisition of data. Beam hardening curve linearisation[8] corrects the projection data by assuming that the sample is composed of just one material. This is a reasonable assumption for samples which are approximately homogeneous, but cannot be used for heterogeneous materials or when a container is present. One suitable group of correction methods are the so-called post-reconstruction methods,[15] which rely on *reprojection* of data to estimate the polychromatic and monochromatic projections. Reprojection uses forward-projection of rays to generate simulated projections from a reconstructed tomogram, based on knowledge about the x-ray spectrum and material properties. For monochromatic data containing no other artifacts, the reprojection and the original data will be the same. However, a reconstruction containing beam hardening artifacts is not, in general, fully consistent with the experimental data, so the reprojected data will be different from the experimental projections. Post-reconstruction methods use the difference between the experimental data and the simulated reprojection as a correction that can be applied to the data. However, even if the x-ray spectrum and detector response is known, one must have a perfect reconstruction of material attenuations. In addition, even for simple samples, the method fails for cylindrical samples where beam hardened projections will reconstruct into a completely consistent tomogram with commensurate beam hardening artifacts.[10] In this work we consider the post-reconstruction method proposed by Krumm *et al*[11] which we'll term *Referenceless Post-Reconstruction Correction*, or the Krumm, Kasperl and Franz (KKF) method. The KKF method allows the correction to operate with no knowledge of the materials or x-ray spectrum, but requires that the sample can be decomposed into homogeneous regions.

3. REFERENCELESS POST-RECONSTRUCTION CORRECTION BY KRUMM, KASPERL AND FRANZ (KKF)

3.1 Outline of the KKF technique

Given normalised projection data, $I(\vec{v},\vec{x})/I_0$, the integrated attenuation, R, can be derived as $R = -\log(I/I_0)$ (according to Eqn. 1); this is reconstructed using standard filtered backprojection methods to generate a tomogram with beam hardening artefacts. The KKF method proceeds using a segmentation of the reconstructed tomogram that distinguishes each material present. Each phase (material) in the segmented image is then reprojected to create a simulated monochromatic projection for each material. Since each reprojection assumes a monochromatic beam, the value of each pixel in the resulting projections is related directly to the path length through the corresponding material. For each material i one can work either in the space of integrated attenuation:

$$R_i(\vec{v},\vec{x}) = \int_0^\infty \mu_i g_i(\vec{x}+\vec{v}s)ds, \tag{2}$$

where $g_i(\vec{r})$ is the binary represesentation of the ith material segmentation at the point $\vec{r}$. Alternatively one can work in I-space:

$$I(\vec{v},\vec{x})/I_0 = e^{-R_i(\vec{v},\vec{x})}. \tag{3}$$

Since there is a projection for each material, N materials will result in projection data consisting of an N-vectors $\vec{R'}(\vec{v},\vec{x})$ for each pixel. One can then construct a scatter plot in $\mathbb{R}^{N+1}$ in which each point corresponds to a single pixel, with the original projection data value of that pixel shown on the vertical axis, while the remaining axes show the reprojection vector for that pixel. Each point therefore represents the experimental attenuation along a single ray, as well as the path length through each material for that ray. If the original beam was monochromatic, the points would lie on a N dimensional hyperplane such that the attenuation was directly proportional to the path lengths:

$$R_{mono}(\vec{v},\vec{x}) = \sum \mu_i R_i(\vec{v},\vec{x}). \tag{4}$$

For data from a polychromatic beam, the points lie on a N-dimensional hypersurface in which attenuation is some function of each material's path length:

$$R_{exp}(\vec{v},\vec{x}) = f(R_1(\vec{v},\vec{x}), R_2(\vec{v},\vec{x}), ...). \tag{5}$$

The KKF method proceeds by first fitting a hyperplane through this polychromatic data, to represent the best monochromatic approximation; this is equivalent to finding the best estimate for the attenuation coefficients μ_i. A hypersurface is fitted through the data and the difference between this and the hyperplane forms the required amount of correction to be applied at each point:

$$R_{corr}(\vec{v},\vec{x}) = R_{exp}(\vec{v},\vec{x}) + f(R_1(\vec{v},\vec{x}), R_2(\vec{v},\vec{x}), ...) - \sum \mu_i R_i(\vec{v},\vec{x}). \tag{6}$$

It should be emphasised that the KKF method works as a *correction*, *i.e.*, it does not transform the data, rather it tries to improve the data. Therefore it is not necessary to achieve complete consistency between the original and projected data, which is an advantage when complete information is not available, and has the potential to make the method robust to errors in assumptions or in segmentation.

As an aside we note that as presented by Krumm *et al*,[11] the method applies the corrections in intensity space (I/I_0) rather than in integrated attenuation space (R) as we show here. The authors do not explain the derivation or motivation for this choice, and (as we show in section 6) there are some tangible advantages to using R-space. In this paper, when we refer to the KKF method we refer to the method as modified to apply in attenuation space.

3.2 Segmentation Considerations

Krumm *et al*[11] demonstrated that this method works excellently for a range of sample types when all the materials (phases), are distinguished from one another accurately. However often the separation of different phases is the goal of μCT analysis and if an accurate segmentation is possible on the original image then a correction is not needed. An inaccurate segmentation used for the correction can result in the introduction of new artefacts. This occurs because, if phase 1 requires a different amount of correction to phase 2, an area of phase 1 which is wrongly identified as phase 2 will be corrected differently to other areas of phase 1 and may no longer resemble them. One must therefore take great care when using correction methods such as this since they can introduce artefacts which are indistinguishable from real features.

The method has a level of robustness to slightly inaccurate segmentations, however an automatic segmentation method is desired, and therefore there is no quality control on the segmentation. Automation has obvious advantages, especially in the case in which data collection and reconstruction both take a very long time. Minimum user input after the process starts allows the entire data collection, reconstruction and correction process to run unaided and continue overnight. Automation also has the ability to give consistent results, unlike more subjective user-led segmentations. The drawback of course is that there is no guarantee on the quality of the segmentation. To avoid the potential loss of fidelity through an inaccurate segmentation, we suggest the grouping of multiple materials which are not easily distinguished into a single phase, as discussed in Section 4.

Certain phases must still be segmented, for example the container and the sample, however automatic segmentation becomes relatively straightforward after we have grouped the materials. The choice of automatic segmentation method is reliant on the type and quality of the images being processed. The use of a watershed segmentation[16] using the gradients as defined by a Sobel filter may be enough to separate the phases. In this paper, because the phases were approximately cylindrical in shape, we used a k-means clustering technique[17] in which the radial distance and grayscale intensity were used to distinguish the phases.

4. CORRECTION ASSUMING A SINGLE MATERIAL SAMPLE

The first sample we consider is experimental data: a 5 mm diameter Bentheim sandstone imaged at 80 kV with a 0.3 mm flat sheet of Aluminium as a filter; this filter is insufficient to alleviate beam hardening for such a sample. Figure 1(a) shows a two dimensional slice of the sandstone, which is shown again in Figure 1(b) with the contrast adjusted to emphasise the cupping artefact in the rock. We propose that the beam hardening can be corrected by treating the entire sandstone sample as a phase made up of air and silicon, rather than performing an internal segmentation to separate the pore space. This means we are assuming that every ray which passes through a certain amount of the air-silicon phase will experience a similar amount of beam hardening.

Figure 1(c) shows the first possible segmentation in which we imagine the rock is convex. The resulting corrected image, seen in Figure 1(d), has corrected the cupping artefact apart from a narrow ring around the sample which is slightly brighter. We explain this by first noting that there is a variance in attenuation for x-rays passing through the same length of segmented phase which is related to the proportion of air and silicon along that path. For rays which pass through only a small amount of the segmented region, *i.e.*, rays which are approximately tangential, the variance in attenuation becomes large compared to value of attenuation. A ray passing through the very edge of the segmented region could actually have passed through 100% air or 100% silicon but will be treated the same, which results in the subtle bright ring seen. The average solid volume fraction is much smaller at the edge of the sample, so that grazing rays pass through much less solid material relative to their path length through the sample than do the "interior" rays. The effect is quickly reduced as the attenuation becomes larger, and so it would be valid to ignore the edge of the sample.

If required, the edge ring artefact can be improved by allowing the segmentation to form a closer fit around the sample, as in Figure 1(e). The correction from this segmentation, seen in Figure 1(f), reduces the problem, but there is still a very subtle darker ring at the edge of the sample. We have shown that the method can work effectively to correct the beam hardening in a multimaterial object by segmenting it as a single material and applying the referenceless post-reconstruction correction technique. However there are constraints on what materials we can group into the same phase, namely that all rays of the same path length through the segmented phase will cause similar beam hardening. This assumption is met if the attenuation of all the grouped materials is similar or if the material is uniformly heterogeneous, as in this case.

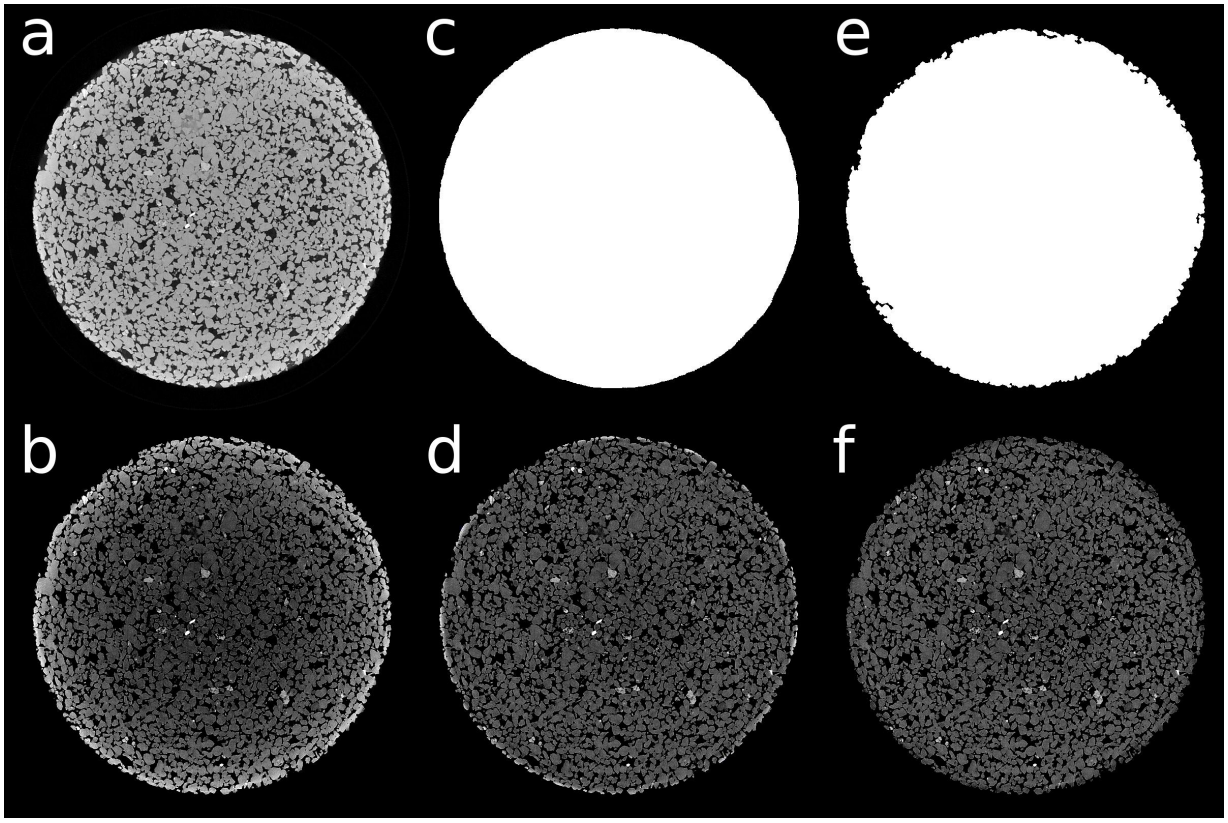

Figure 1. Cross section of a 5 mm diameter Bentheim sandstone sample used to demonstrate the effectiveness of the correction method for different segmentations. The core was imaged using a micro-focus X-ray source with a W target using an 80 kV accelerating voltage and *under*-filtered using 0.3mm Al. (a) Original reconstruction (b) rescaled to emphasise cupping artefact due to beam hardening. (c) Convex segmentation used to designate rock material and (d) its resulting correction. (e) Tighter fit segmentation and (f) its resulting correction

5. CORRECTION OF A SAMPLE IN A CONTAINER

5.1 Correction Assuming Two Materials: Sample and Container

We now consider simulated data representing the case of a quartz rock sample in a highly attenuating titanium container. Figure 2(a) shows a simulation of this scenario, created by adding together multiple reprojections at energies ranging from 20 keV up to 100 keV using silicon and titanium attenuation coefficients obtained from NIST.[18] The specific set of X-ray energies used was {20, 30, 40, 50, 60, 80, 100} keV with a corresponding, normalised spectral-weighting as follows: {0.24, 0.30, 0.18, 0.12, 0.06, 0.06, 0.04}. We see in Figure 2(a) that the beam hardening has caused a bright halo effect around the inside edge of the container. This halo poses a challenge when segmenting the air surrounding the rock, especially if it was to be attempted automatically. The container itself is straightforward to segment, for example by using an automatic k-means clustering on the intensity values. To avoid a segmentation on the interior of the container, we consider it as a single phase. To be clear, we have segmented the image into two regions; the container is the first phase and everything inside the container is grouped and treated as the other phase. The KKF method is then applied and the subsequent correction can be seen in Figure 2(b). The intensity line plot in Figure 2(c) demonstrates that the cupping artefact has been removed in the container phase, signified by a flat line for the grayvalues in that region.

Figure 2(d) shows a rescaled intensity line plot which demonstrates that the cupping in the rock sample itself has been reduced towards the centre, but is still present closer to the container. What's more, this correction does not necessarily make the segmentation of the air any easier since the bright halo around the inside edge of the container is still present. An attempted segmentation of the air could be made and then improved upon by iterative applications of the technique, but each segmentation reinforces itself in the subsequent reconstruction,

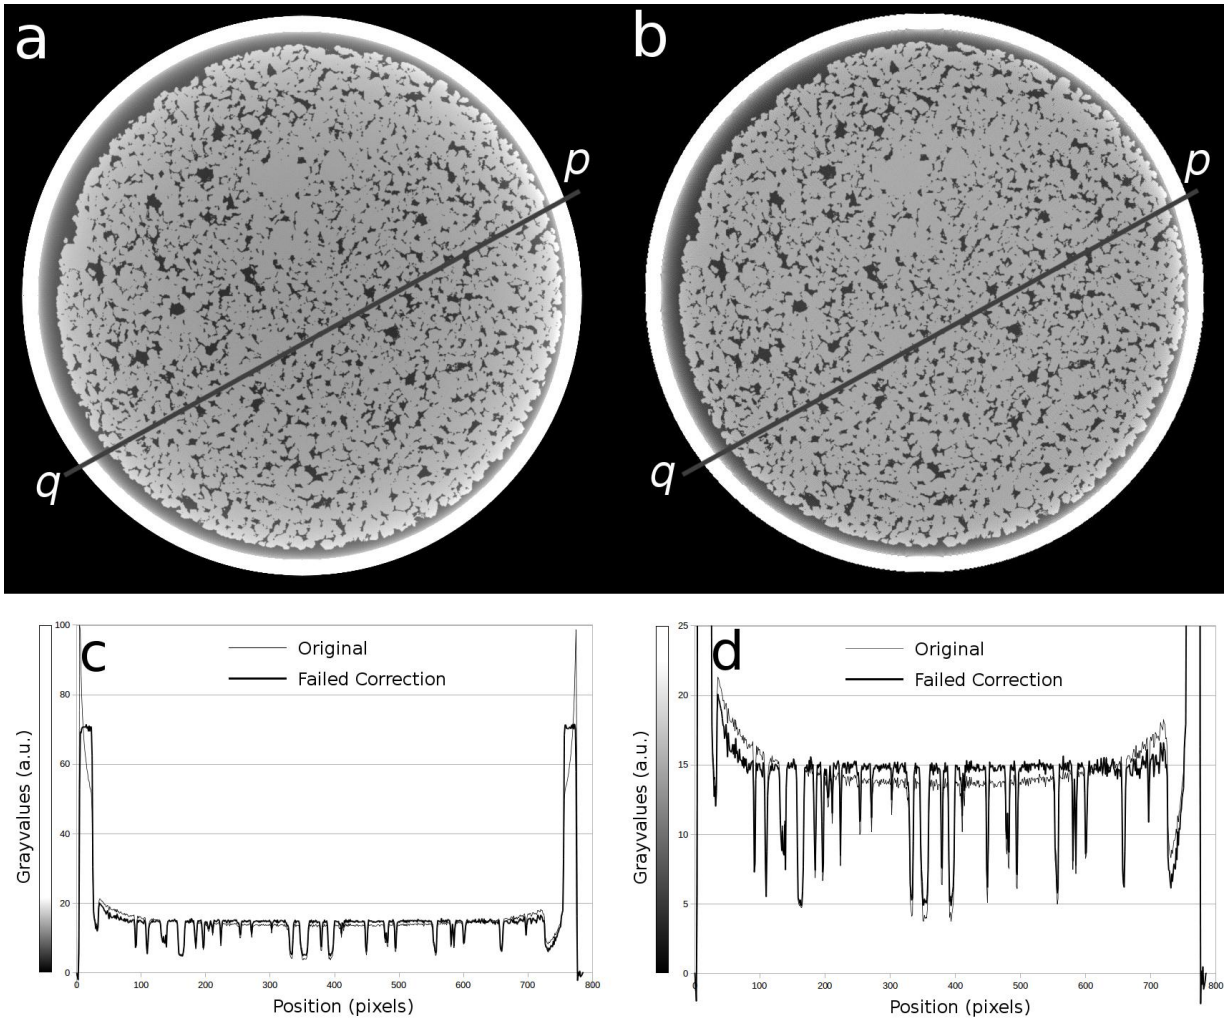

Figure 2. Cross section of a simulated silicon sample in a titanium container. (a) The original reconstructed image with grayvalues rescaled to show the cupping artefact in the sample. (b) Failed two material correction attempt, with the inside of the container treated as a single material. (c) Grayscale intensity line plot showing the full range of values, (d) zoomed in on the area of interest.

and so this iterative method does not guarantee that the segmentation is accurate. In the next section we develop a modification of the KKF method, similar to that performed by Van de Casteele *et al*[9] and particularly Paziresh *et al.*[19] It requires two correction stages, the first of which removes the inner halo artefact caused by the container.

5.2 Container Precorrection Followed by Sample Correction

5.2.1 Container-only Correction

Instead of treating the interior of the container as a single phase, we consider the correction which would be applied if only the container was present, *i.e.*, the same container filled with air rather than a sample. To do this we must approximate what the attenuation would have been if an empty container was imaged by the same polychromatic beam. We therefore have to ignore x-rays which pass through the centre of the container, since they could have been hardened by the sample and will not be representative of an empty container.

Now, if the inside wall of the container is convex, like in the case of a cylinder, then every possible path length through the container, is captured by at least one ray that does not pass through the interior of the container. To put another way, grazing rays can probe all possible path lengths. Therefore, from this container-only data, one can build a complete correction curve for the container material:

$$C_{estimated}(\vec{v},\vec{x}) = f(\mu_c C_{seg}(\vec{v},\vec{x})) \tag{7}$$

Where C is the attenuation due to the container and μ_c is the attenuation coefficient of the container approximated by fitting a straight line through the same set of rays. If we take the difference of these two functions, we get a correction as a function of container path length to apply to the rest of the rays in the projection.

$$\begin{aligned} R_{precorrected} &= R_{exp}(\vec{v},\vec{x}) - \mu_c C_{seg}(\vec{v},\vec{x}) + C_{estimated}(\vec{v},\vec{x}) \\ &= R_{\epsilon}(\vec{v},\vec{x}) + C_{exp}(\vec{v},\vec{x}) - \mu_c C_{seg}(\vec{v},\vec{x}) + C_{estimated}(\vec{v},\vec{x}) \\ &= R_{\epsilon}(\vec{v},\vec{x}) + C_{corr}(\vec{v},\vec{x}) \end{aligned}$$

Where R_ϵ is the attenuation due to material which is not the container, which we note has not been corrected.

To apply this principle, we segment the container and inside in the same way as in Section 5.1 which generates the same point cloud, shown in Figure 3. We ignore the grey data points, which represent rays which have not exclusively passed through the container, and fit a 1 dimensional curve to find $C_{estimated}$ and straight line through the black data points to find $\mu_c C_{seg}$.

We see the result of this correction in Figure 4(a). There are a number of things to note; most importantly that the bright halo around the inside of the container has been removed entirely, allowing for a very straightforward segmentation of the air phase surrounding the rock. There are two interesting aspects of the intensity line profile which we need to explain, the overall drop in intensity and the new *reverse cupping* effect in the rock, seen in our rescaled intensity line plot in Figure 4(c). The overall intensity drop is not significant, and is a result of the correction assuming that there was no attenuation within the container. The reverse cupping can be explained with respect to the implicit assumption we have made by ignoring rays which pass through the inside of the container, which is that the beam hardening occurs only in the container. In reality there is beam hardening occurring in the rock sample as well, which manifests itself in two artefacts. Firstly there is the traditional cupping artefact which causes a darker region in the centre because of the varying amount of rock passed through. Secondly, there is a reverse cupping artefact which causes the centre region to be brighter than the edges. This can be explained by considering rays *pq* and *pr* in Figure 4(a). Inner regions of the sample are probed by only rays like *pq* which, because of their angle of entry, have passed through comparatively less container than rays like *pr*, which probe the outer regions. Therefore the outer regions appear less dense since they are probed by harder beams on average. The relative impact of these two beam hardening artefacts is dependent on the thickness of the container, its attenuation relative to the sample and the fan angle of the

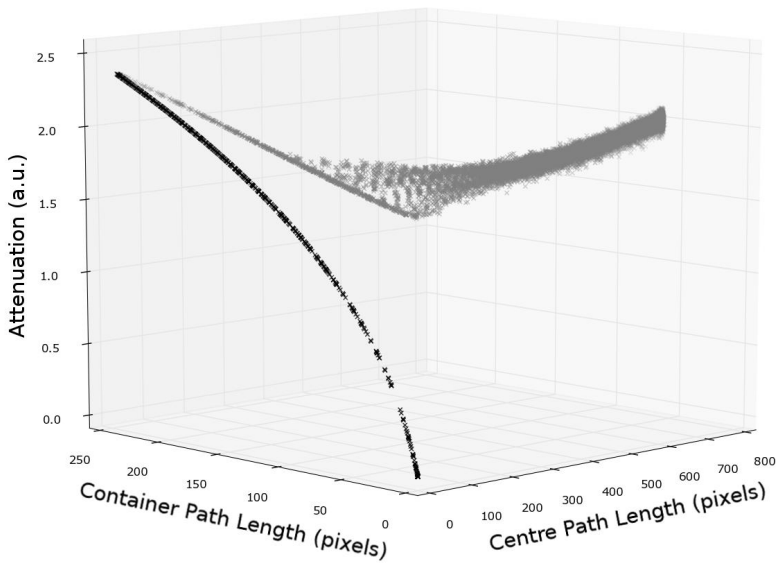

Figure 3. Three dimensional point cloud used in the correction process. The black data points indicate the region fitted for the container-only correction technique

incoming beam. In this case the container has a much higher attenuation compared to the sample and so the reverse cupping effect dominates.

It is worth noting the relationship between this technique and the use of filters during acquisition discussed in Section 2. The attenuation of an x-ray is defined as $R = ln(I/I_0)$ where I is the intensity of a beam with the sample present, and I_0 is the intensity of the beam with no sample present, known as the *clear field*. When a filter is used, its direct effect on the intensity of the beam is removed by the clear field, since identical filtering is present for I and I_0. However, the indirect effect on intensity by altering the beam spectrum, *i.e.*, hardening the beam, is still present.

In a fan beam arrangement with a flat sheet of filtering, rays of higher angles will pass through more material and so the reverse cupping artefact described for the container correction is also present for pre-filtering. The degree to which the cupping and reverse cupping effects influence the sample is again dependent on the attenuation, thickness and fan angle. The reverse cupping effect is not as dramatic for a flat sheet as it is for the container correction because the ray path lengths show less variation for flat filters. Figure 5 depcits these path length variations. So the container correction technique described is analogous to applying filtering and removing the effect with a clear field. The difference is that if the filtering occurs in the field of view, *e.g.*, if we 'filter' with a container, the beam hardening can be corrected for.

5.2.2 Final Sample Correction

The pre-corrected image can now be segmented using the approach from Section 4 such that the rock sample is considered as a single material. This was done automatically by applying a k-means clustering of the intensities and then excluding all regions not connected to the previously segmented container region. The automatic segmentation worked well in this case, but more sophisticated segmentation routines could easily be adopted if required. We can now apply the referenceless post-reconstruction correction algorithm to our newly segmented image. The resulting correction has removed the beam hardening artefacts almost entirely, as demonstrated in Figure 4(b) and intensity line plot in Figure 4(d). We do note the slight drop in intensity seen at the very edges of the intensity line plot, which is explained as in Section 4.

5.3 High Density Inclusions

We have considered so far the case of a fairly uniformly heterogeneous sample which can be treated as one material. However sometimes there are materials within the sample which need to be segmented because of the

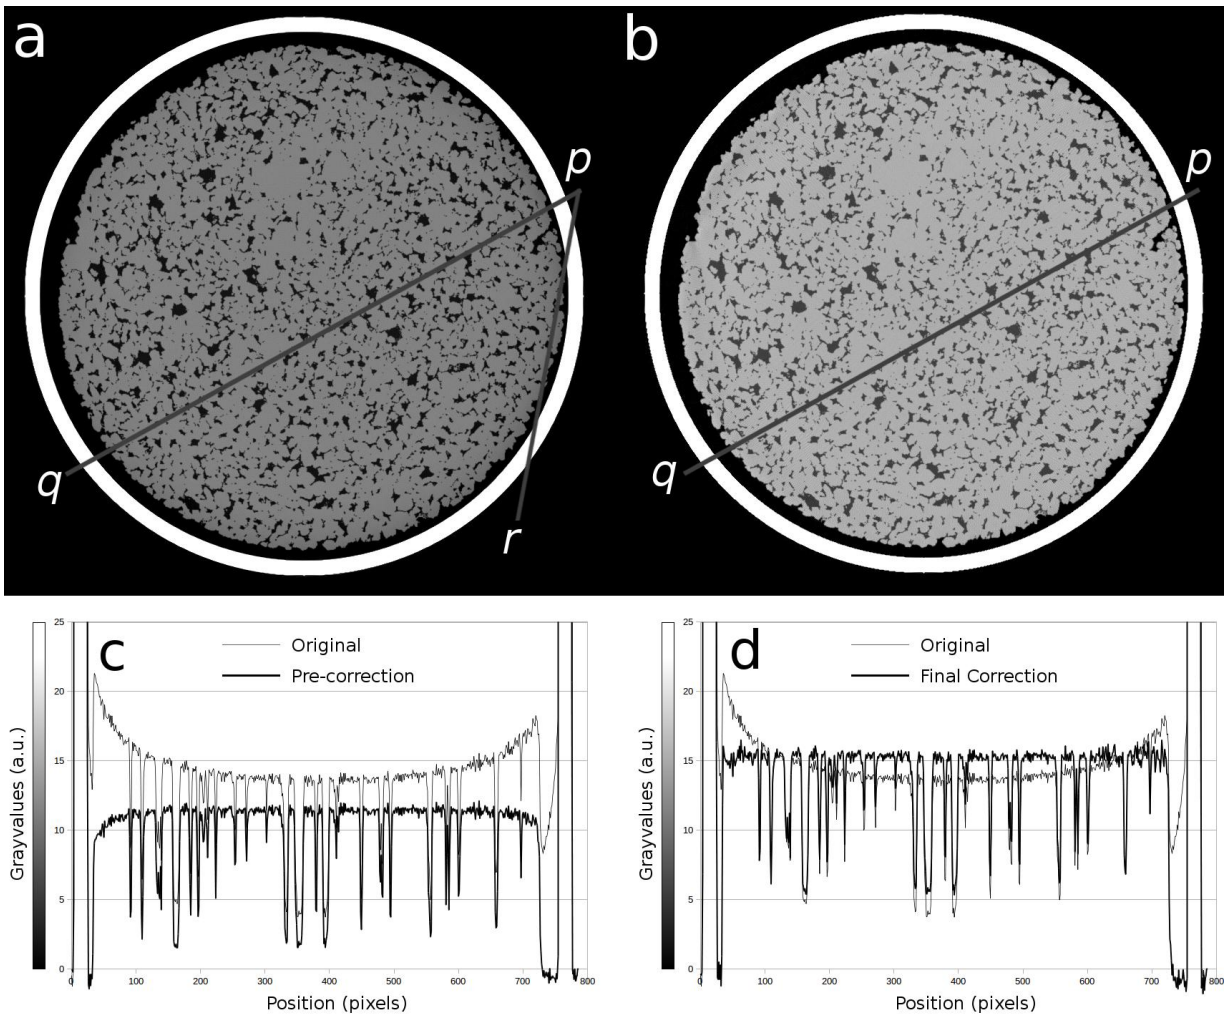

Figure 4. Cross section of corrections applied to the sample in Figure 2(a), again with the contrast rescaled to show the effects in the sample. (a) The resultant reconstruction after only the beam hardening in the container is removed. (b) The final corrected reconstruction after a segmentation of the air phase has been performed. (c) The grayscale intensity line plots of the original and pre-corrected image, exhibiting the reverse cupping effect. (d) The final correction compared to the original image. Note that (c) and (d) have been zoomed in to the sample region and we therefore cannot see the flattening effect on the grayvalues in the container. However the effect in the container region is very similar to that shown in Figure 2(c)

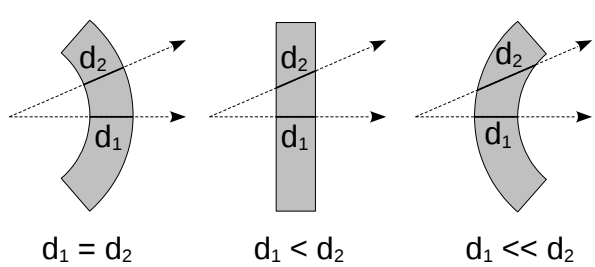

Figure 5. Schematic demonstrating the relative path lengths for x-rays with zero fan angle (d_1) and non zero fan angle (d_2). (a) A spherical filter results in all beams passing through the same amount of filter. (b) A flat sheet filter results in longer path lengths for greater fan angles. (c) A container will result in even longer path lengths for greater fan angles.

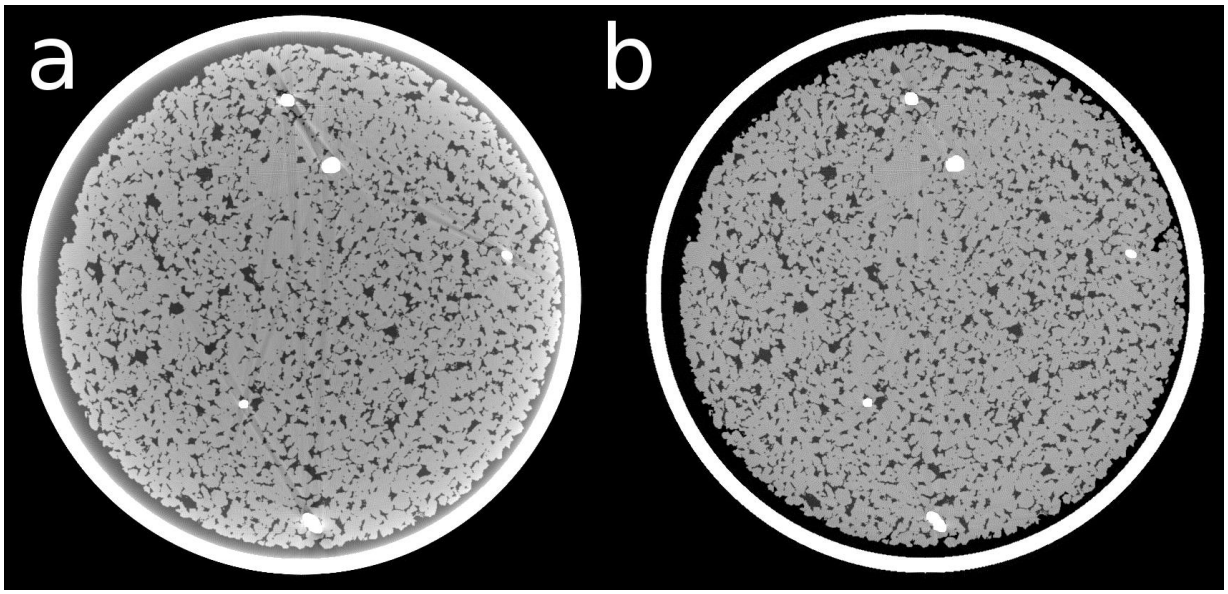

Figure 6. Cross section of a simulated silicon sample in titanium container with highly attenuating iron inclusions. (a) Original image, rescaled to show the effect within the sample. Note the streaking artefacts characterised by dark regions between the bright spots. (b) The final corrected image, using a three material correction, with the streaking artefacts almost completely removed.

obscuring effect they have on the rest of the image, regardless of the drawback of a change in relative contrast between phases. This is especially true of highly attenuating materials which cause streaking artefacts when imaged, for example high density pyrite inclusions in a sandstone. This has been simulated by including 5 iron inclusions in our simulation, which cause significant streaking effects as seen in Figure 6)(a). The high attenuation of the inclusions make them straightforward to separate from the rock matrix, again using a simple k-means segmentation. This sample is now treated as three materials; the container, the rock material segmented after the pre-correction discussed in Section 5.2.1 and the high density inclusions. Applying the correction method with these segmentations almost entirely removes the streaking, as seen in Figure 6(b).

6. CORRECTION SPACE: INTENSITY OR ATTENUATION

As mentioned in Section 3, the KKF correction was applied in attenuation space (R-space), rather than intensity space (I-space) as was done in the paper by Krumm *et al.*[11] If the sample is segmented into its component materials perfectly then the choice of correction space does not matter and both corrections will produce the same result. The major difference between the two methods is the effect the slope of the hyperplane fit has on the correction. In both cases, the slope in each direction determines the average grayscale value for each phase. For the R-space correction, the contrast within a phase remains constant for any slope, *i.e.*, if the pore space is grouped in the same phase as the rock then the average absolute grayscale difference between the rock and the pores will remain approximately constant. In the case of the I-space correction, the contrast within a phase will change to match the contrast change between phases *i.e.*, the overall consistency of contrast is maintained. However this also means there is a 'correct' slope for the R_{mono} fit and if any other plane is fit the results will be degraded, especially with regard to contrast changes within a phase. This means that the R-space correction is more robust to non-accurate segmentations, but comes with the disadvantage (or potential advantage) of the contrast within a phase and between phases being independent. In this paper where we have grouped multiple materials into a single phase and are not concerned about the relative contrast between the container and the sample, the R-space correction was more appropriate to ensure there was no contrast change within the sample phase. However in other cases, such as when all the materials are segmented, the I-space correction potentially has more fidelity to the original relative contrast values of the entire image.

For the pre-correction case in which we correct only the container, we're required to apply the correction in

the R-space. This is because we ignore the R values of rays which do not pass exclusively through the container when finding R_{mono} and R_{poly}, and so when we apply this correction to the rest of the rays there is no guarantee that our correction equation will be defined. In I-space, the correction equation for each ray is analogous to Equation 6:

$$I_{corr}(\vec{v},\vec{x})/I_0 = -\ln(e^{R_{exp}(\vec{v},\vec{x})} + e^{f(R_1(\vec{v},\vec{x}),R_2(\vec{v},\vec{x}),\ldots)} - e^{\sum \mu_i R_i(\vec{v},\vec{x})}) \tag{8}$$

It is possible that for certain rays the logarithm is undefined. This happened especially when there were high density inclusions within the sample, so $e^{-R_{exp}}$ was very small for rays which passed through the high density regions.

7. CONCLUSION

We have shown that the referenceless post-reconstruction correction technique can be effectively applied to reduce beam hardening artefacts of cylindrical samples in containers. It should be noted that, being referenceless, this method cannot extract physically meaningful attenuation values, but serves to enable the identification/segmentation of the various materials present in the sample. We have only considered approximately cylindrical samples here, however, there is no restriction on sample geometry, the method should work equally well on non-cylindrical samples in containers. It was demonstrated that certain materials could be treated as a single phase and still produce good artefact reduction. The correction requires an accurate segmentation of all other phases, which was facilitated by performing an initial pre-correction on the image by correcting the beam hardening exclusively in the container. As such, we have introduced the building blocks for a work-flow which can automatically correct beam hardening artefacts in cylindrical samples.

ACKNOWLEDGMENTS

The authors wish to thank Michael Turner and Andrew Fogden for their assistance and advice in sample preparation and μCT imaging. A.P.S. acknowledges support of a Future Fellowship from the Australian Research Council. A.M.K. acknowledges funding through the DigiCore consortium.

REFERENCES

[1] Ketcham, R. A. and Carlson, W. D., "Acquisition, optimization and interpretation of x-ray computed tomographic imagery: applications to the geosciences," *Computers & Geosciences* **27**(4), 381–400 (2001).

[2] Arns, C. H., Bauget, F., Limaye, A., Sakellariou, A., Senden, T., Sheppard, A., Sok, R. M., Pinczewski, W., Bakke, S., Berge, L. I., et al., "Pore-scale characterization of carbonates using x-ray microtomography," *SPE JOURNAL-RICHARDSON-* **10**(4), 475 (2005).

[3] Wang, Y., Uchida, T., Westferro, F., Rivers, M. L., Nishiyama, N., Gebhardt, J., Lesher, C. E., and Sutton, S. R., "High-pressure x-ray tomography microscope: Synchrotron computed microtomography at high pressure and temperature," *Review of scientific instruments* **76**(7), 073709 (2005).

[4] Lesher, C. E., Wang, Y., Gaudio, S., Clark, A., Nishiyama, N., and Rivers, M., "Volumetric properties of magnesium silicate glasses and supercooled liquid at high pressure by x-ray microtomography," *Physics of the Earth and Planetary Interiors* **174**(1), 292–301 (2009).

[5] Wildenschild, D., Vaz, C., Rivers, M., Rikard, D., and Christensen, B., "Using x-ray computed tomography in hydrology: systems, resolutions, and limitations," *Journal of Hydrology* **267**(3), 285–297 (2002).

[6] Cnudde, V., Cnudde, J., Dupuis, C., and Jacobs, P., "X-ray micro-ct used for the localization of water repellents and consolidants inside natural building stones," *Materials characterization* **53**(2), 259–271 (2004).

[7] Varslot, T., Kingston, A., Myers, G., and Sheppard, A., "Considerations for high-magnification high-cone-angle helical micro-ct," *Proc. SPIE Developments in X-ray Tomography VIII* **8506**, 850614:1–10 (2012).

[8] Herman, G. T., "Correction for beam hardening in computed tomography," *Physics in medicine and biology* **24**(1), 81 (1979).

[9] Van de Casteele, E., Van Dyck, D., Sijbers, J., and Raman, E., "A model-based correction method for beam hardening artefacts in x-ray microtomography," *Journal of X-ray Science and Technology* **12**(1), 43–57 (2004).

[10] Kingston, A., Myers, G., and Varslot, T., "X-ray beam hardening correction by minimizing reprojection distance," *Proc. SPIE Developments in X-ray Tomography VIII* **8506**, 85061D:1–10 (2012).

[11] Krumm, M., Kasperl, S., and Franz, M., "Reducing non-linear artifacts of multi-material objects in industrial 3d computed tomography," *NDT & E International* **41**(4), 242–251 (2008).

[12] Taina, I., Heck, R., and Elliot, T., "Application of x-ray computed tomography to soil science: A literature review," *Canadian Journal of Soil Science* **88**(1), 1–19 (2008).

[13] Alvarez, R. and Macovski, A., "Energy selective reconstructions in x-ray computerized tomography," *Phys. Med. Biol.* **21**, 733–744 (1976).

[14] Van Geet, M., Swennen, R., and Wevers, M., "Quantitative analysis of reservoir rocks by microfocus x-ray computerised tomography," *Sedimentary Geology* **132**(1), 25–36 (2000).

[15] Nalcioglu, O. and Lou, R., "Post-reconstruction method for beam hardening in computerised tomography," *Physics in medicine and biology* **24**(2), 330 (1979).

[16] Beucher, S. and Meyer, F., "The morphological approach to segmentation: the watershed transformation," *Mathematical Morphology in Image Processing (Ed. E. R. Dougherty)* , 433–481 (1993).

[17] MacQueen, J. B., "Some methods for classification and analysis of multivariate observations," *Proc. 5th Berkeley Symposium on Mathematical Statistics and Probability* , 281–297 (1967).

[18] Hubbell, J. H. and Seltzer, S. M., "Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients," *National Institute of Standards and Technology* (1996).

[19] Paziresh, M., Kingston, A., Myers, G., and Latham, S., "Software x-ray beam hardening correction of cylindrical specimens," *Proc. Int. Conf. Tomography of Materials and Structures (ICTMS)* , 187–191 (2013).