

Techniques in helical scanning, dynamic imaging and image segmentation for improved quantitative analysis with X-ray micro-CT

Adrian Sheppard¹, Shane Latham¹, Jill Middleton¹, Andrew Kingston¹, Glenn Myers¹, Trond Varslot^{1,2}, Andrew Fogden¹, Tim Sawkins¹, Ron Cruikshank¹, Mohammad Saadatfar¹, Nicolas Francois¹, Christoph Arns³ and Tim Senden¹

¹Department of Applied Mathematics, Research School of Physics and Engineering, Australian National University, Canberra 0200, Australia [email:Adrian.Sheppard@anu.edu.au]

²Current address: Lithicon Pty Ltd, 72 Northbourne Avenue, Canberra 2600, Australia

³School of Petroleum Engineering, University of New South Wales, Sydney 2052, Australia

Keywords: dynamic tomography, helical cone-beam tomography, micro-computed tomography, image registration, image segmentation

ABSTRACT

This paper reports on recent advances at the micro-computed tomography facility at the Australian National University. Since 2000 this facility has been a significant centre for developments in imaging hardware and associated software for image reconstruction, image analysis and image-based modelling. In 2010 a new instrument was constructed that utilises theoretically-exact image reconstruction based on helical scanning trajectories, allowing higher cone angles and thus better utilisation of the available X-ray flux. We discuss the technical hurdles that needed to be overcome to allow imaging with cone angles in excess of 60°. We also present dynamic tomography algorithms that enable the changes between one moment and the next to be reconstructed from a sparse set of projections, allowing higher speed imaging of time-varying samples. Researchers at the facility have also created a sizeable distributed-memory image analysis toolkit with capabilities ranging from tomographic image reconstruction to 3D shape characterisation. We show results from image registration and present some of the new imaging and experimental techniques that it enables. Finally, we discuss the crucial question of image segmentation and evaluate some recently proposed techniques for automated segmentation.

1. INTRODUCTION

The last decade has seen an explosion in the number and quality of micro-computed tomography (MCT) instruments installed at numerous locations worldwide, capitalising on the availability of high-quality components for the generation and detection of X-rays, as well as the rise of GPU computing that allows full-scale image reconstruction on personal computers. MCT beamlines now exist at all major synchrotrons, including TOMCAT at the Swiss Light Source (SLS), ID19, ID22 and ID15 at the European Synchrotron Radiation Facility (ESRF) and 2-BM (XOR) and 13-BM (GSECARS) at The Advanced Photon Source (APS) in the US, to name but a few. These facilities offer 3D imaging resolution at a range of scales, from 100nm to about 20µm, with some form of X-ray optics required for resolution better than 500nm. The most efficient beamlines can acquire datasets of 2000³ voxels at 2-5 micron resolution in less than one minute.

Laboratory facilities have seen similar advances, with commercial instruments from several vendors also offering resolutions approaching 100nm and datasets of 2000³ voxels obtainable in several hours at resolutions from 2-50µm. The MCT facility at the Australian National University (ANU) has been under constant development since first becoming operational in 2000 and has been used for research into many areas, most notably geology (Arns et al 2005), petrophysics (Arns et al 2004), tissue engineering (Jones et

al 2007), granular materials (Aste et al 2006) and paleontology (Long et al 2006). Three instruments currently in operation: the first instrument using a 200kV reflection-style X-ray source still functions although the original CCD detector has been replaced by a large area amorphous-silicon (a-Si) flat panel detector offering higher quality images, better radiation tolerance and a faster readout rate, while two new instruments have been constructed with reflection-type sources and a-Si panels. The older system has a resolution of about $3\mu\text{m}$ while the newer systems are about $1.5\mu\text{m}$. Since its inception the facility has aimed to derive quantitative results from the imaging program, an emphasis that has driven continual development in both hardware and software. In this paper we describe some recent advancements at the facility.

2. HELICAL-SCANNING CONE-BEAM TOMOGRAPHY

One major constraint in laboratory MCT is beam flux. Laboratory sources generate X-rays from a micron-sized region in a target material, usually tungsten, onto which an electron beam is focussed. If the electron beam intensity is too high ($>2\text{W}/\mu\text{m}^2$) then the target material is vaporized. The maximum allowed X-ray flux is therefore proportional to the area of this region. Since the instrument resolution is proportional to this region's diameter, improving resolution by a factor of 2 causes a factor of 4 reduction in the maximum tolerated beam power. Other types of X-ray source have been proposed to overcome this limitation, but as of 2013 no alternative is sufficiently stable or suitably housed for use in our facility.

Because the X-rays are produced uniformly in all directions from the target, one can maximise the flux passing through the sample and captured at the detector by placing the sample closer to the source. This results in the detector capturing a larger solid angle of X-rays. Unfortunately, circular-scanning cone-beam tomography depends on approximations that are only valid at low cone angles, with the artifacts resulting from this approximation becoming unacceptable above a cone angle of $\sim 5^\circ$. We have implemented theoretically-exact helical-scanning MCT (Varslot et al 2011), based on the Katsevich inversion formula (Katsevich 2002) and believe that the ANU facility is the only high cone-angle MCT instrument. A number of technical obstacles have needed to be overcome while developing high cone-angle helical scanning, we describe these in the following text.

The first issue is that of thermal drift causing relative movement between source and sample. While all effort is made to avoid such movement, our experience has shown that some thermal drift is sufficiently common to be worth correcting. At lower cone-angles one only need be concerned with transverse movements, whereas in high cone-angles slight changes in the source-sample distance cause an appreciable change in magnification. Thermal drift effects are removed by extending the reference-scan drift correction method introduced by Salmon et al (2009) to accommodate magnification changes (Myers et al 2011c).

The second issue is overall system alignment. This is more onerous for helical scanning for which there are two additional geometric alignment parameters as compared to circular scanning, making a total of 7. In particular, the source-sample distance must be known to several microns, in contrast with circular scanning where source-sample distance only affects voxel size. This parameter is very difficult to determine to sufficient accuracy experimentally since the source spot is several hundred microns behind the front face of the X-ray tube. To correct for any potential geometric misalignments, we use an iterative passive autofocus method (Kingston et al 2011, Varslot et al 2012) that finds the 7 helical geometry parameters as those which generate the sharpest tomogram. This technique is particularly effective since it finds the optimum value for the geometric parameters for the overall data acquisition, which may be slightly but significantly different to those measured prior to the imaging.

Another major issue is that of inhomogeneous magnification. At high cone angles, the side of the sample close to the X-ray source sees a much higher geometric magnification than the far side; a factor of 3 variation in magnification is typical. In 1-PI helical scanning, rays encompassing 180° are collected for each voxel, so that each part of the sample is in the radiographic field of view for about one half-turn of the helix (it is not exactly one half-turn due to the fan-angle). Some parts of the sample will be predominantly on the far side during this half turn and consequently see a relatively low magnification, while other parts see a relatively high magnification, resulting in a tomogram where some regions are sharper than others. This is a fundamental problem that has no straightforward solution; we have chosen to use a "double helix"

scanning trajectory (Varslot et al 2012), which largely removes the inhomogeneity in resolution, at a cost of a doubling in the volume of acquired data and consequently a doubling of the reconstruction times.

A final problem that is more serious in high cone-angle imaging is that of secondary radiation, emanating from surfaces inside transmission-type X-ray sources. This problem was discussed extensively in Boone et al (2012); for high cone-angle imaging the secondary radiation can only be reduced to acceptable levels through the placement of a pinhole collimator on the front of the X-ray tube.

Having resolved these issues the facility produces images of significantly higher quality that are captured in significantly shorter times. Nevertheless, for the bulk of the images acquired at the facility the highest practical image quality is desired, so that scan times of 10-20 hours are typical. Figure 1 shows an example of the data acquired using a helical scanning trajectory.




Figure 1: (a) Slice through an Image of a rock core 36mm long and 8mm in diameter, with $5.8\mu\text{m}$ voxel size. The two white squares on the image show the regions zoomed in (b) and (c), which show a field of view of $1.75 \times 1.75\text{mm}$. Acquisition of a dataset of the same field of view and quality by stacking circular tomograms would have taken nearly 10 times longer. From Varslot et al (2011), used with permission.

3. DYNAMIC TOMOGRAPHY

In medical imaging, dynamic tomography is used to remove the unavoidable and undesirable artefacts caused by patient movement. In materials science the problem is reversed: it is precisely the changes that one wishes to capture. In addition, it is usually possible to capture complete static images of the sample before and after the dynamical process has taken place. We have developed techniques for imaging just these changes, for experiment types where the changes from one frame to the next are highly constrained by the physics of the system (Myers et al 2011a).

The technique is particularly well suited to studying the dynamics of immiscible two-phase fluid displacements in porous media. Our approach exploits two physical constraints of this system:

- (a) The solid material is non-porous, so changes can only occur inside the pore space
- (b) The fluids are incompressible, so the change in attenuation when one fluid replaces the other is a step change. Therefore, ignoring partial volume effects, an image of the change-set is a binary image.

This prior information can be used in place of missing data to reconstruct the dynamic fluid configurations from what would otherwise be incomplete measurements. This is very similar to *binary tomography*, where objects are reconstructed from a very small number of projections, and *compressed sensing*, where signals are reconstructed from noisy and incomplete data by making assumptions about the characteristics of the noise. While the above constraints are specific to the two-phase fluid system, to which this technique is applied in Myers et al (2011b), other constraints may be formulated for reactive flow processes and material deformation under applied stress.

Our basic approach to using these constraints is explained in Myers et al (2011a) and consists of formulating the reconstruction as an optimisation problem where the parameters to be optimised are (a) the discrepancy between the experimental data and a re-projection of the reconstructed image, and (b) the spatial localisation of the solution. This optimisation problem is solved using the simultaneous iterative reconstruction technique (SIRT). GPU computing is used to cope with the computational demands of this method.

This algorithm results in order-of-magnitude reductions in the amount of data needed for each tomographic image and enables reconstruction of 512^3 image data from as little as 72 low-quality projections over 360° . Ground-truth experiments have been conducted which demonstrate the effectiveness of the approach. Future work will investigate the difference between sparser projections and more frequent but noisier data. The latter has theoretical advances but can make reconstruction prohibitively computationally expensive. Finally, we note that the method is equally applicable to data acquired at synchrotrons where it has the potential to reduce radiation damage and data storage while also reducing acquisition times




Figure 2: Snapshots of air displacing a NaI-doped water from the pore space of a Bentheimer sandstone sample. Images are 210×210 voxels, taken from 2000^3 datasets. Dynamic tomography has enabled images of this type of fluid displacement to be acquired an order of magnitude faster than with static imaging.

4. IMAGE REGISTRATION AND COMBINING EXPERIMENT WITH IMAGING

One major challenge faced in tomography is dealing with very large data; for example it is not uncommon at the ANU facility to generate images of $2000 \times 2000 \times 10000$ voxels, resulting in datasets approaching 100GB. Timely analysis of these datasets requires that the computation be distributed across cluster-style computers. All the algorithm development in the group has therefore used the Message Passing Interface (MPI), the standard library for distributed memory programming, with almost all algorithms demonstrated to scale well to several hundred processes. While the software that's been developed contains a great many components for the processing of grey-scale and binarised images, only image registration and segmentation are discussed in the current work.





Figure 3: (a) Small region of a slice from a micro-CT image of a complex, clay-rich sandstone, with 3µm voxel size. The clay-filled regions are poorly resolved. (b) The corresponding area imaged later using SEM at 1 µm pixel size, with images down to 5nm resolution also acquired. To prepare the sample for SEM imaging it was resin impregnated and polished down to expose a planar face from within the field of view of the micro-CT image. The slice shown in (a) is a slightly oblique cut through the tomographic volume, identified by 2D-3D registration as matching the face imaged by the SEM. This is an invaluable technique for adding high resolution information to a micro-CT image. Note the pore in the top right corner where material was lost during preparation of the polished surface for SEM imaging.

Image registration, the aligning of one image with another, is a key component for extending the capabilities of an MCT facility. Both 2D-3D and 3D-3D registration has been implemented, using similarity-based optimisation techniques (Latham et al 2008). An example of 2D-3D registration is shown in figure 3. Image registration greatly extends the capability of an MCT facility by allowing perfect alignment of samples before and after experimental procedures, allowing the changes to be quantified, generally by image subtraction. One can take arbitrarily many images separated by arbitrary lengths of time, so the experimental work can involve any number of steps and can be carried out in any location over any length of time. At the ANU facility the experimental work frequently involves multiple immiscible fluid phases such as the crude oil and water phases shown in Figure 4 (see e.g. Kumar et al 2009 and Singh et al 2011). Another very common approach is the use of contrast agents to identify clays, microporosity and organic matter in geomaterials. Still other, ongoing work involves deformation under mechanical loading and reactive flow, particularly dissolution and precipitation associated with CO₂. The described algorithm is for rigid registration and can only be applied to cases of mechanical deformation when a significant volume of the image remains unaltered, as in Derluyn et al (2013).



(a) (b) (c)

Figure 4: (a) 2D slice from tomogram of 5mm Bentheimer sandstone sample. (b) Corresponding slice from registered image after filling the sample with 1.2M NaI brine then displacing the brine with crude oil in a centrifuge. Phase contrast artefacts are prominent. (c) Subtraction of dry from wet image showing fluid locations, with phase contrast effects largely cancelled out since the images in (a) and (b) contain similar levels of phase contrast despite the difference in attenuation contrast. Tomograms are 2000×2000×1600 voxels with a voxel size of 2.8µm. Slices are 225×225 pixels for FOV of 0.5mm.

5. IMAGE SEGMENTATION

Image segmentation - classifying image into discrete "phases" - is a critical pre-requisite for image-based quantitative analysis and modelling. Many of the methods still in use today are labour intensive and suffer from operator subjectivity. Accurate quantification of the efficacy of segmentation methods is difficult since the "correct" segmentation is only known for synthetic datasets, which cannot capture the exact noise or blurring characteristics of real data. A great many segmentation methods have been proposed over the years, most with computer vision or medical imaging in mind, and often untested on micro-CT data. Development and validation of automatic segmentation methods therefore remains an important goal.

In recent years attention has focused on the probabilistic segmentation methods introduced by Geman and Geman (1984) with Markov Random Field (MRF) methods prevalent. These methods use Bayes' theorem to combine a Gaussian noise model for the voxel intensities in each material phase with an MRF term that incorporates a "cost" for inter-phase boundaries. An N -phase segmentation problem is therefore cast as a global optimisation problem with N degrees of freedom for each voxel. Solving the full optimisation problem is not possible, so approximate methods are required. The Iterated Conditional Modes (ICM) algorithm (Besag 1986), was found to be relatively effective for 2D images by Berthod et al (1996) and more recently by Kulkarni et al (2012) for 3D synthetic and micro-CT images.

In isolation, MRF methods require initialization (seeding) with the mean and variance of the voxel intensities for each material phase. In Berthod et al (1996) and Kulkarni (2012) this is done manually, selecting a single homogeneous region for each phase. Numerous techniques have been proposed for automatic initialization (for one review, see Jain 2010); in this work we use a finite Gaussian mixture model (GMM) (Jones and McLachlan 1990), which can be considered an extension to k-means clustering (Jain 2010). The GMM is fitted to the histogram of intensities using the Expectation Maximization algorithm of Jones and McLachlan (1990). However, neither GMM nor MRF methods allow for partial volume effects or blurring, so that voxels containing a mixture of material phases are problematic. Combining GMM with the watershed transform applied to a gradient image remedies this deficiency, since it is reasonable to assume that the local gradient minimum in each watershed region is relatively free of the partial voxel effect.

Therefore we consider two primary segmentation methods. The Thresholding approach (T) classifies each voxel by its intensity alone, through maximum likelihood thresholding from the fitted GMM. The Watershed Region Classification (WRC) method classifies an entire watershed region as the same class. The region class is assigned by maximum likelihood using the 27 voxel neighbourhood of the gradient minimum (seed-voxel) in each watershed region. For each Gaussian distribution we calculate the probability that the 27 intensities are drawn from that distribution and assign the class as the distribution giving the maximum result. The only user-selectable parameter for WRC and T is the number of phases for the GMM process. Both methods are applied standalone or used as the initial condition for an MRF method implemented using ICM.

We use the simplest and most common form of regularization in the MRF which simply penalizes the surface area of the segmented image to a degree controlled by a scale factor β , see Eqn. 1. The phase labeling \hat{L} which minimizes the negative log likelihood with a smoothing bias is given by:

$$\hat{L} = \underset{L}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \left[\ln \sqrt{2\pi\sigma_L^2} + \frac{(x_i - \mu_L)^2}{2\sigma_L^2} + \sum_{\{S_i, S_j\} \in C} \beta \gamma(l_{S_i}, l_{S_j}) \right] \right\}, \quad (1)$$

where

$$\gamma(l_{S_i}, l_{S_j}) = \begin{cases} -1 & \text{if } l_{S_i} = l_{S_j} \\ +1 & \text{if } l_{S_i} \neq l_{S_j} \end{cases}$$

and where C is a local neighbourhood, l_{S_i} is the phase label given to the voxel site S_i , N is the number of voxels and μ_L and σ_L^2 are the mean and variance of each phase. The ICM calculates the phase which minimizes the negative log likelihood for each individual voxel given its neighbourhood and then compares the global result of Eqn. 1 with the previous global result until a convergence criterion is reached.

Three 3D synthetic datasets were created to simulate some of the properties of porous material μ CT images. Sample 1 (figure 6) has a high surface area to volume ratio of 0.42 and has a 2 voxel “gap” between each cubic region. It has moderate noise and a small amount of Gaussian blur. Sample 2 (Figure 7) and an additional sample 3 (not illustrated) have less noise and moderate Gaussian blurring. The use of 7 phases in these two samples along with cubic region size allow control of the surface area, with the smaller cubic regions of sample 2 having a normalised surface area of 0.32 compared to sample 3’s 0.16. They do not contain any “gap” between cubic regions and contain the full range of phase-boundaries. To complement these synthetic datasets, a micro-CT image of a partially saturated unconsolidated sand (sample 4, shown in Figure 9) has also been tested with the sample containing three physical phases: oil, contrast enhanced brine and sand grains.

Accuracy was measured in all synthetic data cases by comparing each voxel value of the control segmentation with the equivalent voxel of the segmentation under question. In the case of the real data a manual segmentation and the tomogram has been compared visually with the resulting images. When ICM was used, it was applied multiple times with a wide range (0.01 to 50) of β values to produce multiple segmentations. For the synthetic datasets, the best β was determined to be the one which gave the segmentation with the lowest error rate. For the sandstone tomographic image, the best β was determined by visual inspection. Figure 5 shows the results of the application of the GMM method to the histograms of datasets 1, 2 and 4.

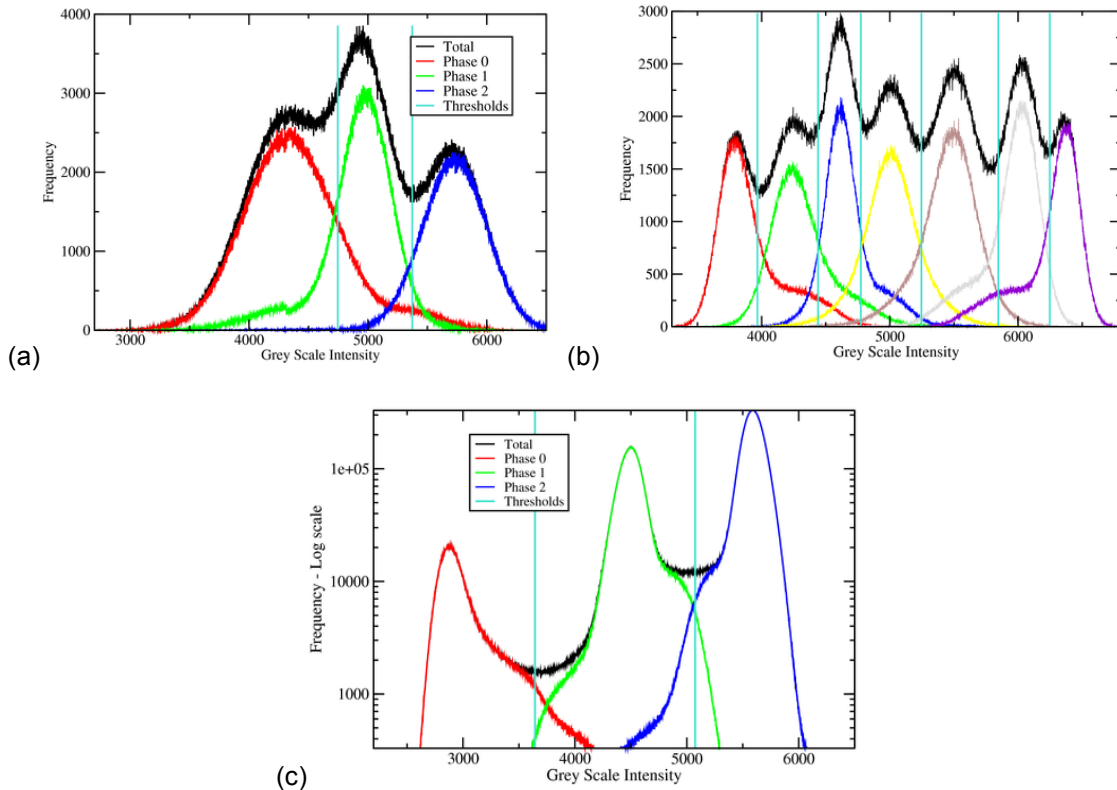


Figure 5: (a,b,c) Histograms of voxel intensities (in black) for the samples 1, 2 and 4 respectively. The red, green, blue, etc. lines show the distribution of intensities in each phase after the Gaussian mixture model (GMM) has been used for watershed region classification (WRC). The thresholds determined from the GMM are shown as vertical lines.






Figure 6: Performance of segmentation methods applied to sample 1 and shown on a 250x300 pixel subset from the 3D data set. a) grey scale image; b) control segmented image from which the gray scale image was derived; c) WRC (90% voxels correct); d) WRC with ICM post-processing with optimal β value of 0.1 (96% correct); e) Threshold (86% correct); f) T plus ICM post-processing with optimal $\beta=0.1$ (96% correct). In this case ICM made a notable improvement to both images, but only when an optimal or near-optimal value was chosen for β .






Figure 7: Performance of segmentation methods on sample 2, shown on a 350x450 pixel subset from the 3D image. a) Synthetic grey scale image; b) control segmented image; c) WRC (95% voxels correct); d) WRC plus ICM with optimal $\beta=5$ (94% correct); e) T (70% correct); f) T plus ICM with $\beta=0.1$ (90% correct).

The results for a small 2D slice from sample 1 are shown in figure 6. ICM improved both methods considerably, but was very sensitive to the β value as seen in figure 8. Sample 2 (Figure 7), was well segmented by WRC with ICM making no improvement to the WRC segmentation, but improving T considerably. An additional data set (not shown), which contains less surface area, yielded a WRC accuracy of 98% while the T was 81% accurate. Both were improved with ICM at $\beta=5$, reaching 99% and 98% respectively. The μ CT data of sample 4 (Figure 9) is relatively sharp with very little noise. For this dataset WRC again performed well, and better by visual inspection than an independently performed manual segmentation. Neither method was significantly improved by ICM.

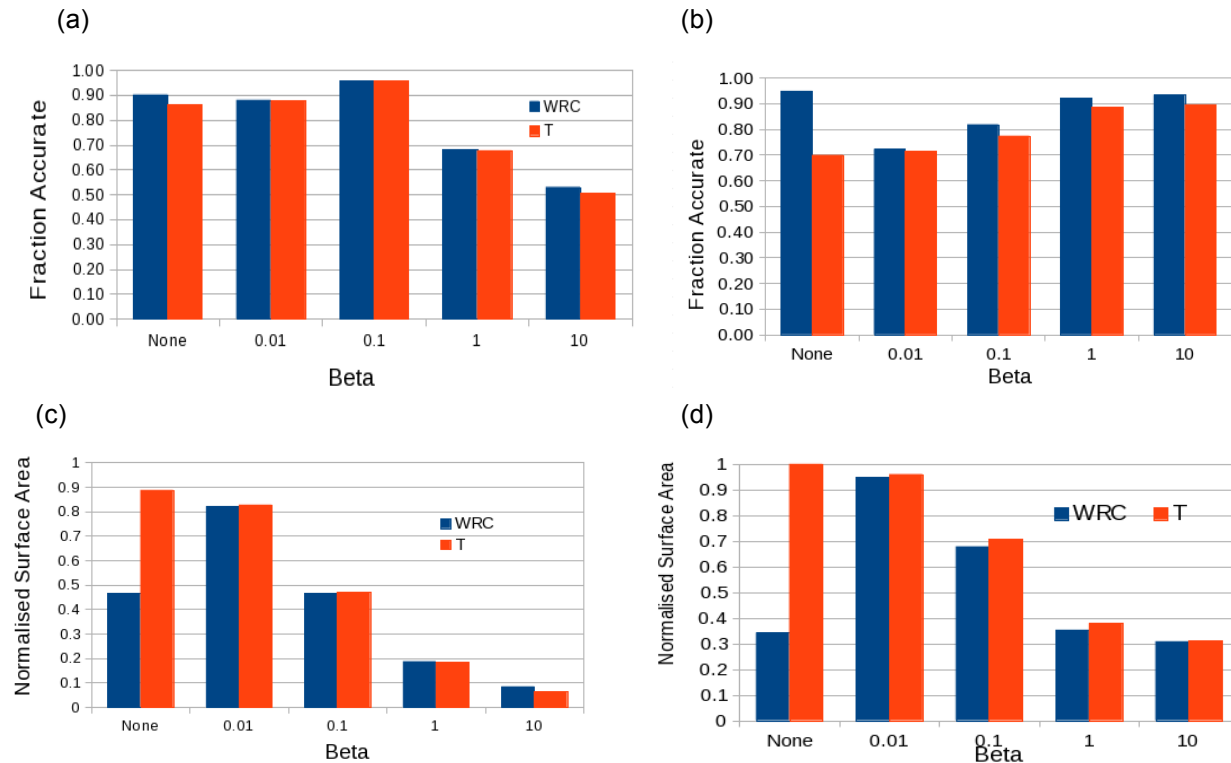


Figure 8: a) segmentation accuracy vs β for sample 1 showing a dramatic drop in accuracy after the optimal value $\beta=0.1$; b) For sample 2, $\beta=0.1$ is far from optimal, with the segmentation accuracy significantly higher for much larger beta; c) and d) normalised surface area vs β for sample 1 and sample 2 respectively, highlighting the sensitivity of the segmentation to this parameter.




Figure 9: 400x350 pixel subset from 3D image of sample 4, an unconsolidated sand containing oil and NaI-doped brine. a) Grey-scale tomogram; b) manual segmentation; c) WRC, the most accurate, particularly on phase boundaries and corners; d) T showing the intermediate phase interrupting the phase boundaries of the other two phases. ICM was tried on this data but even using the optimal β it did not improve the outcome.

6. CONCLUSIONS AND FURTHER WORK

Research into many aspects of the micro-CT imaging pipeline continues at the ANU micro-CT facility. From recent work we can conclude

- The application of high cone-angle helical scanning to micro-CT has thrown up many challenges, which have either been solved or reduced.
- Dynamic tomography has the potential for order-of-magnitude improvements in data acquisition efficiency when acquiring 4D tomographic "movies" of certain dynamic processes.
- Image registration is a powerful tool that allows the combination of micro-CT with other imaging modalities and with *ex-situ* experimental programs that do not suffer the limitations or expense of *in-situ* experiments.
- Automated segmentation remains an elusive goal, with recently proposed methods showing promise but with no single method able to perform well for the wide variety of materials imaged at the ANU facility. Partial volume effects, that become significant in samples with high surface area or sub-resolution features, are particularly difficult to address. For higher quality datasets with fewer partial-volume voxels and good contrast-to-noise all the tested methods can be recommended.

Future work will continue in these areas while also considering issues in beam hardening, region-of-interest imaging and phase contrast effects. In the near future the group expects to construct a micro-CT instrument with sub-micron resolution. It is expected that the severe limits on beam flux generally encountered in this regime will be ameliorated by the use of high cone-angle helical scanning. It is also expected that phase contrast effects will be prominent at high resolutions, providing both a challenge and an opportunity.

7. ACKNOWLEDGEMENTS

We acknowledge the support of the Australian Research Council through Future Fellowship FT100100470 and Discovery Projects DP110102964 and DP110102888, the ANU/UNSW Digital Core Consortium, and of a substantial grant of supercomputer time provided by Australia's National Computing Infrastructure (NCI) National Facility.

8. REFERENCES

- Arns, C., F. Bauguet, A. Limaye, A. Sakellariou, T. Senden, A. Sheppard, R. Sok, V. Pinczewski, S. Bakke, L. Berge, P. Oren, M. Knackstedt, (2005) "Pore-Scale Characterization of Carbonates Using X-Ray Microtomography." *SPE Journal* **10** (4): p. 475-484 . doi:10.2118/90368-PA.
- Arns, Christoph H., Mark A. Knackstedt, W. Val Pinczewski, and Nicos S. Martys, (2004) "Virtual Permeametry on Microtomographic Images." *Journal of Petroleum Science and Engineering* ,**45**(1): p. 41–46.
- Aste, T., M. Saadatfar, and T. J. Senden. (2005). "Geometrical Structure of Disordered Sphere Packings." *Physical Review E* , **71** (6): 61302.
- Berthod, M., Z. Kato, S. Yu, and J. Zerubia (1996), "Bayesian image classification using Markov random fields. *Image and Vision Computing*", **14**(4): p. 285-295.
- Besag, J. (1986). "On the statistical-analysis of dirty pictures". *Journal of the Royal Statistical Society Series B-Methodological*, **48**(3): p. 259-302.
- Boone, M. N., J. Vlassenbroeck, S. Peetermans, D. Van Loo, M. Dierick, and L. Van Hoorebeke. (2012). "Secondary Radiation in Transmission-type X-ray Tubes: Simulation, Practical Issues and Solution in the Context of X-ray Microtomography." *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **661** (1): 7–12.
- Derluyn, H., M. Griffa, D. Mannes, I. Jerjen, J. Dewanckele, P. Vontobel, A. Sheppard, D. Derome, V. Cnudde, E. Lehmann and J. Carmeliet (2013), "Characterizing Saline Uptake and Salt Distributions in Porous Limestone with Neutron Radiography and X-ray Micro-tomography." *Journal of Building Physics* **36** (4): p. 353–374. doi:10.1177/1744259112473947.
- Geman, S. and D. Geman (1984). "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images". *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**(6): p. 721-741.
- Jain, Anil K. (2010). "Data Clustering: 50 Years Beyond K-means". *Pattern Recognition Letters* **31**(8): p. 651–666. doi:10.1016/j.patrec.2009.09.011.
- Jones, A. C, C. H Arns, A. P Sheppard, D. W Hutmacher, B. K Milthorpe, and M. A Knackstedt (2007), "Assessment of Bone Ingrowth into Porous Biomaterials Using Micro-CT." *Biomaterials*, **28** (15): p. 2491–2504.
- Jones, P. N., and G. J. McLachlan.(1990). "Algorithm AS 254: Maximum Likelihood Estimation from Grouped and Truncated Data with Finite Normal Mixture Models". *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **39**(2): p. 273–282.
- Katsevich, A. (2002). "Theoretically Exact Filtered Backprojection-type Inversion Algorithm for Spiral CT." *SIAM Journal of Applied Math*: 2012–2026.
- Kingston, A., A. Sakellariou, T. Varslot, G. Myers, and A. Sheppard (2011). "Reliable Automatic Alignment of Tomographic Projection Data by Passive Auto-focus." *Medical Physics* **38** (2011): 4934.
- Kulkarni, R., M. Tuller, W. Fink, and D. Wildenschild. (2012) "Three-Dimensional Multiphase Segmentation of X-Ray CT Data of Porous Materials Using a Bayesian Markov Random Field Framework". *Vadose Zone Journal*, **11**(1) vzj2011.0082.
- Kumar, M., T Senden, A Sheppard, J Middleton, and M Knackstedt (2010), "Visualizing and Quantifying the Residual Phase Distribution in Core Material." *Petrophysics* **51** (5): p. 1.
- Latham, S., T. Varslot, and A. Sheppard (2008). "Image Registration: Enhancing and Calibrating X-ray Micro-CT Imaging." *Proc. of the Soc. Core Analysts, Abu Dhabi, UAE*.
- Long, J. A, G. C Young, T. Holland, T. J Senden, and E. M.G Fitzgerald. (2006). "An Exceptional Devonian Fish from Australia Sheds Light on Tetrapod Origins." *Nature*, **444** (7116): p. 199–202.
- Middleton, J., S. Latham and A. Sheppard (2013). "Segmentation Method Performance Comparison: an Automated Watershed Method and ICM Post Processing", submitted to the 1st International Conference on the Tomography of Materials and Structures (ICTMS), Ghent, Belgium.

Myers, G.R., A.M. Kingston, T.K. Varslot, M.L. Turner, and A.P. Sheppard (2011a). "Dynamic Tomography with a Priori Information." *Applied Optics* 50 (20): 3685–3690.

Myers, G.R., A. M. Kingston, T. K. Varslot, M. L. Turner, and A. P. Sheppard. (2011b), "Dynamic X-ray Micro-tomography for Real Time Imaging of Drainage and Imbibition Processes at the Pore Scale." *International Symposium of the Society of Core Analysts*, Austin, TX, USA, paper #SCA2011-42,.

Myers, Glenn R., Andrew M. Kingston, Trond K. Varslot, and Adrian P. Sheppard (2011c). "Extending Reference Scan Drift Correction to High-magnification High-cone-angle Tomography." *Optics Letters* 36(24): 4809–4811. doi:10.1364/OL.36.004809.

Salmon, P. L., X. Liu, and A. Sasov (2009) "A Post-scan Method for Correcting Artefacts of Slow Geometry Changes During Micro-tomographic Scans." *Journal of X-Ray Science and Technology* 17(2): 161–174.

Singh, K., R. K Niven, T. J Senden, M. L Turner, A. P Sheppard, J. P Middleton, and M. A Knackstedt. 2011. "Remobilization of Residual Non-Aqueous Phase Liquid in Porous Media by Freeze- Thaw Cycles." *Environmental Science & Technology* 45 (8): 3473–3478.

Varslot, T., A. Kingston, G. Myers, and A. Sheppard. (2011). "High-resolution Helical Cone-beam micro-CT with Theoretically-exact Reconstruction from Experimental Data." *Medical Physics* 38: 5459.

Varslot, T., A. Kingston, G. Myers, and A. Sheppard (2012). "Considerations for High-magnification High-cone-angle Helical micro-CT." In *SPIE Optical Engineering+ Applications*, 850614–850614.