

University of Southampton
Faculty of Engineering and Physical Sciences
Department of Mechanical Engineering

**Development of correlative neutron
and X-ray computed tomography to
study fluid dynamics and structural
deformation at the micro-scale in
plant and soil systems**

DOI: 10.5358/SOTON/T0033

Thomas Clark
ORCID: 0000-0001-6276-7061

March 2022

Thesis for the degree of Doctor of Philosophy
in
Engineering & the Environment

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
Department of Mechanical Engineering

Doctor of Philosophy

Thomas Clark

The interactions between plant roots and soil are an area of active research, particularly in terms of water and nutrient uptake. Since non-invasive, *in vivo* studies are required, tomographic imaging appears an obvious method to use, but no one imaging modality is well suited to capture the complete system. X-ray imaging gives clear insight into soil structure and composition. However, water is very weakly attenuating to X-rays, and biological matter also displays poor contrast. Neutron imaging presents a complementary view where water and biological matter are better distinguished, but the soil minerals are imaged with inferior contrast and resolution in comparison to equivalent X-rays scans. This work aims to develop robust methods for complementary X-ray/neutron tomographic imaging of plant root samples. These should lead to new insight into water and nutrient transport in soil. The primary challenges of this project are: to develop experiments that will meet the requirements of both imaging modalities and the biological requirements of the plant samples and to develop ways to register a pair of reconstructed volume images of samples that have been produced at entirely separate facilities.

This work investigates the use of fiducial markers for point-based registration concerning the material, number and distribution of markers to address the

registration challenge, first with simulation and then experimentally with plant samples imaged using neutrons and X-rays. The neutron scans were collected at the IMAT instrument at ISIS Neutron and Muon Source and the X-ray scans both at X-ray Imaging Centre at the University of Southampton and the I12 beamline at Diamond Light Source. A marker segmentation algorithm designed to automate the registration process is presented and evaluated, as are methods for combining the registered data from the two modalities to optimise the technique and facilitate segmentation, quantification and further analysis.

Contents

Declaration of Authorship	19
Acknowledgements	21
1. Introduction	25
1.1. Introduction	25
1.2. Aims and objectives	26
2. Literature review	29
2.1. X-ray and neutron imaging	29
2.1.1. X-ray attenuation mechanisms	29
2.1.1.1. Photoelectric absorption	30
2.1.1.2. Rayleigh scattering	30
2.1.1.3. Compton scattering	30
2.1.1.4. Pair production	30
2.1.2. X-ray imaging systems	34
2.1.2.1. X-ray sources	34
2.1.2.2. X-ray detectors	37
2.1.3. Neutron attenuation mechanisms	37
2.1.4. Neutron imaging systems	38
2.1.4.1. Neutron sources	38
2.1.4.2. Neutron imaging detectors	40
2.2. Tomographic imaging techniques	41
2.2.1. Time-lapse CT and golden ratio acquisition	43
2.2.2. Artefacts	44
2.2.2.1. Ring artefacts	44
2.2.2.2. Beam hardening	46
2.2.2.3. Scatter	46
2.2.2.4. Noise	47
2.2.2.5. Starvation	47

2.2.2.6. Undersampling	47
2.2.3. Multi-modal imaging	47
2.3. Root and soil imaging	50
2.3.1. Soil	50
2.3.2. Roots	52
2.3.3. Plant root and soil imaging	54
2.3.3.1. X-ray	54
2.3.3.2. Neutron	54
2.3.3.3. Correlative	59
3. Distribution of fiducial markers to optimise the registration of multi-modal CT data	61
3.1. Introduction	61
3.2. Simulation	64
3.3. Materials and methods	71
3.3.1. X-ray and neutron imaging equipment	71
3.3.1.1. IMAT	71
3.3.1.2. μ -VIS HMX	72
3.3.2. Cadmium imaging and registration experiment	72
3.4. Results	76
3.5. Discussion and conclusions	78
4. Borosilicate as a fiducial marker material and a sphere fitting algorithm for fiducial localisation	81
4.1. Introduction	81
4.2. Materials and methods	82
4.2.1. Materials survey	82
4.2.2. I12	85
4.2.3. Borosilicate trials	86
4.3. Image processing	93
4.3.1. Stack concatenation	97
4.3.2. Sphere fitting	98
4.4. Results	101
4.5. Discussion	106
4.6. Conclusions and future work	115

5. Correlative X-ray and neutron imaging for water segmentation and quantification	117
5.1. Introduction	117
5.2. Equipment and methodology	118
5.2.1. Preliminary tests	118
5.2.2. Experimental methods	123
5.3. Segmentation and quantification of water in soil and sand . . .	124
5.3.1. Time series neutron radiography	124
5.3.2. Beam hardening correction	126
5.3.3. Independent processing of complementary modalities .	129
5.3.4. Joint processing of complementary modalities	139
5.3.5. Quantification	146
5.4. Conclusions and future work	149
6. Final conclusions and future work	153
A. Code Listings	157
A.1. Marker simulation	157
A.1.1. fidtest.py	157
A.1.2. fidlib.py	160
A.2. Stack concatenation	164
A.3. Sphere fitting	169
A.3.1. test.py	169
A.3.2. Volume.py	178
A.3.3. SphereBuilder.py	182
A.3.4. CircleHunter.py	185
A.4. Masked reconstruction simulation	188
B. Original figures	193
Bibliography	201

List of Figures

2.1. Photoelectric absorption of an incident X-ray photon. The incident electron transfers its energy to the photoelectron, ejecting it with kinetic energy equal to the difference in energy between the incident photon and the electron binding energy. An electron from the outer shell moves to fill the vacancy in the K shell and the difference in binding energies is consequently released as a characteristic X-ray.	31
2.2. The mass attenuation and mass energy-absorption coefficients of germanium as a function of photon energy. The K-edge can be seen at 1.11×10^{-2} keV. The L-edges occur between 1.216×10^{-3} keV and 1.414×10^{-3} keV. Reproduced from NIST X-Ray Mass Attenuation Coefficients Database [Hubbell and Seltzer, 2004].	32
2.3. Rayleigh and Compton scattering diagram. The top case shows Rayleigh scattering where the photon is absorbed by the electron and another photon is emitted with equal energy but also with a slight change in direction. The lower case shows Compton scattering where the absorbed photon provides enough energy to free an electron and only the remainder of the energy is transferred to the scattered photon.	33
2.4. How the types of photon interaction vary with photon energy and atomic number. Based on figure from Podgorsak et al. [2005].	34
2.5. A typical bremsstrahlung and characteristic radiation spectrum for a tube source. The maximum photon energy is 150 keV which shows that the tube voltage must be 150 kV. A steep decline in intensity can be seen below about 40 keV as the X-rays produced at these low energies are quickly absorbed. Characteristic radiation peaks can be seen in the 50-80 keV range. Based on figure from [PhysicsOpenLab, 2017].	35

List of Figures

2.6. A diagram showing how golden angle based aquisition progressively fills the scan arc with an infinite series of projections. The top figure shows the first four angles in the sequence and the bottom figure shows the first seventeen.	45
2.7. The external structure of a root.	53
2.8. An example slice from [Li and Tang, 2019] showing the contrast between sand, air and water in a high-resolution synchrotron image.	55
2.9. A tomogram of a 12 day old chickpea plant grown in a cylinder of sandy soil and presented in [Moradi et al., 2011]. Sub-figures b, c and d show slices at depths of 4, 6 and 8 cm respectively. .	57
2.10. Time resolved neutron tomography of a lupin root system after the injection of 4 ml of deuterated water (D_2O) from the bottom. Reproduced from [Totzke et al., 2017].	58
3.1. Diagrams showing the three types of registration errors: fiducial localisation error (FLE), fiducial registration error (FRE) and target registration error (TRE). Based on Figure 1 of [West et al., 2001].	64
3.2. Potential marker locations for the simulation to test.	65
3.3. The test points used to evalutate the target registration error througout the volume and give an estimate of TRE_M	66
3.4. The change in target registration error as the number of fiducial points is increased.	66
3.5. The average contribution of a single marker to the target registration accuracy as the number of markers increases.	67
3.6. The optimal fiducial marker distributions found by the simulation for $N = 10$ (top) and $N = 16$ (bottom).	68
3.7. A comparison of the optimal marker distributions to the mean, minimum and maximum TRE_M values found in 1500000 random trials. (A) (Top) the full range of N values. (B) (Bottom) $N = 12 - 16$	70
3.8. A high-level IMAT imaging system schematic.	72
3.9. Diagram illustrating the intended marker layouts for samples 1, 5 and 7.	74

3.10. Four quartz sample tubes with cadmium fiducial markers attached by aluminium tape. The green and white tape marks the bottom of the 60 mm IMAT field of view to aid in sample alignment. From left to right, the marker configurations are 1, 3, 5, 7.	75
3.11. A slice and partial volume render from the neutron scan of sample 7.	76
3.12. A slice and partial volume render from the X-ray scan of sample 8.	77
3.13. Sample 5 had the worst registration accuracy despite having the most fiducial markers. This figure shows how an artefact prevents accurate segmentation (segmentation shown in red) of a cadmium piece, creating a large FLE for the corresponding fiducial marker.	78
3.14. A slice from sample 8 with the X-ray data on the left and in red and the neutron data on the right and in green. This slice shows the match of a fiducial marker, the accuracy of the registration and the complementarity of the modalities.	79
4.1. A demonstration of the intended marker distribution. The markers are located in arcs near the top and bottom of the sample with minor variations in their height to ensure that they do not contribute attenuation in the same slices.	87
4.2. A sample positioned in front of the detector at I12.	88
4.3. The experimental setup at IMAT. Two samples are placed in front of the detector on the multi-axial rotation stage.	89
4.4. Attempts to reconstruct a slice from Small Soil 5 with three different centre of rotation values. In (a), where the centre of rotation is 205, the marker at the top of the sample is well resolved - besides the artefacts - however, the tape around the outside of the tube and the inner radius of the sample tube are not. Once the centre of rotation has reached 209 in (c), the inner edge of the tube is reconstructed well with precise edges, but the marker is not and has split into two.	92

List of Figures

4.5. The slice from Small Soil 5 shown in figure 4.4, this time reconstructed assuming a total rotation of 367.031°. The whole sample appears to be correctly reconstructed with the new parameters.	96
4.6. The cross-section of a sphere is always circular, regardless of orientation.	99
4.7. The images show the results of the Hough transform applied to the same slice before (left) and after (right) the application of a median filter with a kernel size of 7. Each circle found by the transform is marked with a red centre and a green circumference.	100
4.8. Slices from the X-ray (left) and neutron (pre-injection) (right) scans of sample LargeSoil3 showing one of the fiducial markers. The marker is easily distinguishable in each modality, primarily by shape in the X-ray image but also by the strength of attenuation in the neutron modality due to the boron dopant. It can be seen that the types of artefact present around the markers in figures 3.11, 3.12, 3.13 and 3.14 have been eliminated.	101
4.9. An example slice from the X-ray scan of sample LargeSoil4 shows markers that have not been kept within the field of view and cannot accurately be located and used in registration or subsequent analysis.	103
4.10. Two markers placed side by side on a sample. While the two markers are easily distinguished in the X-ray image (left), the scattering artefacts make it difficult for them to be separated by the DBSCAN algorithm in the neutron data (right). Consequently, these two markers were found by hand for the registration but discarded when the algorithm was rerun for analysis of the registration.	108
4.11. Two nominally matching slices from different scans of sample SmallSoil3 overlaid. The pre-injection scan is coloured red, and the post-injection scan is cyan. The difference in the marker positions between the two scans suggests that the sample has been rotated around a point close to the centre of the sample.	110
4.12. FRE between the two neutron scans plotted against the FLE.	111
4.13. The FRE between each pair of scans.	112

4.14. Average FRE across all scans for each number of markers used in the registration.	113
5.1. A series of radiographs showing the progressive injection of 1 ml of water into the soil at a rate of 0.05 ml/min. From left to right these images show the sample at the beginning of the injection, then after 8, 12 and 20 minutes respectively (to the nearest minute).	120
5.2. The difference between two radiographs showing the change in water and deuterated water between the two acquisitions. The left sample shows water injected into dry soil, and the right shows deuterated water injected into a wet soil sample, displacing water further down the tube. The dark regions show the water, and the light represent the deuterated water.	121
5.3. The first and last radiographs taken of LargeSoil4 after deuterated water was injected into the sample. (a) The radiograph taken immediately after the injection of deuterated water. (b) The final radiograph, taken 2 hours and 24 minutes after the injection. (c) The difference between the first and last radiographs.	125
5.4. The ROI from the summed difference projections for the sample LargeDrySoil.	127
5.5. The attenuation from figure 5.4 plotted against path length. . .	128
5.6. Neutron attenuation through the fluid in the SmallWetSand sample as a function of path length.	129
5.7. Slice 4402 from the sample LargeWetSand.	131
5.8. The absolute difference between the pre- and post-injection images of slice 4402.	132
5.9. Slice 8027 from the sample LargeWetSand.	133
5.10. Taking the difference between the pre- and post-injection images for slice 8027 of LargeWetSand. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.1.	134

List of Figures

5.11. X-ray data for slice 4402 of LargeWetSand. 5.11a shows the data after registration. It is converted to a binary image using a threshold to produce 5.11b, which can then be utilised for further analysis.	135
5.12. The difference (before and after the injection of deuterated water) image of slice 4402 shown before (a) and after (b) a mask is applied.	136
5.13. The difference (before and after the injection of deuterated water) image of slice 8027 shown before (a) and after (b) a mask is applied. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.2.	137
5.14. A demonstration of the impact of poor registration. These slices were taken from sample LargeDrySand where the fiducial registration error between the post-injection neutron scan and the other scans was approximately 1 mm.	138
5.15. Slice 1851 from the sample SmallSoil4.	139
5.16. The difference between slice 1851 of the pre- and post-injection neutron scans of SmallSoil4. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.3.	140
5.17. The difference between the pre- and post-injection scans of slice 1851 from SmallSoil4 with an X-ray mask overlaid. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.4.	141
5.18. The multi-modal data processing approach applied to this data set. The processing from sample to registered tomograms is covered in section 4.3. This chapter presents all subsequent processing from that step.	143
5.19. The simulated slice: wet, dry and binarised	144
5.20. The reconstructed slices after masking. The left image shows the benchmark result where the simulated pre- and post-injection data was reconstructed separately, and the mask applied to the difference between the reconstructions. The right image shows the new method where the difference data was reconstructed with iterative masking.	145

5.21. The error maps of the reconstructed slices after masking. The left image shows the benchmark result where the simulated pre- and post-injection data was reconstructed separately, and the mask applied to the difference between the reconstructions. The right image shows the error produced with the new method where the difference data was reconstructed with iterative masking. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.5.	145
5.22. Slice 4402 from the sample LargeWetSand reconstructed used the SIRT algorithm with an X-ray mask. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.6.	146
5.23. The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with an X-ray mask. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.7. . .	147
5.24. The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with the X-ray mask applied afterwards. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.8.	148
5.25. A calibrated slice (4724) with a calibration bar to show the relationship between the amount of fluid in a voxel (in picolitres) and the grey level.	150
B.1. Taking the difference between the pre- and post-injection images for slice 8027 of LargeWetSand.	194
B.2. The difference (before and after the injection of deuterated water) image of slice 8027 shown before (a) and after (b) a mask is applied.	195
B.3. The difference between slice 1851 of the pre- and post-injection neutron scans of SmallSoil4.	196
B.4. The difference between the pre- and post-injection scans of slice 1851 from SmallSoil4 with an X-ray mask overlaid.	197

List of Figures

B.5. The error maps of the reconstructed slices after masking. The left image shows the benchmark result where the simulated pre- and post-injection data was reconstructed separately, and the mask applied to the difference between the reconstructions. The right image shows the error produced with the new method where the difference data was reconstructed with iterative masking.	198
B.6. Slice 4402 from the sample LargeWetSand reconstructed used the SIRT algorithm with an X-ray mask.	198
B.7. The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with an X-ray mask.	199
B.8. The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with the X-ray mask applied afterwards.	200

List of Tables

2.1. Soil types as defined by particle diameter, reproduced from [Taiz and Zeiger, 2002].	51
3.1. The fiducial marker arrangements. FRE measurements and TRE_M values for each of the sample tubes that were scanned, as well as their products which are taken as an indicator of TRE. A representative selection of these are illustrated by a diagram in figure 3.9 and a photo of some of the prepared sample tubes is shown in figure 3.10.	74
4.1. Neutron and X-ray transmission values for a range of materials that were considered for markers. All the transmission values were based on a 1 mm path length. Only transmission is calculated for borosilicate because it is a mixture not a single element or compound.	83
4.2. The detector properties for each imaging camera and optics module at I12 in EH1. For EH2, the field of view for M1 and M2 are extended to 46×20 and 20×20 , respectively [Connolley et al., 2020].	85
4.3. Sample conditions. The particle size classifications are given in Table 4.3b.	86
4.4. The over-rotation measured for each neutron scan. The samples are split based on whether they were scanned on rotation stage B (4.4a) or rotation stage C (4.4b) of the multi-axial rotation stage.	94

List of Tables

4.5. The number of fiducial markers attached to each sample compared to the number of markers found both before and after registration by the sphere fitting algorithm. Any markers found in some but not all of the X-ray, pre-injection neutron and post-injection neutron scans were discounted from the last two columns as it is only markers found in all modalities that can contribute to the analysis shown subsequently.	104
4.6. The FLE of each sample based on the two neutron scans. The final three columns show the FLE components along each axis. (All values in voxels.)	105
4.7. The FRE values found between each scan for each sample. (All values in μm .)	106
4.8. Axial average components of FRE for each pair of scans. (All values in voxels.)	109

Declaration of Authorship

I, *Thomas Clark*, declare that this thesis titled *Development of correlative neutron and X-ray computed tomography to study fluid dynamics and structural deformation at the micro-scale in plant and soil systems* and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this university;
- Where any part of this thesis has been previously submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- Where I have consulted the published work of others, this is always clearly attributed;
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- Parts of this work have been published as:
 - Burca, G., Nagella, S., Clark, T., Tasev, D., Rahman, I.A., Garwood, R.J., Spencer, A.R.T., Turner, M.J. and Kelleher, J.F., 2018. Exploring the potential of neutron imaging for life sciences on IMAT. Journal of Microscopy, 272(3), pp.242-247.

Declaration of Authorship

- Clark, T., Burca, G., Boardman, R., Blumensath, T., 2020. Correlative X-ray and neutron tomography of root systems using cadmium fiducial markers. *Journal of Microscopy* 277(3), pp. 170-178.

Signed:

Date:

Acknowledgements

I want to thank my supervisors at the University of Southampton: Drs Thomas Blumensath and Richard Boardman. Your wise guidance, warm friendship and practical help have seen me through the twists and turns of this project.

Working across multiple sites with both internal and external supervision has been challenging at times, particularly as the scope of the project and the type of work has evolved. I am incredibly thankful for the fantastic support I have received from Dr Genoveva Burca as my supervisor at IMAT. Your experience and guidance have been invaluable to this project. It has been a pleasure to work alongside you on each of my visits to Harwell.

This work was proposed and initiated by Dr Sam Keyes, whose guidance in getting me started and setting me off in the right direction has made this work possible.

Prof. Ian Sinclair has served as my internal examiner for each milestone during this project. This thesis is undoubtedly far stronger for your insightful critiques and suggestions.

Throughout this project, I have enjoyed spending a considerable amount of time with the μ -VIS X-ray imaging centre at Southampton and received friendly assistance and encouragement not just from those linked with my work but also from the whole team and many students.

I collected the X-ray data that supports chapters 4 and 5 of this thesis at the I12 beamline at Diamond Light Source. I'm very grateful to Dr Oxana Magdysyuk and Kaz Wanelik for their generous support and assistance. They even gave up sleep to keep our experiment running through the night. I also want to thank Ronan Smith and Keiran Ball, who accompanied me on that trip to share the shifts of late-night sample changes. Best of luck to you both in your ongoing research.

It has been a great blessing to work for Christ Church Southampton for the last year while I have completed the writing of this thesis. I have learned a lot this year that I wouldn't have done otherwise. I am full of thanks to the

Acknowledgements

whole team for showing incredible generosity and flexibility as I have struggled to balance two very different jobs. It has provided a delightful contrast to be working in such a close-knit team to serve people alongside the many hours sitting at my desk writing.

On a similar note, I want to acknowledge my housemates for their patience with me as I have increasingly worked from home and shut myself away with my thesis.

Finally, I want to thank both Judith and my family for their loving support and encouragement as I have slowly completed this work.

Nomenclature

μ Linear attenuation coefficient

ψ_g Gravitational potential

ψ_p Hydrostatic potential

ψ_s Osmotic potential

ψ_w Water potential

Σ Macroscopic cross-section

σ Microscopic cross-section

σ_c Absorption cross-section

σ_e Elastic scattering cross-section

σ_i Inelastic scattering cross-section

d_k Distance from a point of interest to the k^{th} axis

f_k RMS distance of all fiducial points to the k^{th} axis

I Attenuated beam intensity

I_0 Incident beam intensity

N Number of fiducial points

n Material density

x Path length through the sample

D Aperture diameter

L Aperture to sample distance

1. Introduction

1.1. Introduction

The human race depends upon plants to provide the vast majority of its food, either directly or indirectly. As the earth's population increases, a corresponding growth in crop yield is required. Current estimates state that crop production will need to double by 2050 to keep pace with projected population growth [Tilman et al., 2011], a target that current trends are insufficient to meet [Ray et al., 2013]. Climate change will make this all the more difficult, mainly through reduced water availability and the drive to reduce fertiliser usage [Atkinson et al., 2019].

The *green revolution* is the name given to a period in the mid-twentieth century when several scientific advancements, including the introduction of fertilisers and genetic modification, led to a tremendous gain in crop yields in a relatively short period [Bishopp and Lynch, 2015]. The green revolution was primarily centred on manipulating the portion of the plant that is visible above the ground, and roots were largely overlooked. However, the root system is central to plant functions such as water and nutrient uptake, anchorage and interaction with symbiotic organisms, so it has been recognised that root growth and development could be an area that is not yet fully exploited to maximise crop yield [Herder et al., 2010]. It has been suggested that the deployment of crops with more efficient water and nutrient uptake due to improved root traits could lead to a second green revolution and help to address the worldwide challenge of food security [Lynch, 2007].

Non-invasive, *in vivo* studies of plant roots present a challenge traditionally addressed through rhizoboxes (containers that force 2D growth conformation with transparent windows for observation) or transparent, artificial growth media. Although widely used, these methods generally result in root systems that vary considerably from those grown in natural soil conditions [Kuijken et al., 2015]. X-ray computed tomography is the primary method for 3D imaging

1. *Introduction*

of root systems grown in soil but is not without limitations [Mairhofer et al., 2011]. While X-rays show the soil structure well, they are not an ideal tool for imaging water distribution in particular since there is very little contrast between the water, plant roots and any other biological soil constituents. None of these hydrogen-rich constituents is imaged clearly, and as a result, it is difficult to differentiate between them in the image data produced.

Neutron imaging offers a solution as many imaging techniques are similar, but the mechanisms by which neutrons interact with matter are very different. Hence different elements, in particular a number of light elements such as hydrogen, provide strong contrast in a neutron image [Kaestner et al., 2017]. Many experiments have shown neutron imaging to be well suited to showing water dynamics, where X-ray imaging would have struggled [Menon et al., 2007, Matsushima et al., 2009, Moradi et al., 2011, Warren et al., 2013, Zarebanadkouki et al., 2013, Totzke et al., 2017]. Neutron imaging has its limitations too. It is slower, less readily available and produces images of lower resolution when compared to equivalent X-ray techniques. Just as X-ray imaging cannot provide a good representation of all the materials in a plant sample, neutron imaging does not show the soil minerals and structure well.

The value of multi-modal methods that image samples using both X-ray and neutron techniques has been identified and begun to be exploited in crop science [Kim et al., 2013, Kaestner et al., 2016, Mawodza et al., 2021] and more widely [Mannes et al., 2015, Ziesche et al., 2020]. Although the requirements and challenges of such multi-modal studies are well understood, many of these experiments rely on methodologies and data processing schemes specific to the instrument or experiment. There is currently no common framework of tools and best practices for correlative X-ray and neutron imaging.

1.2. **Aims and objectives**

This work aims to develop multi-modal methods that combine X-ray and neutron imaging techniques to provide a more complete scan of a root-soil system. These methods would allow better observation and study of water and nutrient uptake and mucilage production within plant samples. The two main challenges in achieving this are image acquisition and image analysis. First, a sample must be prepared and imaged to meet the requirements of both the X-ray and neutron imaging instruments and the biological needs of the

plant. Second, image processing techniques must be developed to allow the two scans of a sample to be registered and combined accurately in a way that facilitates further analysis. The methods developed in this work are designed to be generalised beyond studies of plants to provide a framework and tool-set for correlative X-ray and neutron imaging more widely.

The first phase of this work is to develop general methods for collecting and registering complementary datasets. We built upon this during the second phase of the work, using methods for combining and enhancing multi-modal datasets. These techniques were tested using an investigation into water fluxes within soils.

Chapter 3 presents a study of fiducial markers as a method for registration. Simulations were conducted to observe the impact of the number and placement of fiducial markers on registration accuracy, and the principles shown by the simulations were then tested in practice. Chapter 4 builds upon this by considering the material used for fiducial markers and presents and evaluates a marker segmentation routine designed to allow automated registration of correlative datasets. Chapter 5 concerns subsequent processing of these correlative datasets after registration to allow the potential of complementary modalities to be maximised and facilitate segmentation, quantification and further analysis.

2. Literature review

2.1. X-ray and neutron imaging

Tomography, and transmission imaging more generally, depends on the fact that a beam of penetrating radiation will be attenuated when travelling through matter. The radiation intensity exponentially decreases as it passes through matter due to the loss of particles by absorption and scattering, and the attenuation can be mapped across an area to produce an image. While the mechanisms for absorption and scattering and some of the nomenclature used to describe these properties are different between X-rays and neutrons, the principles of beam attenuation are common to both X-ray and neutron beams.

The Lambert-Beer law gives the total attenuation of a monochromatic beam through a homogeneous object (eq. 2.1), where I is the attenuated beam intensity, I_0 is the incident beam intensity, μ is the linear attenuation coefficient and x is the path length through the object. The attenuation coefficient is the sum of the attenuation coefficients for each mechanism of attenuation. This equation is based on the assumptions that all scattered radiation is removed from the beam, that the attenuation coefficient does not vary with x and that the incident beam contains radiation at a single energy level so that the energy dependence of the attenuation coefficient can be ignored.

$$I = I_0 e^{-\mu x} \quad (2.1)$$

2.1.1. X-ray attenuation mechanisms

Typically the attenuation mechanism that applies to a given photon is dependent on its energy relative to the binding energy of the electron with which it interacts.

2. Literature review

2.1.1.1. Photoelectric absorption

If a photon interacts with an inner shell electron and the photon's energy is greater or equal to that of the electron binding energy, then the X-ray may be absorbed due to the photoelectric effect, as shown in figure 2.1. In this case, the photon gives up all its energy, and the excited electron is ejected from the atom with kinetic energy equal to the difference between the incident photon energy and the electron binding energy. An electron fills the gap left by the ejected electron from an outer shell with less binding energy, producing a characteristic X-ray with energy equal to the difference in binding energy between the two shells [Seibert and Boone, 2005].

For photoelectric absorption to take place, the energy of the incident photon must be at least as great as the electron binding energy. This results in a sudden jump in the mass attenuation of an element where the energy equals the K-shell electron binding energy, and photoelectric absorption becomes possible. This is known as a K-edge. Similarly, other edges can be seen at the other electron shell binding energies. Figure 2.2 shows the mass attenuation coefficient of germanium as an example of this.

2.1.1.2. Rayleigh scattering

If a photon's energy is significantly less than the binding energy of the electron it interacts with, then the photon may be scattered with no loss of energy. This scattering occurs by the electron absorbing the photon and immediately returning to its original energy level by emitting a photon of equal energy in a slightly different direction [Seibert and Boone, 2005]. This process is also known as coherent or elastic scattering, and is illustrated in figure 2.3.

2.1.1.3. Compton scattering

If the energy of a photon is much higher than an electron's binding energy, then Compton (inelastic) scattering may occur. Figure 2.3 shows how, when energy is absorbed from the photon by the electron, both the electron and a scattered, lower energy photon are emitted [Seibert and Boone, 2005].

2.1.1.4. Pair production

For photons with very high energies (>1.022 MeV) it is possible for the photon energy to be converted into mass through pair production. A positron and

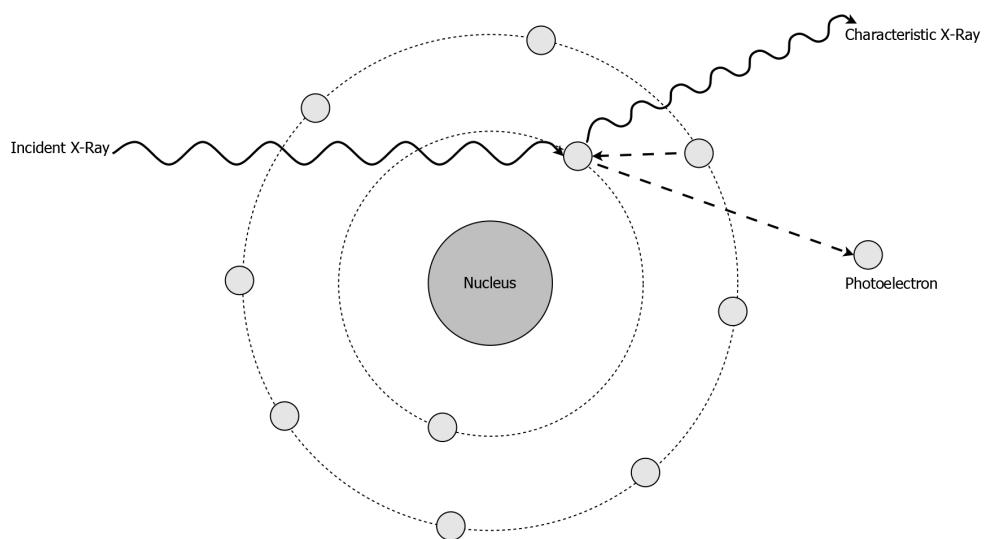


Figure 2.1.: Photoelectric absorption of an incident X-ray photon. The incident electron transfers its energy to the photoelectron, ejecting it with kinetic energy equal to the difference in energy between the incident photon and the electron binding energy. An electron from the outer shell moves to fill the vacancy in the K shell and the difference in binding energies is consequently released as a characteristic X-ray.

2. Literature review

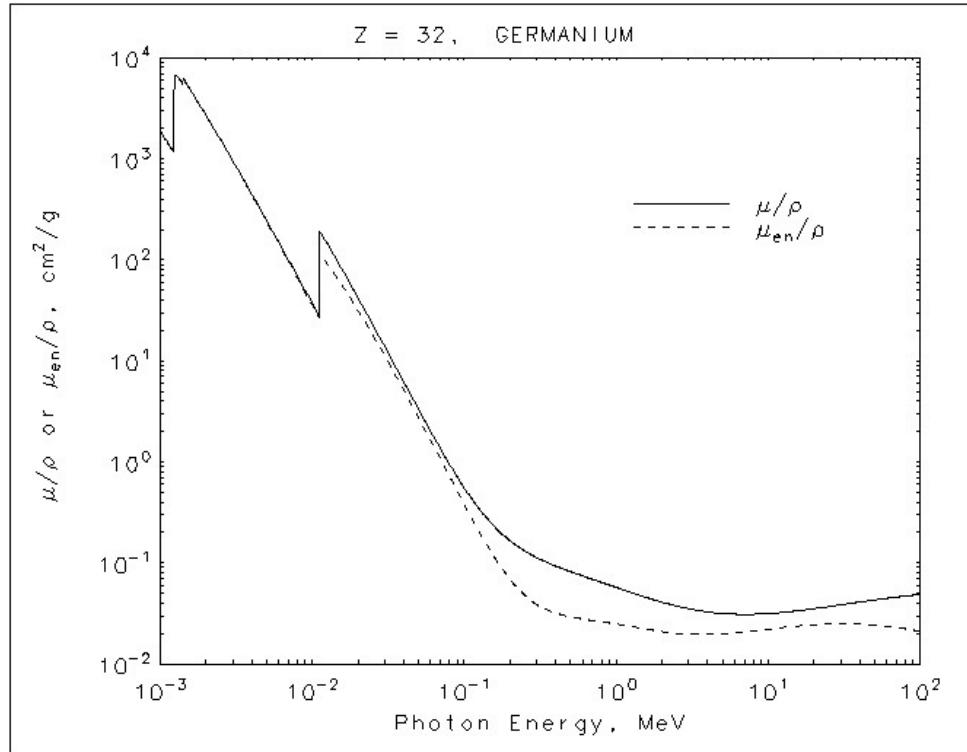


Figure 2.2.: The mass attenuation and mass energy-absorption coefficients of germanium as a function of photon energy. The K-edge can be seen at 1.11×10^{-2} keV. The L-edges occur between 1.216×10^{-3} keV and 1.414×10^{-3} keV. Reproduced from NIST X-Ray Mass Attenuation Coefficients Database [Hubbell and Seltzer, 2004].

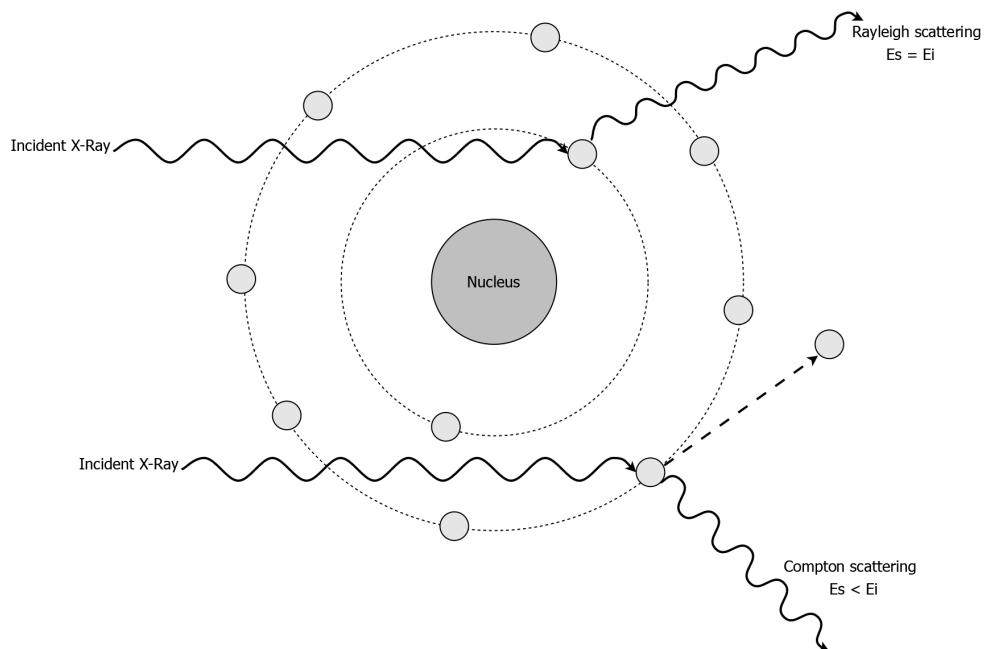


Figure 2.3.: Rayleigh and Compton scattering diagram. The top case shows Rayleigh scattering where the photon is absorbed by the electron and another photon is emitted with equal energy but also with a slight change in direction. The lower case shows Compton scattering where the absorbed photon provides enough energy to free an electron and only the remainder of the energy is transferred to the scattered photon.

2. Literature review

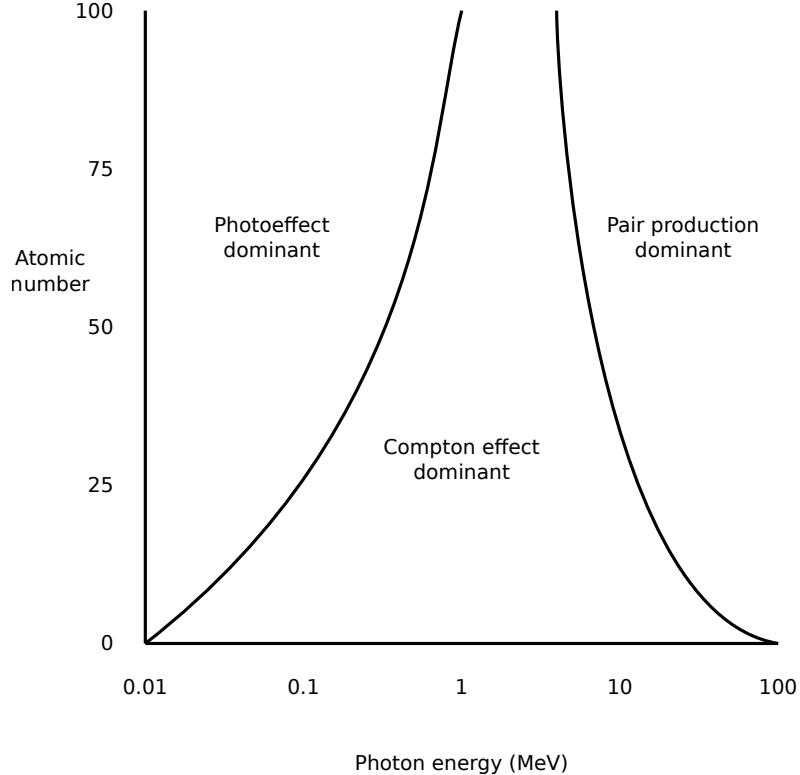


Figure 2.4.: How the types of photon interaction vary with photon energy and atomic number. Based on figure from Podgorsak et al. [2005].

electron pair are ejected. Eventually, the electron will lose its energy and the positron will collide with an electron producing two photons of 511 keV.

In a typical imaging configuration, the photoelectric effect dominates at low energies until Compton scattering takes over at a higher energy that is material dependent. Pair production is only significant for significantly higher energies and materials with a high atomic number. The precise energies where these mechanisms occur are material dependent [Seibert and Boone, 2005]. Figure 2.4 shows how the dominant interaction type varies with photon energy and atomic number.

2.1.2. X-ray imaging systems

2.1.2.1. X-ray sources

Lab-based imaging systems commonly use X-ray tube sources to produce the necessary beam by accelerating electrons and colliding them with a target to

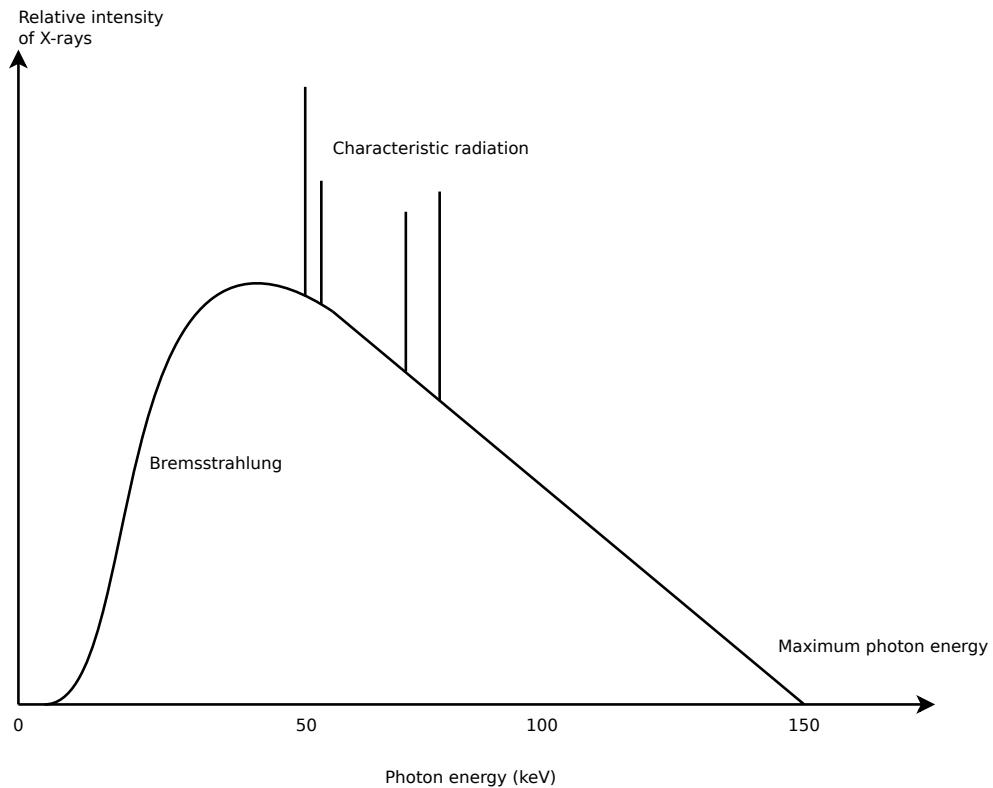


Figure 2.5.: A typical bremsstrahlung and characteristic radiation spectrum for a tube source. The maximum photon energy is 150 keV which shows that the tube voltage must be 150 kV. A steep decline in intensity can be seen below about 40 keV as the X-rays produced at these low energies are quickly absorbed. Characteristic radiation peaks can be seen in the 50-80 keV range. Based on figure from [PhysicsOpenLab, 2017].

2. Literature review

generate Bremsstrahlung radiation [Reed, 2005].

In X-ray radiography, the power of a source is given as a kVp value. kVp stands for kilovolts peak and is the tube voltage between the anode and cathode. The electron volt (eV) is the unit that describes the energy of the particles. As electrons are accelerated across the tube, they gain the same energy as the tube potential, so $kVp = keV$. As electrons collide with the target, they slow down, and some of their kinetic energy is converted to X-rays. The electrons interact with the nuclei within the target, so the energy lost (and emitted X-ray energy) is dependent on how close the electron is to the nucleus with which it interacts. If an electron hits a nucleus directly, all its energy is transferred, so the X-ray produced will have energy equal to the tube potential. However, it is more likely that the electron will pass near the nucleus, and a lower proportion of its energy will be transferred. A continuous spectrum of energy is produced. A higher proportion of lower energy X-rays is produced, partly because electrons are less likely to hit nuclei than to pass by at a distance and partly because electrons can interact more than once but lose energy with each interaction. The output spectrum diminishes at particularly low energies. This is because low energy X-rays are more easily attenuated and hence do not travel very far through matter [Reed, 2005]. Typically the mean keV will be about a third of the kVp, but this is dependent on the target material and can be increased using filters that remove lower energy photons from the beam. Figure 2.5 shows a typical output spectrum from a tube source. Provided the tube potential is high enough, characteristic radiation can be produced by a similar mechanism to the photoelectric absorption described above. In this case, however, it is an incident electron rather than a photon that frees an electron from the atom.

In a synchrotron-based system, electrons are accelerated until they are close to the speed of light travelling around a storage ring. As the beam's path is bent, the electrons lose energy in the form of electromagnetic radiation. This radiation is highly collimated and has a brightness many orders of magnitude higher than a tube source could produce [Kinney and Nichols, 1992]. The high flux produced allows for far more subtle variations in absorptivity to be detected. In addition to this, synchrotron radiation can be tuned to a narrow energy band using a monochromator which removes artefacts associated with energy dependent absorption (2.2.2) [Landis and Keane, 2010].

2.1.2.2. X-ray detectors

Initially, X-ray detectors could not measure incident X-rays directly but by measuring their interaction products. Early detectors took advantage of the ability of X-ray radiation to ionize gasses. An array of high pressure xenon chambers with alternating anodes and cathodes between the chambers will produce a current when incident X-rays release electrons from the xenon atoms [Buzug, 2008]. The low probability of photoelectric absorption can be improved by increasing the pressure and the height of the chambers.

Modern detectors either use direct semiconductor systems or indirect scintillator based systems. Direct detectors use photoconductive materials such as silicon doped with lithium to capture incident X-ray photons and convert them to electron hole pairs via the photoelectric effect. A bias voltage draws the electrons and holes to the corresponding electrodes, generating a current proportional to the intensity of the incident radiation [Chotas et al., 1999].

An indirect detector consists of a scintillator and a photon detector. The scintillator converts incident X-rays to light that the photon detector can then measure. The choice of scintillator material is critical since it determines both the quantum efficiency and the time constant of the conversion process. To provide a fast fluorescence decay and a high quantum efficiency, ceramic materials made of rare earth oxides such as gadolinium oxysulphide ($\text{Gd}_2\text{O}_2\text{S}$) are used [Buzug, 2008].

An anti-scatter collimator grid is often added to the detector to remove scattered X-ray quanta from the detected beam since these would reduce the image contrast. This comes at a cost however to the detector fill factor and the spatial resolution [Carlton and Adler, 2012].

Direct or indirect flat panel detectors can be produced to take advantage of semiconductor production techniques by forming a matrix of detector elements on a single substrate. This results in a detector with a high fill factor, good dynamic range and a linear response [Buzug, 2008].

2.1.3. Neutron attenuation mechanisms

The extent to which neutrons interact with nuclei of a given isotope is expressed as a cross-section.

Microscopic cross-sections are defined as the likelihood of a particular interaction between an incident neutron and a target nucleus but can be thought

2. Literature review

of as the effective target area a nucleus presents to an incident neutron. The standard unit for the microscopic cross-section (σ) is the barn, which equals 10^{-28} m^2 .

A macroscopic cross-section (Σ) can be derived from the microscopic cross-section and the material density (n) in nuclei per m^3 by equation 2.2 [Anderson et al., 2009].

$$\Sigma = n\sigma \quad (2.2)$$

The three primary interactions in neutron imaging are elastic (σ_e) and inelastic (σ_i) scattering and absorption through radioactive capture (σ_c). Generally, a neutron scattering reaction occurs when a target nucleus emits a single neutron after a neutron - nucleus interaction but in an elastic scattering reaction there is no net energy transfer into the nucleus. In an inelastic scattering reaction some energy from the incident neutron is absorbed to the recoiling nucleus and the nucleus remains in an excited state. In a radiative capture reaction, the incident neutron is absorbed and one or more gamma rays are emitted [Anderson et al., 2009].

The total cross-section is calculated as the sum of the cross-sections for each type of interaction (eq. 2.3).

$$\sigma = \sigma_e + \sigma_i + \sigma_c \quad (2.3)$$

Using the total macroscopic cross-section, neutron transmission can be calculated through a sample of depth x using the Lambert-Beer law (eq. 2.1) with the total macroscopic cross-section as the linear attenuation coefficient.

2.1.4. Neutron imaging systems

2.1.4.1. Neutron sources

The neutron beam required for an imaging instrument is typically produced either by a reactor source or an accelerator based source. Each source type will offer different beam characteristics and facilities generally design instruments that will take advantage of the particular properties of their source. Nuclear reactors produce neutrons through the fission process where a nuclear fuel such as uranium-235 is split into a pair of smaller nuclei and two more neutrons. In a

2.1. X-ray and neutron imaging

typical research reactor, one neutron will be used to continue the chain reaction and one is available for external use. More modern research reactors such as FRM-II in Munich [Neuhaus and Petry, 2007] use very compact reactor cores with very high enrichment which allows for high neutron flux to be available to the surrounding instruments. It became apparent in the 1990s that further increasing the power of research reactors to generate higher flux for instruments would be prohibitively expensive and that future development of higher-flux sources would use accelerators [Lawler, 1998].

Accelerator based sources can produce neutrons either by accelerating electrons and using the Bremsstrahlung photo-neutron reaction or by using the spallation process where protons hit a solid target to generate neutrons. Electron accelerators are comparatively inexpensive to construct, however the large amount of heat produced with each neutron severely limits the performance capabilities [Anderson et al., 2009]. Consequently, spallation sources that make use of higher powered proton accelerators are more common. High energy protons can create large numbers of “spalled” neutrons from a collision with heavy nuclei. For example, a 1 GeV proton is capable of producing approximately 25 neutrons from a lead target, with heat deposition in the target of about half of the proton beam power [Anderson et al., 2009]. This results in an order of magnitude less heat to be dissipated when compared to a fission reaction producing equivalent flux. Most accelerator-based neutron sources are pulsed. This allows far higher instantaneous power and neutron flux since the target can cool between pulses given a sufficiently low duty cycle. Furthermore, short-pulse spallation sources allow time-of-flight (TOF) measurements to take place. In the TOF method, the speed of a neutron is measured by timing its flight from the source to the detector. This information can be used to calculate the energy spectrum of the neutrons.

Both accelerators and reactors produce neutrons with higher energy levels than imaging applications can use effectively. Moderators are used to slow neutrons from energies in the MeV range to one electron-volt or lower. Reactors use moderators that reduce the neutron energy through inelastic scattering. This requires a material containing light elements, such as H_2O or D_2O (heavy water). The moderator may be surrounded by “reflector” materials that will scatter some fast neutrons back into the moderator. The moderators for an accelerator based source typically vary more as there are more parameters that can be varied to optimise the source for particular applications [Anderson et al.,

2. Literature review

2009]. Neutrons need to undergo a number of scattering events before they reach thermal equilibrium. To maintain the short pulse needed for TOF techniques the moderator needs to be kept small. This means that there will be some under-moderated neutrons in the beam so most pulsed neutron sources position the moderators tangentially to the target to minimise the flux of high energy neutrons and gammas entering the moderator. The moderated beam will therefore consist of two components: the thermal equilibrium component and the under-moderated epithermal component. The thermal component has a Maxwell distribution of energies around the moderator temperature and the epithermal component is inversely proportional to the energy of the neutrons which have escaped the moderator before thermalisation [Ikeda and Carpenter, 1985]. The thermal component will increase the pulse width with a long decay but this can be reduced at the expense of thermal flux by cooling the moderator to a lower temperature. The pulse width can be further decreased by decoupling the moderator. A decoupled moderator is surrounded by an absorbing material, such as cadmium, except on the side where the beam emerges. This prevents neutrons entering the moderator from the reflector rather than directly from the target since these neutrons would have taken a longer path and would therefore increase the pulse width. This does of course reduce the flux as more neutrons are being removed from the beam.

2.1.4.2. Neutron imaging detectors

Neutron detection relies on the interaction between a neutron's interaction with an atomic nuclei. If neutrons have sufficient energy then they may ionise surrounding material however it is often lower energy neutrons that are of interest for imaging. These must be detected through neutron capture reactions, which release energy as charged particles or gamma rays. The most common converters used are helium-3 and lithium-6 since they have large neutron absorption cross sections and only emit charged particles. Helium-3 detectors have modest resolution but very low sensitivity to gamma rays. They are typically used in neutron scattering instruments. Lithium-6 is a very efficient scintillation converter and is often used in thermal neutron scintillators [van Eijk et al., 2004].

Some neutron detectors operate in an integration mode, where the individual neutron events are not distinguished. This can produce a high frame rate with gating but is not suitable where the time dependence of the signal is

an important parameter, for example in TOF methods. Detectors that allow counting rather than integration of neutron events overcome this limitation and also facilitate binning of measurements. As a result of this, most neutron instruments use counting detectors [Anderson et al., 2009].

A typical white beam detector setup uses a scintillator screen and a cooled CCD-camera which is placed perpendicular to the beam with a 45° mirror and a lens system to focus the light produced by the scintillator onto the camera's CCD. This keeps the camera out of the neutron path to avoid radiation damage, at the expense of light collection efficiency. The camera acts as integrating detector so it is not suitable as a detector for a TOF system. To produce a counting scintillation detector, a photomultiplier is used to detect the light at each pixel location. Avalanche photodiodes are a possible alternative but they produce more noise and are more expensive to produce [Anderson et al., 2009].

Another approach to neutron detection is to use a solid-state semiconductor system. If charged particles produced by a neutron capture event enter a biased semiconductor device then a pulse will be generated. The device can then be connected to fast readout counting electronics. This type of system could present a number of advantages including small size, low power consumption and a high data rate. Glass microchannel plates (MCPs) accelerate electrons through microscopic channels to amplify electronic signals. By doping the glass channels with an appropriate neutron converter, the device can act as a neutron detector. These offer very high performance as counting neutron detectors with excellent spatial and time resolution making them an ideal choice for TOF applications [Anderson et al., 2009].

2.2. Tomographic imaging techniques

Computed tomography is a technique that expands transmission imaging to 3D by combining information from a series of 2D radiographs recorded as the sample rotates about a single axis relative to the source and detector. In systems where rotating the sample is impractical, such as medical scans, this is achieved by rotating the source and detector around the sample. The reconstructed 3D volume is represented as a stack of slices, each representing a plane perpendicular to the axis of rotation.

First, projection images are collected by taking a radiograph at each angle as the sample rotates through either 180° or 360° . Generally, it can be assumed

2. Literature review

that the attenuation of a beam through a sample along a given path will be the same in either direction, hence a 180° rotation is sufficient. When this assumption is not true, for example, if the sample scatters the beam and multiple scattering interactions bring radiation in and out of a path at different points, a 360° rotation will generally give a better image as it is less susceptible to artefacts. The addition of artefact correction processing can reduce the difference in performance between the 180° and 360° scans [Apostolopoulos et al., 2008]. To meet Nyquist sampling requirements and avoid undersampling artefacts the number of projections collected needs to be at least equal to the width of the detector in pixels multiplied by 2π . Where the sample doesn't fill the width of the detector number can be reduced by replacing the width of the detector with the maximum width of the sample in pixels. This is particularly useful with in neutron imaging where the scan times are particularly long and with samples that are significantly taller than they are wide.

Once collected, the projection data must be converted into slices using a reconstruction algorithm. A range of reconstruction algorithms exist, but the most common is filtered back projection (FBP), where a discretised version of the inverse Radon transform is used to reconstruct the volume analytically. The attenuation measured by each pixel in the detector is assumed to be constant over the entire ray path through the sample. By summing all the paths with appropriate rotations corresponding to the projection angles, a 2D image is formulated from the 1D line integrals. However, this image will be blurred until a ramp filter is applied to remove low frequencies [Hsieh et al., 2009].

Iterative algorithms are more computationally intensive but aim to improve on filtered back-projection, in particular by reducing noise and streak artefacts. They fall broadly into two categories: algebraic and statistical methods [Geyer et al., 2015]. The development of these has been primarily driven by medical imaging, where the aim is to achieve sufficient image quality while minimising the patient's dose. Reconstruction methods that can better handle the noise and other artefacts associated with shorter scan times support this aim and have been studied as long as CT has been in use. Despite pre-dating FBP, it is only as the availability of computational power has grown that these iterative methods have become viable and gained widespread use, taking over from the computationally simpler FBP [Willemink and Noël, 2019].

Iterative reconstruction algorithms use both projection and back-projection steps to reconstruct the sample then evaluating its projections to minimise the error between the projections collected and the projections that would be expected given the reconstructed data. In algebraic methods such as ART (algebraic reconstruction technique) [Gordon et al., 1970] and SIRT (simultaneous iterative reconstruction technique) [Gilbert, 1972], this is done by forming a system of linear equations to be solved [Penczek, 2010]. Statistical methods such as maximum likelihood methods assign weights to the data based on the statistical uncertainty (noise) and often seek to model the physical system to give better results in the forward projection step [Geyer et al., 2015].

2.2.1. **Time-lapse CT and golden ratio acquisition**

Time-lapse imaging aims to scan dynamic processes, either with radiography or tomography, by taking a series of scans of a sample as it changes. Depending on the speed of the process, there can be challenges in obtaining good quality tomographies of dynamic processes because reconstruction methods generally assume that a sample is invariant while all the projections are being scanned for a single tomography. As a result, artefacts can occur as a sample moves if the dynamic process is not sufficiently slow compared to the rate at which scans are completed [Kaestner et al., 2016]. Increasing the acquisition rate can help reduce artefacts. However, this is at the cost of increased noise and reduced resolution. Neutron imaging particularly struggles with time-series imaging as the acquisition times are so much longer. A sample that could be scanned in under a minute at an X-ray synchrotron beamline would likely take between four and eight hours to collect an equivalent neutron scan, depending on the particular beamline parameters and desired image quality.

One technique that can help overcome these limitations is to change the acquisition sequence for the projections to one that will have a time averaging effect during reconstruction. A range of acquisition schemes have been proposed and evaluated [Mueller et al., 1997, Guan et al., 1998], but the most common alternative to sequential acquisition is to generate angular increments between projections by the golden ratio. Golden ratio acquisition is a flexible approach since it provides an infinite series of angles that can be stopped at any point, and it outperforms other acquisition schemes such as random access or prime number decomposition [Kohler, 2004]. Using this acquisition scheme the step size between each projection is defined as the scan arc (360° or 180°)

2. Literature review

multiplied by the golden ratio ϕ , giving an angle of approximately 222.5° for a 360° scan. Figure 2.6 shows the first four and the first seventeen projections to be collected to illustrate the way that the sequence produced by this scheme produces an infinite sequence of angles.

$$\phi = \frac{\sqrt{5} - 1}{2} \approx 0.618$$

A scan that does not have enough projections will suffer from sampling artefacts. Using this method, however, a scan can be continued indefinitely. The number of projections for each volume (and therefore the time interval) can be selected after the acquisition to create a balance between avoiding motion artefacts while maximising time resolution and minimising artefacts due to under-sampling. In this case, iterative reconstruction techniques may be preferred to FBP since they are less sensitive to under-sampling artefacts and allow the use of additional regularisation [Kaestner et al., 2016].

2.2.2. Artefacts

A range of artefacts can be produced in a tomographic imaging system where physical effects cause the image produced to deviate from the sample it is meant to capture. These artefacts are detrimental to image quality and can reduce the accuracy of subsequent measurements taken from the reconstructed volume. In the following section, the common artefacts that occur in scans produced for this project are discussed.

2.2.2.1. Ring artefacts

A faulty or non-linear pixel in the detector can lead to a consistent dark or bright spot in the radiographs. When reconstructed, this will produce a ring in the affected slice. Ring artefacts can often be reduced or removed using image processing either before or after reconstruction takes place, often by identifying and removing consistent lines from the sinograms [Boin and Haibel, 2006]. Other solutions include sample shuttling, where the sample is moved slightly within the beam to ensure that the faulty detector pixel does not maintain a constant position relative to the centre of rotation.

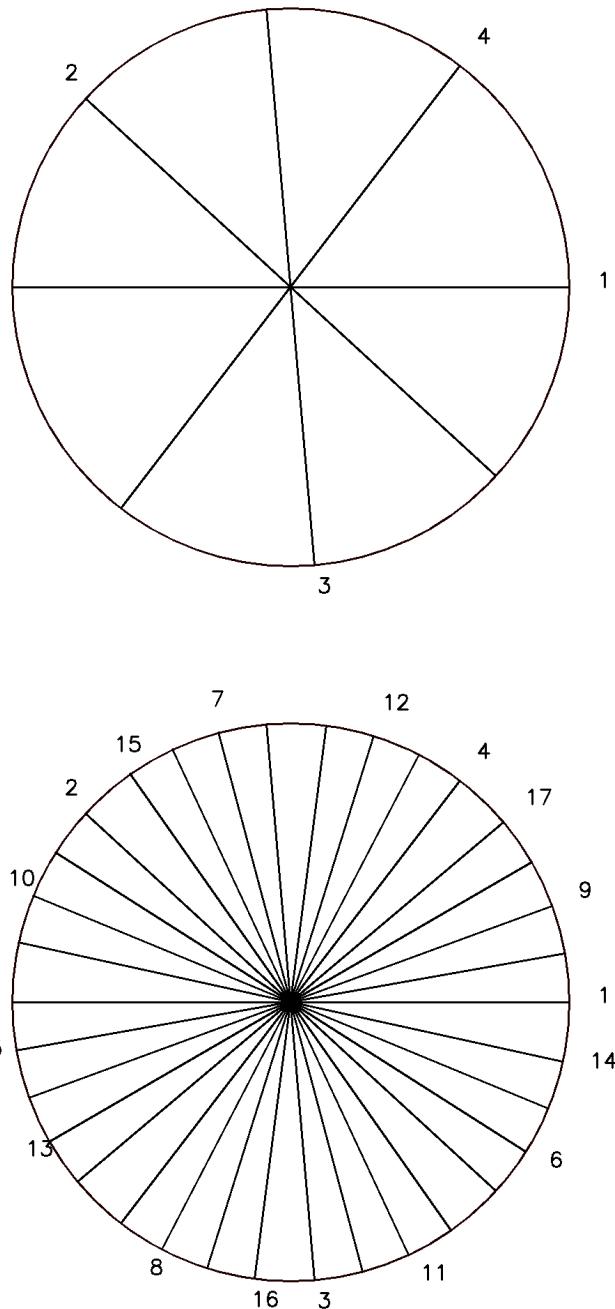


Figure 2.6.: A diagram showing how golden angle based acquisition progressively fills the scan arc with an infinite series of projections. The top figure shows the first four angles in the sequence and the bottom figure shows the first seventeen.

2. Literature review

2.2.2.2. Beam hardening

Beam hardening occurs when a polychromatic beam passes through a highly attenuating sample, and some of the lower energy components of the beam are lost. This results in a beam of lower overall intensity but higher average energy leading to two visual artefacts. Firstly, highly attenuating objects appear more strongly attenuating at their boundaries and less attenuating in their centres. Secondly, dark streaks can appear between highly attenuating objects. Filtering the beam as it leaves the source removes lower energy components before the beam reaches the sample. This filtering reduces beam hardening but also the flux, which can harm the signal to noise ratio. The absence of low energy particles can be detrimental to contrast in weakly attenuating materials. A range of software methods can also be applied to the scan data to reduce beam hardening artefacts. These fall broadly into the categories of linearisation methods and iterative methods [Zhao and Li, 2015]. Another approach used for beam hardening correction is dual-energy CT. In this case, the sample is scanned twice with different beam energies, and this information simulates a monochromatic scan, which would be free from beam hardening [Boas and Fleischmann, 2012].

2.2.2.3. Scatter

Scattering artefacts occur where a photon or neutron has a scattering interaction with the sample. Instead of passing through the sample or being absorbed, the photon or neutron changes direction and energy. This can result in it arriving at a different pixel in the detector. This can manifest as reduced contrast in the data or streaks, similar to those seen with beam hardening. This can be addressed either by seeking to prevent scattered radiation from arriving at the detector or by using software methods to correct for the effects of scattering [Schörner, 2012]. In the first case, the amount of scattered radiation reaching the detector can be reduced by increasing the sample to detector distance. This improvement comes at the cost of both resolution and signal to noise ratio, however. Another common approach is to place an anti-scatter grid in front of the detector. This grid prevents off-axis radiation from entering each pixel of the detector [Kyriakou and Kalender, 2007]. Software approaches aim to avoid the reduced contrast and signal to noise ratios introduced by the physical mitigation techniques and instead adjust the data to account for the

scatter [Zhao et al., 2016]. Some methods use additional measurements taken with a beam stop array to measure the scatter so that it can be removed [Ning et al., 2004], while others work purely with the uncorrected image [Zhao et al., 2016].

2.2.2.4. Noise

The detection of X-rays or neutrons is a stochastic process, and therefore susceptible to Poisson noise due to statistical error if the detector count is low. Whether achieved by increased beam intensity or scan time, a higher count will improve the signal-to-noise ratio but may not be viable. FBP reconstruction assumes accurate projection data and fails to take the noise into account. Iterative reconstruction methods use statistical models to account for the noise and therefore give better results with noisy data [Boas and Fleischmann, 2012].

2.2.2.5. Starvation

If part of a sample is too highly attenuating for the beam to penetrate, very little radiation reaches the detector. Depending on the sample geometry, this may be particular to a subset of angles. The results of this low count at the detector will be particularly noisy projections [Barrett and Keat, 2004]. Once reconstructed, this manifests as streaks through the paths of high attenuation. As a consequence of noise in the image, starvation artefacts are more severe when FBP is used for reconstruction. However, these artefacts can be reduced substantially with additional data processing [Mori et al., 2013, Chang et al., 2014].

2.2.2.6. Undersampling

If the number of projections used for tomography is too low, then this can lead to aliasing. Typically, this occurs because a larger angular interval is chosen to minimise the scan time or the sample's dose [Barrett and Keat, 2004]. As in the case of noise, this is a problem that FBP is particularly susceptible to, and iterative methods can provide significant improvements.

2.2.3. Multi-modal imaging

Multi-modal imaging seeks to overcome the weaknesses of an imaging modality by using a second modality to provide complementary information. Kaestner

2. Literature review

et al. suggest two cases where bimodal imaging is desirable. In the first case, the sample is static, but some features are difficult to identify due to the low contrast between the materials — the second case concerns samples with a dynamic component. One modality can identify the static structure while the other captures the displacements [Kaestner et al., 2016].

The samples considered in this work combine both these cases as there is little variation in X-ray attenuation between the water, roots and organic soil content, but there is also interest in the transport and uptake of water through the soil and roots, which calls for dynamic studies.

The NEUTRA and ICON instruments at the Paul Scherrer Institut, Switzerland, are the first neutron imaging beamlines to install X-ray imaging equipment, allowing complementary X-ray and neutron images to be produced *in situ* without the sample being moved and delays being introduced between acquisitions using different modalities [Kaestner et al., 2016]. At NEUTRA, the X-ray source can be placed in the line of the neutron beam to give a nearly identical geometry. The same detector can be used for both modalities with a change of scintillator screen. This configuration removes the need for registration since the alignment of the two systems is similar [Mannes et al., 2015]. ICON uses a different arrangement where the X-ray beam crosses the neutron beam at an angle that can be set between 80° and 135° to minimise cross talk and shadowing of either modality [Kaestner et al., 2017]. This arrangement is designed to allow simultaneous acquisition of both modalities to ensure the best possible match between data if the sample can change. In this case, registration of the two modalities is required, and the data produced so far suggest that registration cannot be based on common local features. The available calibration data of the configuration and the common sample position between the two modalities means that an accurate initial guess can be used, so global optimisation strategies have resulted in successful registration [Kaestner et al., 2017].

The bi-modal nature of these two instruments means that combined X-ray and neutron imaging is now a standard option at PSI. The combined installations mean that issues related to the sample being moved between facilities or time passing between the X-ray and neutron scans are entirely avoided. Furthermore, registration is either avoided or simplified. The research output of this facility shows the value of this instrument configuration. Similar instrument designs have subsequently been developed at NIST [LaManna et al.,

2.2. Tomographic imaging techniques

2017] and ILL [Tengattini et al., 2020]. These instruments follow a similar pattern of adding an X-ray tube source and detector to an imaging beamline at a reactor or spallation source and optimising beam geometry, detector choice and shielding to ensure minimum crosstalk and similar beam geometries for pixel-wise comparison.

This work evaluates how facilities can achieve similar results by gathering bimodal data in separate instruments. Such methods would enable a wide range of experiments at a more extensive range of facilities that cannot produce similar bi-modal instrument configurations due to space or cost restrictions and allow further performance enhancements such as those achieved by using synchrotron X-ray imaging techniques rather than tube-based sources.

Prototype systems are also being developed to achieve bimodal X-ray and neutron imaging with low energy electron accelerators rather than research reactor or spallation neutron sources [Yu et al., 2021]. As discussed in section 2.1.4.1, electron accelerator sources are comparatively inexpensive to construct but are limited in their performance by the heat produced with each neutron due to the photo-neutron reaction. They do, however, produce both X-rays and neutrons from a single source that can be distinguished by a single detector using time-of-flight techniques which makes them suitable for bimodal imaging [Yu et al., 2021].

Magnetic resonance imaging (MRI) is another three-dimensional imaging technique that could be used in conjunction with X-ray or neutron imaging. Unlike X-ray and neutron imaging, it does not rely on the attenuation of a penetrating beam but instead applies pulsed magnetic field gradients to the sample and measures the nuclear spin density in a target region [Carlson, 2006]. In particular, this technique is sensitive to hydrogen (at least in liquid and gaseous forms), although other nuclei with unpaired nuclear spins can also be investigated [Carlson, 2006].

MRI is useful for studying plant hydraulics on different length scales within intact plants [Van As, 2006]. Ferromagnetic and paramagnetic materials in the soil result in a poor signal to noise ratio for MRI and mean that the technique is generally limited to carefully selected growth media such as glass beads or pretreated sand [Moradi et al., 2011]. This means that the experimental design constraints for MRI require the samples to be prepared in a way that is further from the natural growth methods of the plants being studied, rendering the technique less attractive for studies focussing on roots and soil.

2. Literature review

Another modality that could be considered but has been ruled out of this project's scope is positron emission tomography (PET). Garbout et al. show that PET can be used in conjunction with X-ray computed tomography to visualise and quantify the uptake of carbon. The X-ray data showed the soil and root structure, and the PET scanner was used to observe carbon uptake, detecting the radioisotope ^{11}C [Garbout et al., 2012]. In general, the contrast between materials is somewhat similar using X-ray imaging and PET, so it is best suited to cases where X-rays observe the static structure and PET captures the displacements of a dynamic component.

2.3. Root and soil imaging

While the study of roots and soil motivates this project, it is also true that the main body of this work is focused on developing techniques to better study plants and soil rather than a direct study into them. With that in mind, a summary of the necessary background is presented here, followed by an account of the increasingly sophisticated experiments designed to study plants and soil through X-ray and neutron imaging.

2.3.1. Soil

Water content and the rate of water movement in soils is dependent on soil type and structure [Taiz and Zeiger, 2002]. Soils can be categorised by properties such as soil chemistry and organic content. However, soils are primarily categorised by their particle size when considering the hydraulic properties, as shown in table 2.1. In heavily watered soil, the water percolates down by gravity through the spaces between particles, displacing or trapping air in the soil. Water may fill the entire channel or act as a film, adhering to the surface of soil particles. Sandy soils have large spaces or channels between particles, whereas clay soils typically have much smaller channels. As a result, water drains quickly through sandy soils, remaining only at the particle surfaces or interstices between particles. In clay soils, the smaller channels do not allow the water to drain so freely, although clay particles may aggregate into 'crumbs' with the aid of organic matter, which helps improve soil aeration and infiltration of water.

The moisture-holding capacity of soils is called the field capacity and is defined as the water content of the soil after it has been saturated and excess

2.3. Root and soil imaging

Table 2.1.: Soil types as defined by particle diameter, reproduced from [Taiz and Zeiger, 2002].

Soil	Particle diameter (μm)	Surface area per gram (m^2g^{-1})
Coarse sand	2000 - 200	< 1 - 10
Fine sand	200 - 20	< 1 - 10
Silt	20 - 2	10 - 100
Clay	< 2	100 - 1000

water has been allowed to drain away. Clay with a high humus content has a high field capacity and might retain 40% water by volume. In contrast, sandy soils typically retain 3% water by volume.

Water potential (ψ_w), the relative tendency of water to move from one area to another, is the sum of osmotic potential (ψ_s), hydrostatic potential (ψ_p) and gravitational potential (ψ_g). In soil or plant cells, the gravitational potential is typically so negligible that it is not included. Osmotic potential is generally very low in soils except when the soil contains a substantial concentration of salts. For wet soils, the hydrostatic potential is close to zero. As the soil dries out, however, it decreases to give a significant negative pressure. This is a result of the cohesive and adhesive forces in water that produce capillary action.

In general, water predominantly moves through soils by a bulk flow driven by a pressure gradient, although some water movement is contributed by diffusion of water vapour [Taiz and Zeiger, 2002]. As plants absorb water from the soil, they deplete the soil of water near the surface of the roots. This depletion reduces hydrostatic pressure in the water near the root surface and establishes a pressure gradient in the soil. Water moves through the interconnected pore spaces by bulk flow down the pressure gradient towards the root surfaces.

The rate of water flow in soils is dependent on two factors: the size of the pressure gradient in the soil and the hydraulic conductivity of the soil. Soil hydraulic conductivity measures the ease with which water moves through the soil and varies with soil type and water content. As the soil's water content (and hence the water potential) decreases, air moves into the vacated soil channels restricting water to the channel's edges, reducing the hydraulic conductivity. In arid soils, the water potential may fall below the permanent wilting point. The permanent wilting point is defined where the potential is so low that the plants cannot regain the necessary internal hydrostatic pressure

2. Literature review

to maintain necessary physiological processes, even if all the water loss ceases. Therefore the water potential of the soil is less than or equal to the osmotic potential of the plant. As a result, the permanent wilting point is not a unique property of the soil but depends on any plants present.

2.3.2. Roots

Close contact with a high surface area is required for a root to absorb water from the soil. The growth of root hairs into the soil maximises this area. Root hairs are extensions of the root epidermal cells (the outermost layer of cells) that grow in the region of maturation and significantly increase the surface area of the root. Root hairs make up more than 60% of the surface area of some roots. Figure 2.7 illustrates the external structure of a root.

Roots grow through cell division in the apical meristem or region of cell division at the tip of the root. The root cap protects this area. The cells that make up the root cap detach from the root as it grows and are thought to modify the surrounding soil conditions to help the plant further [Graham et al., 2006]. Above the region of cell division is the region of elongation. Here the cells grow, particularly in the vertical direction, and begin to develop specialisations. The zone of maturation contains developed cells with their different roles.

Water enters the root most readily in these young regions, particularly the root hair zone. More mature regions of the root may have a protective hydrophobic layer of tissue called an exodermis or hypodermis that prevents water from entering [Taiz and Zeiger, 2002].

The close contact between soil and root surface is easily broken when the soil is disturbed. As a result, newly transported seedlings and plants require protection from water loss for a few days after transportation. Soon, new growth re-establishes soil-root contact, and the plant can better withstand water stress.

The root cap secretes a gel called mucilage or mucigel, which surrounds the root tip. This mucilage lubricates the root, helping it to move through the soil. It also helps control the rhizosphere (the immediate environment surrounding a root) by preventing the roots from drying, protecting the root from toxic substances and enhancing the absorption of water and minerals [Kutschera-Mitter et al., 1998, Graham et al., 2006].

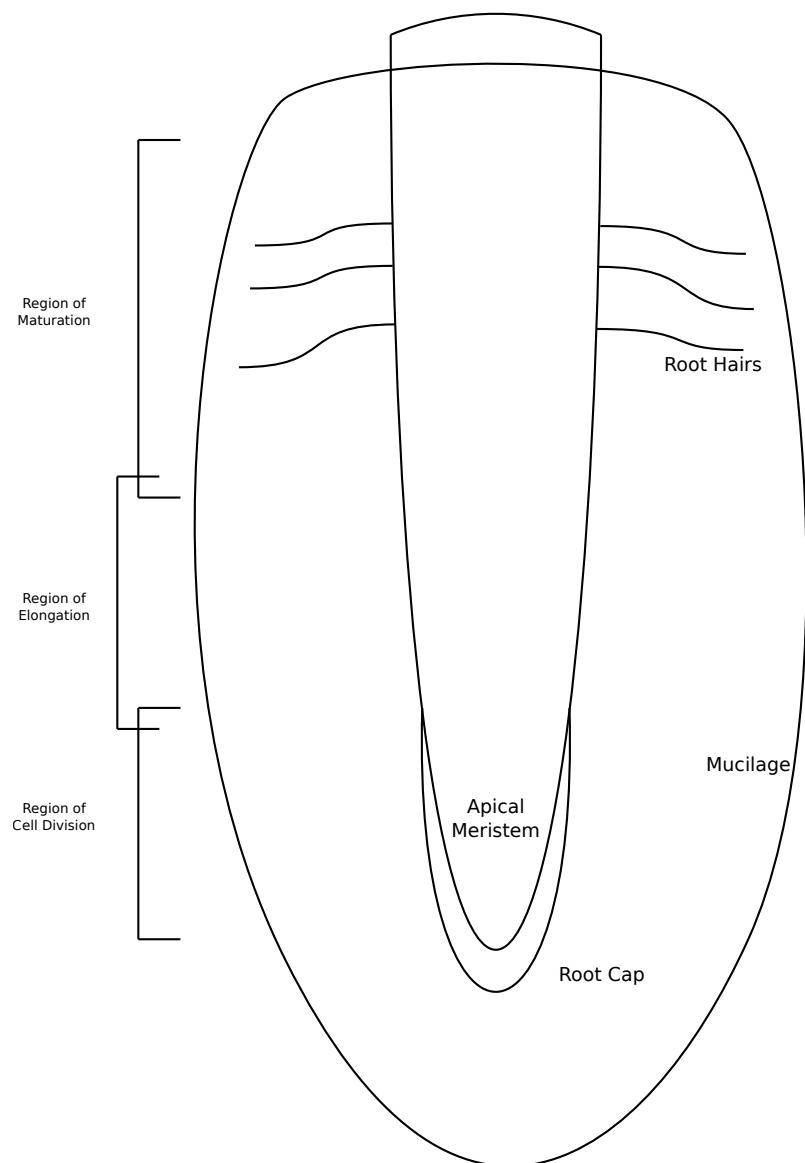


Figure 2.7.: The external structure of a root.

2. Literature review

2.3.3. Plant root and soil imaging

2.3.3.1. X-ray

The development of high-resolution synchrotron imaging has allowed detailed and accurate modelling of roots at ever-decreasing scales. In 2013, Keyes et al. demonstrated that individual root hairs could be seen and used for image-based modelling [Keyes et al., 2013]. Following this, the application of time-resolved imaging allowed Keyes to observe the growth of a root over a day with one hour time steps. Digital volume correlation could observe the full-field soil deformation to give more information [Keyes et al., 2016]. In this work, Keyes also observed that X-ray imaging might not suit all potential applications and called for similar investigations using neutron imaging in conjunction with X-rays.

Koebernick et al. imaged two genotypes of Barley (with and without roots hairs) and was able to observe differences in the rhizosphere pore structures [Koebernick et al., 2017]. Helliwell et al. investigated the interactions between soil texture, bulk density and plant species and their impact on root growth and soil densification using X-ray CT [Helliwell et al., 2019]. Each of these studies shows the use of X-ray tomography to study soil and root structure. In studies that use X-rays as the sole imaging modality but seek to investigate hydraulic properties, however, it remains common to use other methods in conjunction with imaging that depend on the destructive extraction of roots; as in Rabbi et al.'s investigation into the way plant roots redesign the rhizosphere to alter its structure and consequently its water dynamics [Rabbit et al., 2018].

Li and Tang's investigation probed the ability of high-resolution synchrotron X-ray imaging to segment water from unsaturated, compacted soil samples [Li and Tang, 2019]. They showed that they could achieve good water characterisation with glass beads or natural sand by using contrast media to enhance the air-water contrast (fig. 2.8). Using clay, however, aggregates formed, and in this case, it remains impossible to observe the thin water films and bridges between particles.

2.3.3.2. Neutron

In 2007, Menon et al. demonstrated that neutron radiography could be used to study plant roots. To ensure that the samples were thin enough to allow neutron penetration, the plants were grown in rhizoboxes. Some of the plants

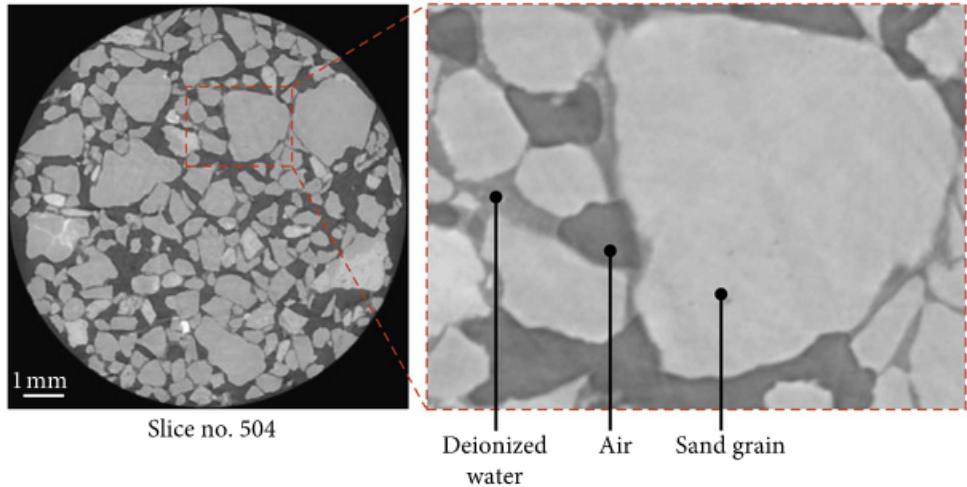


Figure 2.8.: An example slice from [Li and Tang, 2019] showing the contrast between sand, air and water in a high-resolution synchrotron image.

were grown in sand, while others were grown in soil. Menon observed that the plants grew better in the soil, but the sand offered better contrast with the roots for imaging. It was also shown that a reduction in root growth could be observed where the soil was contaminated by boron and zinc [Menon et al., 2007].

Although there is evidence to show that water can be identified in some X-ray images [Li and Tang, 2019], the requirements this adds to the experimental design are stringent and not typically compatible with scanning a root system *in situ*. Water has a far greater attenuation coefficient for neutrons which allows for better image contrast. Matsushima et al. took advantage of this to demonstrate that water flow velocity can be estimated in plants using cold neutron imaging. A radiograph was produced every minute to track the progress of water through the plant. Pulses of D_2O provided contrast with the water since they allow better neutron penetration. As a result of the changes in attenuation, the pulses could be traced, and their velocity estimated [Matsushima et al., 2009].

Neutron tomography offers a three-dimensional visualisation of a sample and can be used to show soil water content around root systems. This technique was used by Moradi to show an increase in water content towards root surfaces and therefore suggest that roots may modify the hydraulic properties of soil

2. Literature review

to improve uptake in dry conditions. This work demonstrated the advantages of tomography over radiography for the study of roots and soil, claiming to set new benchmarks for spatial resolution and quantification of water content distribution [Moradi et al., 2011]. Slices from one of the scans are shown in figure 2.9.

Subsequent experiments show that neutron imaging and its ability to map water are well suited to the study of water transport within plants and soil. For example, Defraeye et al. attempted to use neutron imaging to produce quantitative water measurements of leaf water flow and transpiration in real-time. It was shown that differences in water content and distribution could be quantified, but dehydrated scans would be needed to give absolute measurements as the dry leaf matter contribution could not be corrected without them [Defraeye et al., 2014].

Zarebanadkouki et al. applied similar methods to quantify the flux of water within roots by tracing pulses of D_2O through the soil and separate roots. The development of new methods for these experiments led to new information about the root behaviour, for example, that water uptake is not uniform along a root [Zarebanadkouki et al., 2013], and informed more accurate mathematical models of water transport within root systems [Zarebanadkouki et al., 2014].

Initially, this type of dynamic experiment was only conducted using neutron radiography as tomography was considered too slow [Warren et al., 2013]. The techniques were subsequently expanded into 3D as high-speed neutron tomography techniques were developed, offering a temporal resolution of three minutes. This increased speed came primarily through the use of on-the-fly tomography, which, while well established for X-rays, had not been utilised in neutron tomography beyond exploratory trials [Dierick et al., 2005] due to its high flux requirements and low spatial resolution. Zarebanadkouki et al. discovered that this technique offered sufficient accuracy to quantify the changes of D_2O across the root tissue and concluded that fast neutron tomography was a critical technique that would allow studies to uncover more information to better model water transport properties across root tissue [Zarebanadkouki et al., 2015].

Totzke et al. subsequently managed to reduce neutron tomography acquisition times to 10 s, demonstrating that the resulting motion blur and reductions to contrast and signal to noise ratio do not produce a sufficient decrease in image quality to prevent the roots being segmented and pulses of deuterated

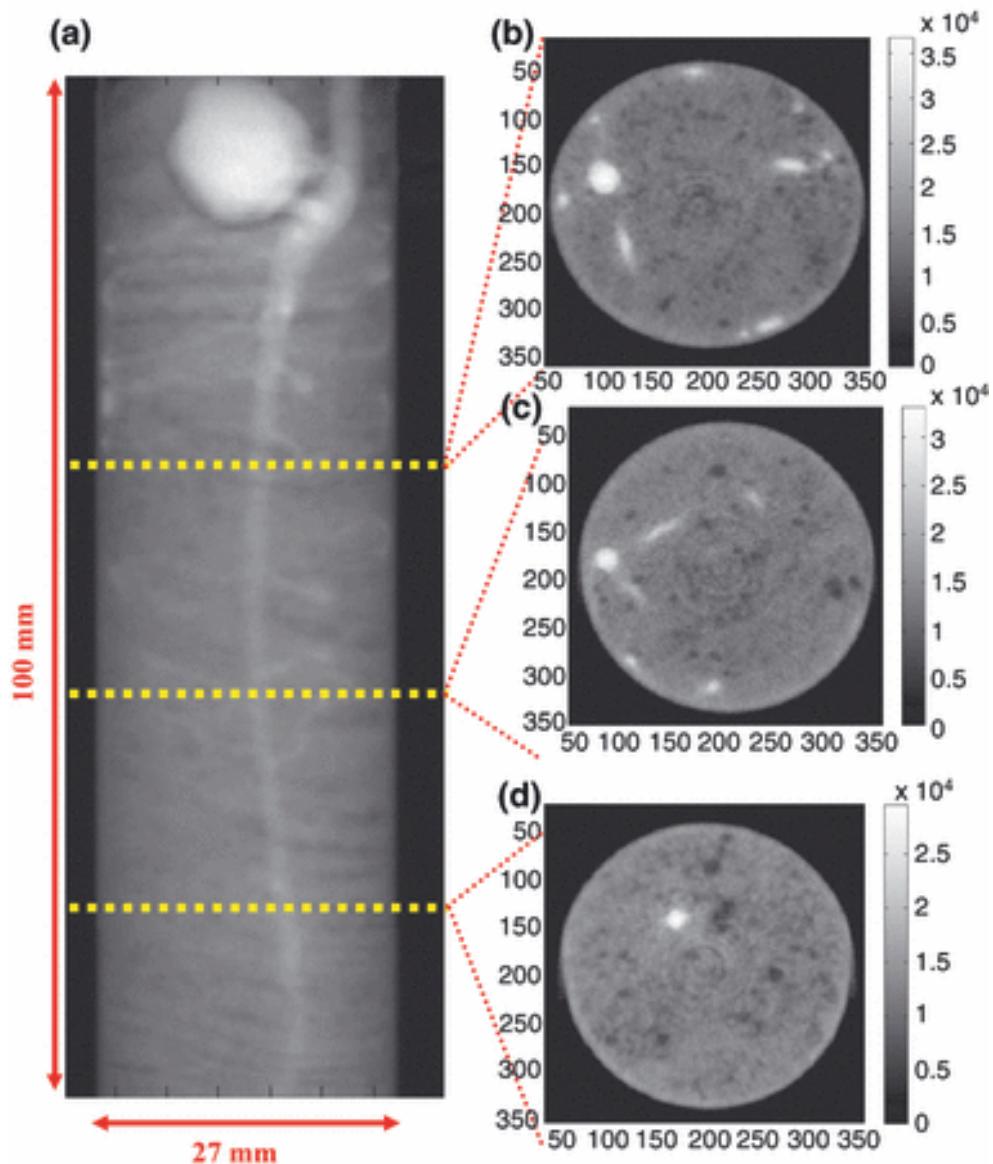


Figure 2.9.: A tomogram of a 12 day old chickpea plant grown in a cylinder of sandy soil and presented in [Moradi et al., 2011]. Sub-figures b, c and d show slices at depths of 4, 6 and 8 cm respectively.

2. Literature review

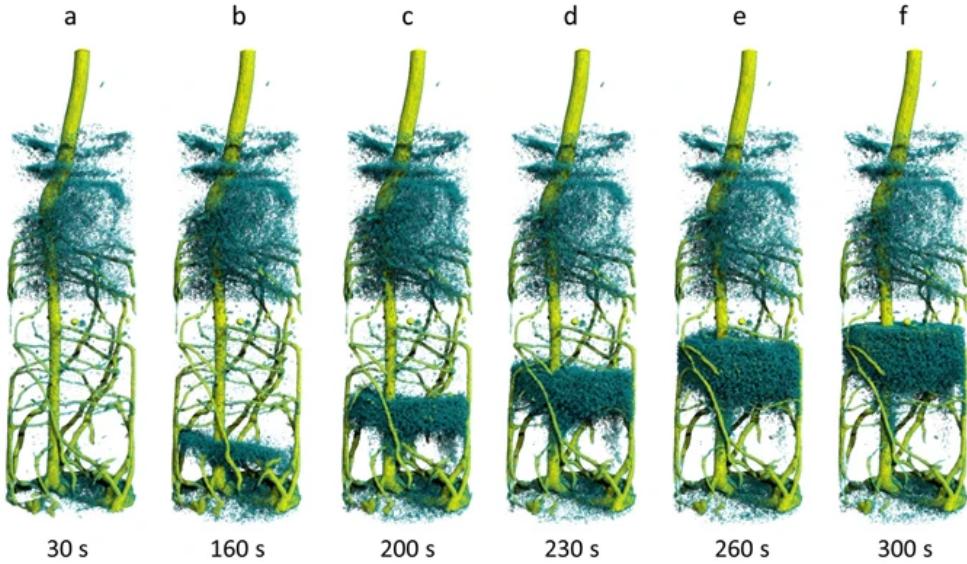


Figure 2.10.: Time resolved neutron tomography of a lupin root system after the injection of 4 ml of deuterated water (D_2O) from the bottom. Reproduced from [Totzke et al., 2017].

water being traced through the system [Totzke et al., 2017]. The accelerated acquisition time is achieved by reducing the exposure time per projection as the sample is continuously rotated. Several experimental factors were considered to ensure the best possible image quality while maintaining this speed. Perhaps most significantly, this experiment was conducted at the CONRAD-2 beamline of the BER II research reactor [Kardjilov et al., 2016], which has a neutron flux at the sample position approximately two orders of magnitude greater than that achieved at PSI in [Zarebanadkouki et al., 2015]. Furthermore, the neutron flux was enhanced by adjusting the pinhole to allow for a high flux neutron beam, even at the cost of resolution. In addition to this, pixel binning was applied to the data to enhance the signal to noise ratio, again at the cost of resolution. Throughout the work, the trade-offs between speed/high flux and resolution were fine-tuned to ensure the resolution was sufficient to capture the features of interest to the study and no more. This work demonstrates the experimental capabilities to allow water dynamics to be imaged in time-resolved 3D at a far higher temporal resolution (fig. 2.10).

2.3.3.3. Correlative

Despite the progress being made in crop science with both X-ray and neutron imaging, there has not yet been a significant uptake of correlative methods. Where multi-modal imaging with X-rays and neutrons was discussed in section 2.2.3, some of the studies cited did investigate this area with roots or soil as an explicit application, but these can essentially be considered proof-of-concept experiments where the main focus of the work is on the new technique rather than on the crop science results. A recent investigation by Mawodza et al. shows that interest in these methods remains high, but the complexities associated with multi-modal imaging can lead investigators to conduct separate X-ray and neutron imaging experiments within a study rather than attempting to correlate them [Mawodza et al., 2021]. This illustrates the need for a greater understanding of the benefits and challenges of multi-modal imaging and standard methodologies and tools to simplify the design of these experiments.

3. Distribution of fiducial markers to optimise the registration of multi-modal CT data

3.1. Introduction

Image registration is the process of putting two images of the same thing into spatial alignment. In general, this could mean combining images taken at different times or from different viewpoints, but in multi-modal analysis, this means integrating images of the same sample from different sources. Registration techniques aim to compensate for differences in machine factors such as resolution and orientation, image factors such as translation, rotation or scaling and object factors if the sample can change between images [Cosman, 2012].

Typical image registration frameworks consist of a method for extracting information to be matched between the reference and target images, a class of transformations capable of aligning the images, a search strategy for selecting transformations to test and a similarity metric to determine the accuracy of a given transformation.

To align CT volume data, we require a class of transforms that allows rotation, scaling, and transformation while preserving the shapes - the lines, angles, and relative distances - present in the data. A rigid body transform will strictly preserve the shape but will not allow scaling. An affine transform is more general and will allow all the fundamental transformations we require (translation, rotation and scale); however, it also allows shear, which can distort shapes in the data. In our case, therefore, we shall ensure that there is no shear component included in the transform.

The affine transformation can be represented as a matrix multiplication of a point represented by a vector. Each transform, such as scale or rotate, can be represented as a matrix, and matrix multiplication can combine these

3. Distribution of fiducial markers for the registration of multi-modal CT data

individual steps to give an overall affine transformation. This representation requires all the transforms to be linear. Translation, however, is not a linear transform. To overcome this, we must define the problem using homogeneous coordinates, which add one to the dimensions of the matrix [House and Keyser, 2016]. The following matrices represent the fundamental affine transformations in 3D:

$$Translate : \begin{bmatrix} 1 & 0 & 0 & \Delta x \\ 0 & 1 & 0 & \Delta y \\ 0 & 0 & 1 & \Delta z \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad Scale : \begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$Shear : \begin{bmatrix} 1 & h_{xy} & h_{xz} & 0 \\ h_{yx} & 1 & h_{yz} & 0 \\ h_{zx} & h_{zy} & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$Rotation \text{ about } x \text{ axis} : \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta_x & -\sin \theta_x & 0 \\ 0 & \sin \theta_x & \cos \theta_x & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$Rotation \text{ about } y \text{ axis} : \begin{bmatrix} \cos \theta_y & 0 & -\sin \theta_y & 0 \\ 0 & 1 & 0 & 0 \\ \sin \theta_y & 0 & \cos \theta_y & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$Rotation \text{ about } z \text{ axis} : \begin{bmatrix} \cos \theta_z & -\sin \theta_z & 0 & 0 \\ \sin \theta_z & \cos \theta_z & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Having settled upon a class of transforms to be used, the question of how to extract information from the target and reference volumes must be addressed.

The similarity metric should, in turn, be easily derived from this. Although some cases of combined X-ray and neutron imaging have successfully used mutual information in the sample to achieve registration [Kim et al., 2015], in general, it has been shown that, due to the complementary nature of the modalities, there is no guarantee that there will be similar local features in corresponding datasets (particularly with multiphase images) and it is, therefore, difficult to find a good registration solution based on common features within a sample [Kaestner et al., 2017]. To overcome this difficulty, fiducial markers can be attached to a sample to aid in registration. A fiducial marker is an object placed within an image to provide guaranteed mutual information and therefore be used as a point of reference. To register volumes using fiducial markers, at least three fiducial points must be selected in both the reference and target images so that the registration parameters can be found and a suitable transform applied.

Inevitably, there will be error in the registration process. The majority of the existing literature examining this has been produced by the medical imaging community where the primary registration challenge is to compare subsequent scans of the same patient and look for changes. Maurer et al. defined three types of error that can occur when using fiducial points for registration [Maurer et al., 1997]:

- Fiducial localisation error (FLE): the average error in locating the position of the fiducial points (Fig. 3.1A).
- Fiducial registration error (FRE): the root mean square (RMS) error between corresponding fiducial points after registration (Fig. 3.1B).
- Target registration error (TRE): the error between corresponding points other than the fiducial points after registration (Fig. 3.1C).

Fitzpatrick and West built on these definitions by providing expressions for the expected FRE and TRE errors in terms of the expected FLE and the set of fiducial points used, where N is the number of points (Eqs. 3.1 and 3.2) [Fitzpatrick et al., 1998].

$$E \left\{ \text{FRE}^2 \right\} \approx \left(1 - \frac{2}{N} \right) E \left\{ \text{FLE}^2 \right\} \quad (3.1)$$

3. Distribution of fiducial markers for the registration of multi-modal CT data

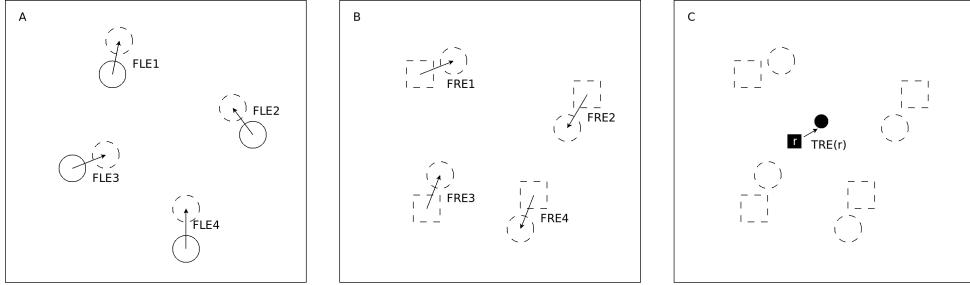


Figure 3.1.: Diagrams showing the three types of registration errors: fiducial localisation error (FLE), fiducial registration error (FRE) and target registration error (TRE). Based on Figure 1 of [West et al., 2001].

$$E \{ \text{TRE}^2(r) \} \approx \frac{E \{ \text{FLE}^2 \}}{N} \left(1 + \frac{1}{3} \sum_{k=1}^3 \frac{d_k^2}{f_k^2} \right) \quad (3.2)$$

$$TRE_M(r) = \frac{1}{N-2} \left(1 + \frac{1}{3} \sum_{k=1}^3 \frac{d_k^2}{f_k^2} \right) \quad (3.3)$$

From this expression, Wang and Song derive an equation (Eq. 3.3) to relate TRE to a particular distribution of markers independently of FLE or FRE. This is achieved by using Eq. 3.1 to substitute FRE for FLE before setting FRE to 1. They propose a deterministic optimisation method for determining the quantity and layout of markers to minimise TRE at a point of interest r [Wang and Song, 2009]. The distance from r to each axis is denoted as d_k , where $k = (1, 2, 3)$ and f_k is the RMS distance of all fiducial points to the k th coordinate axis.

3.2. Simulation

To determine the best configuration of fiducial markers before imaging the samples, a simulation was carried out based on the method presented by Wang and Song [Wang and Song, 2009]. The method was adapted to the case where the sample is cylindrical, and the registration is to be optimised over the whole volume of the sample rather than at a single point of interest.

$TRE_M(r)$ is a dimensionless quantity that relates the expectation value of TRE to that of FRE at a point, r , for a given set of markers M . Potential marker locations were given by a grid of points on the surface of the cylinder

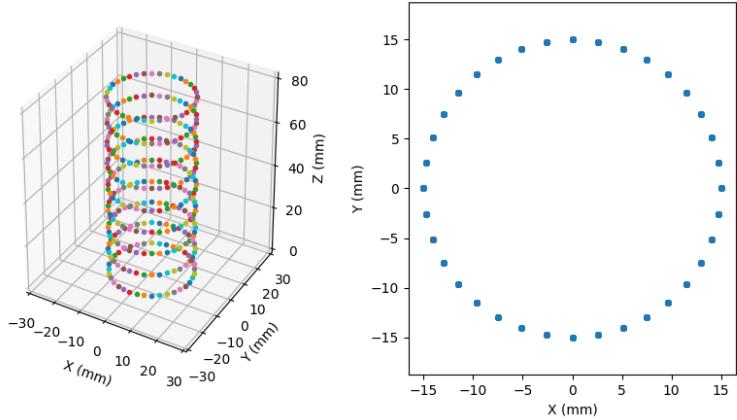


Figure 3.2.: Potential marker locations for the simulation to test.

and these are shown in figure 3.2. By evaluating $TRE_M(r)$ for a set of points evenly distributed throughout the volume and taking the RMS value, an estimate of the TRE_M value is calculated for the whole volume. This means that TRE_M gives a measure of how well two volumes can be registered with a set of markers M independently of the FLE and FRE. The set of points used is illustrated in figure 3.3.

Markers are added to the simulation one by one, minimising the TRE_M each time until the desired number of markers is reached. Each time a marker is to be added, the TRE_M value is calculated for each of the available marker positions in addition to the existing markers from previous steps. The marker at the position which gives the lowest TRE_M is added to the set. Once the desired number of markers is reached, the marker with the highest contribution is repositioned to reduce TRE_M . This step is repeated until the repositioned marker still gives the largest contribution to TRE_M .

The simulation requires the first two points to be provided so that the algorithm can begin to run with the third marker. To ensure that the initial positions of these first two markers do not affect the result, all possible unique combinations (excluding rotationally symmetric equivalents) are tested. The code listing can be seen in section A.1.

The simulation was run for $N = 4, \dots, 16$, where N is the number of fiducial markers, to establish how many fiducial points should be used and how they should be distributed about the sample. Figure 3.4 shows how the value of

3. Distribution of fiducial markers for the registration of multi-modal CT data

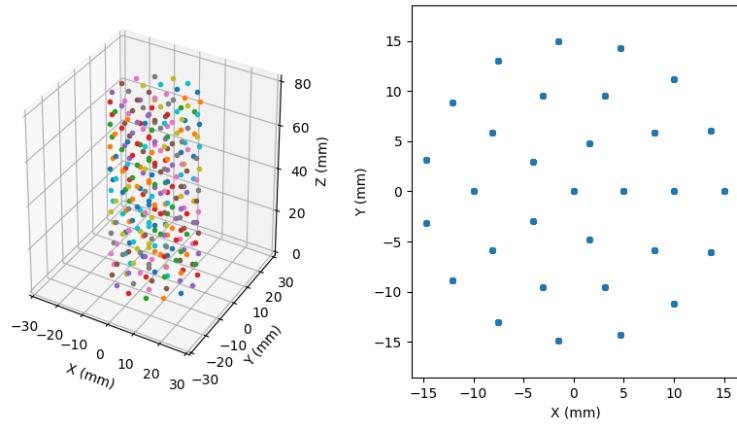


Figure 3.3.: The test points used to evaluate the target registration error throughout the volume and give an estimate of TRE_M .

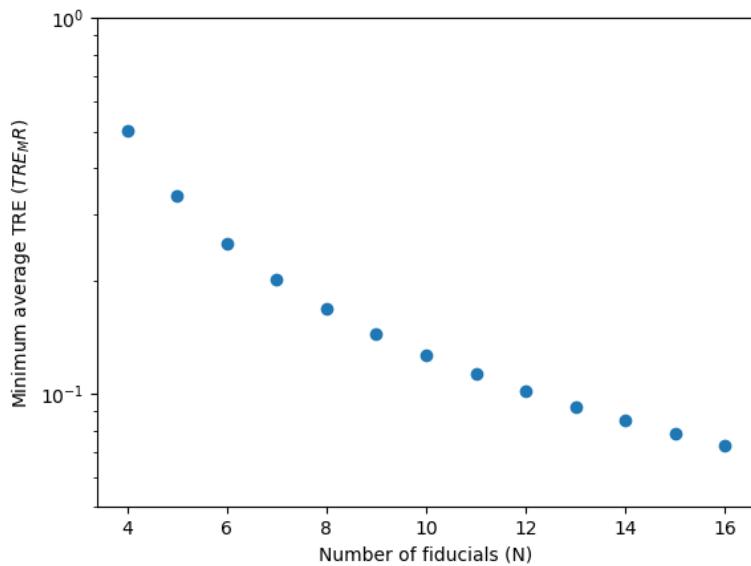


Figure 3.4.: The change in target registration error as the number of fiducial points is increased.

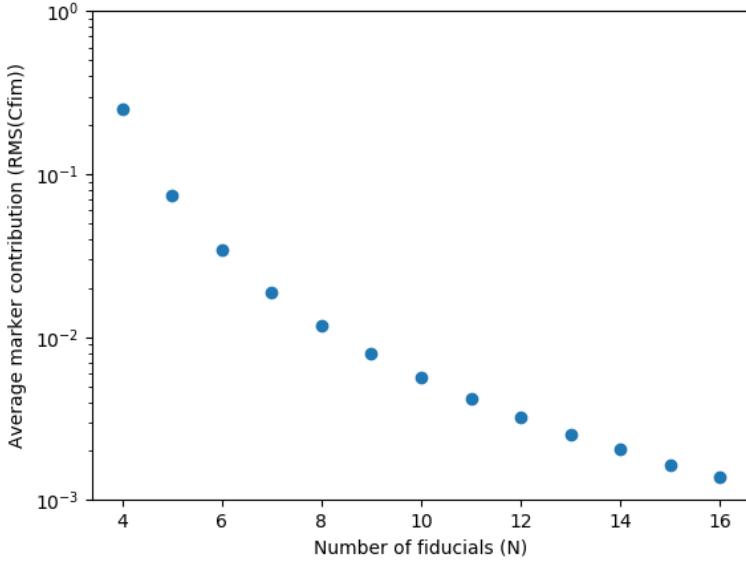


Figure 3.5.: The average contribution of a single marker to the target registration accuracy as the number of markers increases.

TRE_M falls as the number of fiducial markers is increased. The significance of each marker to the accuracy of the registration (Cfim) also falls quickly as N is increased, as shown in figure 3.5. This was calculated by taking the RMS average of the change in TRE_M when each marker is removed from M . This trend confirms that for a larger number of markers, each marker is less significant, so less error is introduced if a marker cannot be accurately located or used for the registration.

Figure 3.6 shows the optimal distributions found by the simulation for $N = 10, 16$. A number of trends can be observed from the distributions of markers for different values of N and applied to the general distribution of fiducial markers around a cylinder.

The first clear pattern is that the markers are distributed evenly between the very top and the very bottom of the sample – maximising the distance between the markers and the centre of the volume. This is a simple principle to apply when attaching markers to the sample and also has practical benefits for the imaging process because it means that the markers can be placed higher and lower than the soil in the sample tube. This means that areas of photon or neutron starvation, that could produce artefacts in the reconstructed data

3. Distribution of fiducial markers for the registration of multi-modal CT data

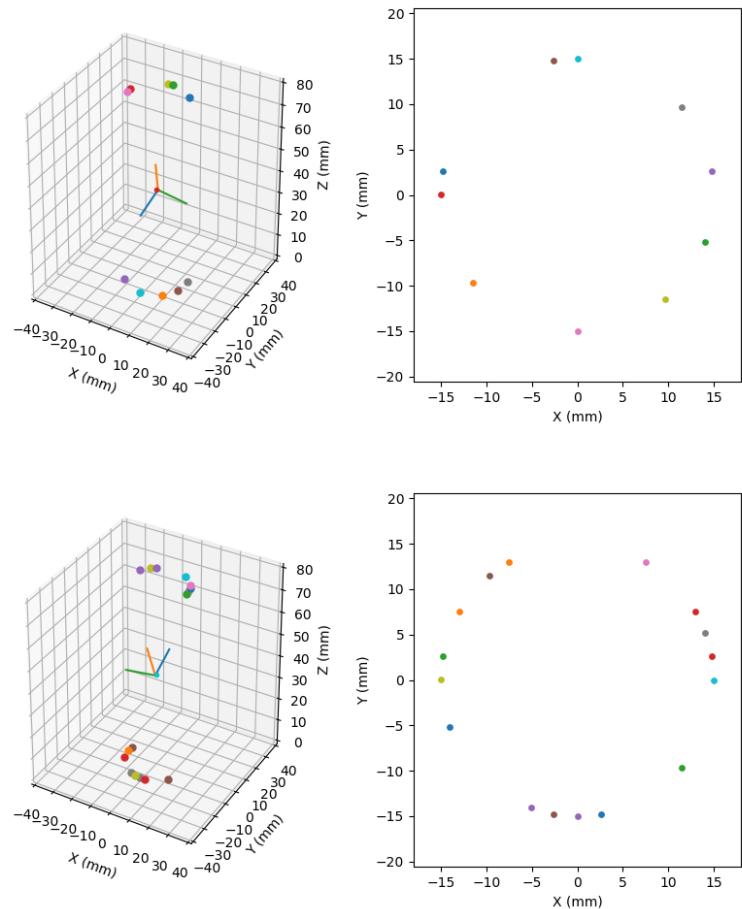


Figure 3.6.: The optimal fiducial marker distributions found by the simulation for $N = 10$ (top) and $N = 16$ (bottom).

3.2. Simulation

due to the markers, can be located away from the region of interest as the soil sample will not fill the very top and bottom of the sample tube.

Another clear pattern is that the markers are placed on a 180° arc at the top and another corresponding arc at the bottom rather than, say, alternating between the top and bottom of the cylinder. The markers are also evenly distributed around the cylinder, ensuring that the centroid of the fiducial points is close to the centre of the sample.

Following these results from the optimisation algorithm, some distributions were chosen to explore these patterns further. By evaluating distributions with an even - but alternating - distribution of markers between the top and bottom, it was found that this generates very similar results. When compared to arcs at the top and bottom, with the same positions in X and Y, the results were equal to 2 significant figures. The minor variations at greater precision did not conclusively show either arrangement to be consistently better than the other for all values of N . It seems reasonable that the simulation gives arcs because it begins optimising with a small number of markers and adds optimised markers iteratively up to the required total rather than starting with N markers and attempting to redistribute them all.

The optimal arrangements found by the simulation were compared to random arrangements to see how significant the differences were. One million five hundred thousand unique, random arrangements were evaluated for each value of N . Figure 3.7 shows the minimum, mean and maximum values from these tests as well as the optimal values (as seen in Figure 3.4). Figure 3.7a shows that a poor arrangement can increase the error by as much as an order of magnitude when very few markers are used. With higher numbers of markers, it becomes clear that the variation between good and bad arrangements becomes insignificant. Figure 3.7b shows the same data in the range where N varies between 12 and 16. It can be seen that the optimal solutions found by the simulation are better than any of the values found in the 1.5 million random arrangements but not by a significant amount. On average, a random distribution of $N + 1$ markers will give better results than the optimal distribution of N markers.

Equation 3.2 states that the expectation of target registration error squared will be proportional to the expectation of the FRE, the number of markers and the distribution of the markers. The simulation results show that with only four markers, the TRE can be reduced to half of the FRE and that it

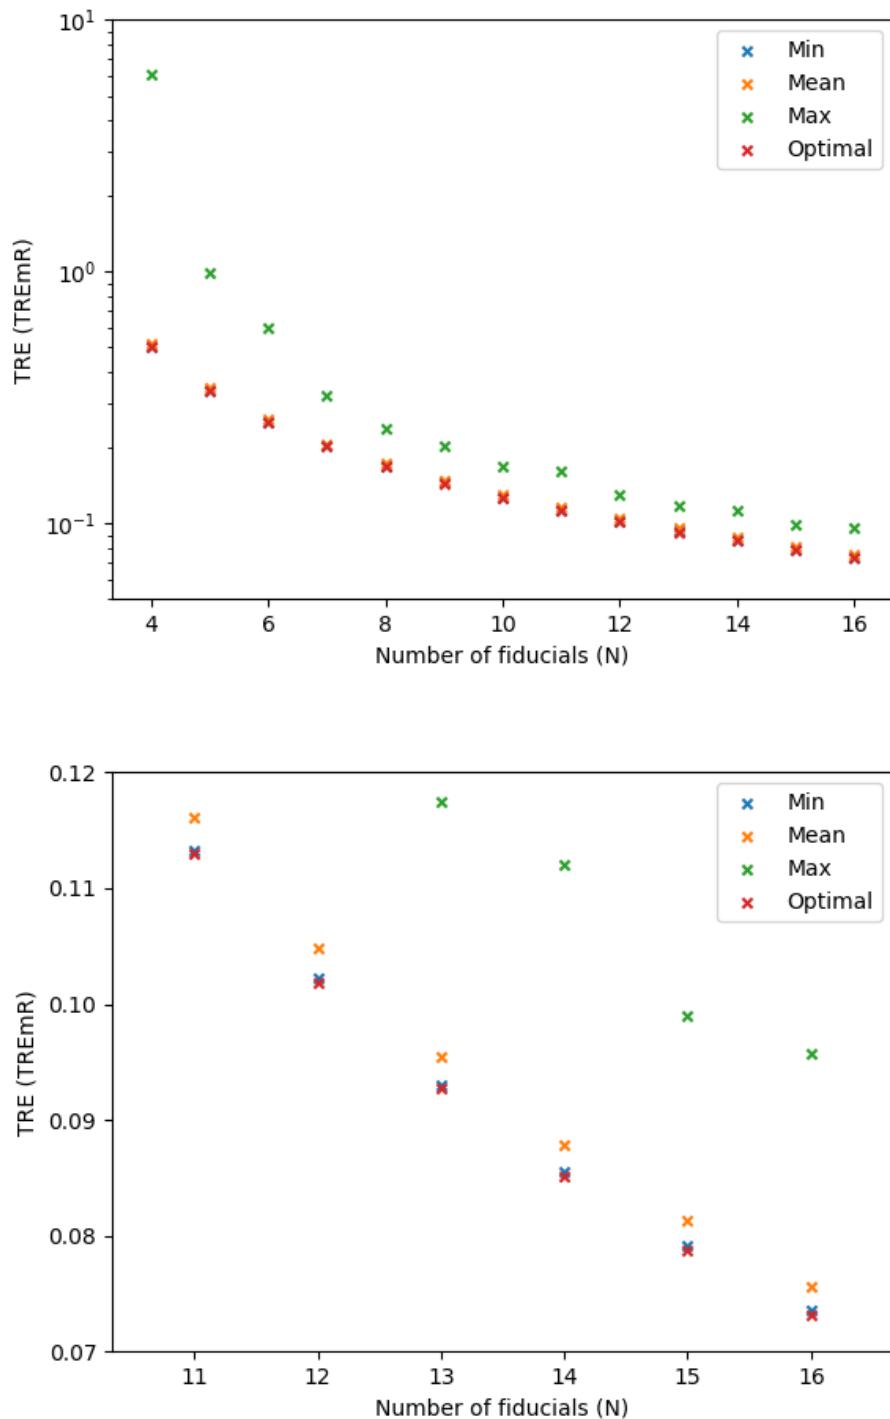


Figure 3.7.: A comparison of the optimal marker distributions to the mean, minimum and maximum TRE_M values found in 1500000 random trials. (A) (Top) the full range of N values. (B) (Bottom) $N = 12 - 16$.

can be reduced to less than a tenth of the FRE by using more than twelve points. Using a high number of markers makes a clear improvement to the result and also allows for the potential loss of a point without seeing a large drop in accuracy.

3.3. Materials and methods

3.3.1. X-ray and neutron imaging equipment

3.3.1.1. IMAT

This project is a pathfinder application for the new IMAT instrument at ISIS Neutron and Muon Source. IMAT is a combined cold neutron imaging and diffraction instrument designed to take advantage of the second ISIS target station to provide neutron radiography, neutron tomography, time-of-flight neutron imaging and spatially resolved diffraction scans [Kockelmann et al., 2013]. Construction of the imaging portion of the instrument was completed in 2016, and since then, IMAT has been running its imaging configuration, with the diffraction system due to be completed in 2022 [Burca et al., 2018]. Figure 3.8 shows a high-level schematic of the instrument.

ISIS TS-2 is a short-pulse source that operates at 40 kW and delivers pulses at a rate of 10 Hz. IMAT uses a cold (18 K), coupled liquid hydrogen moderator to slow the neutrons. A straight super-mirror neutron guide transports the neutrons from the target to the experimental area. Three choppers are placed within the guide to filter the beam. A T0 chopper removes fast neutrons and gamma radiation, then a pair of double-disk choppers define the wavelength band to ensure there is no frame overlap between successive neutron pulses. At the end of the guide is a pinhole selector that allows the aperture diameter (D) to be varied between five values to define different L/D ratios where L (the distance from the aperture to the sample) is 10 m [Burca et al., 2013]. This results in a total flight path of 56 m to the sample. Between the aperture and the sample, the beam travels through a 9 m evacuated flight tube and is shaped by five sets of jaws [Kockelmann et al., 2015]. The sample position has a combined translation and partial rotation system that is rated for up to 1.5 tonnes with a tomography stage for full rotation mounted on top.

IMAT has three detectors that can be changed depending on the application. An MCP camera with a $28 \times 28 \text{ mm}^2$ field of view and a $55 \mu\text{m}$ pixel size that

3. Distribution of fiducial markers for the registration of multi-modal CT data

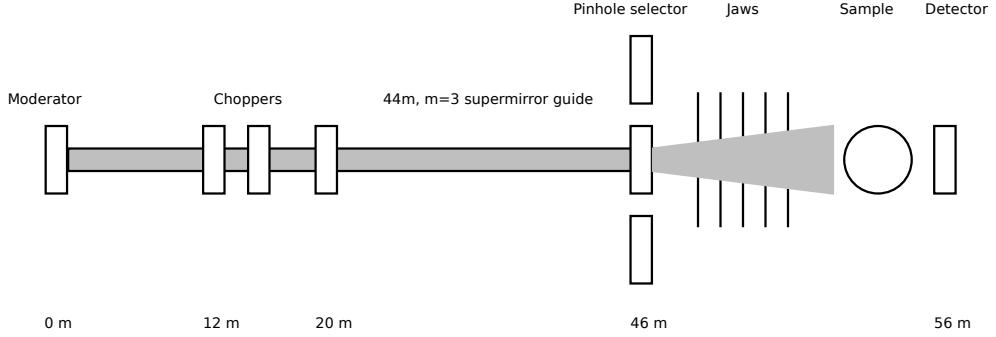


Figure 3.8.: A high-level IMAT imaging system schematic.

is used for TOF measurements [Tremsin et al., 2013]. An optical camera box with a field of view varying between $50 \times 50 \text{ mm}^2$ and $200 \times 200 \text{ mm}^2$ with a Zyla sCMOS 4.2 Plus camera which acts as an integrating detector with a range of scintillators [Finocchiaro et al., 2013]. The final detector is an active pixel sensor using the PImMS-2 CMOS with a $27.7 \times 27.7 \text{ mm}^2$ field of view and a $70 \mu\text{m}$ pixel size which can also be used for TOF applications [Pooley et al., 2017]. In this work, the optical camera was used exclusively except for a few test scans.

3.3.1.2. μ -VIS HMX

The μ -VIS X-ray imaging centre at the University of Southampton is a dedicated centre for X-ray Computed Tomography with a suite of seven scanning systems. Their modified Nikon/X-Tek HMX (225 kVp) [μ -VIS, 2018] was used to produce the X-ray data for this experiment. This is a customised, general-purpose X-ray CT and radiographic inspection system. It can take samples up to 300 mm in height, although this is reduced to approximately 150 mm if the robotic sample exchanger is used. The 225 kVp X-ray source can be configured for high resolution or high flux by using different anodes and uses a PerkinElmer PE1621 flat panel detector to capture the image. The detector is made up of a 2000×2000 matrix of amorphous silicon pixels with discrete gadolinium oxysulphide scintillators [PerkinElmer, 2008].

3.3.2. Cadmium imaging and registration experiment

A set of samples were imaged using both X-ray and neutron tomography and then registered to test the proposed registration scheme and demonstrate the

3.3. Materials and methods

complementarity of the two modalities for further studies.

Cadmium was proposed as a suitable material for fiducial markers because it has a large attenuation coefficient for both neutrons and X-rays, allowing easy segmentation in either modality. An initial X-ray scan, to confirm the suitability of the sample tube and markers, showed that the cadmium produced significant artefacts, primarily due to beam hardening and scattering. It was decided that the artefacts could be reduced to a satisfactory level provided the marker size was minimised and there were sufficient variations in the height of the markers to avoid streaking between two markers in a slice. Smaller cadmium pieces were cut with a maximum thickness of 1 mm, and a new sample tube was scanned to ensure these changes were sufficient.

Once the sample tube and marker configuration were shown to be suitable for imaging, a set of plants were grown and imaged using both IMAT and the HMX to produce a set of complementary volumes that could be used to develop and test registration techniques. Twenty lupin seeds were placed in wet paper towels to germinate. After six days, eight seeds that had begun to sprout were selected and transferred to the sample tubes. This experiment introduced new sample tubes made from boron-free quartz with an inner diameter of 14 mm and a wall thickness of 1.5 mm. Each tube had a single fibreglass wick to draw water for the plant. The wick was surrounded by 1 tsp of sand with particle sizes between 1.18 and 0.6 mm. Soil sieved to the same particle size range was then used to fill the tube, covering the seed around 20 mm from the top. The soil used in this experiment was a sand-textured eutric cambisol collected from a surface plot at Abergwyngregyn, North Wales (53014'N, 4001'W).

Each sample had a different quantity or arrangement of fiducial markers, as listed in Table 3.1. Figure 3.9 shows a diagram of three of these marker arrangements. Four of the sample tubes are then shown in figure 3.10. These marker distributions were selected to allow the simulation results to be compared to measured data. The samples were left to grow for 5 days before the scans began. The neutron scans were conducted first over two days, and the X-ray scans were collected over the following two days.

The neutron scans were conducted using the optical camera with a 135 mm lens, and this gave a field of view of 60×60 mm. A $60 \mu\text{m}$ ZnS/LiF scintillator with a surface area of 90×90 mm was used. The beam was shaped using the 40 mm pinhole, and the jaws were set from 40 mm to 70 mm to match the beam profile to the field of view as closely as can be achieved without

3. Distribution of fiducial markers for the registration of multi-modal CT data

Table 3.1.: The fiducial marker arrangements. FRE measurements and TRE_M values for each of the sample tubes that were scanned, as well as their products which are taken as an indicator of TRE. A representative selection of these are illustrated by a diagram in figure 3.9 and a photo of some of the prepared sample tubes is shown in figure 3.10.

Sample	Fiducial arrangement	FRE	TRE_M	TRE ($FRE \times TRE_M$)
1	5 markers in a vertical line	2.66	0.33	0.89
3	6 markers in arcs at the top and bottom of the FoV	6.49	0.25	1.62
4	8 markers in arcs at the top and bottom of the FoV	2.03	0.17	0.34
5	10 markers in arcs at the top and bottom of the FoV	11.02	0.13	1.38
7	9 markers in arcs at the top, middle and bottom of the FoV	4.33	0.14	0.62
8	9 markers in arcs at the top, middle and bottom of the FoV	2.85	0.14	0.41

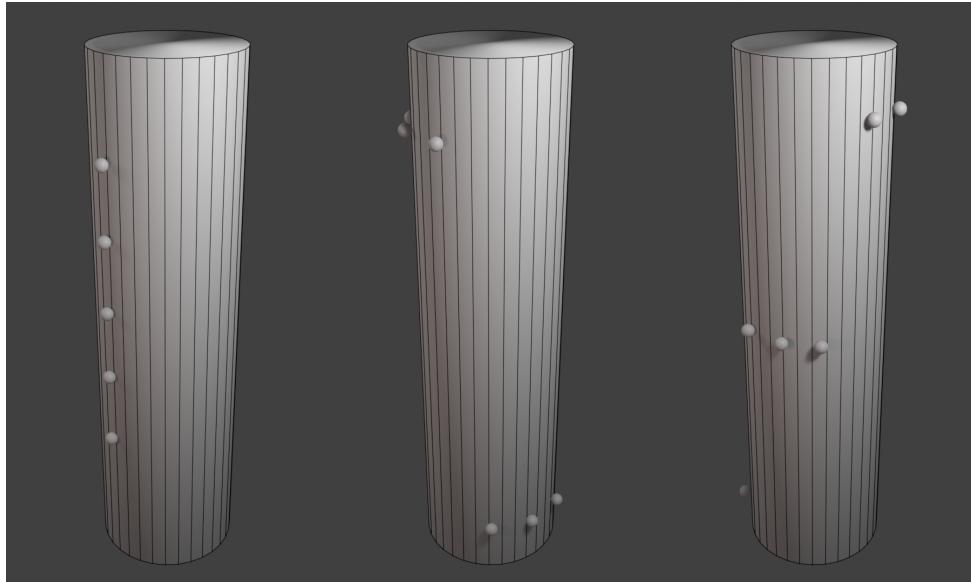


Figure 3.9.: Diagram illustrating the intended marker layouts for samples 1, 5 and 7.

3.3. Materials and methods



Figure 3.10.: Four quartz sample tubes with cadmium fiducial markers attached by aluminium tape. The green and white tape marks the bottom of the 60 mm IMAT field of view to aid in sample alignment. From left to right, the marker configurations are 1, 3, 5, 7.

introducing artefacts. The samples were positioned 15 mm from the detector, and 964 projections were taken with an exposure time of 30 s per projection. The projections were reconstructed with a voxel size of 31 μm , using the filtered back-projection algorithm in Octopus Reconstruction versions 8.9.3.4 and 8.9.4.2 [Dierick et al., 2004]. The X-ray scans were conducted in the HMX at 80 kVp and 87 μA . A total of 1571 projections were taken with four frames and a 500 ms exposure time. The projections were reconstructed using Nikon CT Pro 3D version 2.2.5386.22184 giving a voxel size of 19 μm .

Once the scans were reconstructed, registration was attempted by segmenting the cadmium pieces and taking their centres as fiducial points, before finding and applying the affine transform to best match these points. The cadmium centres were located using a threshold segmentation in conjunction with Fiji's 3D Objects Counter [Schindelin et al., 2012]. An affine transform can be determined to match the two point sets. This is achieved by removing the translation and scaling differences by finding the difference between the centroids of the two points sets and the root-mean-square deviations of the points from their respective centroids. The rotation component is then determined

3. Distribution of fiducial markers for the registration of multi-modal CT data

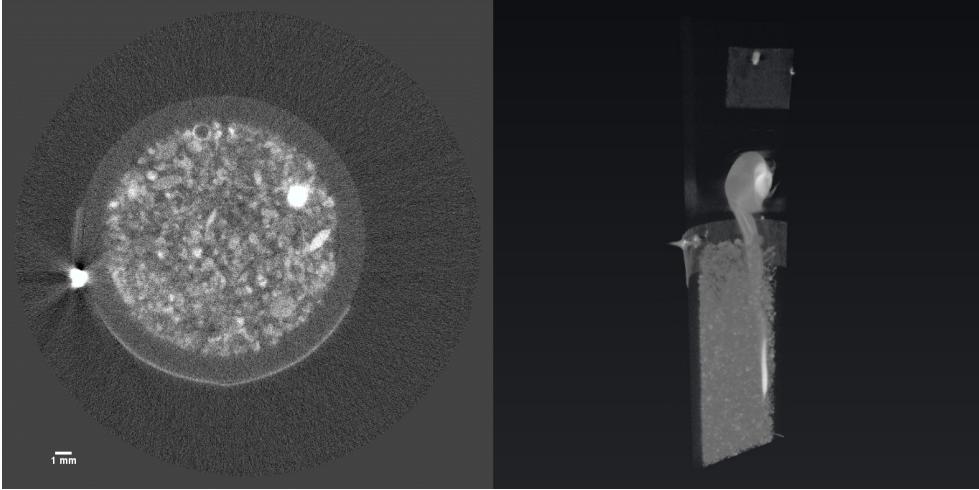


Figure 3.11.: A slice and partial volume render from the neutron scan of sample 7.

using Horn’s algorithm [Horn, 1987]. This gives a closed-form least-squares solution using unit quaternions to represent rotation. The unit quaternion representing the best rotation can be found as the eigenvector associated with the most positive eigenvalue of a symmetric 4×4 matrix composed from the point sets. After the transform was applied, the volumes were then cropped to matching dimensions. In order not to discard the higher resolution data in the X-ray scan, the neutron scan was considered the target volume. As a result, it was scaled up and re-sampled at a higher resolution than IMAT could have achieved to match the resolution of the X-ray data.

3.4. Results

Figures 3.11 and 3.12 show example slices and volumes from the reconstructed data collected using neutron and X-ray imaging, respectively.

Following the registration process, the fiducial markers were then re-segmented in each modality and the new positions compared to give the FRE for each sample. This is assumed to be a reasonable indicator of registration accuracy in the case of an affine transform. The FRE values are shown in Table 3.1. The mean FRE of sample 5 was so much higher than that of the other samples due to one marker, which was misaligned by 18.5 voxels. The mean FRE without taking that marker into account is only 4.5, which is far closer to the

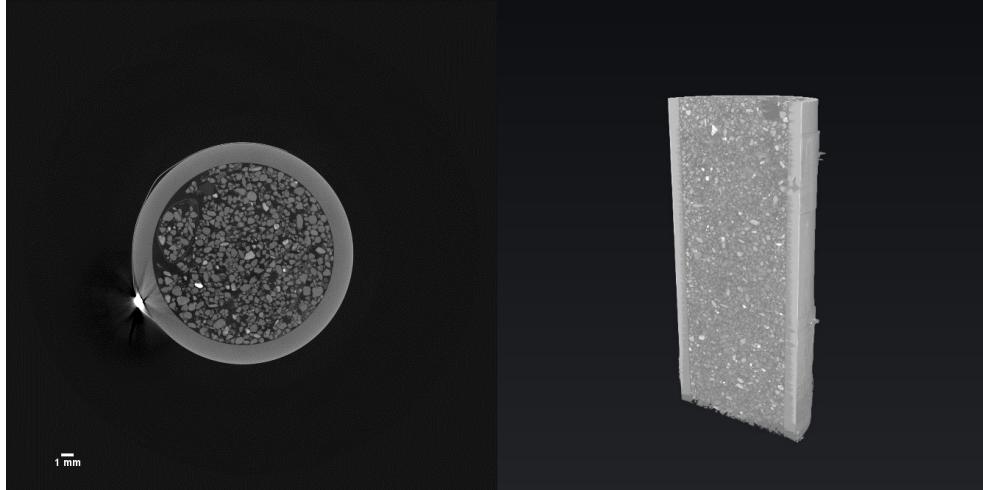


Figure 3.12.: A slice and partial volume render from the X-ray scan of sample 8.

typical FRE values produced by the other samples. This error was introduced as a result of a large FLE for this marker in the X-ray data due to an artefact. Figure 3.13 shows how the threshold segmentation detects an area far greater than that of the cadmium piece due to the artefacts surrounding it. This increase in the volume of the segmented marker pulls the measured centre away from the true value, and perhaps more importantly, away from the position found in the complementary modality to which it is to be matched.

Figure 3.14 shows a slice from sample 8 after registration. In addition to the combined data, each modality is shown separately. The side by side and overlaid comparison clearly shows the differences in contrast and signal to noise ratio between the two modalities but also the accuracy of the registration – particularly when observing the aluminium tape around the outside of the tube. This image also allows the complementarity to be seen clearly. Not only does the neutron data show the plant root with greater contrast than the X-ray data but it shows some of the soil particles that appear in the X-ray data while omitting others. This means that the combined data can be used to infer information about the different materials making up soil particles that could not be distinguished using X-rays alone.

3. Distribution of fiducial markers for the registration of multi-modal CT data

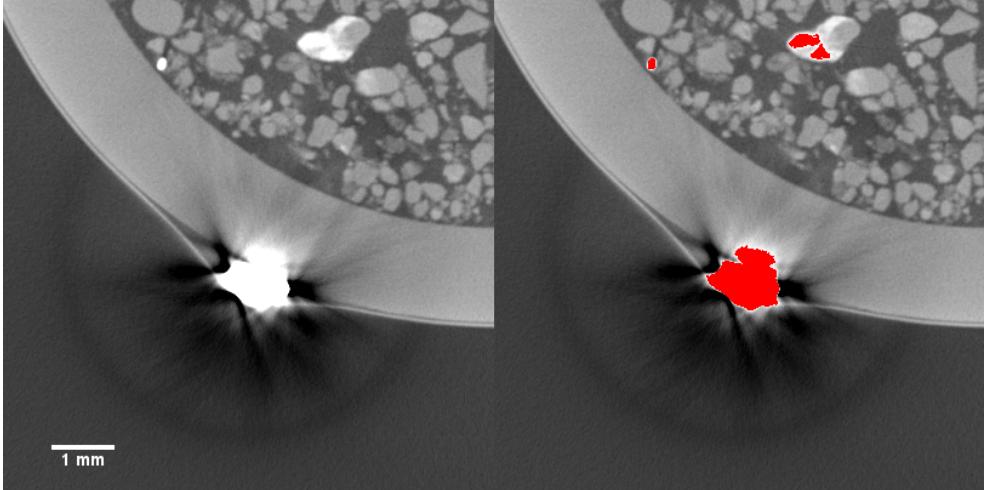


Figure 3.13.: Sample 5 had the worst registration accuracy despite having the most fiducial markers. This figure shows how an artefact prevents accurate segmentation (segmentation shown in red) of a cadmium piece, creating a large FPE for the corresponding fiducial marker.

3.5. Discussion and conclusions

The data collected in this experiment demonstrates the suitability of X-ray and neutron tomography for multi-modal studies, particularly of plant-soil systems. It has been shown that fiducial markers used in conjunction with a registration algorithm allow the data from the two modalities to be registered accurately, overcoming a lack of clear mutual information in the sample. It has also been demonstrated that more information can be taken by combining techniques than could be collected from either technique in isolation. The general target for “sufficient registration accuracy” is to be registered well enough for voxel-wise comparisons. This requires the volumes to be registered to the nearest voxel. This corresponds to a TRE value less than or equal to 0.5 which was achieved for a third of the samples in this experiment.

There is no correlation in these results between the FRE of the samples and the number and distribution of fiducial markers used. For example, sample 1 has the second-lowest FRE but had the worst set of fiducial markers because it had the fewest markers, and their arrangement was close to co-linear. Sample 5 had the greatest number of markers and therefore had the best set of fiducial points, but it showed the worst FRE of all the samples. The product of the FRE and TRE_M values (Table 3.1) gives an estimate for the TRE for each

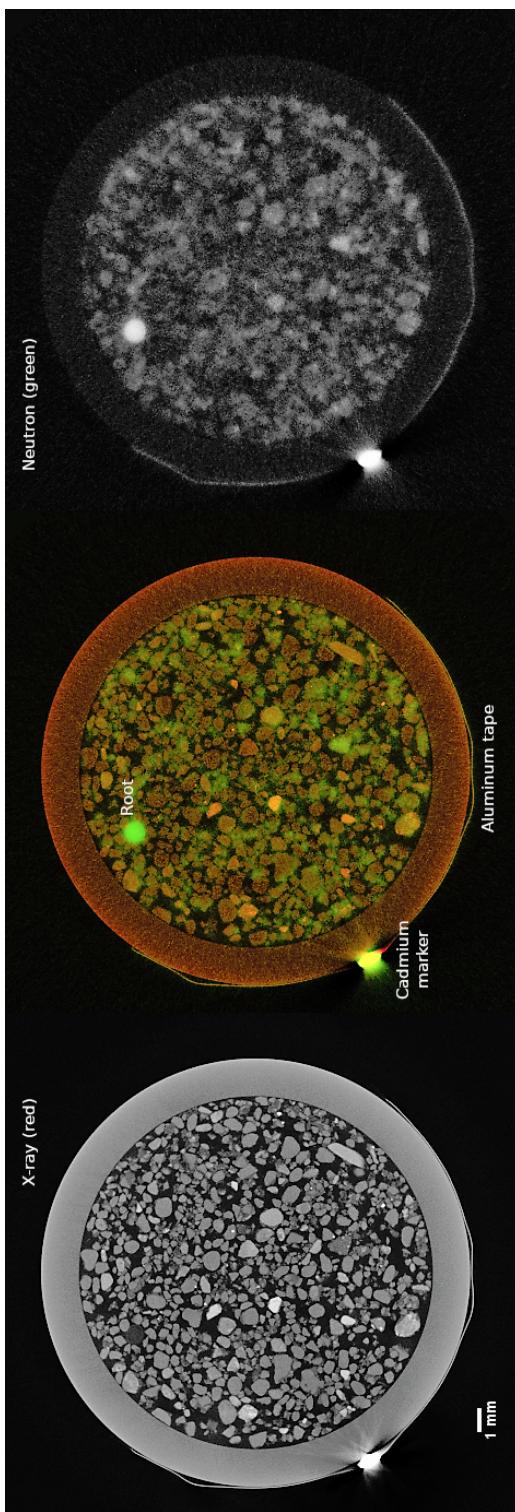


Figure 3.14.: A slice from sample 8 with the X-ray data on the left and in red and the neutron data on the right and in green. This slice shows the match of a fiducial marker, the accuracy of the registration and the complementarity of the modalities.

3. Distribution of fiducial markers for the registration of multi-modal CT data

sample. The variation in FRE is far more significant in this calculation than that of TRE_M , which shows that accurately imaging and locating the markers is more important to the overall registration accuracy than having the markers ideally distributed. As was suggested in the previous section, the registration of some samples could be improved by deliberately ignoring markers with high FLE from the registration process as the error they introduce exceeds the benefit of having an additional marker.

The higher FRE of sample 5 indicates that artefacts in the scans which affect the segmentation of the fiducial markers are the most significant source of error in registration (fig. 3.13). Cadmium's high attenuation introduces artefacts, and these have been shown to reduce the registration accuracy such that the effect of marker arrangement could not be tested in detail. This suggests that although it can be used as a fiducial marker material, it is not ideal. A material that would attenuate both X-rays and neutrons less would result in reduced artefacts, leading to better scan data and more accurate registration because large FLE values as a result of image artefacts are the primary contributor to FRE in the registered data.

In response to these results, further studies were planned to test the suitability of other materials, such as borosilicate as fiducial markers. Movement and other changes in the samples can be observed in the registered results due to the time and travel between the two scans of each sample. In particular, the seedlings begin to droop, and the top layers of soil can move considerably between the two scans. This is addressed in subsequent studies by booking equipment to conduct both scans at the same site with minimal delays between them. Furthermore, these studies integrate synchrotron X-ray imaging which will improve the resolution of the data and allow more complex imaging techniques to be introduced.

4. Borosilicate as a fiducial marker material and a sphere fitting algorithm for fiducial localisation

4.1. Introduction

The previous chapter showed that fiducial markers could be used to facilitate the registration of multi-modal CT datasets. Investigations were carried out to measure how the quantity and distribution of markers could affect the registration accuracy. Although registration was completed for six pairs of CT image data and principles for determining the distribution of markers were determined, it became apparent that the marker localisation could be a far more significant source of registration error, particularly given an unsuitable marker material; in this case cadmium. Cadmium was initially selected because it was highly attenuating to both X-rays and neutrons, which should allow for easy segmentation using a threshold as it would be the most highly attenuating substance in the sample. In practice, it was found that the cadmium pieces were so highly attenuating that significant artefacts were introduced, masking the shape of the cadmium pieces and so contributing significant fiducial localisation error (FLE) to the registration. Furthermore, cadmium is highly toxic and so machining it to produce small pieces was ruled out as prohibitively complex or expensive to do safely.

With this in mind, a more thorough investigation of potential marker materials was completed, and borosilicate markers were chosen to be used for the next experiment. This experiment was conducted using the I12 beamline at Diamond Light Source for high-energy X-ray imaging to complement the neutron imaging data collected at IMAT (as in the previous experiment). I12 gave some significant advantages over the HMX used in the previous experiment, particularly much faster scan times and being on the same site as IMAT so that samples did not have to travel significant distances between scans. Re-

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

ducing the delay between the scans helps to minimise changes in the samples due to plant growth or water uptake. The reduced travel also minimises the exposure of the samples to shaking, which could alter the distribution and packing of the sand or soil particles. While the decrease in attenuation compared to cadmium did successfully remove the negative impact of cadmium on image quality, it presented a more challenging segmentation task for fiducial localisation as the markers no longer stand out as the most highly attenuating material in the samples. As a result, new image processing routines had to be developed for the registration to be completed.

4.2. Materials and methods

4.2.1. Materials survey

A range of materials was initially considered as alternatives to cadmium for producing fiducial markers. The main properties to evaluate are the attenuation of both neutrons and X-rays at relevant energies, the strength of the attenuation and the mechanism of attenuation (absorption being preferable to scattering). High attenuation remained a desirable property, although not so high as cadmium to avoid artefacts. Other factors considered were the material cost, hazards of handling the materials and the ease of obtaining or producing suitable quantities and shapes for markers.

Table 4.1 lists transmission values calculated for a range of elements and compounds considered when evaluating marker material options. For this simulation, the neutron wavelength was taken to be 2.6 Å since this is close to the peak flux wavelength of IMAT. For the X-ray values, 0.07 MeV is used as an approximation to the effective energy of the I12 beamline, although the precise value would depend on the particular filters used. The transmission values are calculated as the proportion of the incident beam that passes through the sample without being absorbed or scattered. The simulation uses the formulas given in 2.1 with X-ray mass absorption coefficients from the NIST standard reference database [Hubbell and Seltzer, 2004] and neutron scattering data from the Periodictable python package [Kienzie, 2020].

Boron and cadmium stand out clearly in the table as the strongest neutron attenuators with absorption approximately thirty times greater than the next most attenuating (cobalt) and no transmission expected. This simulation assumes a single energy level close to the peak. The energy spectrum produced

4.2. Materials and methods

Table 4.1.: Neutron and X-ray transmission values for a range of materials that were considered for markers. All the transmission values were based on a 1 mm path length. Only transmission is calculated for borosilicate because it is a mixture not a single element or compound.

Material	Neutron absorption	Neutron scattering	Neutron transmission	X-ray mass attenuation	X-ray transmission
Li	4.723	0.063	0.620	0.140	0.986
B	144.571	0.683	0.000	0.153	0.985
Al	0.020	0.091	0.989	0.240	0.976
Si	0.012	0.108	0.988	0.272	0.973
Sc	1.592	0.941	0.776	0.534	0.948
Fe	0.314	0.987	0.878	0.900	0.914
Co	4.890	0.509	0.483	0.978	0.907
Cd	168.866	0.301	0.000	4.363	0.646
Silica (SiO_2)	0.007	0.282	0.972	0.222	0.978
Boric Oxide (B_2O_3)	42.789	0.447	0.013	0.171	0.983
Ruby (Al_2O_3)	0.016	0.366	0.963	0.211	0.979
Borosilicate (SiO_2 : 80%, B_2O_3 : 13%, Na_2O : 4%, Al_2O_3 : 3%)	-	-	0.569	-	0.979

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

by IMAT means that there will be some transmission through both boron and cadmium (as the previous chapter shows) due to the higher energy neutrons in the beam. Lithium provides moderate neutron absorption with very low scattering but has very high X-ray transmission and so would not be suitable except in combination with a stronger X-ray absorber. Aluminium and silicon both show relatively high transmission in both modalities, which is why they were considered for sample containers. Iron can quickly be ruled out as its neutron scattering cross-section is larger than its neutron absorption cross-section. This leaves scandium and cobalt the best options from the elements considered as they show significant, but not problematically large, attenuation of both X-rays and neutrons. These both show high neutron activation - cobalt in particular with its decay to nickel with a half-life of 5.27 years, making them less attractive options. The compounds considered include boric oxide, which tempers the high neutron absorption of boron, although not to the degree required, and ruby, which was ruled out quickly since it is primarily a scatterer of neutrons rather than an absorber.

Boron-free quartz had been used for the sample tubes in the previous experiment and so had been shown to appear in both modalities. In the X-ray data, the sample tubes were well resolved, and so quartz would serve well as markers, provided that the segmentation process could be adapted to distinguish shapes as there would be no difference in attenuation compared to the sample tube. The sample tubes attenuated the neutron beam much less, however, so additional attenuation would be desirable. Borosilicate is inexpensive and readily available in various forms, including spheres with a 1 mm diameter. The addition of small amounts of boron would result in more highly attenuating markers that could be easily distinguished in the neutron modality. As mentioned above, the neutron absorption cross-section for boron is high enough to prevent transmission by itself, but as a dopant in a borosilicate marker, this is reduced to a useful 56.8%. This should ensure a sufficient drop in attenuation to avoid neutron starvation artefacts. Having a known spherical shape and size allows for the shape-based approaches to segmentation required for the X-ray modality in particular and also allows artefacts that might occur to be easily removed, which should lower the FLE.

Table 4.2.: The detector properties for each imaging camera and optics module at I12 in EH1. For EH2, the field of view for M1 and M2 are extended to 46×20 and 20×20 , respectively [Connolley et al., 2020].

Camera		PCO.edge		MIRO 310M	
Optical Module	Magnification	Resolution ($\mu\text{m}/\text{pixel}$)	FoV (mm)	Resolution ($\mu\text{m}/\text{pixel}$)	FoV (mm)
M1	0.346	18.53×18.53	46×12	38.10×38.10	46×12
M2	0.82	7.91×7.91	20×12	16.26×16.26	20×12
M3	2	3.24×3.24	8.0×7.0	6.67×6.67	8.5×5.3
M4	5	1.3×1.3	3.3×2.8	2.67×2.67	3.4×2.1

4.2.2. I12

In this experiment, the neutron data were collected at IMAT (3.3.1.1) and the high-energy X-ray data were collected at Diamond Light Source (DLS), using the I12-JEEP (Joint Engineering, Environmental and Processing) beamline. I12 was built on a short straight section of the Diamond storage ring and opened for user experiments in 2009. I12 is a high energy X-ray beamline for imaging, diffraction and scattering, offering two in-line experimental hutches that can receive either a polychromatic or monochromatic X-ray beam that operates at energies of 53-150 keV. I12 uses a 4.2T superconducting wiggler source with 21 full-field periods of 48 mm. The first experimental hutch (EH1) is 51 m from the source and has a maximum beam size of 50 mm (horizontal) by 15 mm (vertical). The second experimental hutch (EH2) is 94 m from the source and allows for a maximum beam size of 94 mm (h) by 28 mm (v) [Drakopoulos et al., 2015]. This means that EH2 is well suited to large samples or experiments with large apparatus.

A variety of detectors are available at I12 to suit a variety of imaging and diffraction experiments. I12 has developed a modular detector system for imaging that allows the choice of scintillators, one of four custom optics modules and either a PCO.edge or MIRO 310M camera. Typically the PCO.edge is used when high resolution is desirable and the MIRO 310M when high-speed imaging is more important. The four optical modules give different options for the trade-off between resolution and field of view, shown in table 4.2.

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

Table 4.3.: Sample conditions. The particle size classifications are given in Table 4.3b.

Sample	Plant	Material	Particle size	Water	Injection
1	N	Sand	Large	N	H ₂ O
2				Y	D ₂ O
3			Small	N	H ₂ O
4				Y	D ₂ O
5		Soil	Large	N	H ₂ O
6				Y	D ₂ O
7			Small	N	H ₂ O
8				Y	D ₂ O
9			Large		
10					
11					
12				Y	D ₂ O
13	Y	Soil	Large		
14					
15			Small		
16					
17					

(b) Particle sizes		
	Small	Large
Soil	< 0.6 mm	0.6 – 1.8 mm
Sand	0.3 – 0.6 mm	0.6 – 1.18 mm

4.2.3. Borosilicate trials

Seventeen samples were prepared for this experiment with a variety of conditions as given in the following table (4.3). Each sample was first scanned using I12 at DLS and was then scanned with IMAT at ISIS over the following five days.

For this experiment, lupin seedlings were grown. These were initially placed on moist paper towels eighteen days ahead of the experiment to begin germination and were transferred to sample tubes filled with soil a week later. The soil was sourced and prepared as described in 3.3.2.

The sample tubes were cylinders made from boron-free quartz, open at both ends so that water could be drawn up through custom-made bases. The inner

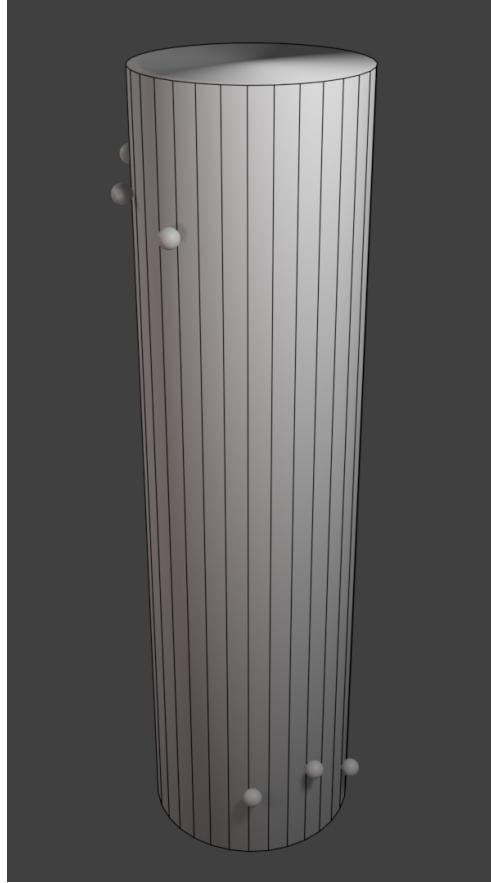


Figure 4.1.: A demonstration of the intended marker distribution. The markers are located in arcs near the top and bottom of the sample with minor variations in their height to ensure that they do not contribute attenuation in the same slices.

diameter was 14 mm, and the wall thickness was 1.5 mm. The height of each tube was 100 mm. A number of fiducial markers were attached to these samples in loose arcs near the top and bottom of each tube as shown in figure 4.1. Variation in the height of these markers was included to avoid multiple markers contributing attenuation to a single slice as this would increase the likelihood of artefacts in the reconstructed data.

Borosilicate spheres with a 1 mm diameter were used as markers for this experiment and were attached to the outside of the tubes using aluminium tape. The markers positions were not precisely aligned and the number of markers varied between samples with an average of 9.4 markers per sample (see table 4.5 for the number of markers on each sample). The beads used

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

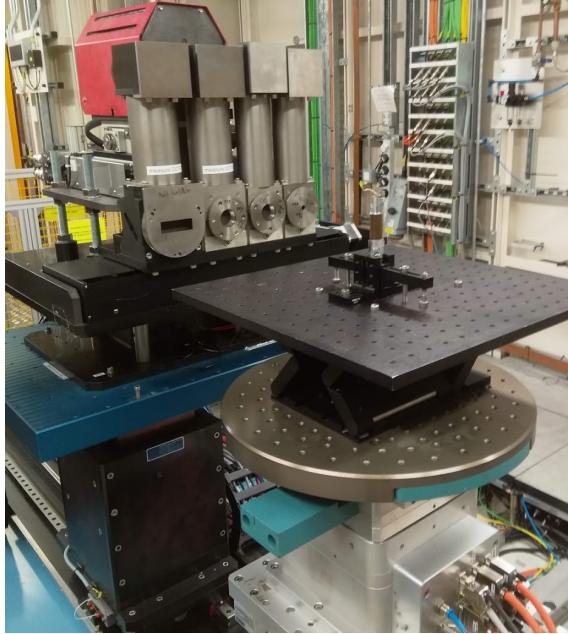


Figure 4.2.: A sample positioned in front of the detector at I12.

had a $\pm 10\%$ size variation, and all calculations were made with the following assumptions of composition: 80% SiO_2 , 10% B_2O_3 and density: 8 g/cm³.

The scans at I12 were conducted in EH1 with the PCO.edge camera and M2 optics giving a voxel size of $7.91 \times 7.91 \mu\text{m}$ and a field of view of $20 \times 12 \text{ mm}$. A 0.004 s exposure was used, and 1800 projections were produced from a 180° scan with a 0.1° step size. In order to get the best possible data from I12's field of view and resolution capabilities, each sample was scanned using five vertical steps. Each different height step had sufficient overlap to recombine the data after reconstruction to create a complete tomogram of the sample.

All of the I12 data was reconstructed at the beamline during the experiment using the in-house Savu reconstruction pipeline [Atwood et al., 2015].

Each sample was scanned multiple times at IMAT. First, a tomography was taken. Then following this, a syringe pump was used to inject 0.25 ml of H_2O or D_2O (depending on the sample) into the sample at a rate of 0.05 ml/min. Once the injection was complete, a series of 720 radiographs were collected with a 10 second exposure to observe the fluid spreading from the injection point through the rest of the sample. After this series of images was recorded, the flow rate of the injected fluid had dropped sufficiently for a second tomography to be taken of the sample.

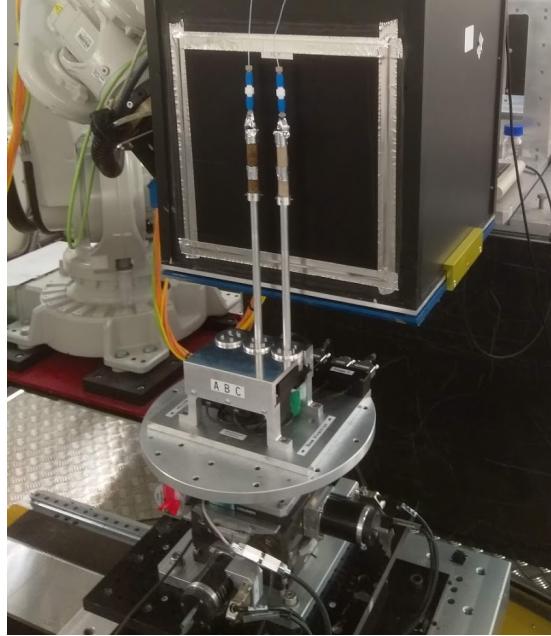


Figure 4.3.: The experimental setup at IMAT. Two samples are placed in front of the detector on the multi-axial rotation stage.

Each tomography at IMAT was conducted over 360° with a step size of 0.6909° , giving 521 projections. IMAT was configured to use the 40 mm pin-hole for the first four samples to be scanned (9, 11, 14 and 17 in table 4.3), but this was changed to 60 mm to give increased exposure for the remaining samples. The pixel size for these scans was $56.7 \times 56.7 \mu\text{m}$. There was a 22 mm sample – detector distance for each scan, and a pair of samples were scanned side by side using the multi-axial rotation stage.

The multi-axial rotation stage is a new piece of apparatus that was first tested at IMAT during the preliminary tests for this experiment (Experiment RB1810843). It allows up to three samples to be placed in the beam simultaneously and each turned on its own, independent axis. This increases the instrument's throughput significantly and therefore helps to overcome one of the limitations of neutron imaging: the comparatively long scan times. The stage is, of course, only useful for samples that require less than half of the width of the field of view, but the tall thin samples of this experiment meet this requirement comfortably.

While it would be possible to fit three samples into the maximum field of view given the 50 mm spacing between the centres of the axis, there are a

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

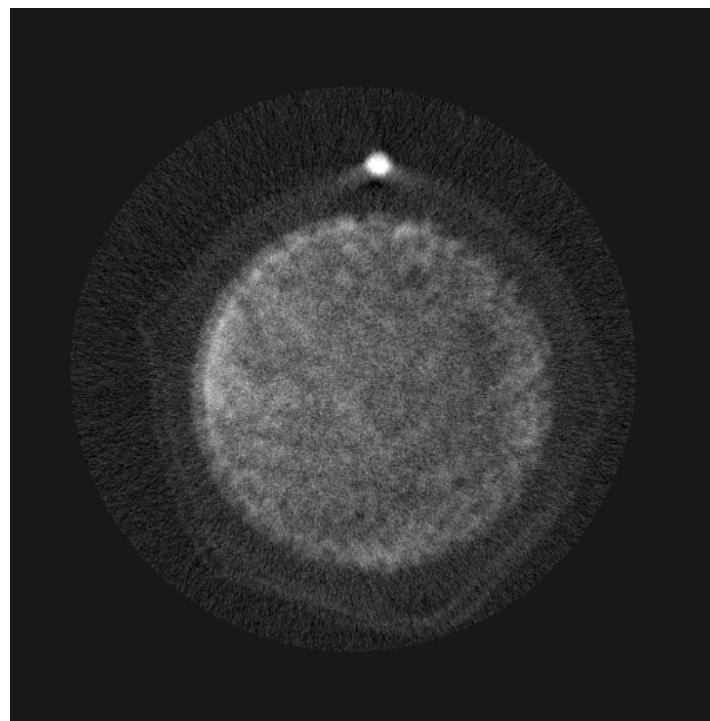
couple of trade-offs associated with this which led to the decision only to use two axes. Firstly, the change in the field of view width requirement would lead to a substantially reduced resolution. Secondly, the IMAT beam profile shows that the neutron flux is greatest at the centre of the beam and reduces with distance from that point [Minniti et al., 2018]. This means that the central sample would receive the neutrons required for the desired image quality more quickly than the two samples closer to the sides. Either longer scan times would be required for these to receive a greater dose, or lower image quality would have to be accepted. Given these limitations, it was decided that scanning two samples at a time with each remaining close to the centre of the beam and receiving a similar dose to each other in a shorter time was the preferred option. This results in each sample receiving slightly lower flux than a sample in the centre of the beam, but this can be offset with a small increase in exposure time. This increase in the scan time is negligible compared to the time saved by scanning two samples simultaneously.

The IMAT data from this experiment was reconstructed using Octopus as in the previous experiment. Initially, there were problems as a set of reconstruction parameters could not be found to reconstruct the whole volume accurately. Figure 4.4 illustrates the problem with a slice from Small Soil 5. Sub-figure (a) shows how 205 is not the correct centre of rotation for the slice due to the evident doubling that is visible at the inner edge of the sample tube and aluminium tape running around the outside of the sample. The image steadily improves through sub-figures (b) and (c) as the centre of rotation is changed first to 207, then to 209. The marker, however, clearly becomes less well resolved, splitting into two separate overlapping images. This type of problem where different parts of a slice were best reconstructed with different centres of rotation was found in a range of slices in a number of the datasets. Generally, if the centre of rotation drifted over the range of slices in a stack, this would indicate a tilt of the rotation axis relative to the detector. Octopus has CoR and tilt correction implemented to address this case. However, a change of the centre of rotation between areas within a single slice is not commonly encountered, so further investigation was required.

There did not appear to be any sample movement of the course of the scan, but when the radiographs at 0° and 360° were compared, it was found that they did not match. It was discovered by taking the difference between the first radiograph at 0° and the last 20 radiographs of the rotation that the

4.2. Materials and methods

(a) CoR: 205



(b) CoR: 207

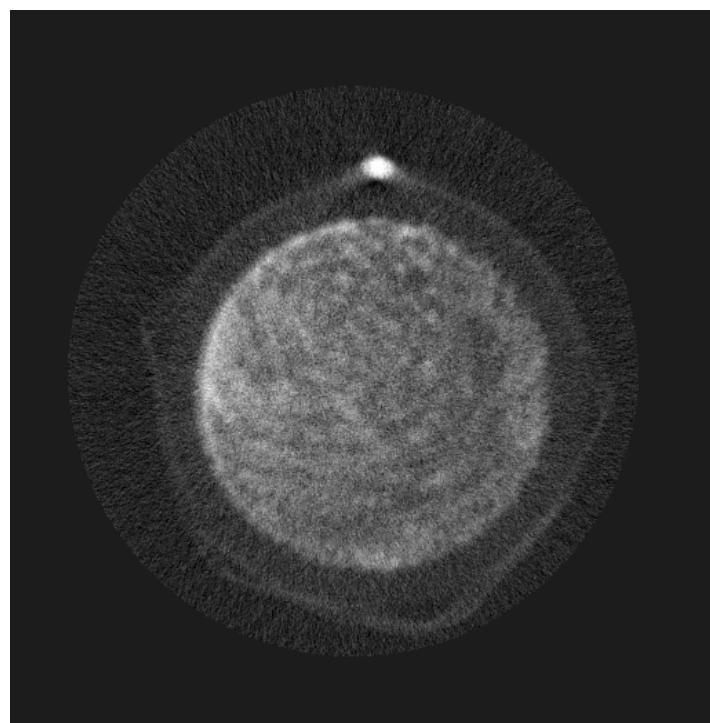


Figure 4.4.

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

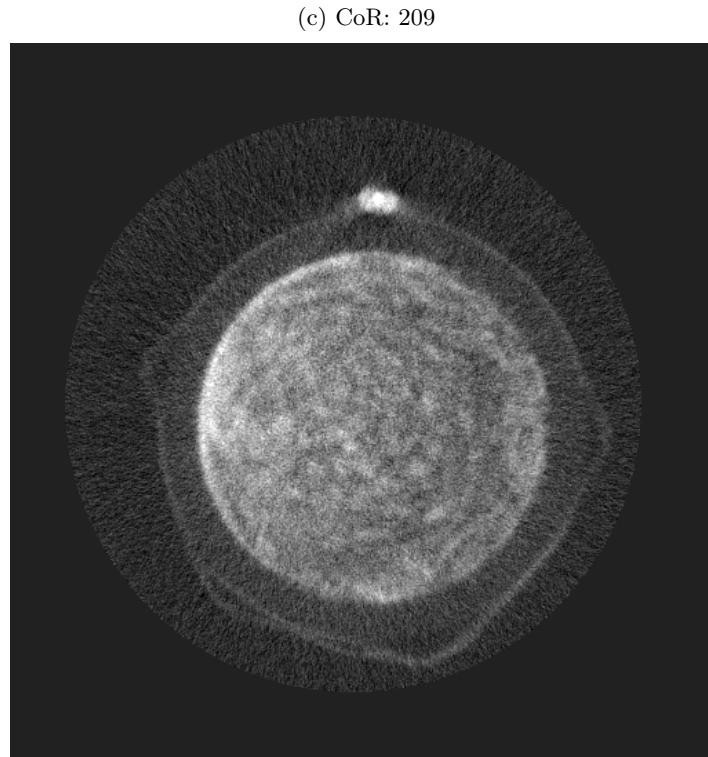


Figure 4.4.: Attempts to reconstruct a slice from Small Soil 5 with three different centre of rotation values. In (a), where the centre of rotation is 205, the marker at the top of the sample is well resolved - besides the artefacts - however, the tape around the outside of the tube and the inner radius of the sample tube are not. Once the centre of rotation has reached 209 in (c), the inner edge of the tube is reconstructed well with precise edges, but the marker is not and has split into two.

sample consistently reached 360° a few projections earlier than expected and over-rotated. The error in the rotation varied from scan to scan, with the error ranging from negligible to 3.9%, which can result in 360° scans being taken over 370°. Since the IMAT log data only recorded the nominal values rather than the true values, the rotation angles for each sample had to be estimated through analysis of the data and these estimated values fed back into the reconstruction algorithm.

Table 4.4 shows the number of projections early each sample reached 360° along with the corresponding angular step size and total rotation. From this data, we see that while neither axis was perfect, Rotation stage C was far more accurate than rotation stage B. It is notable that there is not any consistency in the error between the pre- and post-injection scans of each sample which shows that the inconsistencies occurred without human handling of the samples or rotation stage.

Once the rotation values were calculated from the number of projections it took for the sample to reach 360°, this information could be fed back into the reconstruction algorithm. This immediately solved the problem shown in figure 4.4 and led to consistently improved reconstructions. Figure 4.5 shows the slice from figure 4.4 reconstructed with the new parameters. This solution assumes that the error introduced by the rotation stage is introduced consistently throughout the scan and is not introduced in a few large steps. It would be very difficult to confirm this by looking at the data, but the improvements seen with the approach described suggest that the assumption is valid.

4.3. Image processing

The reconstructed data from I12 was rescaled to best use the dynamic range before the data type was converted from 32 bit real to 8 bit int. This reduced the scale of the data by a factor of four, making it far more practical to work with in all subsequent steps. As mentioned above, the samples were scanned in vertical steps, meaning that five partial scans of each sample required concatenation. While this task is commonly performed by hand, it is a repetitive and time-consuming process, and therefore it makes sense to automate the process when working with a large amount of data. Subsection 4.3.1 presents the algorithm used for this task. Once the I12 data had been rescaled and

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

Table 4.4.: The over-rotation measured for each neutron scan. The samples are split based on whether they were scanned on rotation stage B (4.4a) or rotation stage C (4.4b) of the multi-axial rotation stage.

Sample	(a) Rotation stage B					
	Pre-injection			Post-injection		
	Extra Projections	Step Size	Total Rotation	Extra Projections	Step Size	Total Rotation
Large dry sand	12	0.706°	368.471°	9	0.702°	366.316°
Large wet sand	11	0.705°	367.750°	13	0.707°	369.194°
Large dry soil	11	0.705°	367.750°	12	0.706°	368.471°
Large wet soil	12	0.706°	368.471°	6	0.698°	364.186°
Large soil 2	9	0.702°	366.316°	11	0.705°	367.750°
Large soil 4	11	0.705°	367.750°	10	0.703°	367.031°
Small soil 2	14	0.709°	369.921°	11	0.705°	367.750°
Small soil 3	11	0.705°	367.750°	11	0.705°	367.750°
Small soil 5	10	0.703°	367.031°	12	0.706°	368.471°

Table 4.4.

(b) Rotation stage C

Sample	Pre-injection			Post-injection		
	Extra Projections	Step Size	Total Rotation	Extra Projections	Step Size	Total Rotation
Small dry sand	2	0.692°	361.385°	2	0.692°	361.385°
Small wet sand	1	0.691°	360.691°	3	0.694°	362.081°
Small dry soil	1	0.691°	360.691°	2	0.692°	361.385°
Small wet soil	2	0.692°	361.385°	2	0.692°	361.385°
Large soil 1	3	0.694°	362.081°	2	0.692°	361.385°
Large soil 3	2	0.692°	361.385°	3	0.694°	362.081°
Large soil 5	1	0.691°	360.691°	2	0.692°	361.385°
Small soil 4	0	0.690°	360.000°	1	0.691°	360.691°

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

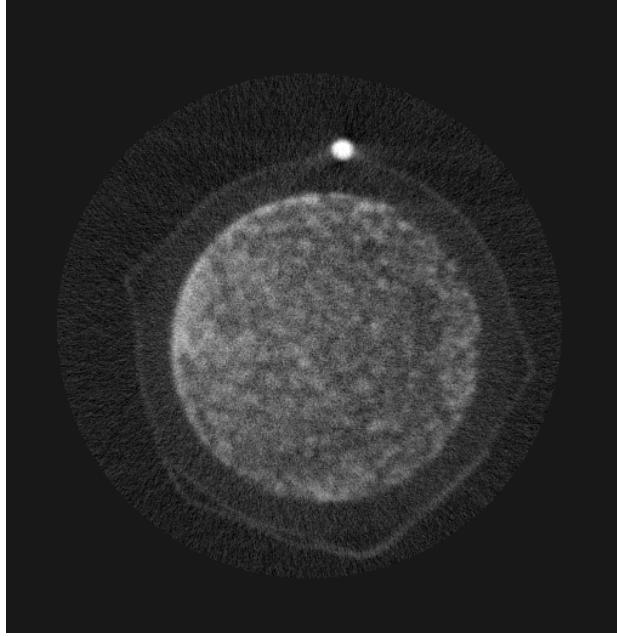


Figure 4.5.: The slice from Small Soil 5 shown in figure 4.4, this time reconstructed assuming a total rotation of 367.031° . The whole sample appears to be correctly reconstructed with the new parameters.

concatenated, with the excess data from overlapping regions removed, its total size was reduced to approximately 1 TB making the volumes far easier to use in subsequent steps.

Registration of the X-ray and neutron data depends on the accurate segmentation of the fiducial markers in each modality. Using a threshold-based segmentation method as in the previous experiment (3.3.2) was not an option with the borosilicate markers since they do not stand out as the most highly attenuating objects in the data. This experiment used 161 markers compared to the 47 used in the cadmium marker experiment, and the larger datasets produced at I12 were also too big to be analysed by the 3D object counter built into Fiji. These factors mean that precisely locating each of the markers by hand would not be a time-efficient option. An algorithm was written to find the markers in an automated, scalable, and repeatable way by finding spheres within volume data. The design and implementation of this tool are given in subsection 4.3.2. Once the markers had been located, registration could be completed using the same method given in 3.3.2.

4.3.1. Stack concatenation

The algorithm used to automate stack concatenation consists of the following steps:

1. Find and remove the overlapping region
2. Calibrate the levels to correct differences in brightness and contrast.
3. Concatenate the calibrated datasets

The algorithm is written with the assumption that the rotational axes of the volumes are aligned correctly and share a common orientation so that the volumes only need to be adjusted by translation in one dimension. The algorithm is applied recursively to concatenate more than two stacks, adding a different dataset to the previously concatenated data with each iteration.

To remove the overlapping region, the algorithm takes the last 20% of one stack and a user-specified proportion of the other. It then examines these regions to find the pair of slices that most closely match. The closeness of a match is evaluated by taking the difference between a pair of slices and returning the standard deviation of the result. Using the standard deviation rather than the difference allows the match between two slices to be evaluated without absolute grey levels becoming a dominant factor rather than the relative differences. Tests were also conducted with an edge detector applied after the difference was taken to see if this would enhance the results further, but this did not improve the performance of the algorithm.

This method for evaluating the closeness of a match between a pair of slices is used repeatedly to find a matching pair of slices. First, every fifth slice in the final portion of one stack is compared to a fixed slice in the second stack. The slices that most closely match from this step are passed as initial values to a minimisation function that keeps the slice in the second stack fixed. This gives the offset between a slice in one stack and the corresponding slice in the other. Finally, this information is passed to another optimisation function that takes the offset and looks for the best match without keeping one dataset fixed to a single slice. The output of this process is the pair of slices in the overlapping regions of the two datasets that most closely match.

Once this is established, the overlapping portions of the stacks are cropped to these points so that the stacks can be concatenated. There may be differences in the grey levels of the two volumes, though, so calibration is carried out

4. *Borosilicate markers and a sphere fitting algorithm for fiducial localisation*

using the matching slices that have been identified before the concatenation is applied.

The calibration of the levels between the datasets is conducted upon a set of representative sample areas within the slices from each dataset that is to become adjacent. Random sample areas of 10 pixels by 10 pixels are evaluated in each of the slices. If the range of grey levels is greater than 0 and the standard deviation is below 20 for each of the slices, then the region is added to the list of sample areas to use for the calibration. Once twenty such regions have been selected, a line of best fit is found for the twenty ordered pairs of grey levels, giving a gradient and an intercept to relate the two slices. The intercept is added to the levels of the slice to be calibrated, and then the levels are multiplied by the gradient resulting in a calibrated pair of datasets.

The full stack concatenation code is included in appendix A.2.

4.3.2. **Sphere fitting**

The current literature for 3D object recognition show that progress is being made in the use of machine learning and computer vision. This has been driven in part by medical imaging [Kawakami et al., 2020] but also other applications such as baggage scanning at airports [Yang et al., 2021]. Unfortunately for this work, these methods depend on training convolutional neural networks with large amounts of data (of the order of 10,000 scans) which is well beyond the scope of this work. More generalised approaches that do not require large quantities of previous scan data but instead are based on the identification of primitive shapes such as spheres appear to have been neglected in favour of these powerful but data hungry techniques. One helpful approach that can be taken from this research however is the method of considering 3D data at scale by considering 2D slices along each axis and then recombining that analysis as in the slice-and-fuse strategy presented by Yang et al. [Yang et al., 2021].

With this in mind, a custom sphere fitting algorithm was produced that would take advantage of more mature 2D shape recognition techniques and extend them to 3D without invoking the additional overhead and requirements of machine learning approaches. The algorithm is based upon the fact that any cross-section of a sphere is always a circle, regardless of the orientation, as shown in figure 4.6. A table listing the centres and radii of circles found within the volume can be produced by looking through 2D slice data along each axis. Spherical objects will appear as lists of circles in successive slices

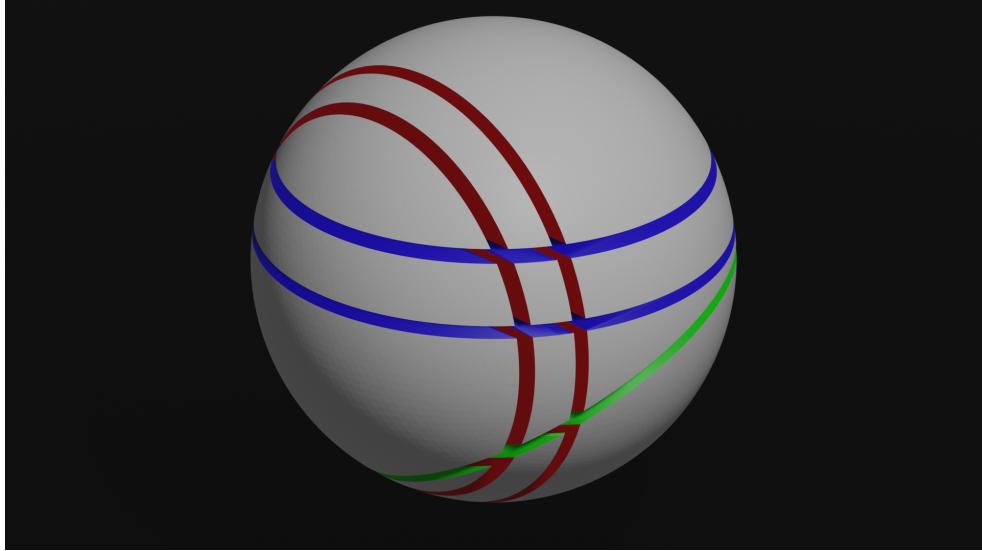


Figure 4.6.: The cross-section of a sphere is always circular, regardless of orientation.

with a common centre point and a radius that increases then decreases with a maximum at the third component of the sphere's centre point. This pattern should be observable in each axis, with the common values for the centre of the sphere being given by the circle centres and the maximum radius.

In practice, data is never free of noise and errors, so an algorithm's implementation must be resilient to inaccurate or missing data. In this case, that could mean that circles are not found in every slice that contains markers or that circles are found with incorrect centres or radii. To overcome this, once the circle centres are found using the Hough circle transform (implemented in OpenCV [Bradski, 2000]), the list is then passed to the DBSCAN clustering algorithm, which is part of the Scikit-learn package for machine learning in Python [Pedregosa et al., 2011]. The DBSCAN algorithm looks for areas of high density which in this case means circles with centres that are close together. By looking for local clusters of circles, the circles corresponding to specific markers can be identified and averaged to estimate the centre coordinates of the marker. Provided that a cluster contains circles along all axes, the value for the centre of the sphere in each axis is calculated without using the circles that are in a plane perpendicular to the axis. This prevents the calculated centre point from being skewed if the circles found in a particular axis for a given marker are not evenly distributed along that axis.

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

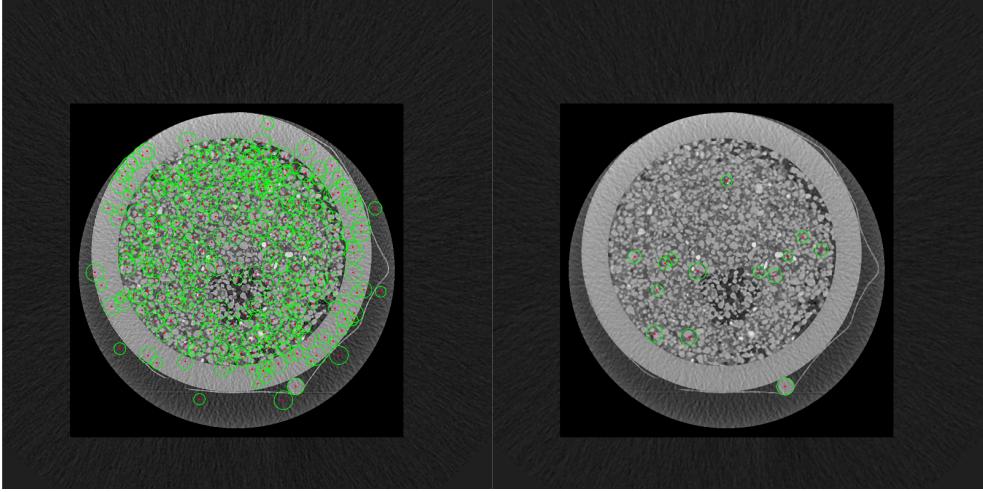


Figure 4.7.: The images show the results of the Hough transform applied to the same slice before (left) and after (right) the application of a median filter with a kernel size of 7. Each circle found by the transform is marked with a red centre and a green circumference.

The different modalities each have different scales and resolutions. As a result of this, parameters for both the Hough circle transform and the DBSCAN clustering algorithms require tuning for each modality. To get the best performance from the Hough circle transform, it is important to ensure that the sensitivity is high enough not to miss any markers. This leads to several false positives, but these can be reduced substantially by applying a median filter to smooth the image. Figure 4.7 shows the Hough transform applied to a slice from an X-ray dataset before and after a median filter is applied. Both identify the marker and some false positives, but it is clear that the filter reduces the number of extra circles. Furthermore, all the remaining circles are within the sample tube. Many of these extra circles are filtered out by the DBSCAN algorithm as it looks for clusters of circles that are close to each other.

The algorithm can still lead to false positives if other objects are found in the data that are approximately circular in at least one axis and have a similar radius to the markers. In this experiment, the sand particles, in particular, were commonly identified as false positives. To remove the majority of these, markers were filtered based on their distance from the centre of the sample. Any spheres found within the inner radius of the sample tube were discarded, and the remainder were kept. After this step, the number of false positives remaining was typically below ten per scan, and so they were removed by

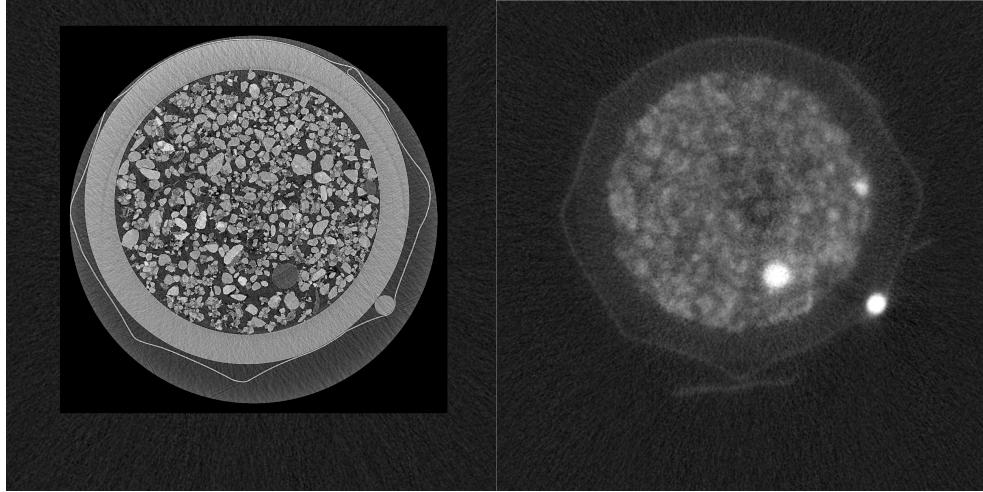


Figure 4.8.: Slices from the X-ray (left) and neutron (pre-injection) (right) scans of sample LargeSoil3 showing one of the fiducial markers. The marker is easily distinguishable in each modality, primarily by shape in the X-ray image but also by the strength of attenuation in the neutron modality due to the boron dopant. It can be seen that the types of artefact present around the markers in figures 3.11, 3.12, 3.13 and 3.14 have been eliminated.

hand. See A.3 for the full code listing.

4.4. Results

In this experiment, borosilicate fiducial markers were used with the aim of reducing artefacts to enable more accurate registration compared to the results presented in Chapter 3. Figure 4.8 shows example slices from the sample LargeSoil3, which contain one of the markers and demonstrate that the use of borosilicate has been successful in this regard as the markers are very clearly distinguished and the kind of artefacts seen around the markers in figures 3.11, 3.12, 3.13 and 3.14 are significantly reduced. Some scattering artefacts are still present in the neutron modality, especially when multiple markers appear in a slice (see figure 4.11). These no longer distort the shape of the markers to the detriment of the segmentation, however, with one exception discussed in section 4.5.

While it is known that the X-ray scans will have better image quality than the neutron data, both in terms of resolution and signal to noise ratio, it is

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

helpful to quantify this for these particular datasets. The X-ray scans consistently have a point spread function (PSF_v) of 3 voxels when measured between the sample tube and air. The PSF_v varies more between scans and even slices with the neutron data but is typically in the range of 20 to 25 voxels.

With no clear difference in attenuation between the markers and the sample tube in the X-ray scans, this experiment called for new methods of segmenting the markers. Given the increased number of markers and the importance of repeatability, an automated system for marker segmentation was developed, as described in section 4.3.2. The markers were found twice using this new tool: once to inform the registration and then a second time after the registration had been applied to evaluate the accuracy of the registration. As the registration process included changes to the scale of the image and re-sampling, the sphere fitting algorithm needed to be re-tuned to look for a different range of radii before being applied to the registered data sets.

The accuracy of this new tool introduces a new source for error in the marker-based registration, so its performance must be considered. The algorithm did not find all of the markers placed on the samples. One of the reasons for this was poor sample alignment during a small number of the scans as not all of the markers were kept within the field of view at I12 where the field of view was not much bigger than the sample radius with the markers included. Figure 4.9 shows an example case of this where a marker has not been kept within the field of view, rendering it useless.

There were limitations of the algorithm, however, which led to further markers not being identified accurately. These are discussed thoroughly in section 4.5. Table 4.5 shows the number of markers that were attached to each sample as well as the number found and used in registration and found after registration to evaluate the registration accuracy. It is worth noting that markers the algorithm missed before registration were found by hand and added to ensure the best possible registration. The number of markers found after registration does exclusively count those found by the automated system. In both cases, any markers found in one or two scans of a sample but not all three were discarded from all subsequent steps, be that registration or further analysis, and are not counted in table 4.5.

The FLE values found for each sample are presented in table 4.6. The FLE value is taken as the RMS difference between the centres of corresponding

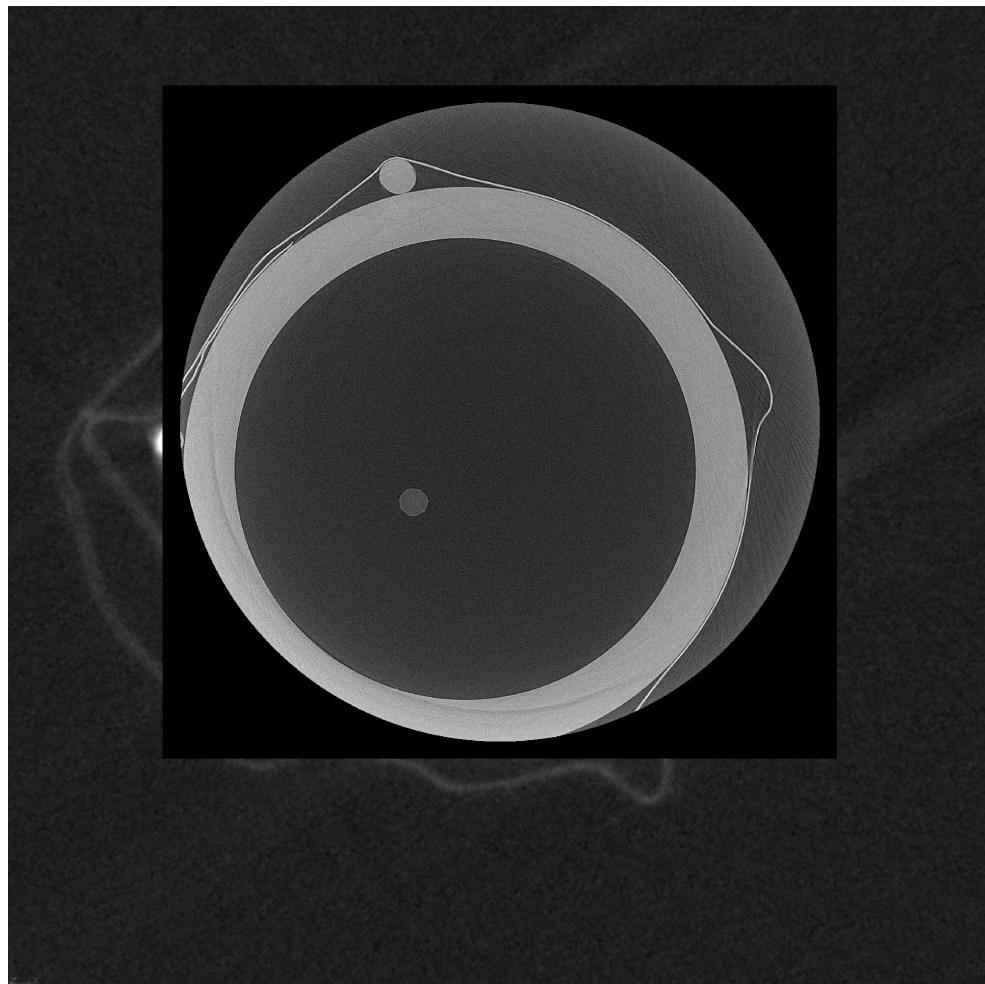


Figure 4.9.: An example slice from the X-ray scan of sample LargeSoil4 shows markers that have not been kept within the field of view and cannot accurately be located and used in registration or subsequent analysis.

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

Table 4.5.: The number of fiducial markers attached to each sample compared to the number of markers found both before and after registration by the sphere fitting algorithm. Any markers found in some but not all of the X-ray, pre-injection neutron and post-injection neutron scans were discounted from the last two columns as it is only markers found in all modalities that can contribute to the analysis shown subsequently.

Sample	Number of Markers	Number of markers used in registration	Number of Markers found after registration
Large dry sand	7	6	6
Large wet sand	8	6	6
Large dry soil	8	4	4
Large wet soil	10	8	7
Small dry sand	6	5	5
Small wet sand	9	8	4
Small dry soil	9	9	7
Small wet soil	10	7	4
Large soil 1	10	7	6
Large soil 2	10	8	8
Large soil 3	11	10	10
Large soil 4	10	8	2
Large soil 5	10	7	7
Small soil 2	11	7	6
Small soil 3	12	9	10
Small soil 4	11	11	5
Small soil 5	9	8	4
Total	161	128	101

4.4. Results

Table 4.6.: The FLE of each sample based on the two neutron scans. The final three columns show the FLE components along each axis. (All values in voxels.)

Sample	FLE	FLE (X)	FLE (Y)	FLE (Z)
Large dry sand	19.2	8.3	9.4	0.5
Large wet sand	21.3	9.3	9.1	0.5
Large dry soil	14.8	5.7	7.0	0.6
Large wet soil	35.8	12.1	15.4	0.7
Small dry sand	3.4	1.4	1.8	0.7
Small wet sand	6.8	3.2	2.3	0.3
Small dry soil	5.5	2.3	2.2	0.3
Small wet soil	4.9	1.8	1.8	0.7
Large soil 1	2.4	0.6	0.6	0.9
Large soil 2	6.5	2.2	2.7	0.5
Large soil 3	9.5	3.9	3.1	0.4
Large soil 4	29.1	10.7	10.8	0.4
Large soil 5	3.3	1.4	1.1	0.3
Small soil 2	29.5	11.6	10.1	0.6
Small soil 3	39.2	14.0	13.7	0.6
Small soil 4	11.0	3.1	4.8	0.5
Small soil 5	16.0	7.1	6.8	1.1

markers in the two neutron scans. The three subsequent columns give the components of the FLE along each axis. Since the X-ray data was not used in the calculation of FLE, markers that were successfully located in both neutron scans but not in the X-ray data were not discarded at this stage. This data cannot be directly compared to the FRE values without conversion since they are given in voxels, not absolute distance, and the data were rescaled as part of the registration process.

Table 4.7 shows the FRE found for each sample between each pair of scans. Each column compares a different pair of registered scans: first the two neutron scans, then each neutron scan to the X-ray scan. The FRE is calculated as the RMS difference between corresponding markers in the two scans after

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

Table 4.7.: The FRE values found between each scan for each sample. (All values in μm .)

Sample	FRE (N diff)	FRE (ND)	FRE (NW)
Large dry sand	995	306	961
Large wet sand	395	339	251
Large dry soil	775	491	595
Large wet soil	1012	391	935
Small dry sand	545	536	96
Small wet sand	240	247	64
Small dry soil	465	429	123
Small wet soil	335	374	99
Large soil 1	408	417	172
Large soil 2	416	332	186
Large soil 3	421	351	214
Large soil 4	247	74	222
Large soil 5	297	291	105
Small soil 2	558	280	481
Small soil 3	708	474	522
Small soil 4	329	350	169
Small soil 5	481	433	165

registration is complete. The post-registration voxel sizes have been used to convert these values from voxels to absolute distance values.

4.5. Discussion

This experiment introduced a new algorithm for the segmentation of spherical objects to be used as fiducial markers in the registration of complementary datasets of different modalities. This algorithm was generally successful in finding the markers placed on the samples for this experiment, but there were some that it could not locate successfully in all three scans (X-ray, neutron pre-injection, neutron post-injection). One hundred and sixty-one markers were

attached to the samples in total, yet due to poor sample alignment, not all of these were kept within the field of view. Figure 4.9 showed an example of this. When the markers were being located the first time to be used to inform the registration, markers that were missed by the automated system were added by hand where possible. In total, this led to 128 markers being used and 33 that could not be seen in all the scans being ignored. When the registration was complete, the algorithm was run again, without any markers being added by hand. This time 101 markers were found, a success rate of 78.9%. It is worth noting that this is very heavily dependent on some parameters that must be tuned. These include the strength of the filtering applied to the datasets, the minimum and maximum radii of the circles, the sensitivity parameters of the Hough circle transform and the EPS and count parameters of the DBSCAN clustering algorithm.

Image artefacts due to poor beam transmission through the markers were not an obstacle to marker localisation in this experiment, as they were in the previous chapter when cadmium markers were used. Figure 4.10 shows the one identified case where artefacts do prevent markers from being identified. Two markers were placed side by side on the sample. In the X-ray image, this presents no problem, but due to the scattering in the neutron modality, the clustering algorithm could not separate the markers into two separate clusters.

Large Soil 4 stands out as a particularly poor result in table 4.5 with only two markers found in all modalities by the algorithm. This occurred because of poor image quality, particularly in the top half of the neutron scans. It is suspected that the sample may not have been fitted securely enough and so moved slightly over the course of the scan.

As described above, the FLE values in table 4.6 are calculated by taking the RMS value of the differences between the positions of corresponding markers in the two neutron scans before registration. This is valid if it can be assumed that the markers are in the same place in both scans, which is to say, no registration is required between these two scans. This should be a valid assumption because they are subsequent scans of the same sample in the same instrument, with no human intervention between the two. The difference should therefore be indicative of the accuracy with which the markers were found.

A component of this error may be due to inaccuracy of the multi-axial rotation stage as discussed in 4.2.3. If this were the case, it would lead to a rotation of the sample between the two scans. This would affect the position

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

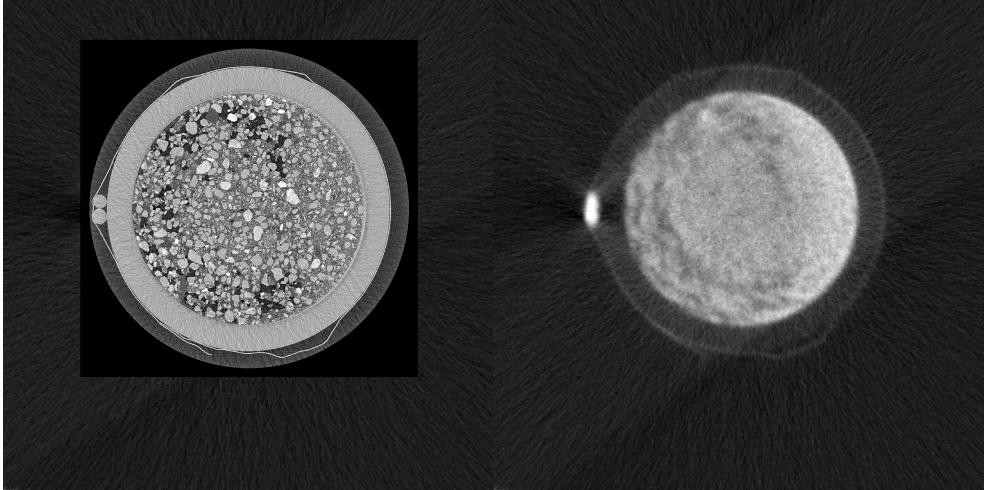


Figure 4.10.: Two markers placed side by side on a sample. While the two markers are easily distinguished in the X-ray image (left), the scattering artefacts make it difficult for them to be separated by the DBSCAN algorithm in the neutron data (right). Consequently, these two markers were found by hand for the registration but discarded when the algorithm was rerun for analysis of the registration.

of the markers in the XY plane but not in the Z direction. If this were the dominant factor in the FLE values, therefore, it is to be expected that the FLE would consistently have a significantly smaller component in the vertical direction.

The data in table 4.6 supports the idea that this type of error may be present. The FLE (Z) values are consistently far lower than the other two components and show less variance. The X and Y components vary over a wide range and appear to be directly proportional. All these are factors to support the idea that there may be rotation between the two scans and a visual inspection of the datasets confirms this, as shown in figure 4.11, where the difference in the marker positions is best explained by a rotation of the sample.

Given that inconsistent, unintended rotation of the sample appears to be the dominant factor contributing to the large X and Y components of the FLE, it is assumed that the Z FLE component is a better indication of the accuracy of the sphere fitting algorithm for the segmentation of the markers. Except for SmallSoil5, these values are all less than one voxel which indicates that the markers are found to a good degree of accuracy. The unintended rotations

Table 4.8.: Axial average components of FRE for each pair of scans. (All values in voxels.)

	X	Y	Z
N diff	12.2	10.5	5.7
ND	14.1	13.1	18.4
NW	5.9	6.5	17.7

between neutron scans, although detrimental to the calculation of FLE, should all be removed by the registration step, meaning that this error term is not carried forward to the FRE or subsequent analysis.

Potential errors in the stack concatenation step of the processing could be another factor that contributes to inaccurate registration. However, these, while technically contributing to FLE, would not be seen until the FRE is evaluated because they only apply to the X-ray scans, and these could not be used in the evaluation of FLE. If errors were to be introduced by the stack concatenation algorithm, it is reasonable to suppose that these would primarily be in the vertical direction as the underlying assumption of the concatenation algorithm is that the scans are aligned in the XY plane, and adjustments are only made in the Z direction. This error would show itself through a higher FRE component in the Z direction where either neutron scan was compared to the equivalent X-ray scan than when the neutron scans were compared to each other. Table 4.8 shows the average component of FRE along each axis between each pair of scans. Like the FLE in table 4.6, the Z component of the FRE between the two neutron scans is far lower than the X and Y components, indicating that the primary error, in this case, is likely to be in the rotation or translation of the sample in the XY plane. For the FRE data where each neutron scan is compared to the X-ray scan, the Z component is the largest component of error, particularly with the post-injection scan, which has a comparatively low error in the XY plane. This suggests that there is a greater error in the localisation of the markers in the Z plane of the X-ray modality. Small errors in the stack concatenation process provide a very plausible explanation for this, although it cannot fully account for the difference in the XY plane between each neutron scan and the X-ray scan.

Figure 4.12 shows the FRE between the two neutron modalities plotted against the FLE. If the two neutron scans were being directly registered to

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

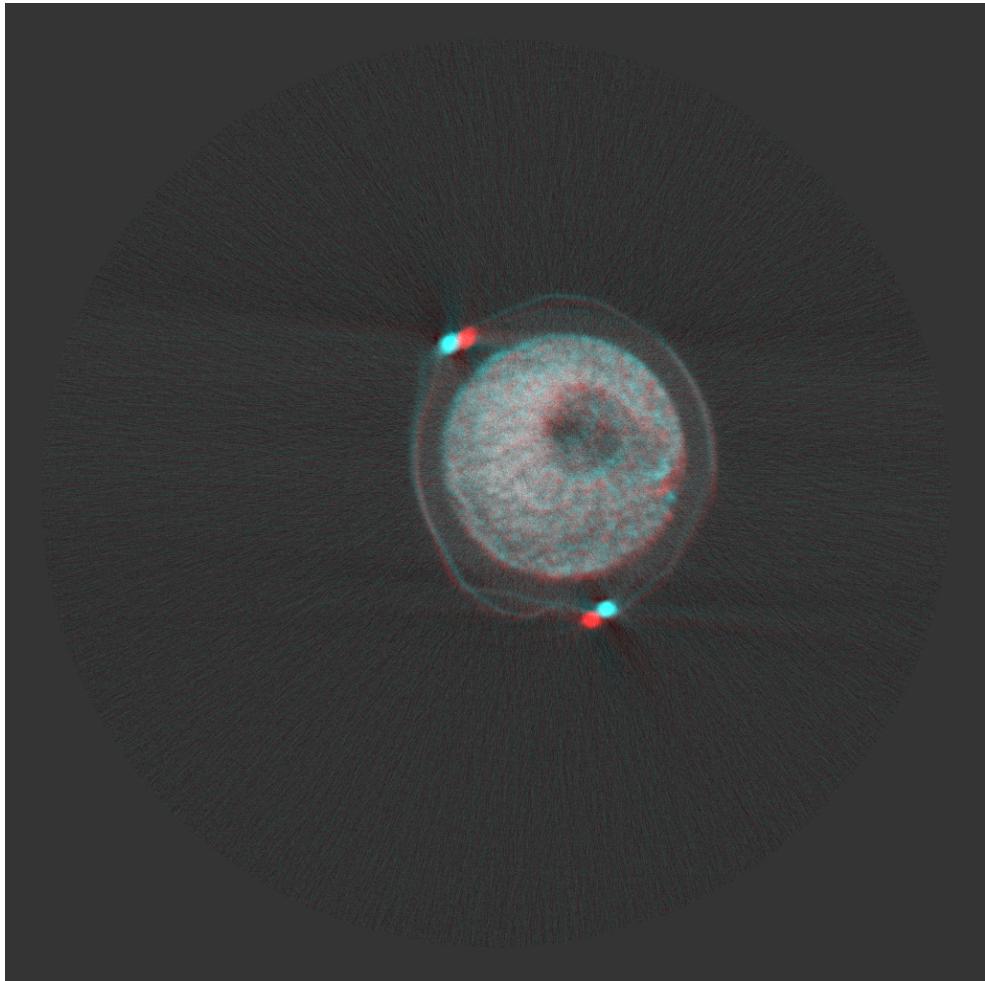


Figure 4.11.: Two nominally matching slices from different scans of sample SmallSoil3 overlaid. The pre-injection scan is coloured red, and the post-injection scan is cyan. The difference in the marker positions between the two scans suggests that the sample has been rotated around a point close to the centre of the sample.

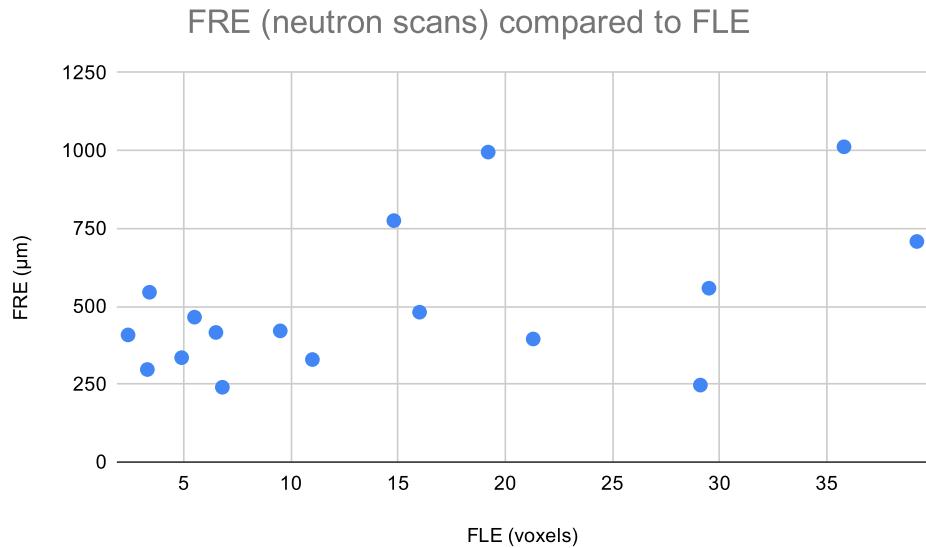


Figure 4.12.: FRE between the two neutron scans plotted against the FLE.

each other, the FLE and FRE should be directly proportional. In this case, however, the neutron scans are both registered to the X-ray scan, and the FLE of the X-ray scan is not considered, so other sources of error will be present and unaccounted for. This explains why, although the FLE and FRE appear to show some correlation, the coefficient of determination is low (0.264).

Figure 4.13 gives a graphical representation of the data in table 4.7. This allows the typical FRE values and the relationships between the scans to be seen more easily. The average FRE (the average distance between registered markers across all samples and modalities) is 394 μm . LargeDrySand and LargeWetSoil are the two samples with the greatest registration errors. For both of these samples, the FRE between the dry (pre-injection) neutron scan and the X-ray scan are close to the average values, but the two FRE values that include the wet (post-injection) neutron scan are far higher. This implies that there are errors in the marker localisation of these samples in the wet neutron scan in particular since the error between the other two scans is so much lower. These two scans skew the average FRE value, and it is interesting to note that the average without these two scans included is far closer to the median FRE value of 351 μm .

In more than three-quarters of the samples, the FRE is greatest between the

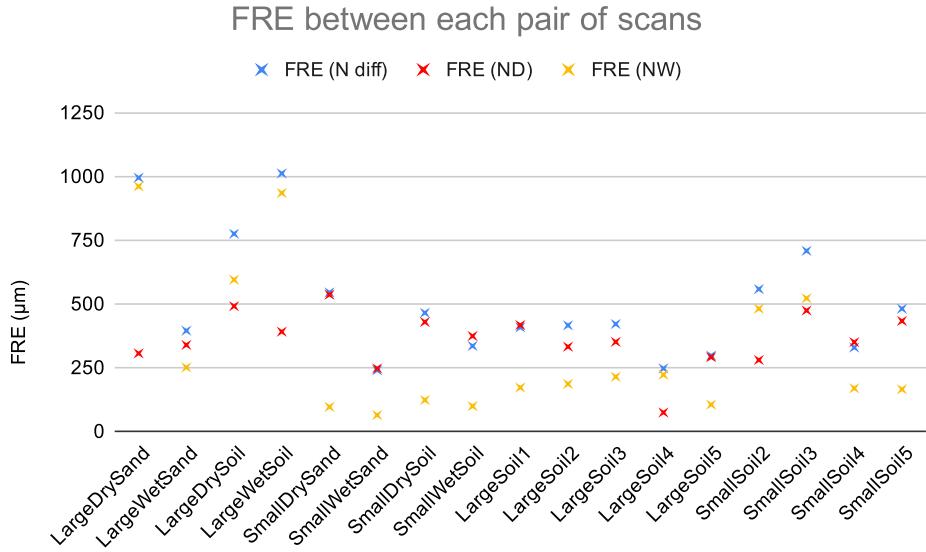


Figure 4.13.: The FRE between each pair of scans.

two neutron scans. This suggests that the evaluation of the marker positions is more precise in the X-ray modality. This is unsurprising due to its superior resolution and signal-to-noise ratio. Interestingly, the average FRE between the wet neutron scan and the X-ray scan is approximately 45 μm lower than that of the dry neutron and X-ray scans, even with the poor results from LargeDrySand and LargeWetSoil. This suggests that the markers were found more accurately in the second neutron scan of each sample after fluid had been injected into the sample. Although some samples had H_2O injected and others D_2O , the fluid used does not appear to correlate to which samples were more accurately registered either before or after the fluid was injected. All the samples that had better registration before the injection were scanned on rotation stage B, although not all samples scanned on rotation stage B were more accurately registered before the injection. Although the correlation is observed, it is not clear how the rotation stage could have any impact on the relative accuracy of marker segmentation between two successive scans.

Overall, the registration accuracy appears to be independent of the different sample parameters (sand or soil, particle size, H_2O or D_2O injection). This is to be expected, provided that the sample does not produce artefacts that extend beyond the outer radius of the sample tube.

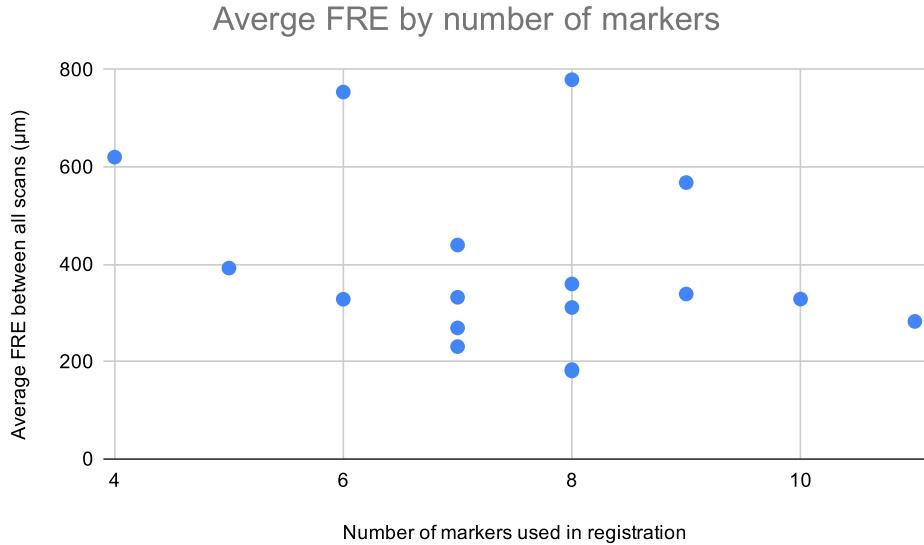


Figure 4.14.: Average FRE across all scans for each number of markers used in the registration.

It was shown in section 3.2 that in the ideal case, the registration accuracy should improve as the number of fiducial markers increases. One of the conclusions of Chapter 3 was that poor marker localisation could overcome this as the addition of inaccurately located markers decreases registration accuracy. Figure 4.14 shows the relationship between the number of markers used in registration (the number of markers found post-registration to evaluate FRE was different for many samples) and the average FRE between each pair of scans for a sample.

It can be seen that the inverse proportionality between the number of markers is not the dominant factor in the relationship between the FRE and the number of markers, and this suggests that the accuracy with which markers can be located in a scan remains the largest source of error in the registration process.

In the experiment with cadmium markers in Chapter 3, the FRE varied between 2.03 and 11.02 voxels. The voxel size after registration was 19 μm , so the FRE can be expressed as between 38.6 μm and 209.4 μm . This is substantially better than the median FRE value of this experiment (351 μm). The registration accuracy of this experiment is substantially lower, despite the significant reduction in artefacts brought about by the use of borosilicate

4. Borosilicate markers and a sphere fitting algorithm for fiducial localisation

markers. Steps have been added to the image processing pipeline, and it is likely that the combined influence of these has contributed to the substantial increase in registration error. Firstly, the automated stack concatenation step applied to the X-ray scans may have introduced errors in the scan data before any registration steps are undertaken, immediately moving markers away from their true position. The higher contribution of FRE from the Z-axis in all cases that include X-ray data would appear to indicate this type of error is present. Secondly, the new algorithm for locating markers may have lower accuracy than the methods used in the previous experiment, which involved more human judgement.

To narrow down the potential sources of error further, the method for locating markers described in section 3.3.2 was applied to the neutron data for LargeWetSoil. Once a threshold value had been found to isolate the markers, the 3D Objects Counter was used to determine the centre of each marker. The neutron markers found with this method typically agreed to within 2.5 voxels of the positions found by the sphere fitting, and the most significant component of this was generally in the vertical direction. The few markers that did not agree to within 2.5 voxels showed lower sphericity and hence could not be located by the sphere fitting method. Unfortunately, this marker segmentation method could not be applied to the X-ray data since the borosilicate did not show a sufficient difference in contrast to the sample tube for a threshold to segment the markers. Also, the file size is more than twenty-eight times larger than the neutron data, and even with significant cropping, this could not be made to fit within the limit of the 3D Objects Counter.

These new marker positions were used to calculate a new set of transforms to register the two neutron scans to the X-ray data. When the new positions of the markers were calculated using these new transforms, the FRE between the two datasets fell to approximately 20% of its previous value. Furthermore, the FRE for the post-injection scan fell from 17.2 voxels to 6.3 voxels, bringing it to a similar value to that of the pre-injection FRE. These calculations show that the sphere fitting algorithm does produce small errors in the locations of the markers it finds and that when these erroneous values are used to inform the registration transformations, they can be scaled up to have a significant impact on the registration accuracy.

4.6. Conclusions and future work

Borosilicate offers clear advantages over cadmium as a fiducial marker material, primarily because it does not produce the significant artefacts seen in both X-ray and neutron scans of cadmium, but also because it is more readily available, can be obtained in a useful form and is far less hazardous to work with. The minor difference in attenuation between the boron-free quartz used for sample tubes and the borosilicate markers meant that new algorithms had to be developed to allow markers to be found and used to inform registration parameters. A new sphere fitting algorithm was developed and used to complete the registration for a collection of seventeen samples. The algorithm was partially successful but was shown to have limitations, in particular lower accuracy with markers of lower sphericity. These errors led to a notable decrease in the accuracy of the registration despite the improvements in the quality of the scans.

There are a variety of options that could be considered to improve upon this registration performance in future experiments. One option would be to reconsider the marker material used to find a material that could be adequately segmented using a threshold-based method. However, the list of potential candidates is short and would likely be expensive, have lower availability, and require work to shape into suitable markers. Another alternative is to work on improving the segmentation algorithm to ensure more accurate results. While it may be difficult to improve the algorithm significantly while remaining independent of the sample, there is scope for interesting algorithms that use additional information, for example, by finding the outer edge of the sample tube and using that to mask both the tube and its contents. Another potential solution is to do a preliminary scan of a collection of borosilicate markers before they are attached to any sample. This scan could be used to measure the sphericity of each bead and therefore allow beads that are less likely to be located accurately to be identified and removed from the experiment at the sample preparation stage.

5. Correlative X-ray and neutron imaging for water segmentation and quantification

5.1. Introduction

This chapter is concerned with the analysis opportunities afforded by the data gathered for the previous chapter. In particular, the benefits associated with the use of complementary neutron and X-ray data for crop science. As in the previous chapters, the methods discussed here are presented with recognition of their applicability towards a broader range of applications.

The work presented in this chapter aims to consider how best to segment and quantify the water concentration within the samples given the three scans collected from each sample. Using a syringe pump to inject known quantities of water into the sample at a controlled rate opened many possibilities for measuring water dynamics and led to the collection of pre- and post-injection tomograms. Using the difference between these two scans makes segmentation of the injected fluid simple in theory but only works if beam hardening can be ignored or corrected. A method to correct beam hardening introduced by the injected fluid by considering the difference between the two scans is presented in section 5.3.2.

The X-ray data can be used to achieve accurate soil segmentation at a higher resolution than that of the neutron data, so a method is proposed to utilise this information as a mask to enhance the reconstruction of the neutron difference images further. Section 5.3.3 discusses the influence of this X-ray based mask when applied to the reconstructed and registered neutron data. In section 5.3.4, however, this process is adapted to use the mask to enhance the neutron data through iterative masked reconstruction. The simulation presented shows that this method should give better results, and a significant difference is observed in the experimental data.

5.2. Equipment and methodology

5.2.1. Preliminary tests

During the planning phase for the experiment that produced the data for this chapter and the previous one, five days of IMAT beam-time were allocated for preliminary tests of the new equipment and techniques proposed (Experiment RB1810843). These tests included the first trial of the multi-axial rotation stage, discussed in the previous chapter (4.2.3), a syringe pump and a golden angle acquisition scheme. These additions allow more advanced techniques to be introduced at IMAT, facilitating new experiments.

Since a previous experiment (RB1830047) allowed a large amount of data to be collected primarily to develop registration techniques and correlative methods with X-rays, the focus of this trial was to gather data that would fully demonstrate IMAT's ability to quantify gradients of fluid distribution within porous media.

All the samples used in this experiment were prepared in boron-free quartz tubes of height 100 mm, inner diameter 14 mm, and outer diameter 17 mm. Cadmium markers were attached to the samples as in chapter 3 to ensure the samples were equivalent to those used in other experiments, although no registration was applied to this trial data. (The borosilicate markers discussed in Chapter 4 were not introduced until after this trial data was collected.) This experiment used two sand fractions: Sand B was made up of particles with diameters between 1.18 mm and 0.6 mm. Sand D contained particles with diameters between 300 μm and 150 μm . Two gels were used in this experiment and were produced by dissolving sodium alginate in water. Gel A contained 1 g/ml, and Gel B contained 2 g/ml.

The beamline was configured for $L/D = 166.7$ (where L , pinhole sample distance, was 10 m and D , pinhole diameter, was 60 mm). The CMOS detector was used with an 85 mm lens and an 80 μm ZnS/6LiF scintillator, giving a 120×120 mm field of view and a pixel size of 56 μm . The CMOS detector is based on an Andor camera used with an exposure time of 30 seconds. All the tomography measurements taken during this experiment used an angle step size of 0.6923° giving 521 projections unless otherwise specified. Two samples were imaged at a time using the multi-axial tomography stage since three samples would require a reduction in resolution and could produce artefacts due to the inhomogeneity of the beam.

5.2. Equipment and methodology

This experiment introduced the use of a syringe pump at IMAT for the first time. The pump allowed accurate fluxes of fluids such as water, deuterated water and more viscous gels into samples with precise control of the volume and rate of delivery and the fluid injection location. The syringe pump used in these experiments was an SP210cZ from WPI. It was connected to the IBEX control system used at ISIS to allow it to be controlled through the IMAT instrument dashboard and scripts. Thin tubes connected the syringes in the pump to the needles inserted into the samples.

The first measurements taken were tomograms of two dry sand samples. One sample contained Sand B, the other contained Sand D and a needle connected to a syringe in the syringe pump. Following this acquisition, 1 ml of water was poured onto the sample containing Sand B, and 0.1 ml was injected into the sample containing Sand D and the tomography was then repeated. Following this, the process was repeated with a further 2 ml added to Sand B and a further 0.1 ml injected into Sand D before a final tomography of these samples. These results demonstrated that water could easily be identified within the images produced and exposed flaws in the experimental method that could be addressed before taking further measurements.

One observation from these first results was that water that is being added to a sample should be added very slowly to avoid unnatural disturbances to the soil structure, such as lifting portions of the sand on pockets of water. Another problem that was quickly identified and fixed was the need for a fixture that minimised the movement of the syringe needle during tomography to avoid any sample disturbance during rotation.

The following configuration used a tube of Sand B and a tube of soil, each wetted to a 10% water content by mass. Three tomographies were conducted with these samples. Between each tomography, 0.05 ml of deuterated water was injected to demonstrate the contrast that allows deuterated water to be seen within a wet sample and subsequently used to determine a flux. This test was unsuccessful because the needles were not attached to the tubes carrying the deuterated water sufficiently well and the fluid leaked before reaching the samples. It was also observed that the samples contained a large amount of water and therefore produced considerable scattering that could result in significant beam attenuation and prevent the deuterated water from being imaged accurately if the concentration was too high. A real-time radiography experiment was conducted to investigate these problems further. Two tubes

5. Correlative imaging for water segmentation and quantification

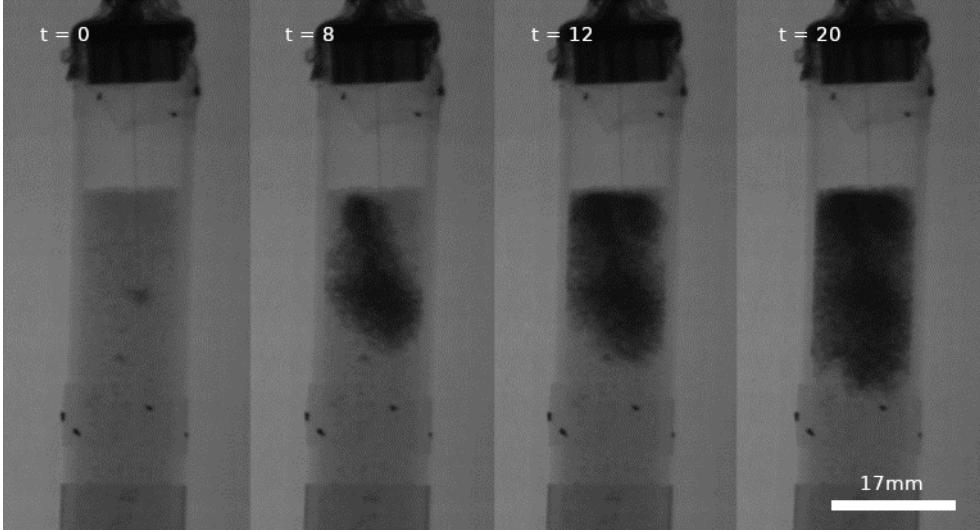


Figure 5.1.: A series of radiographs showing the progressive injection of 1 ml of water into the soil at a rate of 0.05 ml/min. From left to right these images show the sample at the beginning of the injection, then after 8, 12 and 20 minutes respectively (to the nearest minute).

of dry soil were injected with 1 ml of water and deuterated water, respectively, at a rate of 0.05 ml per minute. Radiographs were taken every 30 seconds and showed the progression of the fluid into the samples. A significant difference in contrast between water and deuterated water was observed, and the improvements made to the syringe needle configuration were shown to work. Figure 5.1 presents radiographs that show the water filling the soil over time.

Following this, a series of three pairs of tomographies was attempted again. One sample contained initially dry soil that received two 0.25 ml injections of water between the tomographies. The other sample was kept from the previous radiography experiment and was injected with deuterated water. This sample initially contained 1 ml of water and also received two 0.25 ml injections between tomographies. Figure 5.2 shows the difference between corresponding radiographs from the initial tomography and the second tomography. The injected water and deuterated water can be seen as well as the movement of the water displaced by deuterated water. Following these tomographies, the remaining fluid in the syringes was injected into the samples as real-time radiography took place.

Finally, a set of tomographies was acquired, with gels A and B injected instead of water. They were injected first into dry soil samples using the same

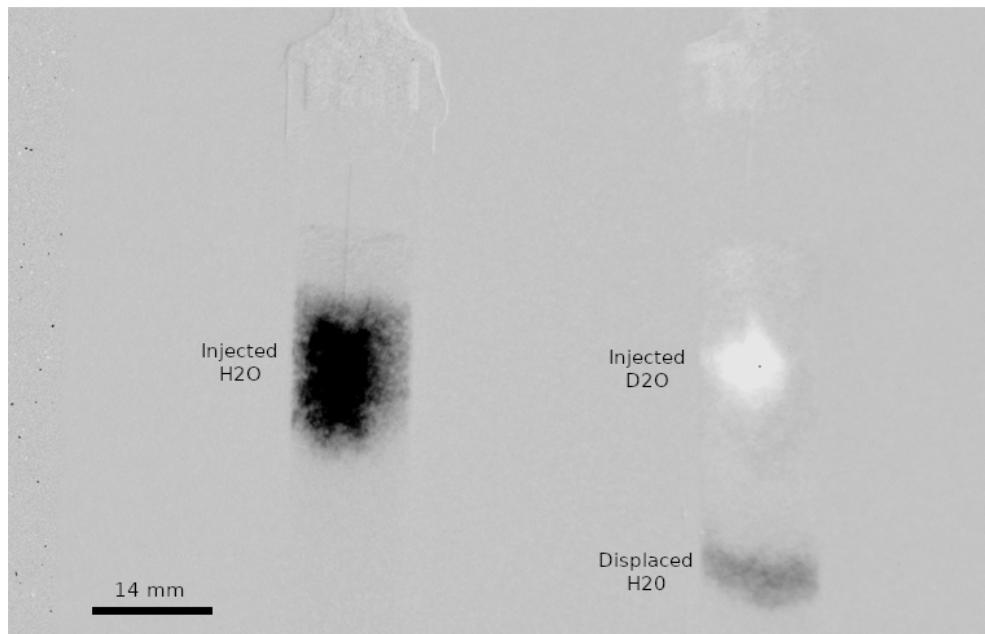


Figure 5.2.: The difference between two radiographs showing the change in water and deuterated water between the two acquisitions. The left sample shows water injected into dry soil, and the right shows deuterated water injected into a wet soil sample, displacing water further down the tube. The dark regions show the water, and the light represent the deuterated water.

5. Correlative imaging for water segmentation and quantification

three tomography sequence. A 20 minute pause was placed before the second and third tomographies to allow the injected gels to diffuse and settle before imaging to reduce motion artefacts as they sink through the soil. As before, the tomography was followed by real-time radiograph series observing injection of the remaining contents of the syringes.

The gel imaging was then repeated with soil samples that had received 1 ml injections of water 24 hours earlier to wet the soil. It was observed that in the previous gel experiment, the fluid had continued to move throughout each scan, so further changes in technique were introduced to reduce motion artefacts by minimising the scan duration. The exposure time for each projection was reduced from 30 seconds to 10 seconds, cutting the scan time by nearly 3 hours. In addition to this, the projection acquisition order was changed to use golden ratio decomposition (2.2.1). This meant that the sample was turned further between each projection, lengthening the scan time slightly but causing a reduction in motion artefacts after reconstruction since the projections are collected out of order, and the position of moving objects in the sample can be averaged. It also allows the number of projections used to reconstruct the tomogram to be determined at the reconstruction stage. This means that the trade-off between signal to noise ratio and motion artefacts can be evaluated after the experiment is complete to find the optimum balance between image quality and temporal resolution.

At the end of the experiment, the final sample was left in place for some tomography speed tests. First, a sequential acquisition scan of 521 projections with 1 second exposure was collected, followed by a golden angle decomposition scan with 360 projections each using a 5 second exposure to inform further any future experiments that might prioritise high-speed tomography over the signal to noise ratio, either for temporal resolution or sample throughput. The impact of the golden angle acquisition sequence on the scan time can be calculated by subtracting the total exposure time by the total scan time and dividing by the number of projections collected. The data from this experiment shows that, on average, 5.74 seconds were added to the collection time of each projection in addition to the exposure time in the sequential scan. In the golden angle scan, an average of 14.98 seconds was required per projection in addition to the exposure time. This significant time increase is best explained by the greater sample rotation required but may also have contributing factors in the IMAT control system configuration.

The significance of this difference to the overall scan time is inversely proportional to the exposure time. The difference between the two acquisition schemes shows that the golden angle scan will always take longer for a given number of projections and exposure time. This difference will be greater for scans with lower exposure times, which are more likely where temporal resolution is critical and golden angle acquisition is considered. If the desired number of projections is known, a golden angle scan will take longer and must offset this disadvantage with the time averaging of projections. Depending on the speed of the dynamic process being scanned, this may or may not be worthwhile. The significant advantage of golden angle scans is the flexibility afforded when the ideal number of projections is unknown. Provided that sufficient projections are collected, the temporal resolution can be improved upon at the reconstruction stage by discarding projections from the start or end of the scan until the impact on the reconstructed image quality is deemed too great. This could result in a shorter scan and better time resolution than a sequential scan where the number of projections is fixed in advance.

For experiments where several scans are taken of similar samples, the best approach could be to begin with a golden angle scan to determine a suitable number of projections and then revert to sequential acquisition with that number of projections. This approach allows the optimum number of projections to be used and minimises the overall scan time by sacrificing the option for time averaging over the projections of each scan.

The significant time costs associated with collecting each projection (by either acquisition scheme) mean that to get the best ratio of total exposure time to total scan time, it is optimal to use fewer projections of greater duration. Attempting to reconstruct from too few projections leads to undersampling artefacts (2.2.2), limiting this trade-off. Iterative reconstruction methods give better results with undersampled data than FBP, so it would be interesting to do a comparative study into whether the image quality attainable in a given scan duration can be improved by using fewer projections but a longer overall exposure time.

5.2.2. Experimental methods

The full experimental details for the collection of the data discussed in the rest of this chapter are given in 4.2.3, but a summary of the samples and scans will be given here. Each neutron scan consisted of a tomography, followed by either

5. Correlative imaging for water segmentation and quantification

water or deuterated water injection, followed by time series radiography and then a second tomography. Samples were prepared with sand and soil, and each of these materials had samples prepared with larger and smaller particle sizes (table 4.3b). Pairs of samples were prepared with each of these four growth media but without a seedling. One of each pair was initially scanned dry and then received an injection of water, and the others were prepared wet and were injected with deuterated water. Ten plant samples were prepared, five in each type of soil. These were kept wet during preparation to ensure plant growth and so were injected with deuterated water. For the complete sample list, see table 4.3. The work discussed here begins after all the data from these scans had been reconstructed and registered using the methods described in section 4.3.

5.3. Segmentation and quantification of water in soil and sand

5.3.1. Time series neutron radiography

The scan duration of a neutron tomography is far too high to capture the movement of injected fluids within a sample. The movement of the fluid through the sample would instead lead to motion artefacts if a tomogram were immediately collected. After each injection, a series of radiographs were collected instead. Using radiography rather than tomography improves the time resolution by two orders of magnitude at the expense of volumetric data. For each sample, 720 images were collected with a 10 second exposure time. The total time between the injection and the acquisition of the second tomogram was around 2.5 hours for each sample.

Figure 5.3 shows the first and last radiographs taken of LargeSoil4 after the deuterated water was injected. The first radiograph shows a bright region within the soil where the deuterated water has just been injected. Over the course of the scans, this area fades away to more closely match the surrounding area as the injected fluid diffuses through the sample. A slight reduction in the neutron attenuation of the surrounding area can be seen in the final radiograph, where the deuterated water has displaced the more strongly attenuating water from the sample. The difference image also shows that the seedling has drooped further over the scan duration.

5.3. Segmentation and quantification of water in soil and sand

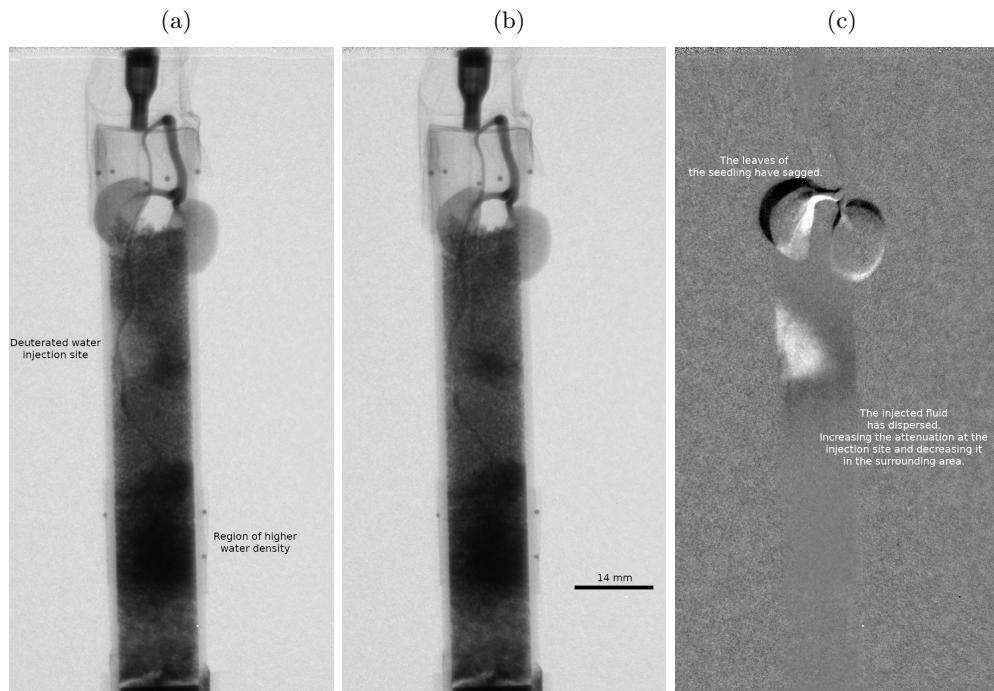


Figure 5.3.: The first and last radiographs taken of LargeSoil4 after deuterated water was injected into the sample.

- (a) The radiograph taken immediately after the injection of deuterated water.
- (b) The final radiograph, taken 2 hours and 24 minutes after the injection.
- (c) The difference between the first and last radiographs.

5. Correlative imaging for water segmentation and quantification

Using these techniques, a great deal of information can be gained, such as the rate that the fluid diffuses into the sand or soil and the depth it reaches in a given time. Accurate quantification, however, remains difficult in only two dimensions. The clearest solution to this is to use time resolved neutron tomography as discussed in section 2.3.3.2, but other techniques such as comparison to X-ray tomography data could provide interesting avenues for further analysis and research without placing such stringent requirements on the imaging capabilities of the neutron instrument.

5.3.2. Beam hardening correction

To further improve the image data analysis, beam hardening correction factors for the water in the neutron scan can be calculated from the projection data. The injected fluid can be isolated by taking the difference between the corresponding pre- and post-injection radiographs of a sample. Assuming that, on average, the water is mixed evenly within the soil or sand, the known geometry of the sample tube can be used to calculate the path lengths through the slice at different points. Without beam hardening effects, the neutron attenuation should be linear with path length. If there is noticeable beam hardening, the attenuation increase should diminish as the path length increases.

Figure 5.4 shows an image produced from the normalised projection data for LargeDrySoil to allow for this analysis. The difference has been taken between the pre- and post-injection radiographs leaving just the injected water. The projections were then summed and cropped to a region of interest that shows only the water. The path length can be calculated as a function of the horizontal position since the inner diameter of the sample is known. This allows the graph in figure 5.5 to be plotted, which shows the maximum attenuation as a function of path length. The attenuation stays low for values with a short path length then rises as the path length increases. While the path length is below 4 mm, other elements contribute to the attenuation, but above this point, the relationship between path length and attenuation appears linear. The linearity indicates that the injected water causes no significant beam hardening in this dataset.

Figure 5.6 shows the same analysis applied to a different sample: SmallWet-Sand. In this case, the results do not show the predictable linear increase of attenuation with path length, but neither do they appear to show the curve expected in the case of beam hardening. Instead, this sample serves to illustrate

5.3. Segmentation and quantification of water in soil and sand

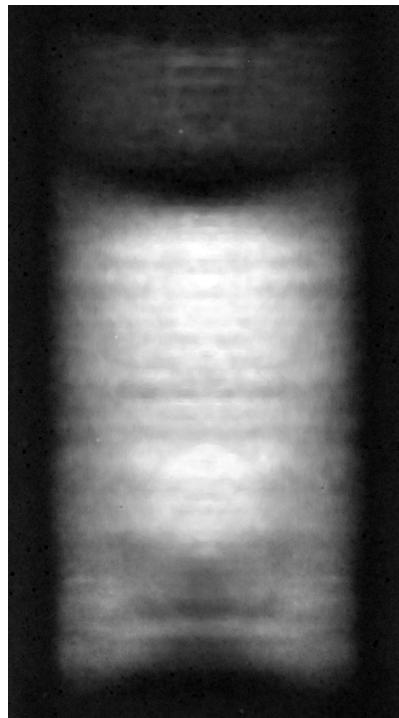


Figure 5.4.: The ROI from the summed difference projections for the sample LargeDrySoil.

5. Correlative imaging for water segmentation and quantification

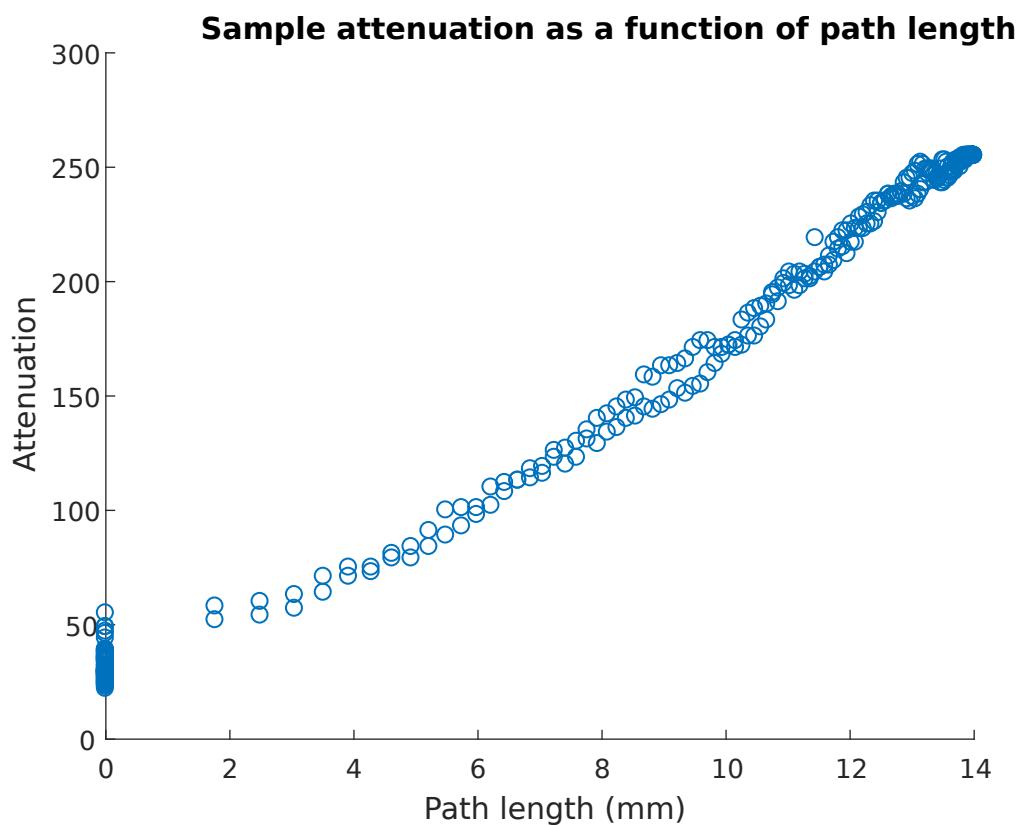


Figure 5.5.: The attenuation from figure 5.4 plotted against path length.

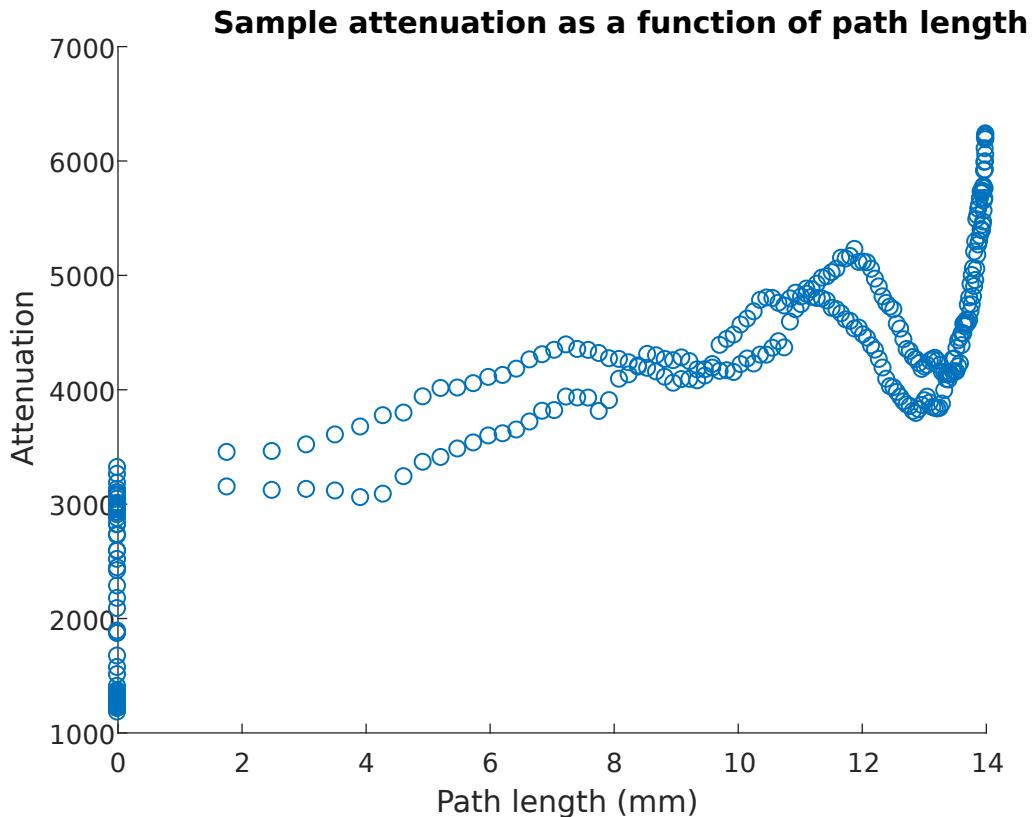


Figure 5.6.: Neutron attenuation through the fluid in the SmallWetSand sample as a function of path length.

how this method can fail when the fluid is unevenly distributed throughout the sample. There are peaks and troughs in the curve and a more significant difference between the two data points at any given path length. This mismatch shows that the image data does not have a vertical line of symmetry at the centre of rotation.

5.3.3. Independent processing of complementary modalities

By comparing the neutron tomograms collected before and after the injection of a known quantity of fluid into a sample, not only can the fluid and its distribution be observed but also quantified. The addition of the complementary X-ray data not only introduces new information but improves the accuracy of measurements taken from the neutron data.

This will be demonstrated first with the LargeWetSand sample because the

5. Correlative imaging for water segmentation and quantification

large sand is well resolved in the neutron data as well as the X-ray, and the registration error of the wet sample is less than half that of the dry equivalent. Figure 5.7 shows a representative slice from the sample. 5.7a shows the slice from the initial tomogram, and 5.7b contains the same slice after the injection of deuterated water has taken place. From this comparison, it can be seen that the two slices are generally well-matched; the registration has aligned the sample well between the two images, and the same sand particles can be seen in the same position in each image. The clear difference that can be seen is the addition of material filling the pore spaces between sand grains towards the top of the sample in 5.7b. Despite nominally being a wet sample where deuterated water is to be injected where water is already present, there does not appear to be any water in the pre-injection image. The fact that the injected deuterated water appears as a brighter (and therefore more highly attenuating) material in the post-injection scan confirms that it has replaced air rather than water. This is no surprise since the large channels between grains give the sand a very low field capacity, and it is to be expected that over 95% of the water (by volume) would drain away through the sand. Another feature of note is the bright particles that can be seen in both scans, most notably in the bottom left of the sample. These indicate that a small proportion of the sand grains is composed of materials far more highly attenuating to neutrons than the rest of the grains.

Figure 5.8 shows the absolute difference between the pre- and post-injection version of slice 4402 shown in figure 5.7. This removes all the information that is common to both images leaving only the differences and therefore displaying them more clearly. The injected fluid remains and is shown with greater contrast by the reduced visibility of the sand and sample tube. The edges of both the sample tube and the sand grains remain faintly visible, although not in the perpendicular direction, which indicates a very slight error in registration.

In some slices, imaging differences are exposed by this, for example, in slice 8027, which is shown in figure 5.9. In this slice, there is less fluid to be seen in the post-injection scan, but the slice does show one of the fiducial markers and also shows how some of the sand grains are more strongly attenuating to neutrons, and so appear far brighter than the others. Figure 5.10 shows the difference between the pre- and post-injection scans and, as in the previously shown slice, most of the sand and the sample tube are removed from the image.

5.3. Segmentation and quantification of water in soil and sand

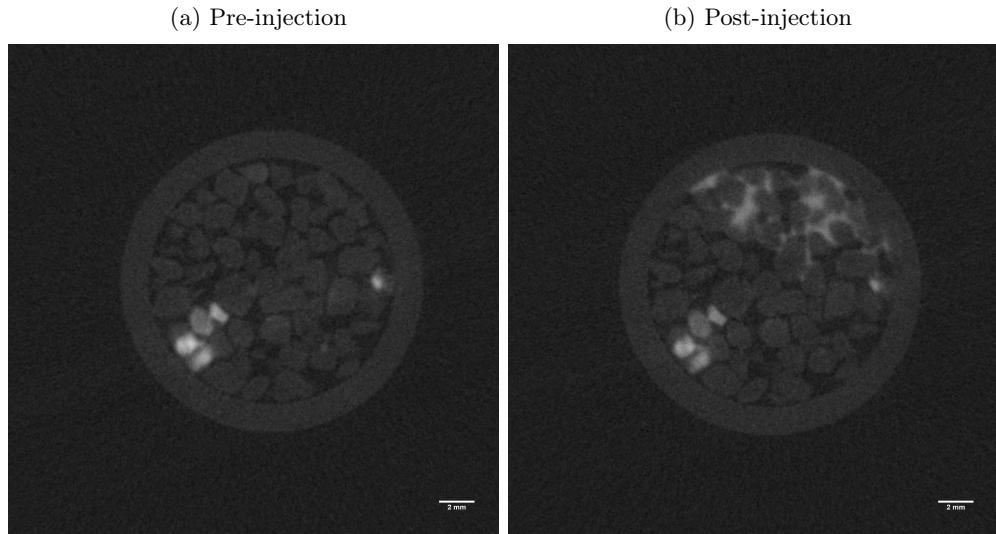


Figure 5.7.: Slice 4402 from the sample LargeWetSand.

Faint arcs can still be seen on some parts of the sample tube edge that show a similar misalignment in the registration step, but interestingly there appears to be greater variation in the marker and the bright sand grains between the two images. It is to be expected that the materials in the sample that attenuate most highly have the highest grey level and so show greater range when the difference is taken. It is interesting that the more strongly attenuating grains of sand remain in the difference image but also gain bright and dark lines on their lower edge. This does not appear to be a product of beam hardening or similar artefact mechanisms since there are no noticeable cupping or streak artefacts around these grains. Similarly, the fiducial marker has a bright glow at the edge, which partially corresponds to the observable misalignment but also suggests that there is a small discrepancy between the two scans. This may be a result of different artefacts in the two scans, as some streaking is visible, particularly in the post-injection slice where the injected deuterated water has increased the attenuation of the sample.

In the X-ray scan, the sand is shown at a far higher resolution, but the water (where present) is not shown. Figure 5.11 shows the X-ray data for slice 4402 of LargeWetSand, which corresponds to the neutron data in figures 5.7 and 5.8. Figure 5.11a shows the original data after registration, but further processing is then applied to prepare it for subsequent analysis. First, the slice is binarised using a threshold then the binary image is cleaned using

5. Correlative imaging for water segmentation and quantification

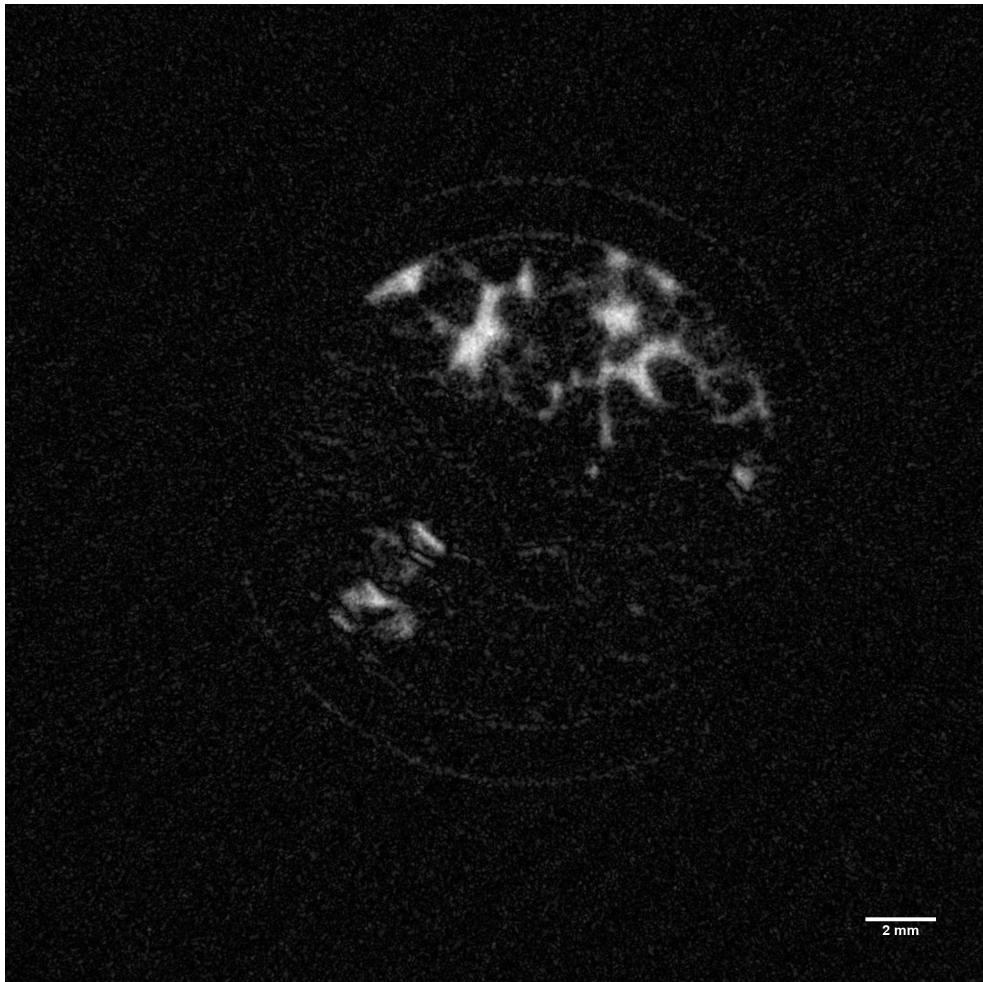


Figure 5.8.: The absolute difference between the pre- and post-injection images of slice 4402.

5.3. Segmentation and quantification of water in soil and sand

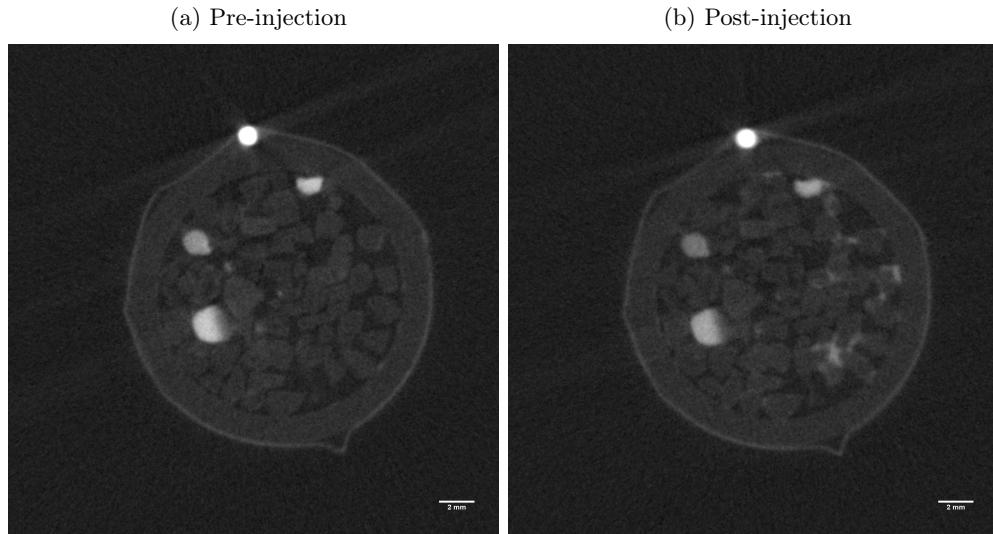


Figure 5.9.: Slice 8027 from the sample LargeWetSand.

erosion and dilation steps to remove small spots and holes from the image. At this point, further analysis could be conducted on the X-ray data to learn more about the sand or soil, and the X-ray data can be used to enhance and therefore improve the analysis of the neutron data.

In the first case, this X-ray data could be used to measure a number of properties of the growth media, including particle sizes, channel widths between particles or the particle-to-air ratio. This could be achieved by making a circular selection just inside the inner ring of the sample to remove all data apart from the growth medium and applying a separation algorithm to ensure adjacent grains are not counted as one before particle analysis tools can be applied to the data.

The simplest way to combine the X-ray data with its corresponding neutron data is to use the binarised volume as a mask. This enhances the contrast further and offers some correction for the neutron imaging limitations. The difference image of slice 4402 from figure 5.8 is shown again in figure 5.12. First in its original form and then with the binary X-ray image overlaid as a mask. The mask is generally very well matched to the grains visible in the neutron data and does a good job of removing most of the material that is not the injected fluid in the pore spaces. Seeing the binarised X-ray and neutron data combined draws attention to the noise that is present in the neutron data. There are some discrepancies in the shapes of the grains between

5. Correlative imaging for water segmentation and quantification

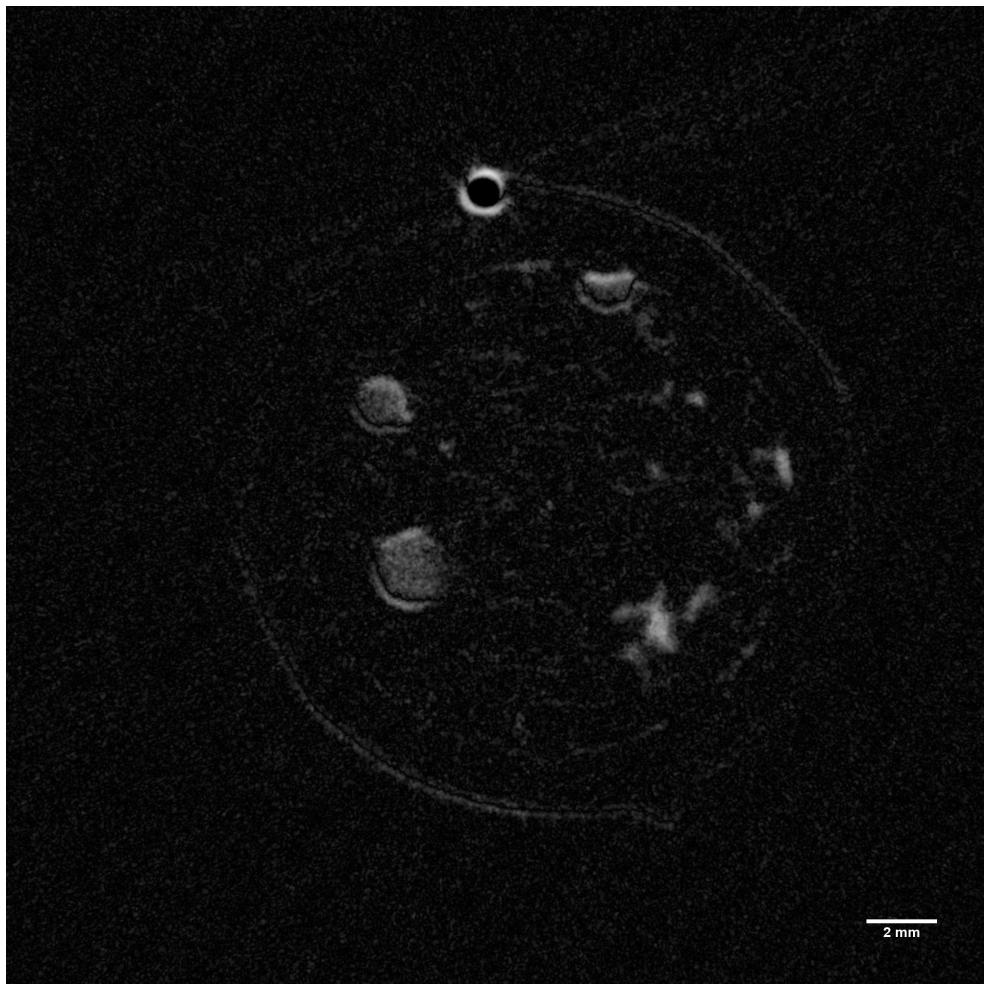


Figure 5.10.: Taking the difference between the pre- and post-injection images for slice 8027 of LargeWetSand. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.1.

5.3. Segmentation and quantification of water in soil and sand

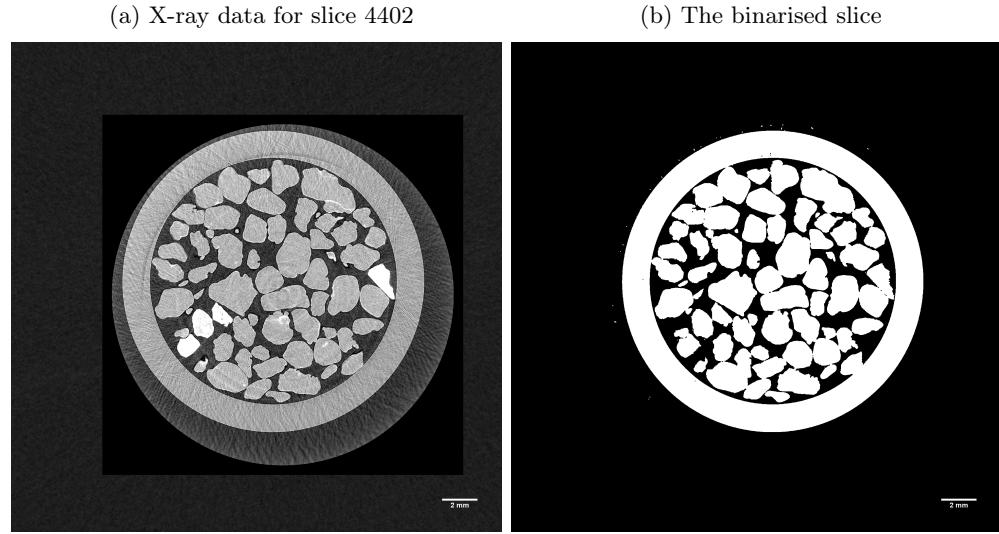


Figure 5.11.: X-ray data for slice 4402 of LargeWetSand. 5.11a shows the data after registration. It is converted to a binary image using a threshold to produce 5.11b, which can then be utilised for further analysis.

the two modalities. The close match of the sample tube between the two modalities suggests that this is not a sign of poor registration, and therefore two possibilities remain. Either some of the grains have shifted as the sample has been transferred between the two beamlines, or the neutron image is not only less precise but also less accurate in resolving the edges of the grains. This is plausible because even the deuterated water scatters neutrons, although to a lesser extent than ordinary water, and so there will be some imaging artefacts causing higher attenuation of the fluid to spread into the grains in the reconstructed image.

Figure 5.13 shows the addition of a similar mask to slice 8027 of the sample. This gives a similar result but particularly illustrates how the mask can remove the more highly attenuating grains within the sand that would otherwise show up with the fluid in the neutron difference image and impact the accuracy of subsequent measurements taken from the volume data.

The impact of small registration errors has been shown in the preceding figures. These errors have not been sufficient to impede the analysis discussed so far, though they could certainly affect the quantitative measurements yet to be taken. Some of the samples were registered with far greater errors (as

5. Correlative imaging for water segmentation and quantification

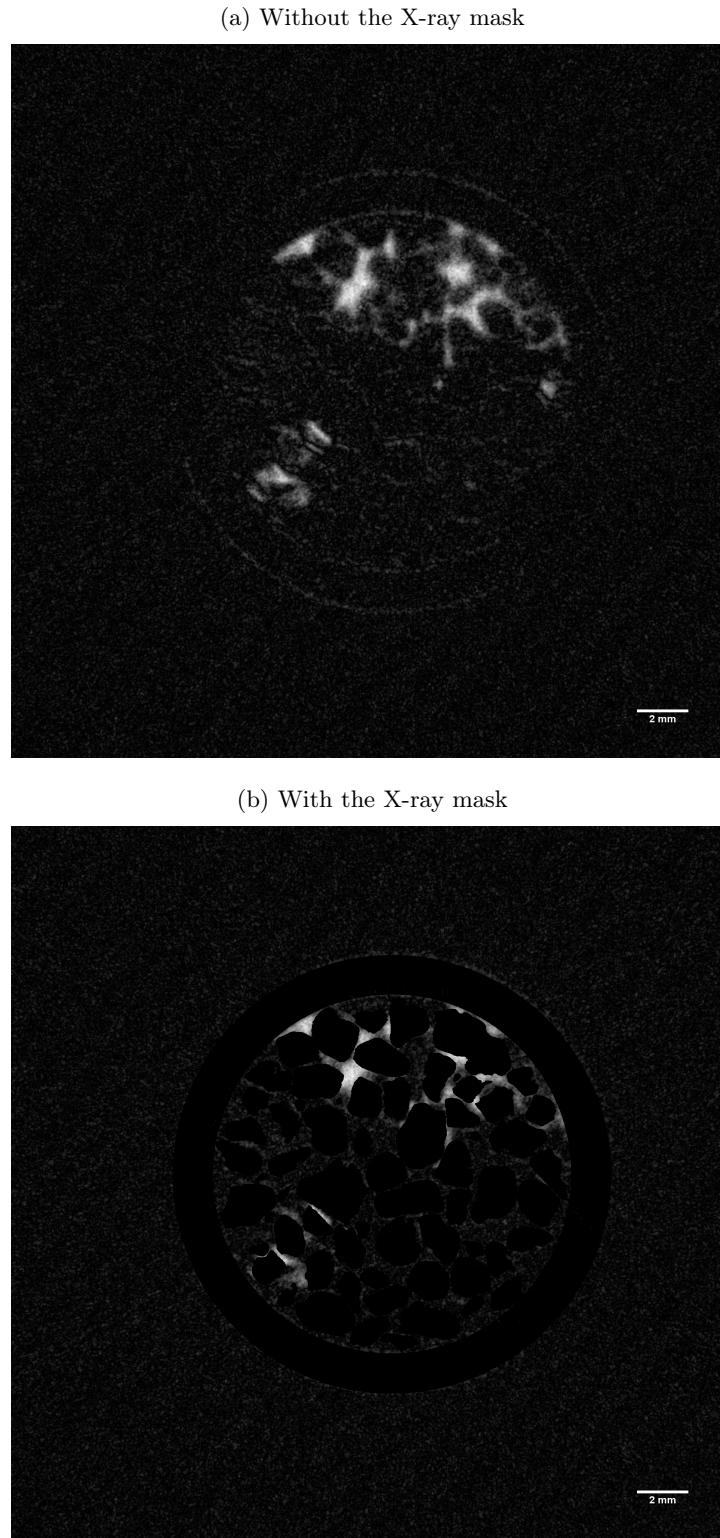
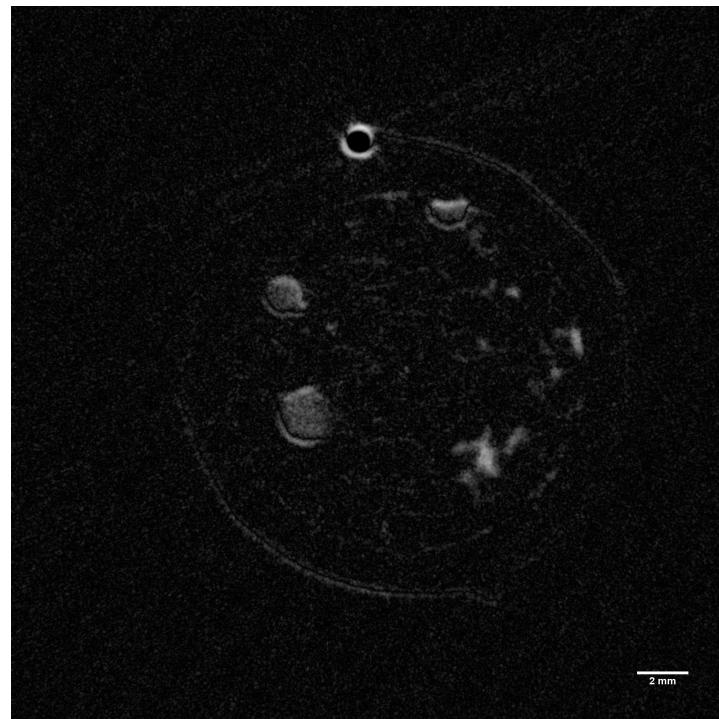


Figure 5.12.: The difference (before and after the injection of deuterated water) image of slice 4402 shown before (a) and after (b) a mask is applied.

5.3. Segmentation and quantification of water in soil and sand

(a) Without the X-ray mask



(b) With the X-ray mask

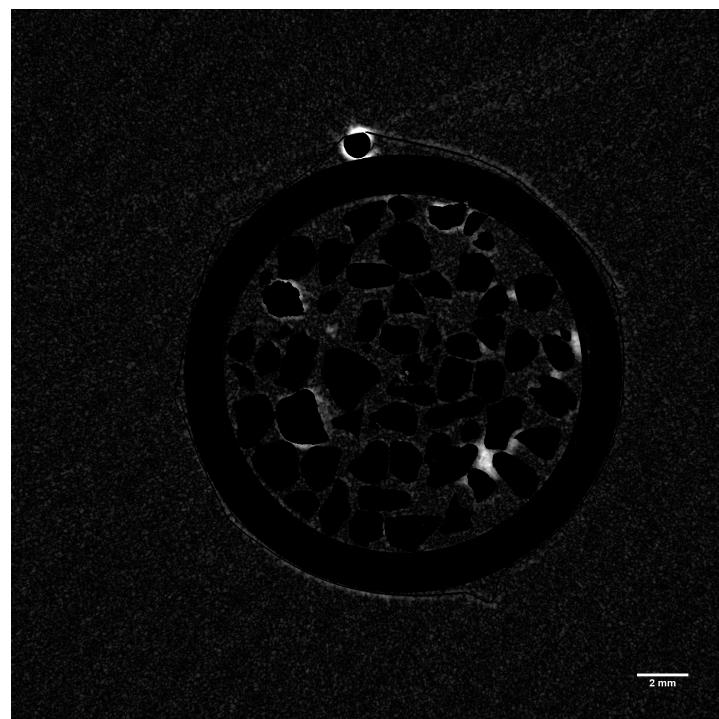


Figure 5.13.: The difference (before and after the injection of deuterated water) image of slice 8027 shown before (a) and after (b) a mask is applied. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.2.

5. Correlative imaging for water segmentation and quantification

(a) A slice taken from above the sand showing the misalignment of two fiducial markers, (b) A slice showing the poor alignment of the X-ray mask to the neutron data.

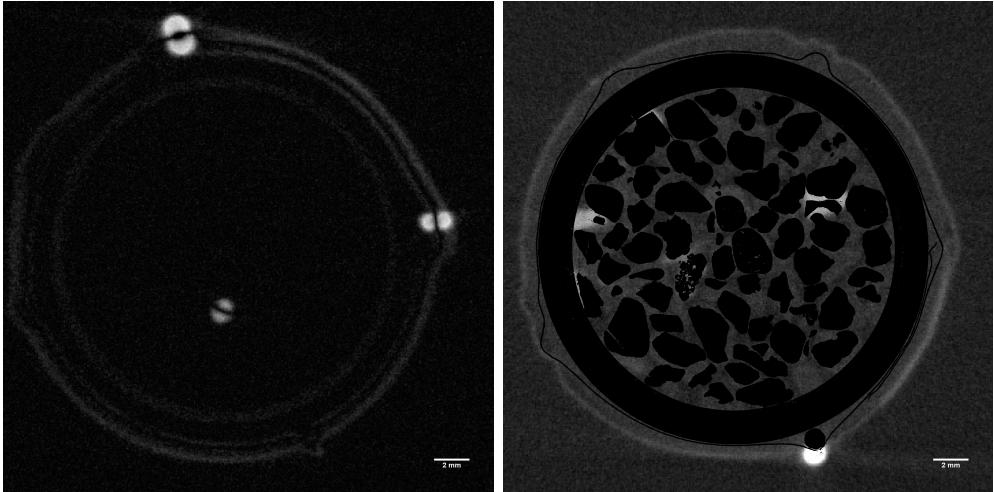


Figure 5.14.: A demonstration of the impact of poor registration. These slices were taken from sample LargeDrySand where the fiducial registration error between the post-injection neutron scan and the other scans was approximately 1 mm.

shown in figure 4.13), and these do prevent any meaningful results from being drawn from the analysis presented. Figure 5.14 illustrates this with slices from LargeDrySand, which along with LargeWetSoil, showed some of the greatest fiducial registration errors in this experiment. This illustrates the need for consistently accurate methods of registration. The pragmatic solution in this instance is to redo the registration of the samples by hand if the registration error exceeds a given threshold that would mark them as useless for further analysis otherwise.

In the case of the LargeWetSand sample shown in the examples so far, the sand grains can be identified in the neutron data as well as the X-ray data. This doesn't hold true in the samples more generally. Figure 5.15 shows the pre- and post-injection versions of slice 1851 from the sample SmallSoil4. This sample contains a seedling grown in soil and the smaller particle size of the growth medium combined with the substantially higher water content. As a result, it is far harder to make out the soil grains in the neutron data. That being said, the increase in water concentration between the scans can be seen by eye and even in the scan taken after the injection, the inside of the

5.3. Segmentation and quantification of water in soil and sand

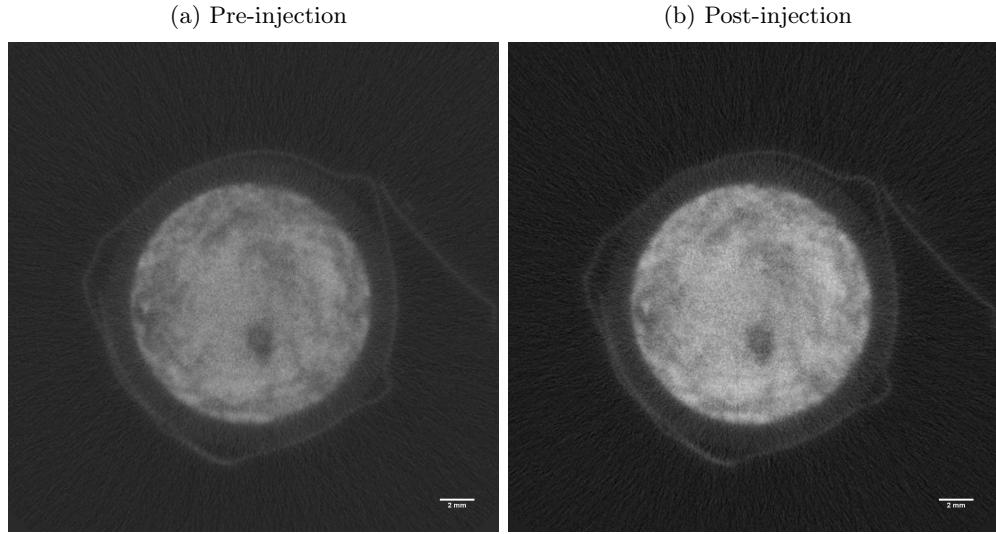


Figure 5.15.: Slice 1851 from the sample SmallSoil4.

sample tube does not appear as a homogeneous region solely dominated by the attenuation of water. The difference between the two images is shown in figure 5.16. The sample shows an even scattering of dark spots throughout the sample that were not visible in either scan alone. This would appear to be the soil particles showing as a lack of change between the two scans. The addition of an X-ray mask in figure 5.17 shows that the soil particles observed in the X-ray data are far larger than the dark spots seen in figure 5.16. It is likely that neutron scattering from the water, and to a lesser extent the deuterated water, in the sample is causing the soil particles to appear far smaller than they otherwise would, blurring the neutron information. This is supported by the difference that can be seen between the blurred edge at the inner wall of the sample tube in figure 5.16, which appears far sharper when the X-ray mask is added to the neutron data.

It can also be seen with this sample that the sample tube is removed far more effectively in the differential image than in the examples from the LargeWet-Sand sample, primarily due to the more accurate registration of this sample.

5.3.4. Joint processing of complementary modalities

We propose an alternative approach to combining X-ray and neutron data for analysis to use both modalities to inform the reconstruction of the neutron

5. Correlative imaging for water segmentation and quantification

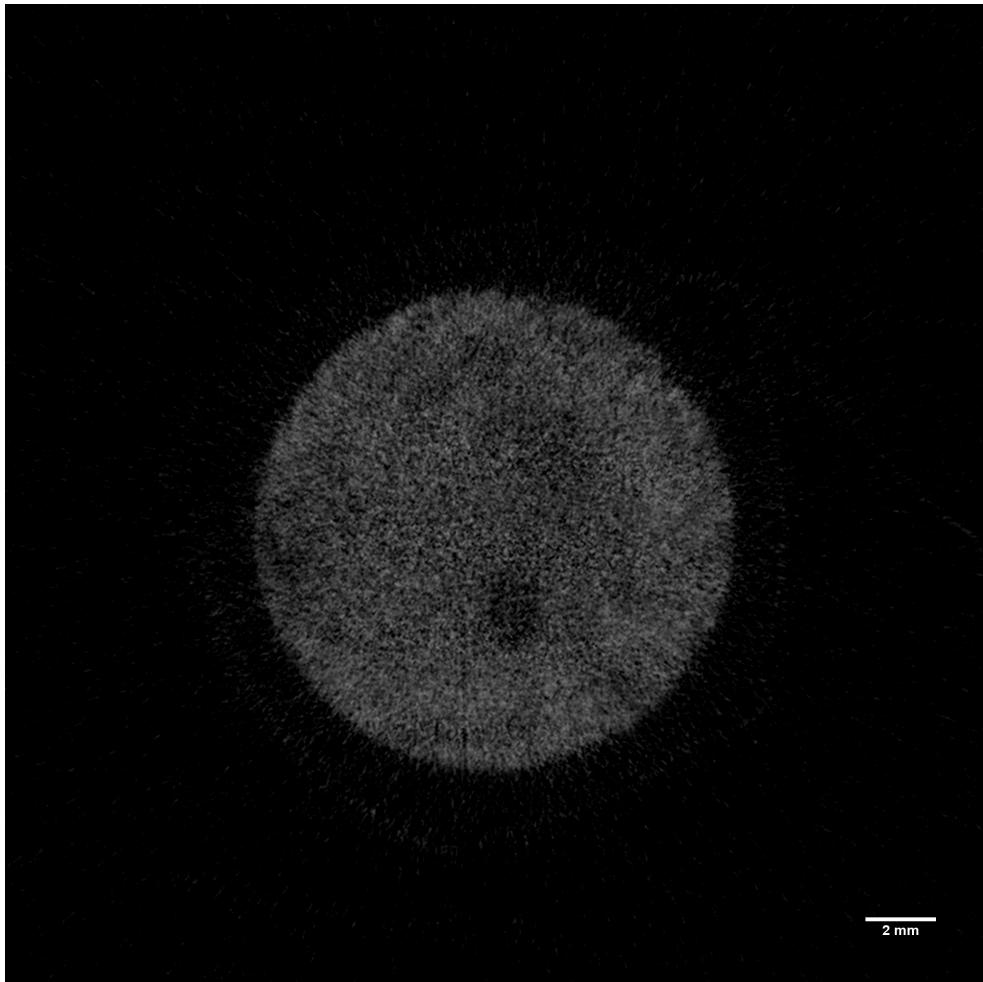


Figure 5.16.: The difference between slice 1851 of the pre- and post-injection neutron scans of SmallSoil4. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.3.

5.3. Segmentation and quantification of water in soil and sand

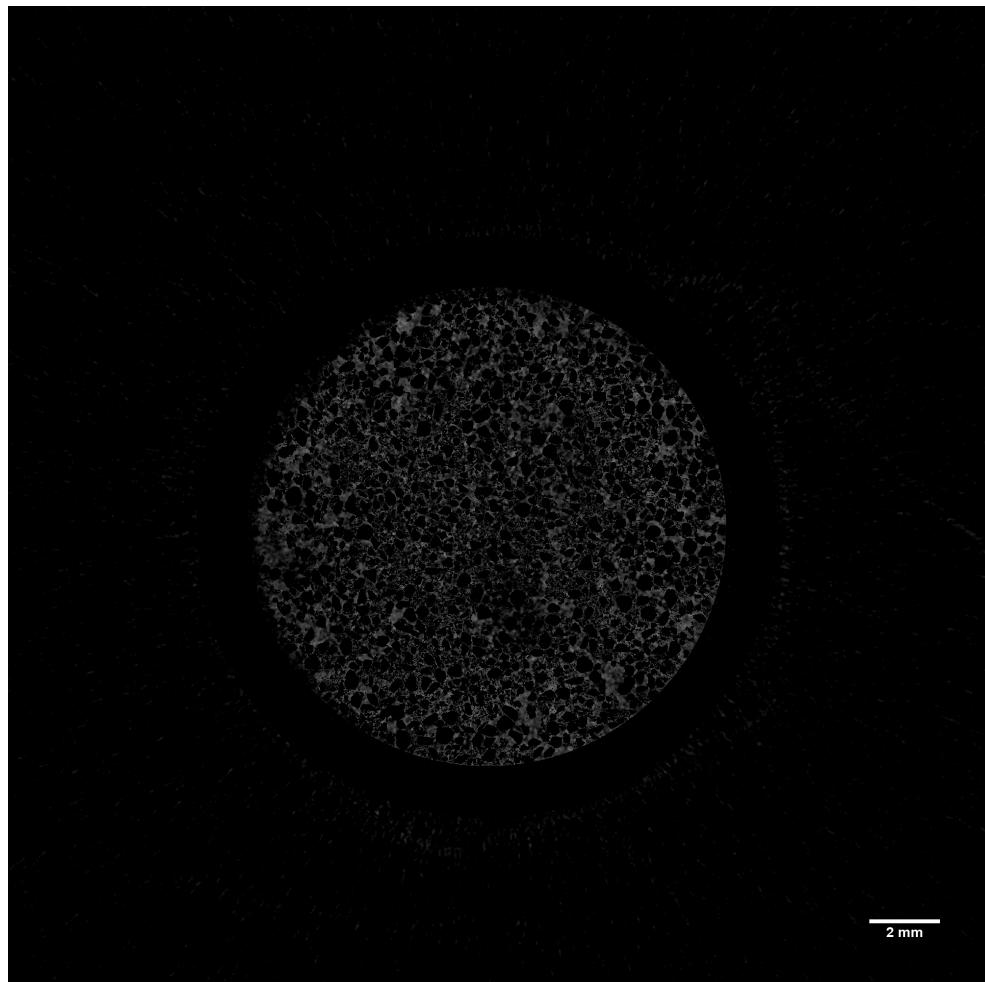


Figure 5.17.: The difference between the pre- and post-injection scans of slice 1851 from SmallSoil4 with an X-ray mask overlaid. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.4.

5. Correlative imaging for water segmentation and quantification

projection data. For this method, a model-based iterative reconstruction algorithm was applied to neutron data with the addition of a mask produced from the corresponding X-ray data. This new method uses the difference between the pre- and post-injection neutron scans to isolate the injected fluid content and reconstruct it with greater precision through the addition of the mask. Figure 5.18 shows the complete processing pipeline from data collection to the analysis presented here.

To reconstruct the neutron data using additional information from the corresponding X-ray data presupposes that the two modalities are well aligned. The methods developed so far all assume registration is applied to the volume data after reconstruction rather than the projection data. In order to use the registration tools already developed and avoid having to register the projection data, a slice from the reconstructed data is used as the input to the registration algorithm. An iterative reconstruction algorithm operates by alternately back-projecting then forward-projecting the data between the reconstructed image and the projection images [Willemink and Noël, 2019]. By inputting data that has already been reconstructed and registered and removing the back-projection step from the first iteration, new requirements for the registration can be avoided. The X-ray mask can also be entered in the reconstructed image space and ensures that the locations not selected by the mask will have a fixed value of 0 throughout the iterative reconstruction process.

Since there is no trusted benchmark data to evaluate this method against, its impact was first tested through simulation and then applied to real data. A simulated slice was created with wet and dry versions to represent the pre- and post-injection scans and a corresponding binarised mask, such as could be created from X-ray data (fig 5.19). So that the simulated neutron slice reflects the lower image quality of neutron data compared to X-ray data, the simulated neutron slice is scaled down by a factor of ten then returned to its original dimensions before Poisson noise is applied. By forward projecting, these slices can be converted to sinograms in preparation for the subsequent reconstruction tests.

A benchmark result is produced from the simulated scan data by reconstructing the pre- and post-injection data and then taking the difference between the reconstructed results applying the mask afterwards to isolate the injected water content. The new algorithm is then implemented by reconstructing the difference between the pre- and post-injection scans with the

5.3. Segmentation and quantification of water in soil and sand

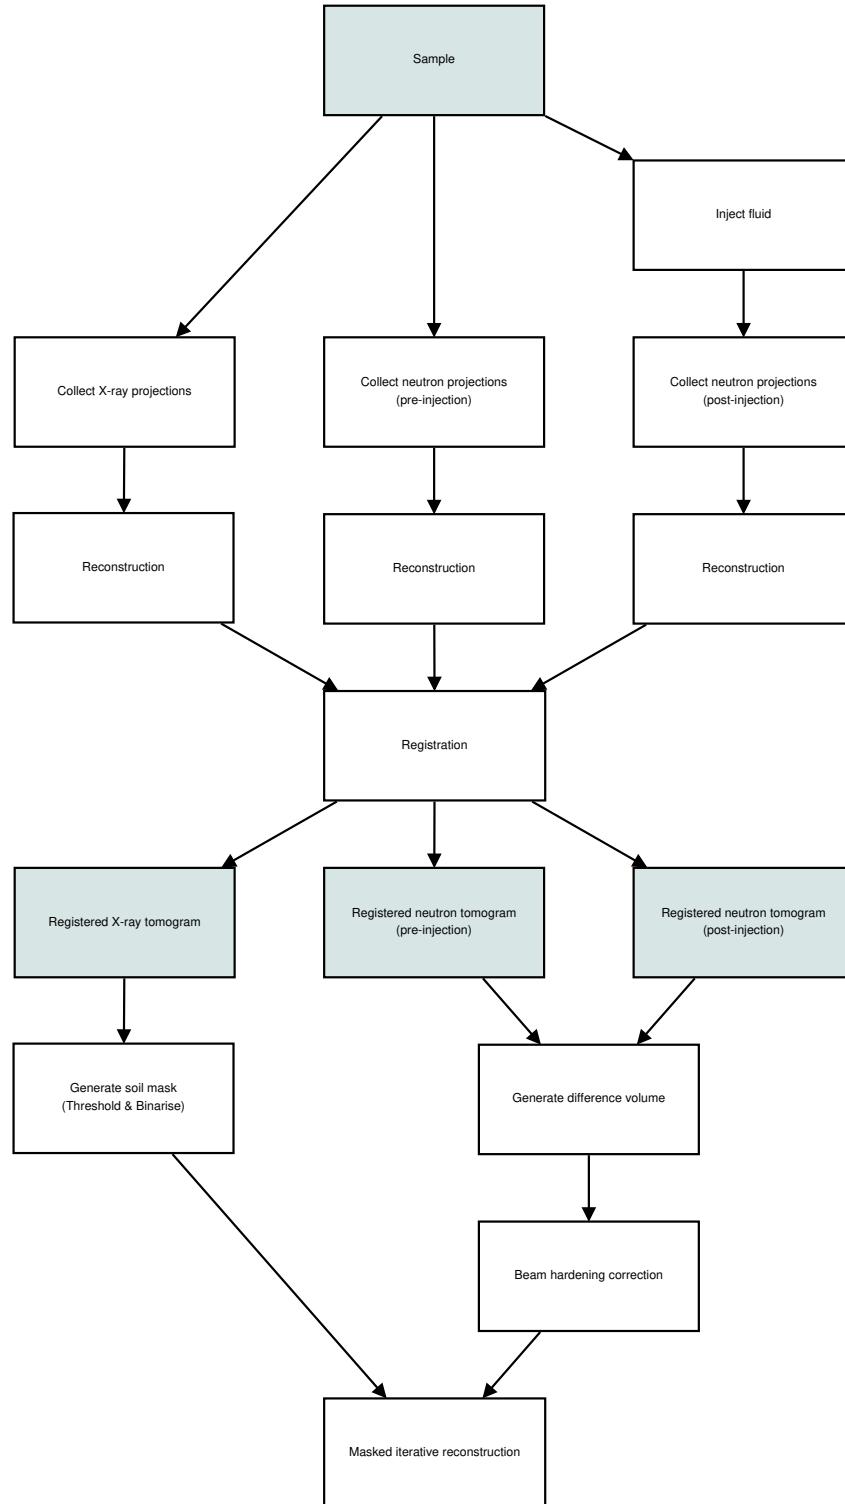


Figure 5.18.: The multi-modal data processing approach applied to this data set. The processing from sample to registered tomograms is covered in section 4.3. This chapter presents all subsequent processing from that step.

5. Correlative imaging for water segmentation and quantification

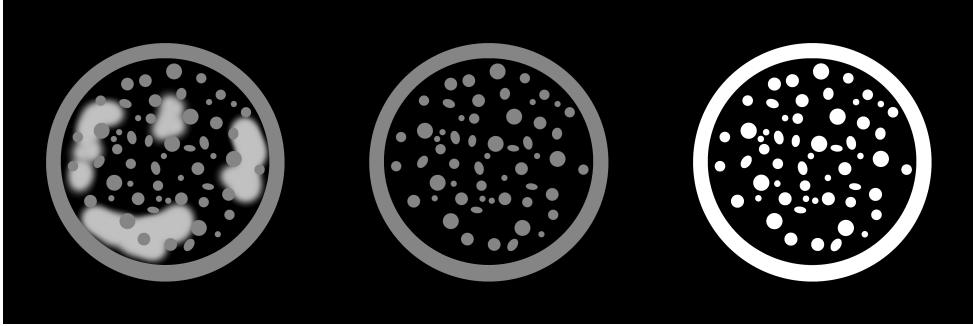


Figure 5.19.: The simulated slice: wet, dry and binarised

mask applied at each step of the iterative reconstruction. Figure 5.20 shows the resulting slices produced by each of these approaches.

To compare the quality of these results, they are subtracted from the original simulated slice (at full resolution), which gives an error map. Figure 5.21 shows the error maps produced by each method. The benchmark result shows rings around each marker, resulting in a far higher error than that of the new algorithm being evaluated. Taking the norm of the error maps makes this a quantitative result. The error in the benchmark test is found to be 13.6 and in the new method, 7.8. The simulation data suggests that the new method gives more precise results for accurate reconstruction and segmentation of the injected fluid, so the method was subsequently tested with real data. The full code for the simulation is included in A.4.

For comparability with the previous section, this method will be demonstrated with slice 4402 from the sample LargeWetSand. Figure 5.7 shows the neutron data that was used as an input to the reconstruction algorithm, and figure 5.11b shows the mask. As in the case of the simulation, the reconstruction was completed using the 2D CUDA implementation of the SIRT algorithm [Gilbert, 1972] from the ASTRA Toolbox [Van Aarle et al., 2016].

Figure 5.22 shows both the pre- and post-injection images for slice 4402 after reconstruction with 150 iterations of the SIRT algorithm and the X-ray mask. The noise in the neutron data is significantly reduced by using an iterative reconstruction algorithm, and the mask ensures that only the pore spaces between sand grains are shown in the reconstructed data. The difference between these two images can be taken (figure 5.23), which shows the injected fluid as in figure 5.12b.

Including the mask in the reconstruction using our new method rather than

5.3. Segmentation and quantification of water in soil and sand

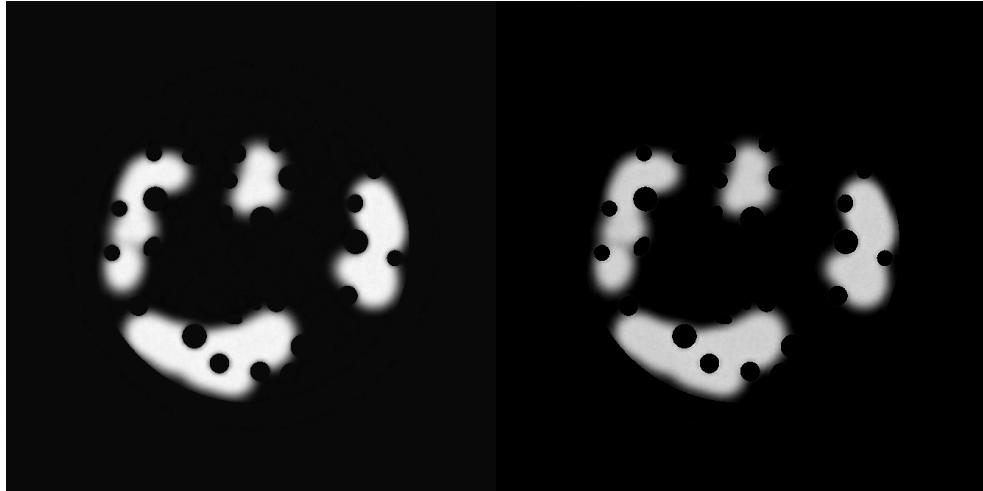


Figure 5.20.: The reconstructed slices after masking. The left image shows the benchmark result where the simulated pre- and post-injection data was reconstructed separately, and the mask applied to the difference between the reconstructions. The right image shows the new method where the difference data was reconstructed with iterative masking.

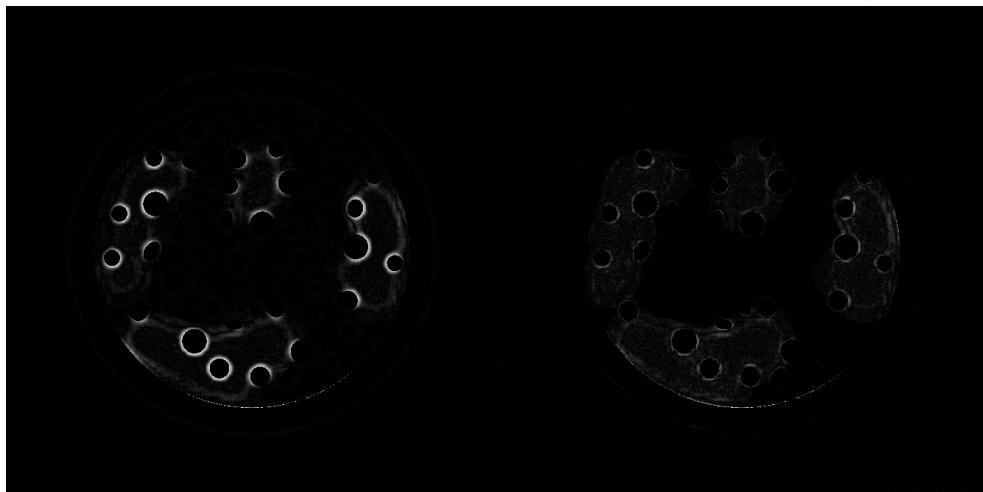


Figure 5.21.: The error maps of the reconstructed slices after masking. The left image shows the benchmark result where the simulated pre- and post-injection data was reconstructed separately, and the mask applied to the difference between the reconstructions. The right image shows the error produced with the new method where the difference data was reconstructed with iterative masking. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.5.

5. Correlative imaging for water segmentation and quantification

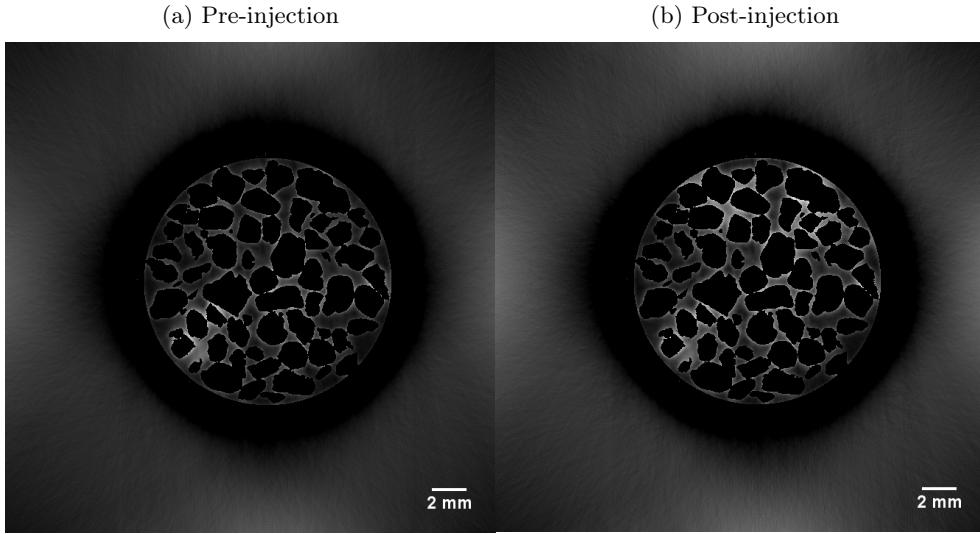


Figure 5.22.: Slice 4402 from the sample LargeWetSand reconstructed using the SIRT algorithm with an X-ray mask. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.6.

applying it afterwards makes a subtle but significant difference. Figure 5.24 shows how the regions of injected fluid appear brighter towards the centre when the mask is not applied until afterwards. In figure 5.23, by comparison, the regions of injected fluid appear more consistent in grey level with brighter edges at the boundaries of the grains. There is also less glow around the mask from the more highly attenuating grains when the mask is applied during reconstruction.

The simulation conducted shows that this novel method of iteratively reconstructing registered data using a mask improves the quality of the neutron data by taking advantage of the higher quality X-ray data. Although this claim is harder to verify in experimental data, the simulation provides strong evidence of the value of this method and the experimental data shows that it has some impact, despite the accuracy of this change being difficult to validate.

5.3.5. Quantification

After applying the techniques described above, one of the most common applications would be to analyse the injected fluid, particularly by quantification. Using the difference between the pre- and post-injection scans makes this very

5.3. Segmentation and quantification of water in soil and sand

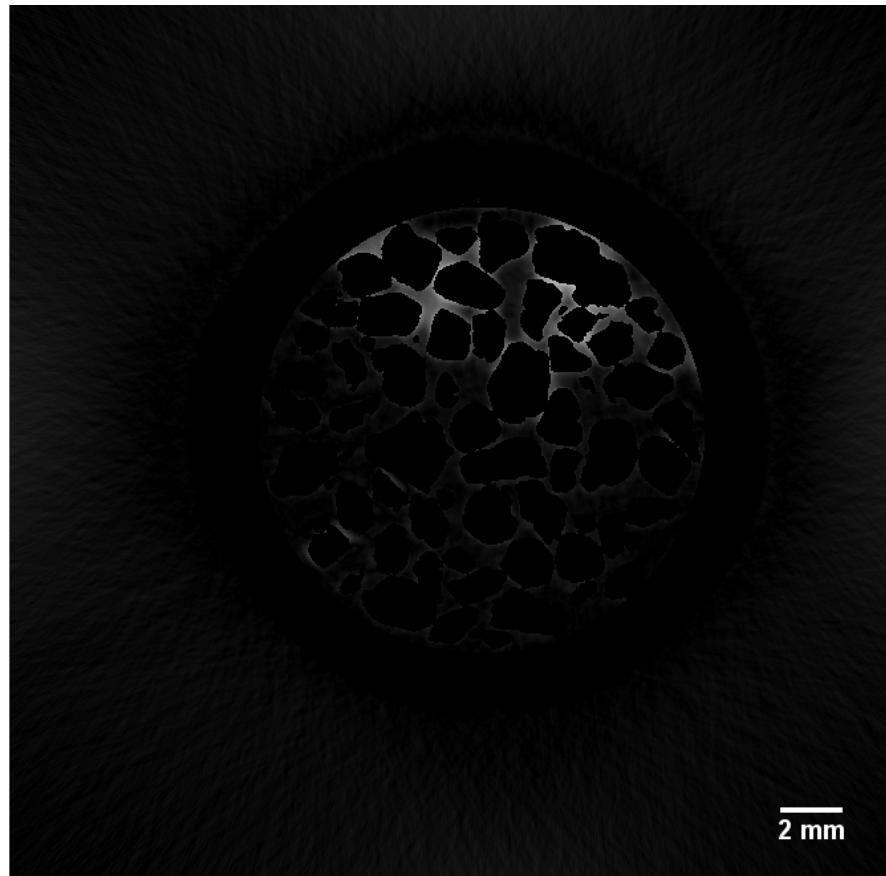


Figure 5.23.: The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with an X-ray mask. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.7.

5. Correlative imaging for water segmentation and quantification

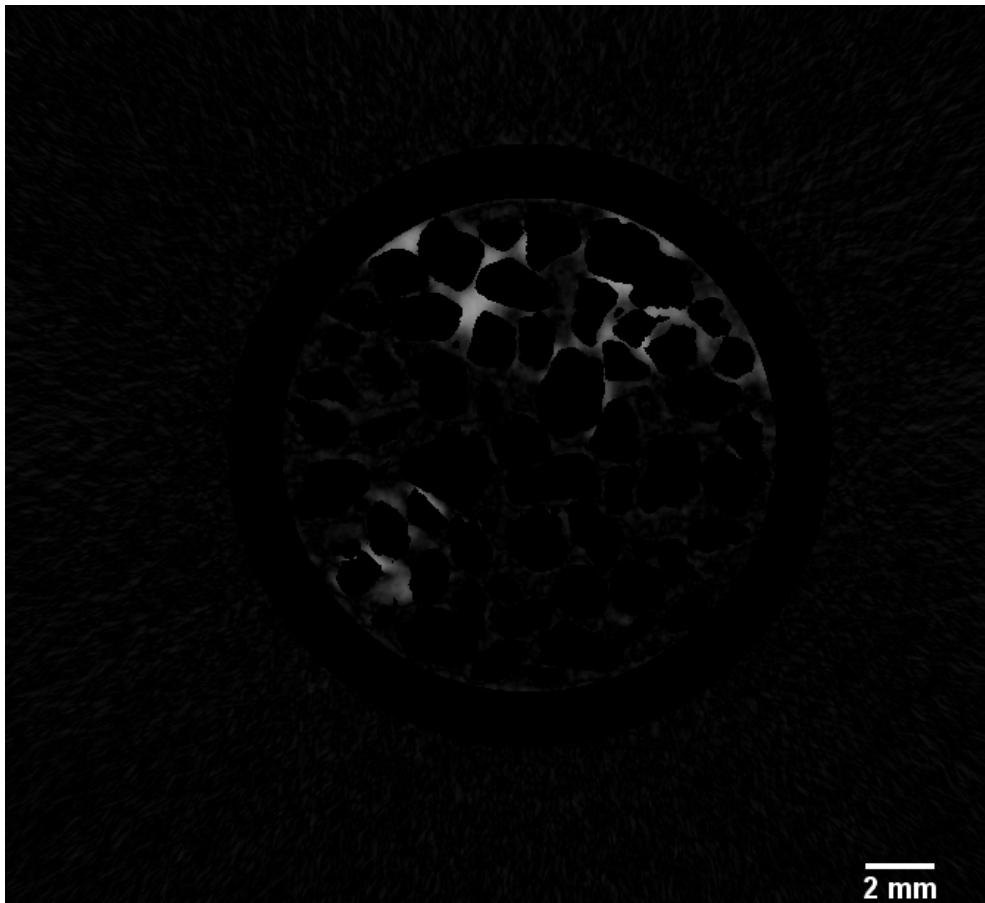


Figure 5.24.: The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with the X-ray mask applied afterwards. The contrast of this figure has been artificially enhanced. The original can be viewed as figure B.8.

simple, and the iterative masked reconstruction method from section 5.3.4 increases the accuracy of the analysis techniques beyond what they would otherwise be.

This section continues to use the LargeWetSand sample to demonstrate the process. The sample received a 0.25 ml injection of D₂O, and this is what is isolated by taking the difference between the pre- and post-injection scans and applying a soil mask based on the X-ray data. All data outside the sample tube can be removed since this information is not helpful to our analysis. By using the histogram analysis tools in Fiji, the number of pixels at each grey level and the total attenuation (sum of all grey levels) can be obtained. If we accept the assumptions that: all 0.25 ml of injected D₂O is within the field of view, all the attenuation in the data is from the injected D₂O, and the attenuation in a voxel is directly proportional to the amount of D₂O it contains, we can infer that the total attenuation of the data corresponds to the full quantity of injected fluid. From this, the volume of fluid in an individual voxel at a given grey level can be calculated using equation 5.1, where the volume of fluid in a voxel V_v is equal to the product of the total volume, V_T , and the grey level, G , divided by the total attenuation, A_T .

$$V_v = \frac{V_T \times G}{A_T} \quad (5.1)$$

Given a pair of calibration values from this formula, a calibration can be interpolated and applied to the data. Figure 5.25 shows a calibration bar added to slice 4724 to illustrate this result.

5.4. Conclusions and future work

This chapter has proposed a new method to evaluate how the data collected in the experiments of the preceding chapter can be processed for subsequent analysis, in particular, a technique that can be applied to take the best advantage of having multi-modal data and multiple neutron scans to get the best possible results. These techniques also demonstrated the importance of good registration, as minor errors in registration often introduced significant errors when these techniques were being tested.

The introduction of a syringe pump at IMAT allows a variety of new experiments to be planned. This work provided the first trial data with the pump

5. Correlative imaging for water segmentation and quantification

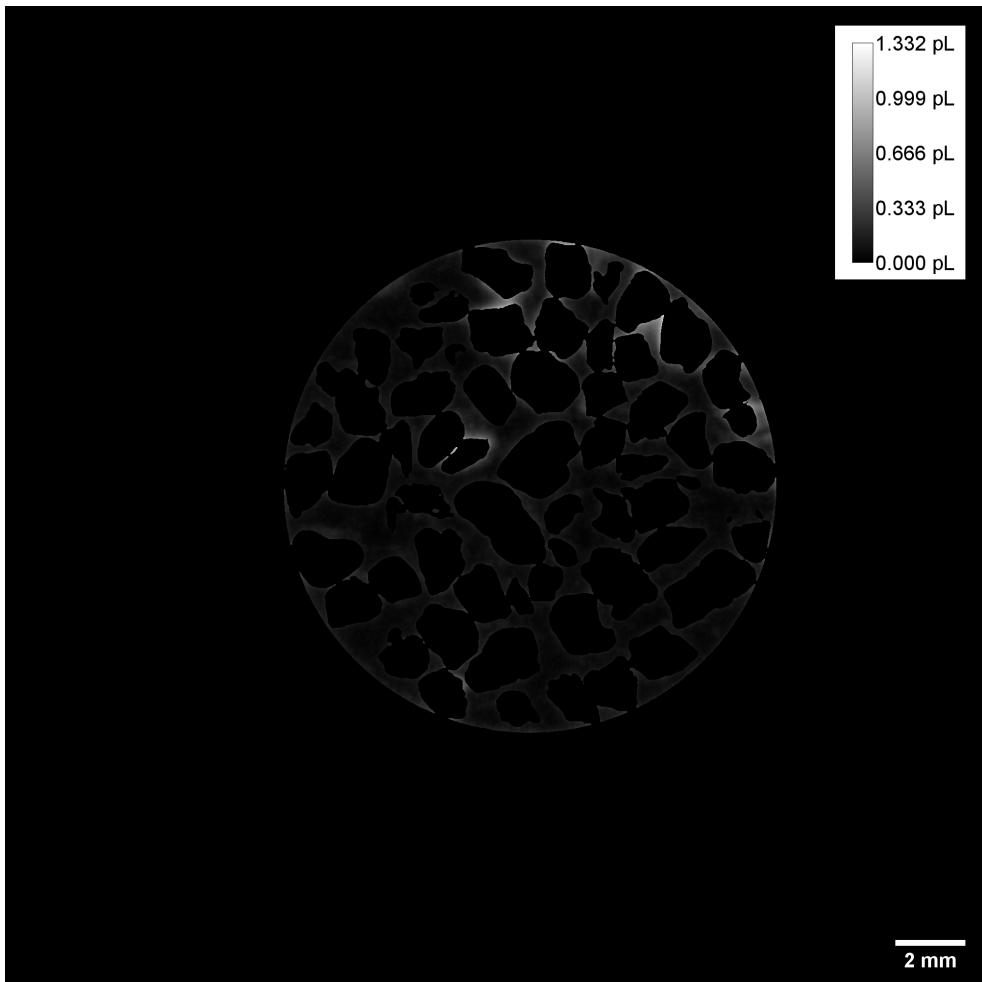


Figure 5.25.: A calibrated slice (4724) with a calibration bar to show the relationship between the amount of fluid in a voxel (in pico-litres) and the grey level.

and demonstrated subsequent processing tools to show the injection of the fluid and its diffusion through a series of radiographs collected over a suitable period. In particular, the power of having pre- and post-injection scans that can be compared to segment a known quantity of fluid from the data easily has been demonstrated. The effectiveness of golden angle acquisition for maximising image quality within limited scan time was evaluated and was found to help determine optimal scan parameters for a given sample but less efficient in a batch of scans after the scan parameters have been chosen. In particular, it is useful when the flow rate of fluid within the sample is unknown.

Having scans with and without the injected fluid makes various processing techniques possible since the segmentation can be completed so quickly and accurately by taking the difference between the scan data. The possibility of taking advantage of this to correct for beam hardening in the injected water was considered. Tests showed that this was only viable when there was high enough saturation of the soil for water to be distributed throughout the sample and where this was the case, little evidence of beam hardening was found. This technique could be helpful to scans with different sample parameters.

Steps were also taken to use the X-ray data to improve the neutron data, introducing it either during or after reconstruction as a mask. The larger sand grains were visible in both modalities, which made them ideal samples for initially testing these processing routines before applying them to the samples with smaller sand or soil that could not be visually verified so easily. The addition of a mask after registration imposed the grain boundaries from the X-ray data onto the neutron data and was shown to help remove the particles allowing better segmentation of the fluid in the pore spaces. The introduction of iterative reconstruction methods resulted in a neutron tomogram with substantially lower noise.

The iterative reconstruction algorithm was introduced because it offered a method for including the X-ray mask data during the reconstruction of the neutron data rather than applying it to the reconstructed volume. This method led to more accurate results: when the mask was applied after reconstruction, the attenuation due to fluid in the pores between the sand grains was highest at the centre of the pores. When the mask was used during reconstruction, the attenuation appeared more evenly distributed within the pores but was highest at the grain boundaries. These results were seen both in experimental and simulated data, and the improved fit with the “true data” that can be known

5. Correlative imaging for water segmentation and quantification

in the simulation gives confidence in the improved accuracy of this method.

The next steps to develop this processing pipeline would be to attempt a more thorough segmentation and quantification routine with appropriate calibration to determine how accurately quantities of fluid can be measured within a sample. More accurate estimates of water concentration could, in turn, allow more accurate beam hardening correction if necessary.

6. Final conclusions and future work

This work has aimed to develop multi-modal methods that can combine X-ray and neutron imaging techniques to scan complex samples that cannot fully be resolved using either technique alone. In particular, root-soil systems, including the water dynamics within. The work was focused, therefore, on developing suitable strategies for image acquisition and image analysis.

Chapter 3 presented initial simulations designed to evaluate the impact of the quantity and distribution of fiducial markers on registration accuracy and then a set of initial trial scans. Strategies for marker distribution were determined from the simulation results. The use of simulation to optimise the placement of fiducial markers rather than evaluate the placement of markers has not been observed in the literature however the findings of this simulation did match well to commonly accepted best practice. The impact of these optimisations on registration accuracy was not seen in practice as the dominant source of registration error was the accurate localisation of the markers. The need for a more suitable fiducial marker material was the most important from a list of lessons to apply to subsequent experiments.

Chapter 4 contains the follow-up experiment that built on the previous results of Chapter 3. Borosilicate was selected as a suitable material for fiducial markers, and a new set of scans were conducted. A large amount of data was generated in this study, making manual processing and registration of the data prohibitively time-consuming. Software tools were therefore developed to (partially) automate the process. In particular a stack concatenation tool and a sphere fitting tool to automatically locate the fiducial markers. The accuracy of these tools and the subsequent registration was then evaluated. The sphere fitting algorithm was shown to be partially successful but with limitations that led to a noticeable decrease in registration accuracy in some instances. This presents an opportunity for future work to build upon it and find ways to ensure that accurate registration of scans does not require lengthy and repetitive manual processing to complete but create an automated tool

6. Final conclusions and future work

that eliminates human error and is scalable for experiments with large numbers of samples to scan.

Chapter 5 assumes the successful collection and registration of correlative X-ray and neutron scans of a sample and discusses subsequent processing strategies. In particular, a processing scheme is proposed to take advantage of a syringe pump. This work was the first to introduce a syringe pump at IMAT, and many future experiments could take advantage of this equipment. The collection of neutron scans before and after the injection of a known quantity of fluid meant that the fluid could easily be isolated in the reconstructed data by taking the difference between the two scans. This dual scanning for easy segmentation enables new approaches to beam hardening correction and simplifies the calibration required for quantification. In addition to this, the use of higher resolution X-ray data to improve the quality of the complementary neutron data was explored. Iterative reconstruction of the neutron data using a mask produced from the X-ray data improves the image quality of the neutron data. This was first demonstrated with simulated data, and similar improvements were seen using actual scan data.

This project has demonstrated many of the complexities involved in registration of complementary CT data produced in different instruments, and the further difficulties in automating the process. In the context of scientific research, such as this project is primarily concerned with, every experiment may be different and experiments are often conducted with low numbers of samples or just a single sample. In these cases, manual registration can be completed and the results of chapters 3 and 4 indicate that this is likely to give a more accurate result as the imperfections of human error are outweighed by the limitations of the automated methods presented in this work. In industrial or medical contexts, larger sets of samples and greater uniformity between scans is to be expected and so in these environments automation of the registration process remains important.

One avenue for further development of automated registration algorithms would be to use machine learning tools either identify markers. There is plenty of work in two dimensional image classification to build upon and the application of machine learning to CT data is already being explored for some medical applications. One of the key requirements for these methods is a large library of training data which does not currently exist for correlative X-ray and neutron CT data and would not be easy to produce. Simulated datasets could

provide a solution, particularly as a supplement to experimental data, but these are difficult to produce with both the accuracy and variability required to properly train an algorithm.

If further work is to be done into registration of complementary CT datasets, it is key that beamlines come up with standardised approaches for applying fiducial points to samples so that a library of datasets can be built up to use for the development of general registration algorithms. This will simplify both experimental design and subsequent data processing, even if that involves manual rather than automated registration. The work presented in chapters 3 and 4 gives a suitable framework for the use of borosilicate fiducial markers that could be used in this capacity.

In the discussion of multi-modal imaging in section 2.2.3, it was noted that there is an increasing trend towards dedicated neutron and X-ray imaging instruments where both modalities can be scanned without having to transfer the sample between scanners. Such an approach has a range of benefits including minimum or no delay between scans for time sensitive samples and processes, no physical disturbance to the sample between scans and either the simplification or avoidance of registration. This work has sought to develop techniques and tools to ensure comparable results can be achieved without using a combined beamline and it has been shown that sub-voxel registration accuracy can be achieved allowing useful comparison of data collected at separate facilities. It is however clear that this can be produced more quickly, easily and reliably by using a dedicated multi-modal beamline and this is likely to remain the case until a reliable tool for at least semi-automated sub-voxel accurate registration can be presented. There remain some advantages to using separate instruments, in particular where specific beamline capabilities are desirable for a given experiment such as monochromatic X-ray imaging which is usually only available at a synchrotron beamline.

The conclusions learned throughout this work and the tools developed in it can inform the experimental design and image processing routines for future experiments that will rely on correlative neutron and X-ray imaging. While individual experiments often have unique requirements, this work offers a framework and toolset that can support the design of subsequent investigations and maximise their potential.

A. Code Listings

A.1. Marker simulation

A.1.1. fidtest.py

```
import numpy as np
from csvReader import writeResultstoCSV, readResultsfromCSV
from fidlib import genPoints, initFidGen, worker, TREmR, Cfim,
    getAverageTREmR

def findbestfidforallN():
    for n in range(Nmin, Nmax+1, 1):
        N = n
        fiducials = pool.map(worker, ((seeds, x, N, fidpoints,
            testpoints)
            for x in range(seeds.shape[0])))
        TREmRnlist = list(map(TREmRlister, fiducials))

        if plotting:
            plt.plotrange(TREmRnlist, N)

        print("N = %d" % N)
        maxIn = TREmRnlist.index(max(TREmRnlist))
        print(TREmRnlist[maxIn])
        poorresults.append(fiducials[maxIn].T)

        minIn = TREmRnlist.index(min(TREmRnlist))
        print(TREmRnlist[minIn])
        results.append(fiducials[minIn].T)

def TREmRlister(x):
    return TREmR(x, testpoints)

def AverageTREmRlister(x):
```

A. Code Listings

```
return getAverageTREmR(x, fidpoints, testpoints)

if __name__ == "__main__":
    tubeheight = 80
    radius = 15
    rings = 4
    anglestep = 10
    heightstep = 10
    global N
    Nmin = 4
    Nmax = 16

    plotting = 1
    quicktest = 0
    readonly = 1
    stats = 0

    # Set plotting to true to produce graphs
    # Set quicktest to true to significantly reduce N,
    # precision and computational time.
    # Set readonly to skip simulation and use csv data

    if plotting:
        import plotting as plt

    if quicktest:
        anglestep = 36
        heightstep = 20
        Nmin = 4
        Nmax = 5

    if (plotting or stats or not readonly):
        fidpoints, testpoints = genPoints(radius, tubeheight,
                                           rings,
                                           anglestep, heightstep)
        if plotting:
            plt.plot(fidpoints, radius, tubeheight)
            plt.plot(testpoints, radius, tubeheight)

    if not readonly:
        import multiprocessing as mp
        nCore = mp.cpu_count()
```

```

pool = mp.Pool(nCore)

results = []
poorresults = []

seeds = initFidGen(fidpoints, tubeheight, heightstep)

findbestfidforallN()

writeResultstoCSV(results, 'results.csv')
writeResultstoCSV(poorresults, 'poorresults.csv')

csvResults = []
poorcsvResults = []
cfimRMS = []

csvResults = readResultsfromCSV(csvResults, 'results.csv',
                                Nmin, Nmax)
poorcsvResults = readResultsfromCSV(poorcsvResults, 'poorresults.csv',
                                    Nmin, Nmax)

if plotting:
    for x in range(0, Nmax-Nmin+1, 1):
        plt.plotfid(csvResults[x], radius, tubeheight, x+Nmin,
                     1)
        plt.plotfid(poorcsvResults[x], radius, tubeheight, x+
                     Nmin, 0)

    cfimList = np.zeros(x+4)
    for i in range(x+4):
        cfimList[x] = Cfim(i, csvResults[x], testpoints)
    cfimRMS.append(np.sqrt(np.mean(np.square(cfimList)))))

TREmRlist = list(map(TREmRlister, csvResults))

plt.plotMinAvgTRE(Nmin, Nmax, TREmRlist)
plt.plotAvgMarkerContr(Nmin, Nmax, cfimRMS)
print(TREmRlist)

if stats:
    if not plotting:
        TREmRlist = list(map(TREmRlister, csvResults))
    print(TREmRlist)

```

A. Code Listings

```

Nset = np.arange(Nmin, Nmax+1)
AvgTREmRlist = list(map(AverageTREmRlister, Nset))
print(AvgTREmRlist)

```

A.1.2. fidlib.py

```

import numpy as np
from numpy import linalg as la

tiny = 0.000001

def findCentroid(fiducials):
    centroid = np.divide(np.sum(fiducials, axis=1), fiducials.
        shape[1])
    return centroid.reshape(centroid.shape[0], -1)

def findInertiaMatrix(fiducials):
    a = np.zeros(shape=(3, 3))
    a[0, 0] = np.sum(np.square(fiducials[[1, 2], :]))
    a[1, 1] = np.sum(np.square(fiducials[[0, 2], :]))
    a[2, 2] = np.sum(np.square(fiducials[[0, 1], :]))
    a[0, 1] = a[1, 0] = -np.sum(np.prod(fiducials[[0, 1], :], axis
        =0))
    a[0, 2] = a[2, 0] = -np.sum(np.prod(fiducials[[0, 2], :], axis
        =0))
    a[1, 2] = a[2, 1] = -np.sum(np.prod(fiducials[[1, 2], :], axis
        =0))
    return a

def perpDist(a, b, c):
    ab = np.repeat(np.subtract(a, b.T), c.shape[0], axis=0)
    cb = np.tile(np.subtract(c, b), (a.shape[0], 1))
    cross = np.cross(ab, cb)
    top = la.norm(cross, axis=1)
    return np.divide(top, la.norm(cb))

def TREm(r, fiducials):
    # not optimised for use with TREmR!!!
    r = np.reshape(np.asarray(r), (1, 3))
    centroid = findCentroid(fiducials)

```

```

inertiaMatrix = findInertiaMatrix(fiducials)
eigval, axes = la.eig(inertiaMatrix)
f = perpDist(fiducials, centroid, (centroid + axes))**2
d = perpDist(r.T, centroid, (centroid + axes))**2
f = np.mean(np.reshape(f, (-1, 3)), axis=0)
df = np.divide(d, f)
TREm = (1 + np.sum(df)/3) / (fiducials.shape[1]-2)
return TREm

def TREmR(fiducials, testpoints):
    centroid = findCentroid(fiducials)
    inertiaMatrix = findInertiaMatrix(fiducials)
    eigval, axes = la.eig(inertiaMatrix)
    pointCs = centroid + axes
    f = perpDist(fiducials.T, centroid, pointCs)**2
    f = np.mean(np.reshape(f, (-1, 3)), axis=0)
    dvals = perpDist(testpoints, centroid, pointCs)**2
    drm = np.mean(np.reshape(dvals, (-1, 3)), axis=0)
    if (np.count_nonzero(f) < 3):
        f[f == 0] = tiny
    df = np.divide(drm, f)
    return (1 + np.sum(df)/3) / (fiducials.shape[1]-2)

def Cfim(i, fiducials, testpoints):
    ifid = np.delete(fiducials, i, axis=1)
    return (TREmR(ifid, testpoints) - TREmR(fiducials, testpoints))
)

def Ccpm(p, fiducials, testpoints):
    for x in np.arange(0, fiducials.shape[1]):
        if(np.all(p == fiducials[:, x])):
            return -5
    pfid = np.column_stack((fiducials, p))
    return (TREmR(fiducials, testpoints) - TREmR(pfid, testpoints))
)

def initCcpm(p, fiducials, testpoints):
    for x in np.arange(0, fiducials.shape[1]):
        if(np.all(p == fiducials[:, x])):
            return -5

```

A. Code Listings

```

pfid = np.column_stack((fiducials, p))
return (TREmR(pfid, testpoints))

def initFidGen(fidpoints, tubeheight, heightstep):
    size = fidpoints.shape[0]
    x = tubeheight//heightstep + 1
    y = fidpoints[0:x]
    repeats = np.arange(size-1, size-(x+1), -1)
    i = np.repeat(y, repeats, axis=0)
    j = np.array([[0, 0, 0]])
    for x in repeats:
        j = np.concatenate((j, fidpoints[size-x:size]), axis=0)
    j = np.delete(j, 0, axis=0)
    seeds = np.concatenate((i, j), axis=1)
    return np.reshape(seeds, (-1, 3, 2), order='F')

def optFid(fiducials, N, fidpoints, testpoints):
    # Add the best possible fiducial points up to N points
    while (fiducials.shape[1] < N):
        if (fiducials.shape[1] == 2):
            potentialFid = np.apply_along_axis(initCcpm, 1,
                                                fidpoints,
                                                fiducials,
                                                testpoints)
        else:
            potentialFid = np.apply_along_axis(Ccpm, 1, fidpoints,
                                                fiducials,
                                                testpoints)
        fiducials = np.column_stack((fiducials,
                                     fidpoints[np.argmax(
                                         potentialFid)]))

    evalFid = np.zeros(N)
    prevMin = -1
    while(1):
        for i in range(fiducials.shape[1]):
            evalFid[i] = Cfim(i, fiducials, testpoints)
        fiducials = np.delete(fiducials, np.argmin(evalFid), 1)
        potentialFid = np.apply_along_axis(Ccpm, 1, fidpoints,
                                            fiducials, testpoints)
        fiducials = np.column_stack((fiducials,
                                     fidpoints[np.argmax(
                                         potentialFid)]))

```

```

    potentialFid)]) )
if np.argmin(evalFid) == prevMin:
    break
prevMin = np.argmin(evalFid)
return fiducials

def genPoints(radius, tubeheight, rings, anglestep, heightstep):
    # arange(start, end, stepsize)
    angle = np.arange(0, 360, anglestep) * (np.pi/180)
    height = np.linspace(0, tubeheight, tubeheight//heightstep+1)

    radpoints = np.stack((radius*np.cos(angle), radius*np.sin(
        angle)), axis=1)
    fidpoints = np.column_stack((np.repeat(radpoints, height.size,
        axis=0),
        np.tile(height, angle.size)))

    testRadii = np.linspace(0, radius, rings)
    testAngles = 360/testRadii[1:]
    testpoints = np.column_stack((np.repeat(np.array([[0., 0.]]),
        height.size, axis=0), height))
    for x in range(rings-1):
        angle = np.arange(0, 360, testAngles[x]) * (np.pi/180)
        angles = np.stack((np.cos(angle), np.sin(angle)), axis=1)
        testpointsXflat = testRadii[x+1] * angles
        testpointsX = np.column_stack((np.repeat(testpointsXflat,
            height.size, axis=0),
            np.tile(height, angle.size)))
    testpoints = np.concatenate((testpoints, testpointsX),
        axis=0)
return fidpoints, testpoints

def worker(arg):
    seeds, i, N, fidpoints, testpoints = arg
    return optFid(seeds[i, :, :], N, fidpoints, testpoints)

def getAverageTREmR(N, fidpoints, testpoints):
    TREmRlist = []
    for i in range(500000):
        good = False    # Test to ensure points are not co-linear

```

A. Code Listings

```
while not good:
    fidset = np.random.choice(len(fidpoints), N, replace=False)
    if not (np.any(np.all(fidpoints[fidset] == fidpoints[fidset][0, :], axis=0)[0:2])):
        good = True
    TREmRlist.append(TREmR(fidpoints[fidset].T, testpoints))
print(min(TREmRlist))
print(max(TREmRlist))
return np.mean(TREmRlist)
```

A.2. Stack concatenation

```
from ij import IJ, ImagePlus, ImageStack
from ij.plugin import ImageCalculator, Concatenator
from ij import WindowManager as WM
from ij.measure import Measurements, Minimizer, UserFunction,
    CurveFitter
from ij.gui import GenericDialog
import random
import datetime

initParams = []

def printTime():
    now = datetime.datetime.now()
    print(now)

def stackSelectionDialog():
    if (WM.getImageCount() < 2):
        print("There should be two image stacks to concatenate.\
              Please open both and try again.")
        exit()
    gd = GenericDialog('Select stacks to concatenate')
    imageTitles = [WM.getImage(id).getTitle() for id in WM.
        getIDList()]
    gd.addChoice("Lower volume", imageTitles, imageTitles[0])
    gd.addChoice("Upper volume", imageTitles, imageTitles[1])
    gd.showDialog()
    if gd.wasCanceled():


```

```

print("User canceled the dialog!")
return

title1 = gd.getNextChoice()
title2 = gd.getNextChoice()
imp1 = WM.getImage(title1)
imp2 = WM.getImage(title2)
return imp1, imp2


def meanDiff(u, v):
    if ((u < 1) or (v < 1)):
        return float('Inf')
    if ((v > 1990) or (u > 1990)):
        return float('Inf')
    a = volume1.getProcessor(u)
    b = volume2.getProcessor(v)
    imp1 = ImagePlus()
    imp2 = ImagePlus()
    imp1.setProcessor(a)
    imp2.setProcessor(b)
    imp3 = ImageCalculator().run("Subtract create", imp1, imp2)
    sumabs = sum([abs(i) for i in imp3.getProcessor().getPixels()])
    return sumabs

def meanDiffsort(x):
    return meanDiff(x, initvol2)

class func1 (UserFunction):
    def userFunction (self, params, x):
        return meanDiff(int(params[0]), int(initParams[1]))


class func2 (UserFunction):
    def userFunction (self, params, x):
        return meanDiff(int(params[0]), int(params[1]))


f1 = func1()
f2 = func2()

```

A. Code Listings

```
def findJoin(vol1, vol2):
    print("Looking for matching slice ...")
    printTime()
    stacksize = vol1.getSize()
    initParam1 = [int(stacksize * 0.6)]
    initvol2 = 200
    initDiff = float('Inf')
    step = 10

    print("Searching for local minima")
    printTime()
    for i in range(int(stacksize * 0.6), stacksize, step):
        meanDiffi = meanDiff(i, initvol2)
        if meanDiffi < initDiff:
            if (i - step == initParam1[-1]) or (i == initParam1[-1]):
                initParam1[-1] = i
            else:
                initParam1.append(i)
            initDiff = meanDiffi

    if len(initParam1) > 3:
        initParam1 = initParam1[-3:]

    m1test = float('Inf')
    m1testresult = []

    print("Optimizing for a fixed slice")
    printTime()
    for i in initParam1:
        global initParams
        initParams = [i, initvol2]
        m1 = Minimizer()
        m1.setFunction(f1, 1)
        m1.setExtraArrayElements(1)
        m1.setParamResolutions([1.0])
        m1status = m1.minimize(initParams, [10])
        m1result = m1.getParams()
        m1value = m1.getFunctionValue()

        if m1value < m1test:
            m1test = m1value
            m1testresult = m1result
```

```

print("Optimizing with 2 free slices.")
printTime()
m2initParams = [m1testresult[0], initvol2]

m2 = Minimizer()
m2.setFunction(f2, 2)
m2.setParamResolutions([1.0, 1.0])
m2status = m2.minimize(m2initParams, [10, 10])

print("Corresponding slices found")
printTime()
m2results = m2.getParams()
vol1end = int(m2results[0])
print("Vol1 end: {}".format(vol1end))
vol2start = int(m2results[1])
print("Vol2 start: {}".format(vol2start))
return vol1end, vol2start

def reduceSubStacks(im1, im2):
    vol1 = im1.getImageStack()
    vol2 = im2.getImageStack()
    vol1end, vol2start = findJoin(vol1, vol2)
    print("Removing overlap")
    printTime()
    IJ.run(im1, "Slice Remover",
           "first={} last={} increment=1".format(vol1end + 1, vol1
           .getSize()))
    IJ.run(im2, "Slice Remover",
           "first=1 last={} increment=1".format(vol2start - 1))

def calibrateLevels(imS1, imS2):
    print("Starting slice calibration")
    printTime()
    vol1 = imS1.getImageStack()
    vol2 = imS2.getImageStack()
    imp1 = vol1.getProcessor(vol1.getSize())
    imp2 = vol2.getProcessor(1)

    regions = 0
    rangeMeans = 0
    vol1Means = []
    vol2Means = []

```

A. Code Listings

```

stdDev = max(imp1.getStats().stdDev, imp2.getStats().stdDev)
print(stdDev)

print("Searching for sample areas")
printTime()
while (regions < 20) or (rangeMeans < stdDev * 3):
    x = random.randint(0, vol1.getWidth() - 10)
    y = random.randint(0, vol1.getHeight() - 10)
    imp1.setRoi(x, y, 20, 20)
    imp2.setRoi(x, y, 20, 20)
    stats1 = imp1.getStatistics()
    stats2 = imp2.getStatistics()
    vol1Range = stats1.max - stats1.min
    vol2Range = stats2.max - stats2.min
    if ((vol1Range < 30) and (vol2Range < 30) and
        (vol1Range != 0) and (vol2Range != 0)):
        regions = regions + 1
        vol1Mean = stats1.mean
        vol2Mean = stats2.mean
        vol1Means.append(vol1Mean)
        vol2Means.append(vol2Mean)
        allmeans = vol1Means + vol2Means
        rangeMeans = max(allmeans) - min(allmeans)
print("Determining and applying correction for selected areas"
      )
printTime()
fitter = CurveFitter(vol1Means, vol2Means)
fitter.doFit(0)
params = fitter.getParams()
print("params (c, m): ({}, {})".format(params[0], params[1]))
IJ.run(impS1, "Add...", "value={} , stack".format(params[0]))
IJ.run(impS1, "Multiply...", "value={} , stack".format(params
[1]))

imp1, imp2 = stackSelectionDialog()
volume1 = imp1.getImageStack()
volume2 = imp2.getImageStack()
print("Suitable stacks have been selected")
reduceSubStacks(imp1, imp2)
calibrateLevels(imp1, imp2)
imp1.setTitle("Corrected levels")
print("Levels corrected: starting calibrated concatenation")
printTime()

```

```

result = Concatenator().concatenate(imp1, imp2, 1)
result.setTitle("Complete")
result.show()
print("Script complete")
printTime()

```

A.3. Sphere fitting

A.3.1. test.py

```

import os
import csv
import numpy as np
from Volume import Volume
from SphereBuilder import SphereBuilder
from CircleHunter import CircleHunter

# import matplotlib.pyplot as plt
# from mpl_toolkits.mplot3d import Axes3D

# sample_tube_inner_rad = 880-900
# X marker_rad = 60-80

sphereBuilder = SphereBuilder()
circleHunter = CircleHunter()

xPath = "../I12Data/"
ndPath = "../RB1920729/Tomodry/"
nwPath = "../RB1920729/Tomo0.25/"
postPath = "../RegisteredVolumes/"
postFPath = "../RegisteredVolumes/FILTERED/"

xCPath = "CircleData/PreRegistration/Xray/"
ndCPath = "CircleData/PreRegistration/NeutronDry/"
nwCPath = "CircleData/PreRegistration/NeutronWet/"
postCPath = "CircleData/Registered/"
postxCPath = "CircleData/Registered/Xray/"
postndCPath = "CircleData/Registered/NeutronDry/"
postnwCPath = "CircleData/Registered/NeutronWet/"

def findCirclesInFolder(path, modality):
    for file in os.listdir(path):

```

A. Code Listings

```

if file.endswith(".raw"):
    if len(file.split("_")) == 3:
        [name, mod, dims] = file.split("_")
        name = name + "_" + mod
    elif len(file.split("_")) == 2:
        [name, dims] = file.split("_")
        dims = dims.split(".") [0]
        dims = dims.split("x")
        dims = list(map(int, dims))
    print(name, file, dims)
    if (name + ".csv") not in os.listdir(postCPath):
        data = Volume(name, path + file, dims, modality)
        data.circles = CircleHunter().findCircles(data)
        data.saveCircleData()
    del data

def findMarkersInFolder(path):
    for file in os.listdir(path):
        if file.endswith(".csv") and (file.count("Markers") == 0):
            data = openCircleData(path + file)
            print(file)
            samples = 10

            if file in ("SmallSoil4_NW.csv", "LargeSoil2_ND.csv"):
                samples = 6
            if file in ("SmallSoil2_NW.csv", "SmallWetSoil_ND.csv",
                        "SmallDrySoil_ND.csv"):
                samples = 2
            if file in ("SmallDrySoil_NW.csv", "SmallDrySand_NW.csv",
                        "SmallSoil4_ND.csv", "SmallSoil2_ND.csv"):
                samples = 1
            if file in ("LargeDrySand_NW.csv"):
                samples = 100

            if "_X.csv" in file:
                samples = 70

            if file in ("SmallWetSand_X.csv", "SmallSoil5_X.csv",
                        "SmallSoil3_X.csv", "LargeWetSand_X.csv",
                        "LargeDrySand_X.csv",
                        "LargeSoil1_X.csv", "SmallDrySand_X.csv"):
                samples = 10

```

```

samples = 50

if file in ("LargeSoil2_X.csv", "LargeSoil3_X.csv",
            "LargeDrySoil_X.csv", "SmallDrySoil_X.csv"):
    samples = 5

if file in ("LargeSoil4_X.csv", "SmallSoil2_X.csv"):
    samples = 30

if file in ("LargeDrySoil_ND.csv", "LargeDrySand_ND.csv",
            "LargeDrySand_X.csv", "LargeWetSoil_X.csv",
            "LargeSoil5_X.csv", "SmallSoil3_X.csv",
            "LargeWetSand_X.csv", "SmallDrySand_X.csv"):
    samples = 50
    markers = sphereBuilder.clusterPoints(data,
                                           samples, 60)
elif file in ("LargeSoil2_X.csv", "LargeSoil3_X.csv",
              "LargeDrySoil_X.csv", "SmallDrySoil_X.csv"):
    markers = sphereBuilder.clusterPoints(data,
                                           samples, 70)
else:
    markers = sphereBuilder.clusterPoints(data,
                                           samples)

with open("ranges.csv") as csvfile:
    allRanges = np.asarray(list(csv.reader(csvfile)))
    indices = np.where(allRanges == file.split(".")
                           [0])
    ranges = np.ndarray.flatten(
        allRanges[indices[0], 1:5]).astype(np.int)

rA = markers[np.logical_and((markers[:, -1] > ranges
                           [0]),
                           (markers[:, -1] < ranges
                           [1]))]
rB = markers[np.logical_and((markers[:, -1] > ranges
                           [2]),
                           (markers[:, -1] < ranges
                           [3]))]

```

A. Code Listings

```
markers = np.vstack((rA, rB))
print(markers)

if file in ("SmallSoil3_NW.csv", "SmallSoil2_NW.csv",
            "LargeSoil5_NW.csv", "LargeWetSand_NW.csv"
            ,
            "LargeSoil2_NW.csv", "LargeDrySand_NW.csv"
            ,
            "LargeDrySand_ND.csv", "SmallSoil3_ND.csv"
            ,
            "SmallSoil2_ND.csv", "LargeSoil5_ND.csv",
            "LargeSoil2_ND.csv", "LargeSoil1_ND.csv",
            "LargeDrySoil_ND.csv", "SmallWetSoil_X.csv"
            ,
            "SmallSoil4_X.csv", "SmallSoil3_X.csv",
            "LargeDrySand_X.csv", "LargeWetSoil_X.csv"
            ,
            "LargeDrySoil_X.csv", "LargeSoil1_X.csv",
            "LargeSoil2_X.csv", "LargeSoil3_X.csv",
            "LargeSoil4_X.csv", "SmallDrySand_X.csv",
            "SmallDrySoil_X.csv", "SmallSoil2_X.csv"):
    markers = sphereBuilder.filterInnerPoints(markers)

if file == "SmallSoil2_ND.csv":
    markers = np.delete(markers, 6, 0)
print(markers)

if file == "LargeWetSand_ND.csv":
    markers = np.delete(markers, 5, 0)
print(markers)

if file == "LargeSoil1_ND.csv":
    markers = np.delete(markers, 11, 0)
    markers = np.delete(markers, 4, 0)
    markers = np.delete(markers, 1, 0)
print(markers)

if file == "LargeDrySoil_ND.csv":
    markers = np.delete(markers, 3, 0)
print(markers)

if file == "LargeDrySand_ND.csv":
    markers = np.delete(markers, 12, 0)
```

```

    markers = np.delete(markers, 11, 0)
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 8, 0)
    markers = np.delete(markers, 6, 0)
    markers = np.delete(markers, 4, 0)
    print(markers)

if file == "SmallWetSoil_X.csv":
    markers = np.delete(markers, 10, 0)
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 5, 0)
    markers = np.delete(markers, 4, 0)
    markers = np.delete(markers, 2, 0)
    print(markers)

if file == "SmallWetSand_X.csv":
    markers = np.delete(markers, 4, 0)
    print(markers)

if file == "SmallSoil5_X.csv":
    markers = sphereBuilder.filterInnerPoints(markers,
                                              550)
    markers = np.delete(markers, 37, 0)
    markers = np.delete(markers, 36, 0)
    markers = np.delete(markers, 35, 0)
    markers = np.delete(markers, 33, 0)
    markers = np.delete(markers, 32, 0)
    markers = np.delete(markers, 30, 0)
    markers = np.delete(markers, 29, 0)
    markers = np.delete(markers, 28, 0)
    markers = np.delete(markers, 27, 0)
    markers = np.delete(markers, 26, 0)
    markers = np.delete(markers, 25, 0)
    markers = np.delete(markers, 24, 0)
    markers = np.delete(markers, 23, 0)
    markers = np.delete(markers, 22, 0)
    markers = np.delete(markers, 21, 0)
    markers = np.delete(markers, 20, 0)
    markers = np.delete(markers, 18, 0)
    markers = np.delete(markers, 17, 0)
    markers = np.delete(markers, 16, 0)
    markers = np.delete(markers, 15, 0)
    markers = np.delete(markers, 14, 0)
    markers = np.delete(markers, 13, 0)

```

A. Code Listings

```
markers = np.delete(markers, 12, 0)
markers = np.delete(markers, 11, 0)
markers = np.delete(markers, 10, 0)
markers = np.delete(markers, 9, 0)
markers = np.delete(markers, 8, 0)
markers = np.delete(markers, 7, 0)
markers = np.delete(markers, 6, 0)
markers = np.delete(markers, 5, 0)
markers = np.delete(markers, 4, 0)
markers = np.delete(markers, 3, 0)
markers = np.delete(markers, 2, 0)
markers = np.delete(markers, 1, 0)
print(markers)

if file == "SmallSoil4_X.csv":
    markers = np.delete(markers, 17, 0)
    markers = np.delete(markers, 16, 0)
    markers = np.delete(markers, 15, 0)
    markers = np.delete(markers, 14, 0)
    markers = np.delete(markers, 12, 0)
    markers = np.delete(markers, 11, 0)
    markers = np.delete(markers, 10, 0)
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 8, 0)
    markers = np.delete(markers, 7, 0)
    markers = np.delete(markers, 6, 0)
    markers = np.delete(markers, 4, 0)
print(markers)

if file == "LargeDrySand_X.csv":
    markers = np.delete(markers, 3, 0)
    markers = np.delete(markers, 1, 0)
print(markers)

if file == "LargeWetSoil_X.csv":
    markers = np.delete(markers, 27, 0)
    markers = np.delete(markers, 25, 0)
    markers = np.delete(markers, 24, 0)
    markers = np.delete(markers, 23, 0)
    markers = np.delete(markers, 22, 0)
    markers = np.delete(markers, 21, 0)
    markers = np.delete(markers, 20, 0)
    markers = np.delete(markers, 18, 0)
    markers = np.delete(markers, 17, 0)
```

```

    markers = np.delete(markers, 16, 0)
    markers = np.delete(markers, 14, 0)
    markers = np.delete(markers, 13, 0)
    markers = np.delete(markers, 12, 0)
    markers = np.delete(markers, 11, 0)
    markers = np.delete(markers, 10, 0)
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 8, 0)
    markers = np.delete(markers, 7, 0)
    markers = np.delete(markers, 5, 0)
    markers = np.delete(markers, 4, 0)
    markers = np.delete(markers, 3, 0)
    print(markers)

if file == "LargeSoil1_X.csv":
    markers = np.delete(markers, 14, 0)
    markers = np.delete(markers, 13, 0)
    markers = np.delete(markers, 12, 0)
    markers = np.delete(markers, 11, 0)
    markers = np.delete(markers, 10, 0)
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 8, 0)
    markers = np.delete(markers, 5, 0)
    markers = np.delete(markers, 0, 0)
    print(markers)

if file == "LargeSoil2_X.csv":
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 7, 0)
    markers = np.delete(markers, 4, 0)
    print(markers)

if file == "LargeSoil3_X.csv":
    markers = np.delete(markers, 3, 0)
    markers = np.delete(markers, 2, 0)
    markers = np.delete(markers, 0, 0)
    print(markers)

if file == "LargeSoil4_X.csv":
    markers = np.delete(markers, 24, 0)
    markers = np.delete(markers, 22, 0)
    markers = np.delete(markers, 21, 0)
    markers = np.delete(markers, 20, 0)
    markers = np.delete(markers, 19, 0)

```

A. Code Listings

```
markers = np.delete(markers, 18, 0)
markers = np.delete(markers, 17, 0)
markers = np.delete(markers, 16, 0)
markers = np.delete(markers, 15, 0)
markers = np.delete(markers, 12, 0)
markers = np.delete(markers, 11, 0)
markers = np.delete(markers, 10, 0)
markers = np.delete(markers, 9, 0)
markers = np.delete(markers, 7, 0)
markers = np.delete(markers, 5, 0)
markers = np.delete(markers, 4, 0)
markers = np.delete(markers, 3, 0)
print(markers)

if file == "LargeSoil5_X.csv":
    markers = np.delete(markers, 13, 0)
    markers = np.delete(markers, 12, 0)
    markers = np.delete(markers, 11, 0)
    markers = np.delete(markers, 10, 0)
    markers = np.delete(markers, 9, 0)
    markers = np.delete(markers, 6, 0)
    markers = np.delete(markers, 5, 0)
    markers = np.delete(markers, 4, 0)
    markers = np.delete(markers, 3, 0)
    print(markers)

if file == "LargeDrySoil_X.csv":
    markers = np.delete(markers, 4, 0)
    markers = np.delete(markers, 2, 0)
    markers = np.delete(markers, 1, 0)
    print(markers)

if file == "SmallDrySoil_X.csv":
    markers = np.delete(markers, 2, 0)
    markers = np.delete(markers, 1, 0)
    print(markers)

if file == "SmallSoil2_X.csv":
    markers = np.delete(markers, 6, 0)
    markers = np.delete(markers, 5, 0)
    markers = np.delete(markers, 4, 0)
    markers = np.delete(markers, 3, 0)
    markers = np.delete(markers, 2, 0)
    markers = np.delete(markers, 1, 0)
```

```

print(markers)

np.savetxt((path + file.split('.')[0] + "Markers.csv")
,
           markers, delimiter=',')

```



```

def openCircleData(filename):
    with open(filename) as csvfile:
        return np.asarray(list(csv.reader(csvfile))).astype(np.
single)

```



```

if __name__ == "__main__":
    # SmallSoil1 has a different size of marker and so requires
    # dedicated radii
    ,
    testData = Volume("SmallSoil1", ndPath +
                      "SmallSoil1_1024x1024x2048.raw",
                      [1024, 1024, 2048], "nd")
    testData.minRadius = 35
    testData.maxRadius = 60
    testData.circles = circleHunter.findCircles(testData)
    testData.saveCircleData()

    testData = Volume("SmallSoil1", nwPath +
                      "SmallSoil1_1024x1024x2048.raw",
                      [1024, 1024, 2048], "nw")
    testData.minRadius = 35
    testData.maxRadius = 60
    testData.circles = circleHunter.findCircles(testData)
    testData.saveCircleData()

    testData = Volume("SmallSoil1", xPath +
                      "SmallSoil1_2560x2560x9265.raw",
                      [2560, 2560, 9265], "x")
    testData.minRadius = 180
    testData.maxRadius = 350
    testData.circles = circleHunter.findCircles(testData)
    testData.saveCircleData()
    # testData.markerLocations = sphereBuilder.clusterPoints(
    #     testData.circles)
    # testData.display

```

A. Code Listings

```

, ,

# findCirclesInFolder(xPath, "x")
# findCirclesInFolder(ndPath, "nd")
# findCirclesInFolder(nwPath, "nw")
# findCirclesInFolder(postPath, "r")
# findCirclesInFolder(postFPath, "r")

, ,

data = openCircleData(ndCPath + "LargeSoil3.csv")
markers = sphereBuilder.clusterPoints(data)

fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')

datax = data[:, 0]
datay = data[:, 1]
dataz = data[:, 2]

ax.scatter(datax, datay, dataz, s=2)
ax.scatter(markers[:, 0], markers[:, 1], markers[:, 2], s=50,
           c='red')
plt.show()
, ,

# findMarkersInFolder(xCPath)
# findMarkersInFolder(ndCPath)
# findMarkersInFolder(nwCPath)
# findMarkersInFolder(postCPath)
findMarkersInFolder(postxCPath)
# findMarkersInFolder(postndCPath)
# findMarkersInFolder(postnwCPath)

```

A.3.2. Volume.py

```

import cv2
import csv
import numpy as np
from math import sqrt

class Volume:
    def __init__(self, name, file, shape, modality):
        self.name = name
        self.shape = shape

```

```

    self.file = file
    self.modality = modality

    self.axis = 2
    self.slice = 0

    if self.modality == "x":
        self.imageScale = 0.3
        self.markerRadius = 65
        self.minRadius = 40
        self.maxRadius = 70

    if (self.modality == "nd") or (self.modality == "nw"):
        self.imageScale = 0.7
        self.markerRadius = 10
        self.minRadius = 6
        self.maxRadius = 12

    if self.modality == "r":
        self.imageScale = 0.3
        self.markerRadius = 45
        self.minRadius = 20
        self.maxRadius = 40

    self.markerLocations = np.empty((0, 3))

def openMarkerRanges(self):
    with open("ranges.csv") as csvfile:
        allRanges = np.asarray(list(csv.reader(csvfile)))
        indices = np.where(allRanges == self.name)
        row = np.ndarray.flatten(allRanges[indices[0], 1:5]).astype(np.int)
    self.ranges = row

def openData(self):
    fd = open(self.file, "rb")
    size = np.prod(self.shape, dtype=np.int64)
    f = np.fromfile(fd, dtype=np.uint8, count=size)
    data = f.reshape(self.shape, order='F')
    self.data = np.swapaxes(data, 0, 1)

def saveCircleData(self):
    if self.modality == "x":
        path = "CircleData/Xray/"

```

A. Code Listings

```
if self.modality == "nd":
    path = "CircleData/NeutronDry/"
if self.modality == "nw":
    path = "CircleData/NeutronWet/"
if self.modality == "r":
    path = "CircleData/Registered/"
np.savetxt((path + self.name + ".csv"), self.circles,
           delimiter=',',)

def openCircleData(self, filename):
    with open(filename) as csvfile:
        self.circles = np.asarray(list(
            csv.reader(csvfile))).astype(np.single
        )

def display(self):
    cv2.namedWindow(self.name)
    cv2.createTrackbar('Slice', self.name, 0,
                      max(self.shape) - 1, self.scroll)
    cv2.createTrackbar('Axis', self.name, 0, 2, self.scroll)
    self.updateDisplay(self.getAxisSlice(self.axis, self.slice
        ))
    cv2.waitKey(0)
    cv2.destroyAllWindows()

def updateDisplay(self, image):
    cimage = cv2.cvtColor(image, cv2.COLOR_GRAY2BGR)

    for circle in self.circles:
        if (circle[-1] == self.axis) and (circle[-2] == self.
            slice):
            cv2.circle(cimage, (int(circle[0]), int(circle[1])
                ),
                       int(circle[2]), (0, 255, 0), 2)
            cv2.circle(cimage, (int(circle[0]), int(circle[1])
                ),
                       2, (0, 0, 255), 3)

    for marker in self.markerLocations:
        if self.axis == 0:
            cv2.circle(cimage, (int(marker[0]), int(marker[2])
                ),
                       5, (255, 0, 0), 5)
        if ((marker[1] - self.slice) < self.markerRadius <
```

```

(marker[1] + self.slice)):
    rad = sqrt((self.markerRadius**2)
                - ((marker[1]-self.slice)**2))
    cv2.circle(cimage, (int(marker[0]), int(marker
[2])), int(rad), (255, 0, 0), 5)
if self.axis == 1:
    cv2.circle(cimage, (int(marker[1]), int(marker[2]))
), 5, (255, 0, 0), 5)
if ((marker[0] - self.slice) < self.markerRadius <
    (marker[0] + self.slice)):
    rad = sqrt((self.markerRadius**2)
                - ((marker[0]-self.slice)**2))
    cv2.circle(cimage, (int(marker[1]), int(marker
[2])), int(rad), (255, 0, 0), 5)
if self.axis == 2:
    cv2.circle(cimage, (int(marker[0]), int(marker[1]))
), 5, (255, 0, 0), 5)
if ((marker[2] - self.slice) < self.markerRadius <
    (marker[2] + self.slice)):
    rad = sqrt((self.markerRadius**2)
                - ((marker[2]-self.slice)**2))
    cv2.circle(cimage, (int(marker[0]), int(marker
[1])), int(rad), (255, 0, 0), 5)

cimage = cv2.resize(cimage, (int(cimage.shape[1]*self.
    imageScale),
                           int(cimage.shape[0]*self.imageScale)),
    interpolation=cv2.INTER_AREA)
cv2.imshow(self.name, cimage)

def getAxisSlice(self, axis, slice):
    if (axis == 0):
        image = self.data[slice, :, :]
        image = np.transpose(image)
    elif (axis == 1):
        image = self.data[:, slice, :]
        image = np.transpose(image)
    else:
        image = self.data[:, :, slice]

```

A. Code Listings

```
    return image

    def scroll(self, dropResult):
        self.slice = cv2.getTrackbarPos('Slice', self.name)
        self.axis = cv2.getTrackbarPos('Axis', self.name)
        self.updateDisplay(self.getAxisSlice(self.axis, self.slice
            ))
```

A.3.3. SphereBuilder.py

```
import numpy as np
from sklearn.cluster import DBSCAN

class SphereBuilder:

    def clusterPoints(self, data, samples, eps=40):
        clusters = DBSCAN(eps=eps, min_samples=samples).fit(data
            [:, 0:3])
        labels = clusters.labels_
        nM = len(set(labels)) - (1 if -1 in labels else 0)
        print(nM)

        markers = np.empty((0, 3))
        for i in range(nM):
            mCluster = data[np.where(labels == i)]

            cForX = mCluster[np.where(mCluster[:, 4] != 1)]
            if cForX.size == 0:
                markerX = np.average(mCluster[:, 0])
            else:
                markerX = np.average(cForX[:, 0])

            cForY = mCluster[np.where(mCluster[:, 4] != 0)]
            if cForY.size == 0:
                markerY = np.average(mCluster[:, 1])
            else:
                markerY = np.average(cForY[:, 1])

            cForZ = mCluster[np.where(mCluster[:, 4] != 2)]
            if cForZ.size == 0:
                markerZ = np.average(mCluster[:, 2])
            else:
                markerZ = np.average(cForZ[:, 2])
```

```

    markers = np.vstack([markers, [markerX, markerY,
                                  markerZ]])

    markers = markers[markers[:, -1].argsort()]
    # markers = markers.astype(int)
    # markers = self.filterInnerPoints(markers)
    return markers

def clusterPointsOldMethod(self, data):
    dataToClusterA = np.squeeze(data[np.where(data[:, 4] == 0
                                              , 0:3:2)])
    dataToClusterB = np.squeeze(data[np.where(data[:, 4] == 1
                                              , 1:3)])
    dataToClusterC = np.squeeze(data[np.where(data[:, 4] == 2
                                              , 0:2)])

    dbA = DBSCAN(eps=12.5, min_samples=4).fit(dataToClusterA)
    labelsA = dbA.labels_
    nA = len(set(labelsA)) - (1 if -1 in labelsA else 0)
    dbB = DBSCAN(eps=12.5, min_samples=4).fit(dataToClusterB)
    labelsB = dbB.labels_
    nB = len(set(labelsB)) - (1 if -1 in labelsB else 0)
    dbC = DBSCAN(eps=12.5, min_samples=4).fit(dataToClusterC)
    labelsC = dbC.labels_
    nC = len(set(labelsC)) - (1 if -1 in labelsC else 0)

    markersA = np.empty((0, 3))
    markersB = np.empty((0, 3))
    markersC = np.empty((0, 3))

    for i in range(nA):
        markersA = np.vstack([markersA, np.append(np.average(
            dataToClusterA[np.where(labelsA == i)], axis=0),
            np.where(labelsA == i)[0].size)])
    for i in range(nB):
        markersB = np.vstack([markersB, np.append(np.average(
            dataToClusterB[np.where(labelsB == i)], axis=0),
            np.where(labelsB == i)[0].size)])
    for i in range(nC):
        markersC = np.vstack([markersC, np.append(np.average(
            dataToClusterC[np.where(labelsC == i)], axis=0),
            np.where(labelsC == i)[0].size)])

```

A. Code Listings

```

print("Plane Markers:")
np.set_printoptions(suppress=True)
# print(markersA)
# print(markersB)
# print(markersC)

markers = np.empty((0, 3))

for row in markersA:
    bIndex = self.closestMatch(row[1], markersB[:, 1])
    cIndex = self.closestMatch(row[0], markersC[:, 0])
    if np.abs(row[0] - markersC[cIndex, 0]) < 4:
        if np.abs(row[1] - markersB[bIndex, 1]) < 4:
            if np.abs(markersB[bIndex, 0] - markersC[
                cIndex, 1]) < 4:
                x = np.average([row[0], markersC[cIndex,
                    0]], weights=[row[2], markersC[
                        cIndex, 2]])
                y = np.average([markersB[bIndex, 0],
                    markersC[cIndex, 1]], weights=[markersB[bIndex,
                        2],
                        markersC[cIndex, 2]])
                z = np.average([row[1], markersB[bIndex,
                    1]], weights=[row[2], markersB[
                        bIndex, 2]])
                marker = [x, y, z]
                markers = np.vstack([markers, marker])

    markers = markers.astype(int)
    # markers = self.filterInnerPoints(markers)
    print(markers)
    return markers

def closestMatch(self, value, array):
    abs_diff = np.abs(array - value)
    return abs_diff.argmin()

def filterInnerPoints(self, markers, dist=397):
    averageCentre = np.mean(markers[:, [0, 1]], axis=0)
    distances = np.linalg.norm(markers[:, [0, 1]] -

```

```

    averageCentre, axis=1)
while np.amin(distances) < dist:
    # print(distances)
    fewerMarkers = np.delete(markers, np.argmin(distances),
        0)
    fAverageCentre = np.mean(fewerMarkers[:, [0, 1]], axis
        =0)
    distances = np.linalg.norm(fewerMarkers[:, [0, 1]]
        - fAverageCentre, axis=1)
    markers = fewerMarkers
    print("removed a false positive:")
print(markers)
return markers

```

A.3.4. CircleHunter.py

```

import numpy as np
import cv2
from multiprocessing import Process, Queue

class CircleHunter:

    def __init__(self):

        self.sliceQueue = Queue()
        self.circleQueue = Queue()
        self.resultQueue = Queue()
        self.nCircleHunters = 1

    def findCirclesInSlice(self, axis, slice, image, modality,
                           minCRadius, maxCRadius):
        if modality == "x":
            circles = cv2.HoughCircles(image, cv2.HOUGH_GRADIENT,
                1, 20,
                param1=50, param2=50,
                minRadius=minCRadius,
                maxRadius=maxCRadius)
        if (modality == "nd") or (modality == "nw"):
            circles = cv2.HoughCircles(image, cv2.HOUGH_GRADIENT,
                1, 20,
                param1=50, param2=20,
                minRadius=minCRadius,
                maxRadius=maxCRadius)

```

A. Code Listings

```
if modality == "r":
    circles = cv2.HoughCircles(image, cv2.HOUGH_GRADIENT,
        1, 20,
                    param1=80, param2=25,
                    minRadius=minCRadius,
                    maxRadius=maxCRadius)

if circles is None:
    pass
else:
    return circles

def addCirclesToList(self, axis, slice, circles):
    if circles is None:
        pass
    else:
        for sublist in circles:
            for circle in sublist:
                if not (circle == [0., 0., 0.]).all():
                    row = np.append(circle, np.array([slice,
                        axis]))
                    self.circles = np.append(self.circles, [
                        row], axis=0)

def findCircles(self, volume):
    self.circles = np.empty((0, 5))
    procs = []

    for i in range(self.nCircleHunters):
        p = Process(target=self.circleFinder_main,
                    args=(self.sliceQueue, self.circleQueue,
                           volume.modality, volume.minRadius,
                           volume.maxRadius)
                    )
        procs.append(p)
        p.start()
    p = Process(target=self.circleLister_main,
                args=(self.circleQueue, self.resultQueue))
    procs.append(p)
    p.start()

    volume.openMarkerRanges()
    volume.openData()
```

```

for a in range(3):
    for b in range(volume.shape[a]):
        # print(a, b, (time.time() - startTime))
        if (a != 2):
            slice = volume.getAxisSlice(a, b)
            self.sliceQueue.put([a, b, slice])
        elif b in range(volume.ranges[0], volume.ranges[1]):
            slice = volume.getAxisSlice(a, b)
            self.sliceQueue.put([a, b, slice])
        elif b in range(volume.ranges[2], volume.ranges[3]):
            slice = volume.getAxisSlice(a, b)
            self.sliceQueue.put([a, b, slice])

for i in range(self.nCircleHunters):
    self.sliceQueue.put('STOP')

self.circles = self.resultQueue.get()
self.circleQueue.put('STOP')
for p in procs:
    p.join()

return self.circles

def circleFinder_main(self, inQueue, outQueue, modality,
                      minRadius, maxRadius):
    for [axis, slice, image] in iter(inQueue.get, 'STOP'):
        print(inQueue.qsize())
        circles = self.findCirclesInSlice(axis, slice, image,
                                           modality,
                                           minRadius, maxRadius
                                           )
        if circles is None:
            pass
        else:
            for sublist in circles:
                for circle in sublist:
                    if not (circle == [0., 0., 0.]).all():
                        row = np.append(circle, np.array([
                            slice, axis]))
                        outQueue.put([row])
    outQueue.put('FINISHED')

```

A. Code Listings

```

def circleLister_main(self , inQueue , outQueue):
    print("circle listener started")
    circles = np.empty((0 , 5))
    i = 0
    for row in iter(inQueue.get , 'STOP'):
        if row == 'FINISHED':
            i = i + 1
            if i == self.nCircleHunters:
                inQueue.put('STOP')
        else:
            print(row)
            row = np.squeeze(np.asarray(row).astype(float).
                astype(int))
            row[[3 , 2]] = row[[2 , 3]]
            if row[4] == 0:
                row[[1 , 2]] = row[[2 , 1]]
            if row[4] == 1:
                row[[0 , 1, 2]] = row[[2 , 0, 1]]
            print(row)
            circles = np.append(circles , [row] , axis=0)
    outQueue.put(circles)
    print("circle listener finished")

```

A.4. Masked reconstruction simulation

```

% Simulation of iterative masked recon
% Setup scan params and geometry
detector_size = 2000;
detector_pixel_size = 1;
n_projections = 522;
angles = linspace2(0, deg2rad(360) , n_projections);

minx = -detector_size/2;
maxx = detector_size/2;
miny = -detector_size/2;
maxy = detector_size/2;

vol_geom = astra_create_vol_geom(detector_size , detector_size ,
    minx , maxx , miny , maxy);
proj_geom = astra_create_proj_geom('parallel' , detector_pixel_size
    , detector_size , angles);

```

A.4. Masked reconstruction simulation

```

% Import simuation slice and mask
sliceWet = imread( 'M:\Data3\CTData\tjc3g13\AstraRecon\Simulation\
    NW.png' );
sliceWet = im2double(sliceWet);
sliceDry = imread( 'M:\Data3\CTData\tjc3g13\AstraRecon\Simulation\
    ND.png' );
sliceDry = im2double(sliceDry);

% Downsample
sliceWetLowRes = imresize(sliceWet, 0.05);
sliceWetLowRes = imresize(sliceWetLowRes, [2000 2000]);
sliceDryLowRes = imresize(sliceDry, 0.05);
sliceDryLowRes = imresize(sliceDryLowRes, [2000 2000]);

% Add noise
scale = 1e10;
sliceWetNoisy = scale * imnoise(sliceWetLowRes/scale, 'poisson');
sliceDryNoisy = scale * imnoise(sliceDryLowRes/scale, 'poisson');

sliceDiff = sliceWetNoisy - sliceDryNoisy;

mask = imread( 'M:\Data3\CTData\tjc3g13\AstraRecon\Simulation\
    SampleSimMask.png' );
mask = imcomplement(im2double(mask));
mask_id = astra_mex_data2d('create', '-vol', vol_geom, mask);

% Create projection data from this
[projNeutronWet_id, projectionsNW] = astra_create_sino_cuda(
    sliceWetNoisy, proj_geom, vol_geom);
[projNeutronDry_id, projectionsD] = astra_create_sino_cuda(
    sliceDryNoisy, proj_geom, vol_geom);
[projNeutronDiff_id, projectionsW] = astra_create_sino_cuda(
    sliceDiff, proj_geom, vol_geom);

% Display a single projection image
% figure, imshow(squeeze(projections(:,:,1))), []

%%%% Reconstruct without mask %%%
% Reconstruct NeutronWet

% Create a data object for the reconstruction
recNW_id = astra_mex_data2d('create', '-vol', vol_geom);

% Set up the parameters for a reconstruction algorithm using the

```

A. Code Listings

```
GPU
cfg = astra_struct('SIRT_CUDA');
cfg.ReconstructionDataId = recNW_id;
cfg.ProjectionDataId = projNeutronWet_id;
% cfg.option.ReconstructionMaskId = mask_id;
cfg.option.MinConstraint = 0;
cfg.option.MaxConstraint = 255;

% Create the algorithm object from the configuration structure
algNW_id = astra_mex_algorithm('create', cfg);

% Run 150 iterations of the algorithm
astra_mex_algorithm('iterate', algNW_id, 150);

% Get the result
recNW = astra_mex_data2d('get', recNW_id);

% Reconstruct Neutron Dry
% Create a data object for the reconstruction
recND_id = astra_mex_data2d('create', '-vol', vol_geom);

% Set up the parameters for a reconstruction algorithm using the
GPU
cfg = astra_struct('SIRT_CUDA');
cfg.ReconstructionDataId = recND_id;
cfg.ProjectionDataId = projNeutronDry_id;
%cfg.option.ReconstructionMaskId = mask_id;
cfg.option.MinConstraint = 0;
cfg.option.MaxConstraint = 255;

% Create the algorithm object from the configuration structure
algND_id = astra_mex_algorithm('create', cfg);

% Run 150 iterations of the algorithm
astra_mex_algorithm('iterate', algND_id, 150);

% Get the result
recND = astra_mex_data2d('get', recND_id);

%%% Reconstruct with mask - water seperately %%%
% Create a data object for the reconstruction
recWaterMasked_id = astra_mex_data2d('create', '-vol', vol_geom);

% Set up the parameters for a reconstruction algorithm using the
```

A.4. Masked reconstruction simulation

```

GPU
cfgM = astra_struct('SIRT_CUDA');
cfgM.ReconstructionDataId = recWaterMasked_id;
cfgM.ProjectionDataId = projNeutronDiff_id;
cfgM.option.ReconstructionMaskId = mask_id;
cfgM.option.MinConstraint = 0;
cfgM.option.MaxConstraint = 255;

% Create the algorithm object from the configuration structure
algWM_id = astra_mex_algorithm('create', cfgM);

% Run 150 iterations of the algorithm
astra_mex_algorithm('iterate', algWM_id, 150);

% Get the result
recWM = astra_mex_data2d('get', recWaterMasked_id);

% Show results post rec mask and iterative mask
recDiff = recNW - recND;
recDiff(~mask) = 0;
figure, imshow(recDiff, []);
figure, imshow(recWM, []);

error = abs(sliceWet - recDiff);
error(~mask) = 0;

errorM = abs(sliceWet - recWM);
errorM(~mask) = 0;

figure, imshow(error, []);
figure, imshow(errorM, []);

norm(error)
norm(errorM)

% Clean up. Note that GPU memory is tied up in the algorithm
% object,
% and main RAM in the data objects.
astra_mex_algorithm('delete', algNW_id);
astra_mex_algorithm('delete', algND_id);
astra_mex_algorithm('delete', algWM_id);
astra_mex_data2d('delete', recNW_id);
astra_mex_data2d('delete', recND_id);
astra_mex_data2d('delete', recWaterMasked_id);

```

A. Code Listings

```
astra_mex_data2d('delete', projNeutronDiff_id);  
astra_mex_data2d('delete', projNeutronWet_id);  
astra_mex_data2d('delete', projNeutronDry_id);
```

B. Original figures

Some of the figures in Chapter 5 have had their contrast stretched so that they are easier to see in this thesis. The original versions of the images are included here.

B. Original figures

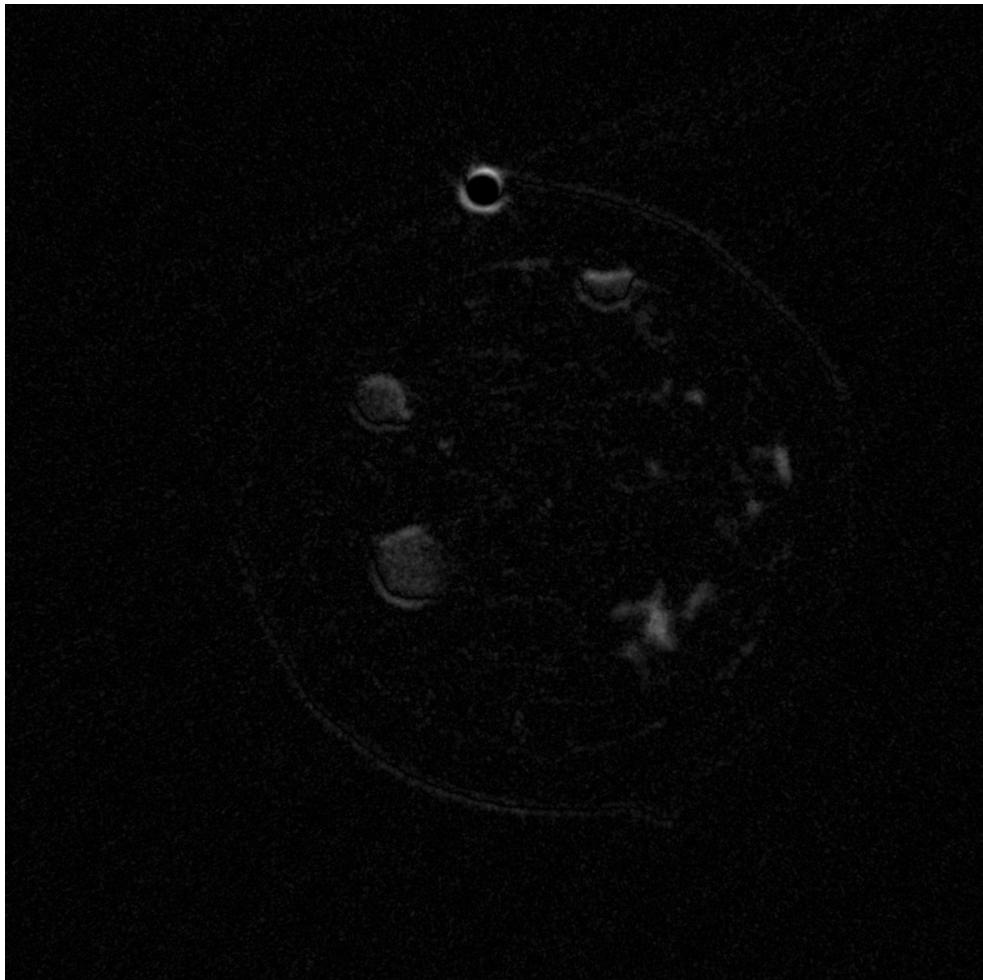
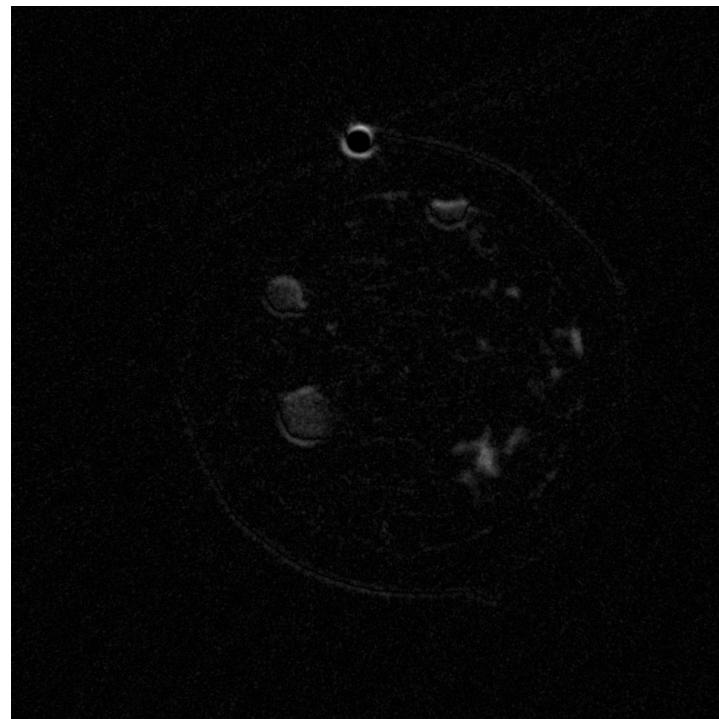


Figure B.1.: Taking the difference between the pre- and post-injection images for slice 8027 of LargeWetSand.

(a) Without the X-ray mask



(b) With the X-ray mask

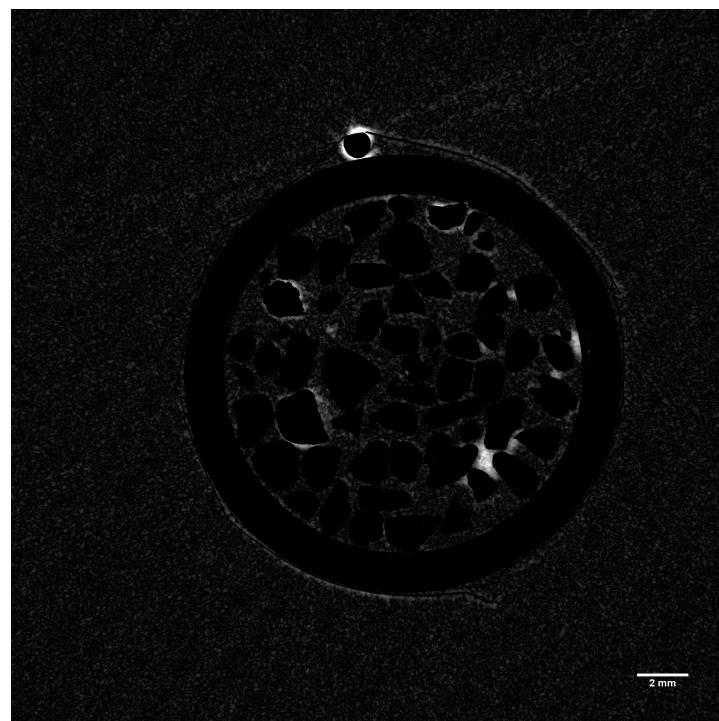


Figure B.2.: The difference (before and after the injection of deuterated water) image of slice 8027 shown before (a) and after (b) a mask is applied.

B. Original figures

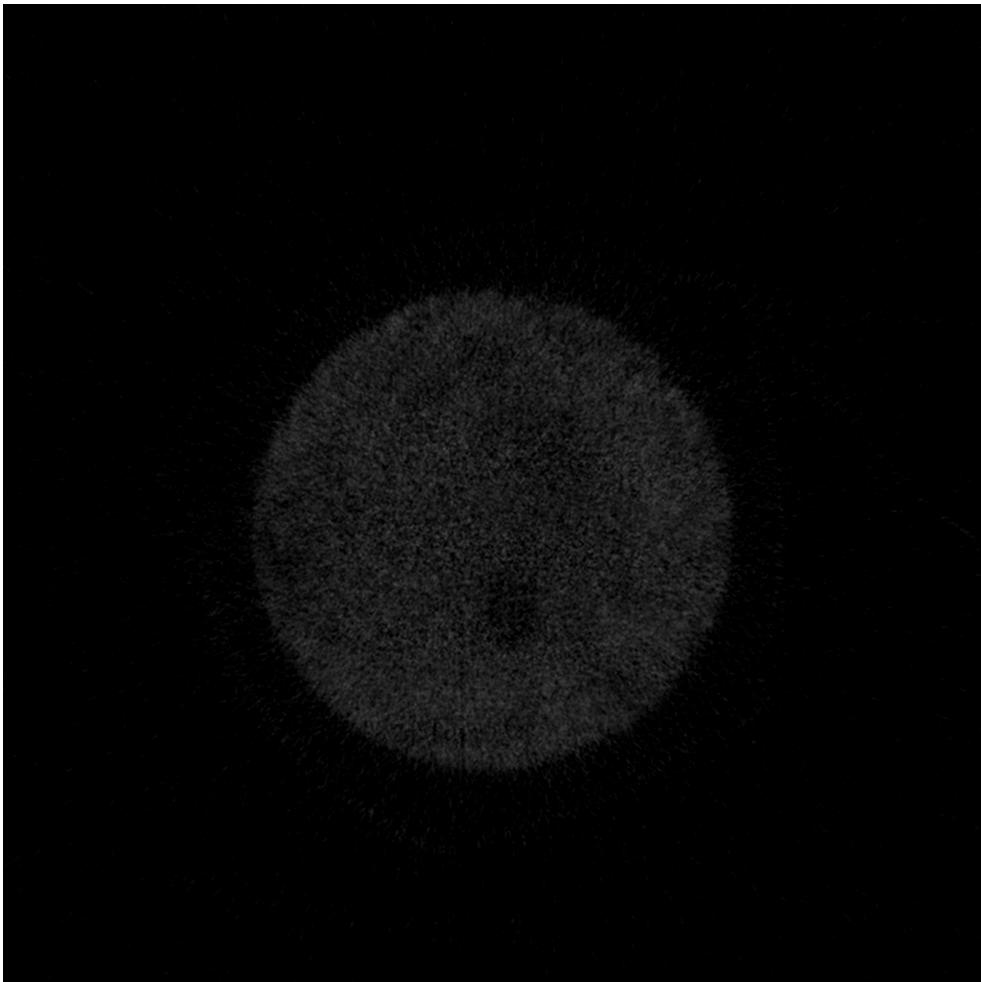


Figure B.3.: The difference between slice 1851 of the pre- and post-injection neutron scans of SmallSoil4.

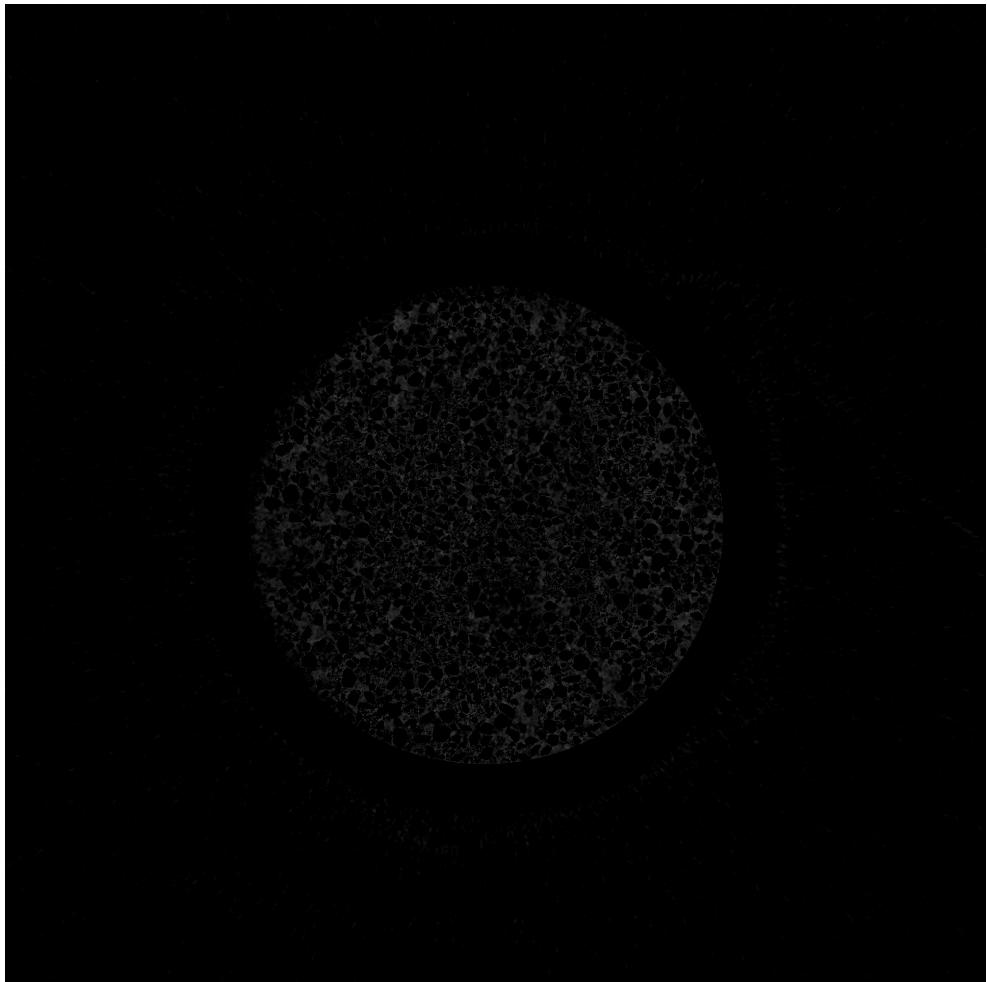


Figure B.4.: The difference between the pre- and post-injection scans of slice 1851 from SmallSoil4 with an X-ray mask overlaid.

B. Original figures

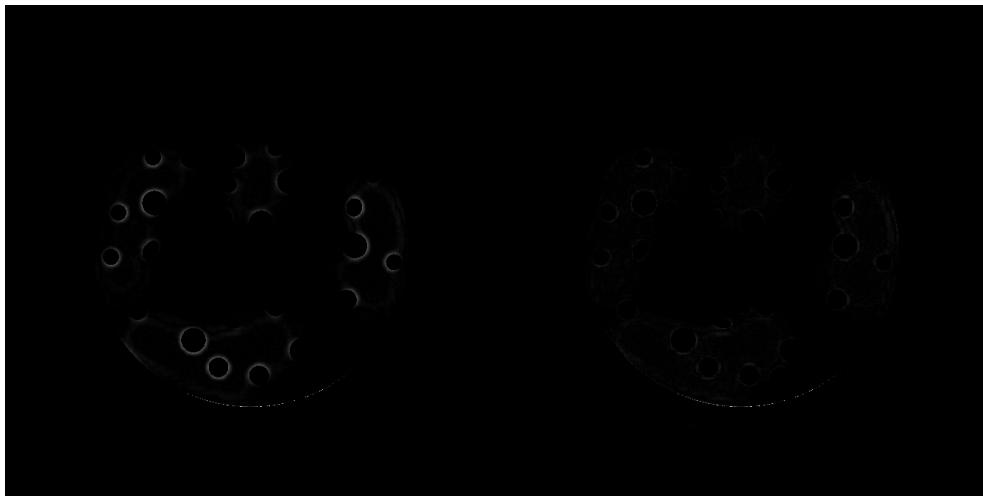


Figure B.5.: The error maps of the reconstructed slices after masking. The left image shows the benchmark result where the simulated pre- and post-injection data was reconstructed separately, and the mask applied to the difference between the reconstructions. The right image shows the error produced with the new method where the difference data was reconstructed with iterative masking.

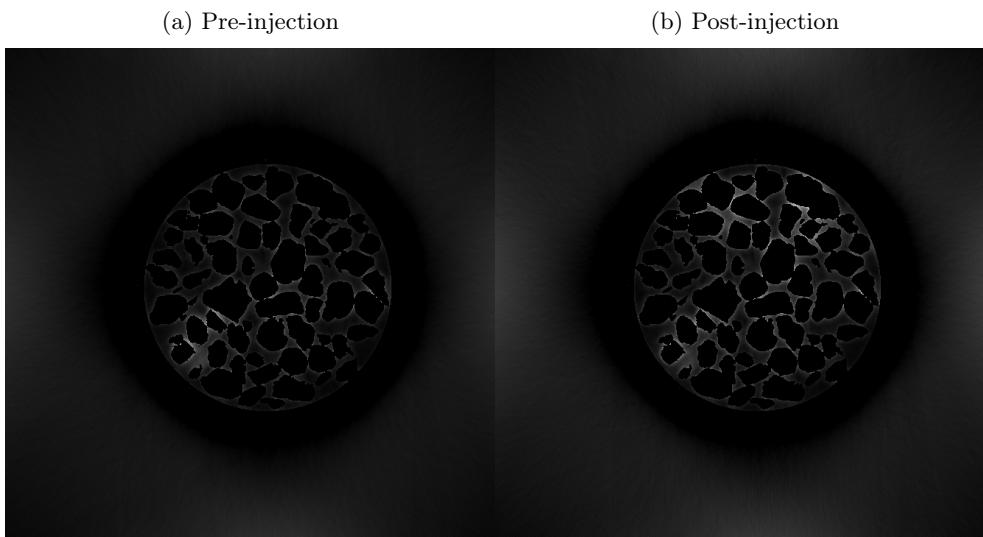


Figure B.6.: Slice 4402 from the sample LargeWetSand reconstructed used the SIRT algorithm with an X-ray mask.

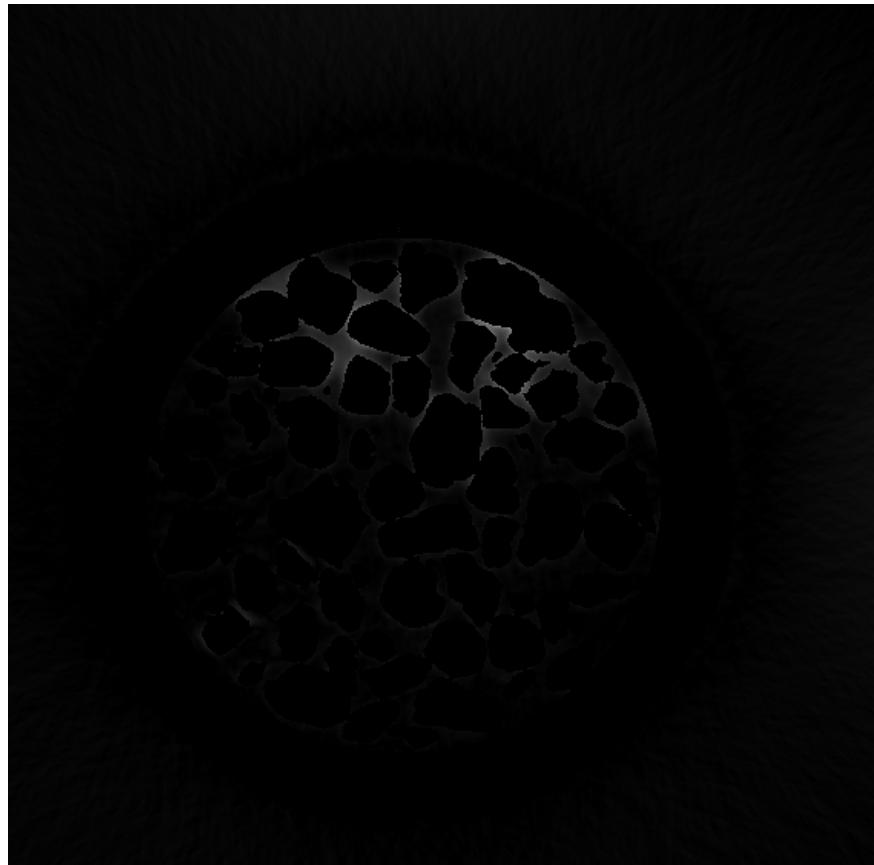


Figure B.7.: The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with an X-ray mask.

B. Original figures



Figure B.8.: The absolute difference between the pre- and post-injection images of slice 4402 when reconstructed using the SIRT algorithm with the X-ray mask applied afterwards.

Bibliography

Ian S. Anderson, Robert L. McGreevy, and Hassina Z. Bilheux. *Neutron imaging and applications : a reference for the imaging community.* Neutron scattering applications and techniques. Springer, New York, 2009. ISBN 9780387786926 0387786929 9780387786933 0387786937 1441946195 9781441946195. URL <http://www.loc.gov/catdir/enhancements/fy1315/2008936128-d.html>.

Dimitris J. Apostolopoulos, Trifon Spyridonidis, Theodoros Skouras, Costas Giannakenas, Christos Savvopoulos, and Pavlos J. Vassilakos. Comparison between 180° and 360° acquisition arcs with and without correction by ct-based attenuation maps in normal hearts at rest. *Nuclear medicine communications*, 29(2):110–119, 2008.

Jonathan A Atkinson, Michael P Pound, Malcolm J Bennett, and Darren M Wells. Uncovering the hidden half of plants using new advances in root phenotyping. *Current Opinion in Biotechnology*, 55:1–8, feb 2019. doi: 10.1016/j.copbio.2018.06.002.

Robert C. Atwood, Andrew J. Bodey, Stephen W. T. Price, Mark Basham, and Michael Drakopoulos. A high-throughput system for high-quality tomographic reconstruction of large datasets at Diamond Light Source. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2043):20140398, 2015. doi: doi:10.1098/rsta.2014.0398. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2014.0398>.

Julia F. Barrett and Nicholas Keat. Artifacts in ct: Recognition and avoidance. *RadioGraphics*, 24(6):1679–1691, 2004. doi: 10.1148/rg.246045065. URL <https://doi.org/10.1148/rg.246045065>. PMID: 15537976.

Anthony Bishopp and Jonathan P. Lynch. The hidden half of crop yields. *Nature Plants*, 1(8), aug 2015. doi: 10.1038/nplants.2015.117.

BIBLIOGRAPHY

F. Edward Boas and Dominik Fleischmann. Ct artifacts: causes and reduction techniques. *Imaging Med*, 4(2):229–240, 2012.

Mirko Boin and Astrid Haibel. Compensation of ring artefacts in synchrotron tomographic images. *Optics express*, 14(25):12071–12075, 2006.

G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.

G. Burca, W. Kockelmann, J. A. James, and M. E. Fitzpatrick. Modelling of an imaging beamline at the ISIS pulsed neutron source. *Journal of Instrumentation*, 8(10):P10001–P10001, 2013. ISSN 1748-0221. doi: 10.1088/1748-0221/8/10/p10001.

G. Burca, S. Nagella, T. Clark, D. Tasev, I.A. Rahman, R.J. Garwood, A.R.T. Spencer, M.J. Turner, and J.F. Kelleher. Exploring the potential of neutron imaging for life sciences on imat. *Journal of Microscopy*, 272(3):242–247, 2018. doi: 10.1111/jmi.12761. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/jmi.12761>.

Thorsten M. Buzug. *Computed tomography: from photon statistics to modern cone-beam CT*. Springer Science & Business Media, 2008. ISBN 3540394087.

William D. Carlson. Three-dimensional imaging of earth and planetary materials. *Earth and Planetary Science Letters*, 249(3):133–147, 2006. ISSN 0012-821X. doi: <https://doi.org/10.1016/j.epsl.2006.06.020>. URL <https://www.sciencedirect.com/science/article/pii/S0012821X0600447X>.

Richard R Carlton and Arlene M Adler. *Principles of Radiographic Imaging*. Cengage Learning, 2012.

Zhiqian Chang, Ruqiao Zhang, Jean-Baptiste Thibault, Ken Sauer, and Charles Bouman. Statistical x-ray computed tomography imaging from photon-starved measurements. In Charles A. Bouman and Ken D. Sauer, editors, *Computational Imaging XII*, volume 9020, pages 113 – 124. International Society for Optics and Photonics, SPIE, 2014. doi: 10.1117/12.2048204. URL <https://doi.org/10.1117/12.2048204>.

Harrell G. Chotas, James T. Dobbins, and Carl E. Ravin. Principles of digital radiography with large-area, electronically readable detectors: A

review of the basics. *Radiology*, 210(3):595–599, 1999. doi: 10.1148/radiology.210.3.r99mr15595. URL <https://doi.org/10.1148/radiology.210.3.r99mr15595>. PMID: 10207454.

T. Connolley, R. Atwood, and O. Magdysyuk. Detectors at I12. December 2020. URL <https://www.diamond.ac.uk/Instruments/Imaging-and-Microscopy/I12/Detectors-at-I12.html>.

P. Cosman. Image registration, 2012. URL <http://www.code.ucsd.edu/pcosman/regis2012.pdf>.

T. Defraeye, D. Derome, W. Aregawi, D. Cantre, S. Hartmann, E. Lehmann, J. Carmeliet, F. Voisard, P. Verboven, and B. Nicolai. Quantitative neutron imaging of water distribution, venation network and sap flow in leaves. *Planta*, 240(2):423–36, 2014. ISSN 1432-2048 (Electronic) 0032-0935 (Linking). doi: 10.1007/s00425-014-2093-3. URL <https://www.ncbi.nlm.nih.gov/pubmed/24923675>.

M. Dierick, J. Vlassenbroeck, B. Masschaele, V. Cnudde, L. Van Hoorebeke, and A. Hillenbach. High-speed neutron tomography of dynamic processes. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 542(1):296–301, 2005. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2005.01.152>. URL <https://www.sciencedirect.com/science/article/pii/S0168900205001944>. Proceedings of the Fifth International Topical Meeting on Neutron Radiography.

Manuel Dierick, Bert Masschaele, and Luc Van Hoorebeke. Octopus, a fast and user-friendly tomographic reconstruction package developed in labview®. *Measurement Science and Technology*, 15(7):1366, 2004. ISSN 0957-0233.

Michael Drakopoulos, Thomas Connolley, Christina Reinhard, Robert Atwood, Oxana Magdysyuk, Nghia Vo, Michael Hart, Leigh Connor, Bob Humphreys, George Howell, Steve Davies, Tim Hill, Guy Wilkin, Ulrik Pedersen, Andrew Foster, Nicoletta De Maio, Mark Basham, Fajin Yuan, and Kaz Wanelik. I12: the Joint Engineering, Environment and Processing (JEEP) beamline at Diamond Light Source. *Journal of Synchrotron Radiation*, 22(3):828–838, May 2015. doi: 10.1107/S1600577515003513. URL <https://doi.org/10.1107/S1600577515003513>.

BIBLIOGRAPHY

V. Finocchiaro, F. Aliotta, D. Tresoldi, R. C. Ponterio, C. S. Vasi, and G. Salvato. The autofocusing system of the IMAT neutron camera. *Review of Scientific Instruments*, 84(9):093701, 2013. ISSN 0034-6748. URL <https://aip.scitation.org/doi/pdf/10.1063/1.4819793>.

J. M. Fitzpatrick, J. B. West, and C. R. Maurer. Predicting error in rigid-body point-based registration. *IEEE Transactions on Medical Imaging*, 17(5):694–702, 1998. ISSN 0278-0062. doi: 10.1109/42.736021. URL <http://ieeexplore.ieee.org/document/736021/>.

Amin Garbout, Lars J Munkholm, Søren B Hansen, Bjørn M Petersen, Ole L Munk, and Radoslaw Pajor. The use of pet/ct scanning technique for 3d visualization and quantification of real-time soil/plant interactions. *Plant and soil*, 352(1):113–127, 2012.

Lucas L. Geyer, U. Joseph Schoepf, Felix G. Meinel, John W. Nance Jr, Gorka Bastarrika, Jonathon A. Leipsic, Narinder S. Paul, Marco Rengo, Andrea Laghi, and Carlo N. De Cecco. State of the art: iterative ct reconstruction techniques. *Radiology*, 276(2):339–357, 2015.

Peter Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of theoretical biology*, 36(1):105–117, 1972.

Richard Gordon, Robert Bender, and Gabor T. Herman. Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. *Journal of theoretical Biology*, 29(3):471–481, 1970.

Linda E. Graham, James M. Graham, and Lee W. Wilcox. *Plant biology*. Pearson/Prentice Hall, second edition, 2006. ISBN 0-13-146906-1.

Huaiqun Guan, Richard Gordon, and Yunping Zhu. Combining various projection access schemes with the algebraic reconstruction technique for low-contrast detection in computed tomography. *Physics in Medicine & Biology*, 43(8):2413, 1998.

Jon R Helliwell, Craig J Sturrock, Anthony J Miller, W Richard Whalley, and Sacha J Mooney. The role of plant species and soil condition in the structural development of the rhizosphere. *Plant, cell & environment*, 42(6):1974–1986, 2019.

BIBLIOGRAPHY

Griet Den Herder, Gert Van Isterdael, Tom Beeckman, and Ive De Smet. The roots of a new green revolution. *Trends in Plant Science*, 15(11):600–607, nov 2010. doi: 10.1016/j.tplants.2010.08.009.

Berthold KP Horn. Closed-form solution of absolute orientation using unit quaternions. *JOSA A*, 4(4):629–642, 1987. ISSN 1520-8532.

Donald House and John C. Keyser. *Foundations of Physically Based Modeling and Animation*, chapter Appendix C: Affine Transformations, pages 335–341. AK Peters/CRC Press, 2016.

Jiang Hsieh et al. Computed tomography: principles, design, artifacts, and recent advances. SPIE Bellingham, WA, 2009.

J. H. Hubbell and S. M. Seltzer. X-ray mass attenuation coefficients. NIST Standard Reference Database 126, July 2004. URL <https://www.nist.gov/pml/x-ray-mass-attenuation-coefficients>.

Susumu Ikeda and John M. Carpenter. Wide-energy-range, high-resolution measurements of neutron pulse shapes of polyethylene moderators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 239(3):536–544, 1985. ISSN 0168-9002. doi: [https://doi.org/10.1016/0168-9002\(85\)90033-6](https://doi.org/10.1016/0168-9002(85)90033-6).

A. P. Kaestner, J. Hovind, P. Boillat, C. Muehlebach, C. Carminati, M. Zarebanadkouki, and E. H. Lehmann. Bimodal imaging at ICON using neutrons and X-rays. *Physics Procedia*, 88:314–321, 2017. ISSN 1875-3892. doi: <https://doi.org/10.1016/j.phpro.2017.06.043>. URL <http://www.sciencedirect.com/science/article/pii/S1875389217300937>.

Anders P. Kaestner, Pavel Trtik, Mohsen Zarebanadkouki, Daniil Kazantsev, Michal Snehota, Katherine J. Dobson, and Eberhard H. Lehmann. Recent developments in neutron imaging with applications for porous media research. *Solid Earth*, 7(5):1281–1292, 2016. ISSN 1869-9529. doi: 10.5194/se-7-1281-2016.

Nikolay Kardjilov, André Hilger, Ingo Manke, Robin Woracek, and John Banhart. CONRAD-2: the new neutron imaging instrument at the Helmholtz-Zentrum Berlin. *Journal of Applied Crystallography*, 49(1):195–202, Feb

BIBLIOGRAPHY

2016. doi: 10.1107/S1600576715023353. URL <https://doi.org/10.1107/S1600576715023353>.

Masashi Kawakami, Kenji Hirata, Sho Furuya, Kentaro Kobayashi, Hiroyuki Sugimori, Keiichi Magota, and Chietsugu Katoh. Development of combination methods for detecting malignant uptakes based on physiological uptake detection using object detection with pet-ct mip images. *Frontiers in Medicine*, 7, 2020. ISSN 2296-858X. doi: 10.3389/fmed.2020.616746. URL <https://www.frontiersin.org/article/10.3389/fmed.2020.616746>.

S. D. Keyes, F. Gillard, N. Soper, M. N. Mavrogordato, I. Sinclair, and T. Roose. Mapping soil deformation around plant roots using in vivo 4D X-ray computed tomography and digital volume correlation. *Journal of Biomechanics*, 49(9):1802–1811, 2016. ISSN 0021-9290. doi: <https://doi.org/10.1016/j.jbiomech.2016.04.023>. URL <http://www.sciencedirect.com/science/article/pii/S002192901630495X>.

Samuel D. Keyes, Keith R. Daly, Neil J. Gostling, Davey L. Jones, Peter Talboys, Bernd R. Pinzer, Richard Boardman, Ian Sinclair, Alan Marchant, and Tiina Roose. High resolution synchrotron imaging of wheat root hairs growing in soil and image based modelling of phosphate uptake. *New Phytologist*, 198(4):1023–1029, 2013. ISSN 1469-8137. doi: 10.1111/nph.12294. URL <https://nph.onlinelibrary.wiley.com/doi/abs/10.1111/nph.12294>.

P. Kienzie. Extensible periodic table. December 2020. URL <https://periodictable.readthedocs.io/en/latest/index.html>.

Felix H Kim, Dayakar Penumadu, Jens Gregor, Nikolay Kardjilov, and Ingo Manke. High-resolution neutron and x-ray imaging of granular materials. *Journal of Geotechnical and Geoenvironmental Engineering*, 139(5):715–723, 2013.

FH Kim, D Penumadu, J Gregor, M Marsh, N Kardjilov, and I Manke. Characterizing partially saturated compacted-sand specimen using 3d image registration of high-resolution neutron and x-ray tomography. *Journal of Computing in Civil Engineering*, 29(6):04014096, 2015.

John H. Kinney and Monte C. Nichols. X-ray tomographic microscopy (XTM) using synchrotron radiation. *Annual review of materials science*, 22(1):121–152, 1992. ISSN 0084-6600.

BIBLIOGRAPHY

W. Kockelmann, S. Y. Zhang, J. F. Kelleher, J. B. Nightingale, G. Burca, and J. A. James. IMAT - A new imaging and diffraction instrument at ISIS. *Physics Procedia*, 43(Supplement C):100–110, 2013. ISSN 1875-3892. doi: <https://doi.org/10.1016/j.phpro.2013.03.013>. URL <http://www.sciencedirect.com/science/article/pii/S1875389213000278>.

Winfried Kockelmann, Genoveva Burca, Joe F. Kelleher, Saurabh Kabra, Shu-Yan Zhang, Nigel J. Rhodes, Erik M. Schooneveld, Jeff Sykora, Daniel E. Pooley, Jim B. Nightingale, Francesco Aliotta, Rosa C. Ponterio, Gabriele Salvato, Dario Tresoldi, Cirino Vasi, Jason B. McPhate, and Anton S. Tremisin. Status of the neutron imaging and diffraction instrument IMAT. *Physics Procedia*, 69:71–78, 2015. ISSN 18753892. doi: 10.1016/j.phpro.2015.07.010.

Nicolai Koebernick, Keith R. Daly, Samuel D. Keyes, Timothy S. George, Lawrie K. Brown, Annette Raffan, Laura J. Cooper, Muhammad Naveed, Anthony G. Bengough, Ian Sinclair, Paul D. Hallett, and Tiina Roose. High-resolution synchrotron imaging shows that root hairs influence rhizosphere soil structure formation. *New Phytologist*, 216(1):124–135, 2017. ISSN 1469-8137. doi: 10.1111/nph.14705. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5601222/pdf/NPH-216-124.pdf>.

Thomas Kohler. A projection access scheme for iterative reconstruction based on the golden section. In *Nuclear Science Symposium Conference Record, 2004 IEEE*, volume 6, pages 3961–3965. IEEE, 2004.

René C.P. Kuijken, Fred. A. van Eeuwijk, Leo F.M. Marcelis, and Harro J. Bouwmeester. Root phenotyping: from component trait in the lab to breeding: Table 1. *Journal of Experimental Botany*, 66(18):5389–5401, jun 2015. doi: 10.1093/jxb/erv239.

L. Kutschera-Mitter, K. M. Barmicheva, and M. Sobotik. *The Importance of Root-Cap Mucilage for Plant And Soil*, pages 673–683. Springer Netherlands, Dordrecht, 1998. ISBN 978-94-011-5270-9.

Yiannis Kyriakou and Willi Kalender. Efficiency of antiscatter grids for flat-detector CT. *Physics in Medicine and Biology*, 52(20):6275–6293, oct 2007. doi: 10.1088/0031-9155/52/20/013. URL <https://doi.org/10.1088/0031-9155/52/20/013>.

BIBLIOGRAPHY

Jacob M LaManna, Daniel S Hussey, E Baltic, and David L Jacobson. Neutron and x-ray tomography (next) system for simultaneous, dual modality tomography. *Review of Scientific Instruments*, 88(11):113702, 2017.

Eric N. Landis and Denis T. Keane. X-ray microtomography. *Materials Characterization*, 61(12):1305–1316, 2010. ISSN 1044-5803. doi: <https://doi.org/10.1016/j.matchar.2010.09.012>. URL <http://www.sciencedirect.com/science/article/pii/S1044580310002706>.

Andrew Lawler. Green light for long-awaited facility. *Science*, 279(5350):470–471, 1998. doi: 10.1126/science.279.5350.470.

Zhong-Sen Li and Lian-Sheng Tang. Using synchrotron-based x-ray microcomputed tomography to characterize water distribution in compacted soils. *Advances in Materials Science and Engineering*, 2019, 2019.

Jonathan P. Lynch. TURNER REVIEW no. 14. roots of the second green revolution. *Australian Journal of Botany*, 55(5):493, 2007. doi: 10.1071/bt06118.

S. Mairhofer, S. Zappala, S. R. Tracy, C. Sturrock, M. Bennett, S. J. Mooney, and T. Pridmore. RooTrak: Automated recovery of three-dimensional plant root architecture in soil from x-ray microcomputed tomography images using visual tracking. *PLANT PHYSIOLOGY*, 158(2):561–569, dec 2011. doi: 10.1104/pp.111.186221.

D Mannes, F Schmid, J Frey, K Schmidt-Ott, and E Lehmann. Combined neutron and x-ray imaging for non-invasive investigations of cultural heritage objects. *Physics Procedia*, 69:653–660, 2015.

U. Matsushima, W. B. Herppich, N. Kardjilov, W. Graf, A. Hilger, and I. Manke. Estimation of water flow velocity in small plants using cold neutron imaging with D2O tracer. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 605(1-2):146–149, 2009. ISSN 01689002. doi: 10.1016/j.nima.2009.01.187.

C. R. Maurer, J. M. Fitzpatrick, M. Y. Wang, R. L. Galloway, R. J. Maciunas, and G. S. Allen. Registration of head volume images using implantable fiducial markers. *IEEE Transactions on Medical*

BIBLIOGRAPHY

Imaging, 16(4):447–462, 1997. ISSN 0278-0062. doi: 10.1109/42.611354. URL <http://ieeexplore.ieee.org/document/611354/?reload=truehttp://ieeexplore.ieee.org/document/611354/>.

Tinashe Mawodza, Manoj Menon, Harriet Brooks, Oxana V. Magdysyuk, Genoveva Burca, and Stuart Casson. Preferential wheat (*triticum aestivum*. l cv. fielder) root growth in different sized aggregates. *Soil and Tillage Research*, 212:105054, 2021. ISSN 0167-1987. doi: <https://doi.org/10.1016/j.still.2021.105054>. URL <https://www.sciencedirect.com/science/article/pii/S0167198721001240>.

M. Menon, B. Robinson, S. E. Oswald, A. Kaestner, K. C. Abbaspour, E. Lehmann, and R. Schulin. Visualization of root growth in heterogeneously contaminated soil using neutron radiography. *European Journal of Soil Science*, 58(3):802–810, 2007. ISSN 1351-0754 1365-2389. doi: 10.1111/j.1365-2389.2006.00870.x.

Triestino Minniti, Kenichi Watanabe, Genoveva Burca, Daniel E. Pooley, and Winfried Kockelmann. Characterization of the new neutron imaging and materials science facility imat. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 888:184 – 195, 2018. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2018.01.037>. URL <http://www.sciencedirect.com/science/article/pii/S0168900218300548>.

Ahmad B. Moradi, Andrea Carminati, Doris Vetterlein, Peter Vontobel, Eberhard Lehmann, Ulrich Weller, Jan W. Hopmans, Hans-Jörg Vogel, and Sascha E. Oswald. Three-dimensional visualization and quantification of water content in the rhizosphere. *New Phytologist*, 192(3):653–663, 2011. ISSN 1469-8137. doi: 10.1111/j.1469-8137.2011.03826.x. URL <http://dx.doi.org/10.1111/j.1469-8137.2011.03826.x>.

Issei Mori, Yoshio Machida, Makoto Osanai, and Kazuhiro Iinuma. Photon starvation artifacts of x-ray ct: their true cause and a solution. *Radiological physics and technology*, 6(1):130–141, 2013.

μ -VIS. Modified 225 kVp Nikon/Xtek HMX, 2018. URL <https://www.southampton.ac.uk/muvvis/about/equipment/hmx.page>.

BIBLIOGRAPHY

Klaus Mueller, Roni Yagel, and J Fredrick Cornhill. The weighted-distance scheme: a globally optimizing projection ordering method for art. *IEEE Transactions on Medical Imaging*, 16(2):223–230, 1997.

J. Neuhaus and W. Petry. Correspondent’s report: Forschungsneutronenquelle Heinz Maier-Leibnitz (FRM II). *Neutron News*, 18(2):13–15, 2007. ISSN 1044-8632. doi: 10.1080/10448630701328356. URL <https://doi.org/10.1080/10448630701328356><https://www.tandfonline.com/doi/pdf/10.1080/10448630701328356?needAccess=true>.

Ruola Ning, Xiangyang Tang, and David Conover. X-ray scatter correction algorithm for cone beam ct imaging. *Medical physics*, 31(5):1195–1202, 2004.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Pas-sos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830, 2011.

Pawel A Penczek. Fundamentals of three-dimensional reconstruction from projections. *Methods in enzymology*, 482:1–33, 2010.

PerkinElmer. *XRD1621NESSeries*, 2008. URL https://www.perkinelmer.co.uk/PDFs/downloads/DTS_16inchDigitalXrayDetectors.pdf.

PhysicsOpenLab. Bremsstrahlung radiation. Online, August 2017. URL <http://physicsopenlab.org/2017/08/02/bremsstrahlung-radiation/>.

Ervin B Podgorsak et al. *Radiation oncology physics*. 2005.

D. E. Pooley, J. W. L. Lee, M. Brouard, J. J. John, W. Kockelmann, N. J. Rhodes, E. M. Schooneveld, I. Sedgwick, R. Turchetta, and C. Val-lance. Development of the "GP2" detector: Modification of the PImMS CMOS sensor for energy-resolved neutron radiography. *IEEE Transactions on Nuclear Science*, 64(12):2970–2981, Dec 2017. ISSN 0018-9499. doi: 10.1109/TNS.2017.2772040.

Sheikh MF Rabbi, Matthew K Tighe, Richard J Flavel, Brent N Kaiser, Chris N Guppy, Xiaoxian Zhang, and Iain M Young. Plant roots redesign the rhizosphere to alter the three-dimensional physical architecture and water dynamics. *New Phytologist*, 219(2):542–550, 2018.

BIBLIOGRAPHY

Deepak K. Ray, Nathaniel D. Mueller, Paul C. West, and Jonathan A. Foley. Yield trends are insufficient to double global crop production by 2050. *PLoS ONE*, 8(6):e66428, jun 2013. doi: 10.1371/journal.pone.0066428.

S. J. B. Reed. *Electron Microprobe Analysis and Scanning Electron Microscopy in Geology*. Cambridge University Press, 2 edition, 2005. doi: 10.1017/CBO9780511610561.012.

Johannes Schindelin, Ignacio Arganda-Carreras, Erwin Frise, Verena Kaynig, Mark Longair, Tobias Pietzsch, Stephan Preibisch, Curtis Rueden, Stephan Saalfeld, Benjamin Schmid, Jean-Yves Tinevez, Daniel James White, Volker Hartenstein, Kevin Eliceiri, Pavel Tomancak, and Albert Cardona. Fiji: an open-source platform for biological-image analysis. *Nat Meth*, 9(7):676–682, 2012. ISSN 1548-7091. doi: <http://www.nature.com/nmeth/journal/v9/n7/abs/nmeth.2019.html#supplementary-information>. URL <http://dx.doi.org/10.1038/nmeth.2019> <http://www.nature.com/nmeth/journal/v9/n7/pdf/nmeth.2019.pdf>.

Karsten Schörner. *Development of methods for scatter artifact correction in industrial x-ray cone-beam computed tomography*. PhD thesis, Technische Universität München, 2012.

J Anthony Seibert and John M Boone. X-ray imaging physics for nuclear medicine technologists. part 2: X-ray interactions and image formation. *Journal of nuclear medicine technology*, 33(1):3–18, 2005. ISSN 0091-4916. URL <http://tech.snmjournals.org/content/33/1/3.full.pdf>.

Lincoln Taiz and Eduardo Zeiger. *Plant physiology*. Sinauer Associates, Sunderland, Mass., 3rd edition, 2002. ISBN 0878938230. URL <http://www.loc.gov/catdir/toc/fy041/2002009242.html>.

Alessandro Tengattini, Nicolas Lenoir, Edward Andò, Benjamin Giroud, Duncan Atkins, Jerome Beaucour, and Gioacchino Viggiani. Next-grenoble, the neutron and x-ray tomograph in grenoble. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 968:163939, 2020.

D. Tilman, C. Balzer, J. Hill, and B. L. Befort. Global food demand and the sustainable intensification of agriculture. *Proceedings of the National*

BIBLIOGRAPHY

Academy of Sciences, 108(50):20260–20264, nov 2011. doi: 10.1073/pnas.1116437108.

C. Totzke, N. Kardjilov, I. Manke, and S. E. Oswald. Capturing 3d water flow in rooted soil by ultra-fast neutron tomography. *Sci Rep*, 7(1):6192, 2017. ISSN 2045-2322 (Electronic) 2045-2322 (Linking). doi: 10.1038/s41598-017-06046-w. URL <https://www.ncbi.nlm.nih.gov/pubmed/28733616>.

A. S. Tremsin, J. V. Vallerga, J. B. McPhate, O. H. W. Siegmund, and R. Raffanti. High resolution photon counting with MCP-timepix quad parallel readout operating at > 1 KHz frame rates. *IEEE Transactions on Nuclear Science*, 60(2):578–585, April 2013. ISSN 0018-9499. doi: 10.1109/TNS.2012.2223714.

Wim Van Aarle, Willem Jan Palenstijn, Jeroen Cant, Eline Janssens, Folkert Bleichrodt, Andrei Dabrovolski, Jan De Beenhouwer, K Joost Batenburg, and Jan Sijbers. Fast and flexible x-ray tomography using the astra toolbox. *Optics express*, 24(22):25129–25147, 2016.

Henk Van As. Intact plant MRI for the study of cell water relations, membrane permeability, cell-to-cell and long distance water transport. *Journal of Experimental Botany*, 58(4):743–756, 12 2006. ISSN 0022-0957. doi: 10.1093/jxb/erl157. URL <https://doi.org/10.1093/jxb/erl157>.

C. W. E. van Eijk, A. Bessière, and P. Dorenbos. Inorganic thermal-neutron scintillators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 529(1):260–267, 2004. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2004.04.163>. URL <http://www.sciencedirect.com/science/article/pii/S0168900204009064>.

Manning Wang and Zhijian Song. Improving target registration accuracy in image-guided neurosurgery by optimizing the distribution of fiducial points. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 5(1):26–31, 2009. ISSN 1478-596X. doi: 10.1002/rcs.227. URL <http://dx.doi.org/10.1002/rcs.227>.

Jeffrey M. Warren, Hassina Bilheux, Misun Kang, Sophie Voisin, Chu-Lin Cheng, Juske Horita, and Edmund Perfect. Neutron imaging reveals internal

BIBLIOGRAPHY

plant water dynamics. *Plant and Soil*, 366(1-2):683–693, 2013. ISSN 0032-079X 1573-5036. doi: 10.1007/s11104-012-1579-7.

Jay B West, J Michael Fitzpatrick, Steven A Toms, Calvin R Maurer Jr, and Robert J Maciunas. Fiducial point placement and the accuracy of point-based, rigid body registration. *Neurosurgery*, 48(4):810–817, 2001. URL <http://ovidsp.tx.ovid.com/ovftpdः/FPDDNCJCBDEDLB00/fs046/ovft/live/gv025/00006123/00006123-200104000-00023.pdf>.

Martin J. Willemink and Peter B. Noël. The evolution of image reconstruction for CT - from filtered back projection to artificial intelligence. *European radiology*, 29(5):2185–2195, 2019.

Anqi Yang, Feng Pan, Vishwanath Saragadam, Duy Dao, Zhuo Hui, Jen-Hao Rick Chang, and Aswin C Sankaranarayanan. Slicenets—a scalable approach for object detection in 3d ct scans. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 335–344, 2021.

Yangyi Yu, Ruiqin Zhang, Lu Lu, and Yigang Yang. The bimodal neutron and x-ray imaging driven by a single electron linear accelerator. *Applied Sciences*, 11(13):6050, 2021.

M. Zarebanadkouki, Y. X. Kim, and A. Carminati. Where do roots take up water? neutron radiography of water flow into the roots of transpiring plants growing in soil. *New Phytol*, 199(4):1034–44, 2013. ISSN 1469-8137 (Electronic) 0028-646X (Linking). doi: 10.1111/nph.12330. URL <https://www.ncbi.nlm.nih.gov/pubmed/23692148>.

Mohsen Zarebanadkouki, Eva Kroener, Anders Kaestner, and Andrea Carminati. Visualization of root water uptake: Quantification of deuterated water transport in roots using neutron radiography and numerical modeling. *Plant Physiology*, 166:487–499, 2014. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4213081/pdf/487.pdf>.

Mohsen Zarebanadkouki, Andrea Carminati, Anders Kaestner, David Mannes, Manuel Morgano, Steven Peetersmans, Eberhard Lehmann, and Pavel Trtik. On-the-fly neutron tomography of water transport into lupine roots. *Physics Procedia*, 69(Supplement C):292–298, 2015. ISSN 1875-3892. doi: <https://doi.org/10.1016/j.phpro.2015.07.037>

BIBLIOGRAPHY

//doi.org/10.1016/j.phpro.2015.07.041. URL <http://www.sciencedirect.com/science/article/pii/S1875389215006513>.

Wei Zhao, Don Vernekkohl, Jun Zhu, Luyao Wang, and Lei Xing. A model-based scatter artifacts correction for cone beam ct. *Medical physics*, 43(4): 1736–1753, 2016.

Yunsong Zhao and Mengfei Li. Iterative beam hardening correction for multi-material objects. *PLOS ONE*, 10(12):1–13, 12 2015. doi: 10.1371/journal.pone.0144607. URL <https://doi.org/10.1371/journal.pone.0144607>.

Ralf F Ziesche, Tobias Arlt, Donal P Finegan, Thomas MM Heenan, Alessandro Tengattini, Daniel Baum, Nikolay Kardjilov, Henning Markötter, Ingo Manke, Winfried Kockelmann, et al. 4d imaging of lithium-batteries using correlative neutron and x-ray tomography with a virtual unrolling technique. *Nature communications*, 11(1):1–11, 2020.