IMAGE REGISTRATION: ENHANCING AND CALIBRATING X-RAY MICRO-CT IMAGING

Shane Latham, Trond Varslot and Adrian Sheppard
Department of Applied Mathematics
Australian National University
Canberra Australia

This paper was prepared for presentation at the International Symposium of the Society of Core Analysts held in Abu Dhabi, UAE 29 October-2 November, 2008

ABSTRACT

A range of 3D image registration techniques have been developed to extend the capabilities of X-ray micro-CT and to allow the calibration of micro-CT data to other microscopic techniques. In contrast to the alignment of 2D images, for which many tools are publicly available, registration of very large 3D images, for which distributed-memory parallel algorithms are required, pose a considerable challenge. We are not aware of available software which is capable of accomplishing this task. To illustrate the diverse capabilities and potential value of our registration techniques, we provide examples of successfully aligned core images.

As a first example, we show 2D back-scattered scanning electron microscopy (BSEM) images of polished thin sections extracted from plugs previously imaged by X-ray micro-CT. BSEM imaging generates nanometre resolution images, and can even provide detailed mineralogical maps through energy dispersive X-ray spectroscopy (EDS). Accurate registration of the 2D microscopic image within the 3D image volume increases the potential for mapping X-ray grey-levels to mineral phases, reveals information about features too fine for X-ray micro-CT and provides an invaluable quality control on the 3D image data.

We also present examples where pairs of 3D micro-CT images are registered. Cores are imaged dry, then removed from the micro-CT apparatus, flooded under different wettability conditions and saturation states, and re-imaged. This process can be repeated with any length of time between image acquisitions. Accurate alignment then allows for the information from all images to be combined, thus adding enormously to the value of both the imaging and the flooding experiments.

INTRODUCTION

Micro-CT imaging and the computational methods for analysing the generated X-ray attenuation data are becoming increasingly popular for characterising many macroscopic properties of porous media. From micro-CT images of dry cores, it is possible to compute properties such as porosity, permeability, conductivity, elasticity, and mercury injection capillary pressure that are in good agreement with laboratory experiment (C. H. Arns et al.

2002, M. A. Knackstedt 2004). In addition to dry-core micro-CT analysis, the imaging of cores that have been flooded with an X-ray blocking liquid allows the study of fluid partitioning in the pore space (Prodanovic et al. 2006, Prodanovic et al. 2007). Flooded core analysis has the potential to enable the determination of key multiphase flow parameters and provide validation data for fluid transport models.

While the study of X-ray Micro-CT tomograms can provide a variety of information regarding pore and rock-matrix structure of a core sample, a complete analysis is often impeded because: (a) X-ray attenuation alone is insufficient to distinguish mineralogy, (b) there exists significant regions in the sample that contain pore/matrix features below the image resolution, (c) there is limited X-ray contrast when imaging fluid distributions and (d) when imaging wet cores, flooding needs to occur quickly and must take place at atmospheric temperature and pressure, since the sample must normally remain resident in the micro-CT apparatus. To overcome (a) and (b), complementary information from other imaging and analysis techniques can be combined with micro-CT to data to reveal the core structure at micro, meso and macro scales (Padhy et al. 2007) as well as mineralogy discrimination. Furthermore, aligning 2D BSEM or thin-section optical images with the corresponding region of the micro-CT 3D images allows the possibility of classifying micro-CT voxels as a particular mineral phase, based on their tomographic grey level. Sub-voxel accurate alignment of the high-resolution 2D images with the micro-CT tomogram also provides a means of making a rapid qualitative assessment of the micro-CT image. Impediment (c) can be overcome by aligning micro-CT images of dry and partially saturated cores, then masking the flooded image with the solid phase obtained from the dry image and classifying voxels belonging to the liquid phase in this masked image to identify the fluid distribution. Similarly, (d) can be avoided through the use of automated registration tools, as it becomes possible to remove samples from the micro-CT apparatus in order to perform flooding. The flooded core can then be imaged at arbitrary times with the resulting flooded and dry tomographic images becoming easily compared when accurately aligned.

In general terms, image registration is the task of bringing into geometric alignment two or more images, of the same scene, taken at different times, from different orientations and/or by different instruments. While solutions to this problem are well represented in the literature for specific problem domains, there is yet to appear a single automated solution which can be reliably applied in all problem domains. There are a variety of different approaches for performing image registration and a recent survey of techniques can be found in (Zitová and Flusser 2003). In this paper, we present a registration algorithm which has been successfully applied to the task of aligning Scanning Electron Microscopy (BSEM) 2D images with micro-CT 3D images, and aligning dry/flooded 3D image pairs. The algorithm can be categorised as a “multi-resolution area-based multi-start optimization” approach and can deal with arbitrary differences in rotational and translational orientation. The C++ implementation is capable of registering large datasets on distributed memory high-performance architectures through the use of the Message

Passing Interface (MPI). To date, micro-CT tomographic images as large as 2×2048^3 bytes have been registered using this implementation.

REGISTRATION ALGORITHM

A common approach (and the approach adopted in this paper) for solving the registration task is to first formulate it as an *optimization problem*. Fundamental to this formulation is a *cost-function* or *distance-metric* which gives a quantitative value indicating the merit of the alignment between two images with respect to a number of transform parameters. These transform parameters define how the two images are overlaid. An optimization algorithm can then be used to discover the transform parameters which give the minimum value of the distance-metric. It is these minimizing-parameters that immediately yield the desired alignment of the two images.

We formulate the optimisation problem using the *correlation coefficient* as the distance-metric and the *similarity* transformation as the deformation, which results in 7 degrees of freedom (3 translational, 3 rotational and 1 isotropic scaling) to be solved by an optimization search strategy. There are many mathematical optimization algorithms for reliably and efficiently discovering local minima, less common are practical optimization methods for reliably finding global minima. Global methods based on annealing, evolution or Monte-Carlo sampling are capable of discovering global minima but do so at the price of a large number of distance-metric evaluations. The appeal of local mathematical optimization methods is fast and accurate convergence (provided the distance-metric satisfies certain conditions), and hence a fewer number of distance-metric evaluations, but the price is that there is no guarantee of discovering a global minimum. To increase the likelihood of local optimization methods finding a global minimum and to further reduce the amount of computation, many registration methods use a multi-resolution approach. The idea is to find transformation parameters which are close to the global minimum using low resolution versions of the original pair of images and then to further perfect the low-resolution transformation parameters using increasingly higher resolutions. In this paper, the multi-resolution search strategy is to use an exhaustive low-resolution search to determine an initial small set of transformation parameters, where at least one element of this initial set of parameters is within the capture radius of a global minimum. These initial sets of transformation parameters are used as starting-solutions for a local numerical optimization method (Powell's method as implemented in Press et al. 2002) to further refine and reduce the transformation parameter sets at increasingly higher resolutions until the final best transformation parameters are determined at the original highest resolution. A more detailed description of the implementation and parallelisation strategy can be found in the appendix.

REGISTRATION RESULTS

In this section, we present results generated by applying the above registration algorithm to align BSEM images with micro-CT images, and to align a micro-CT image of a flooded core sample with the micro-CT image of the corresponding dry core sample.

Alignment of BSEM 2D Images with Micro-CT 3D Images

To obtain high-resolution BSEM images from within the volume imaged by micro-CT, we use the following experimental protocol. Firstly, the sample, usually a cylinder with a diameter of 5mm, is imaged in the micro-CT. The sample is then cut, perpendicular to the cylinder's axis, and one half is impregnated with epoxy. This half is then polished down to form a thin section with a thickness of around 30 microns which is mounted on a glass slide. Clearly this is a destructive process which precludes re-imaging of the sample. Thin-sections are normally imaged using two methods. The first method, BSEM, generates images from the top 1-2 micrometres of the thin-section and all 2D results presented in this work arise from the BSEM method. The second back-illuminated optical microscopy method, as per traditional petrographic analysis, can be used to generate images consisting of the integrated (average) properties over the full 30 micron thickness of the thin-section. Registration of these optical images is beyond the scope of this work, although it is likely our registration method would be capable of aligning these images by using a different distance-metric.

Being a destructive imaging technique, the quality of the 2D BSEM image can be adversely affected by a number of factors including the trapping of air bubbles in the epoxy and the removal of friable core material during polishing. We have observed both these effects in the registered images studied to date. The amount of core material removal appears to be insignificant, however, air bubbles are a significant problem that will affect quantitative characterisation of the pore space.

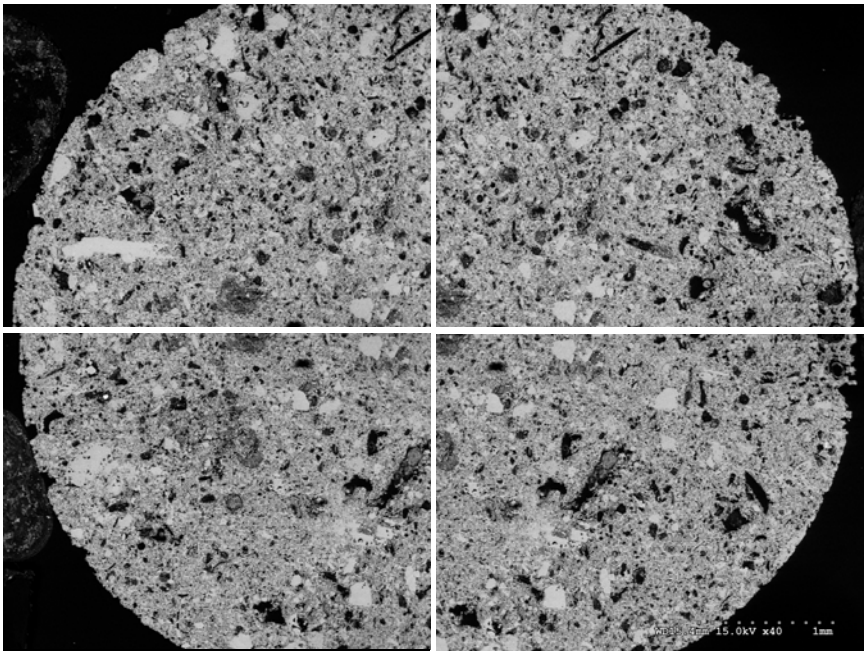

Figure 1: Four "quadrant" BSEM images of carbonate thin section taken at 40× magnification.

Techniques for resolving this problem are beyond the scope of this work. However, it is worth noting that our approach allows for quality control of both the 3D tomographic image and the 2D SEM image, since one can invariably infer the source of discrepancies between the two registered images. It is also possible that fine particles may be dispersed through the pore space during the epoxy filling, however, to date this has not been observed. BSEM images of core sample thin sections are obtained at a number of resolutions, (for the samples shown here, magnifications factors were ≈40×, ≈450× or ≈1000×). For the lowest magnification, a series of BSEM scans is performed so that the entire area of the core-sample thin section is imaged. Figure 1 shows four "quadrant" BSEM images taken at ≈40× magnification of a carbonate thin section with pixel size of 1.24 microns. These lowest resolution BSEM 2D images were individually pre-processed to remove distortion, which is introduced during the BSEM image acquisition.

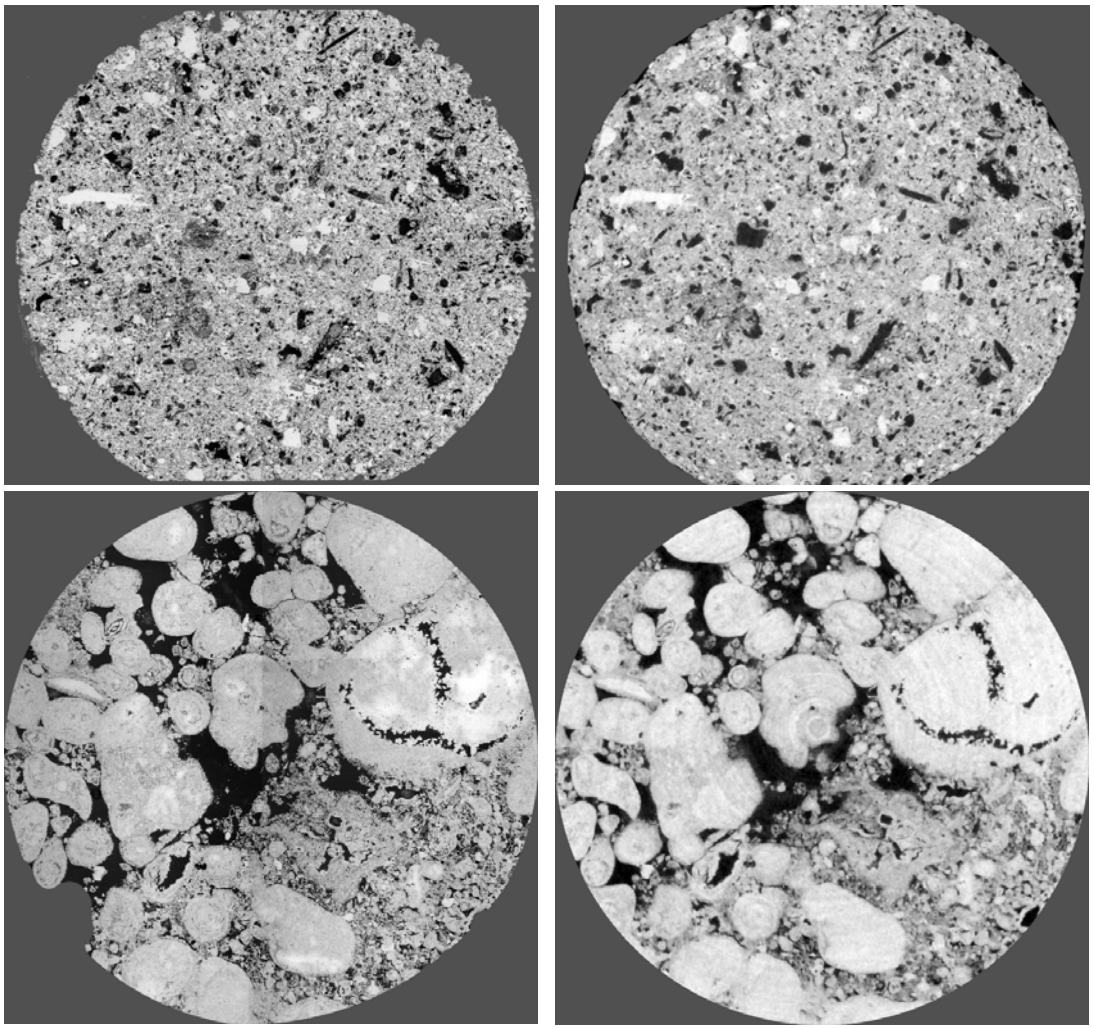

Figure 2: BSEM images of carbonate samples (left column) and registered region of corresponding micro-CT 3D image (right column).

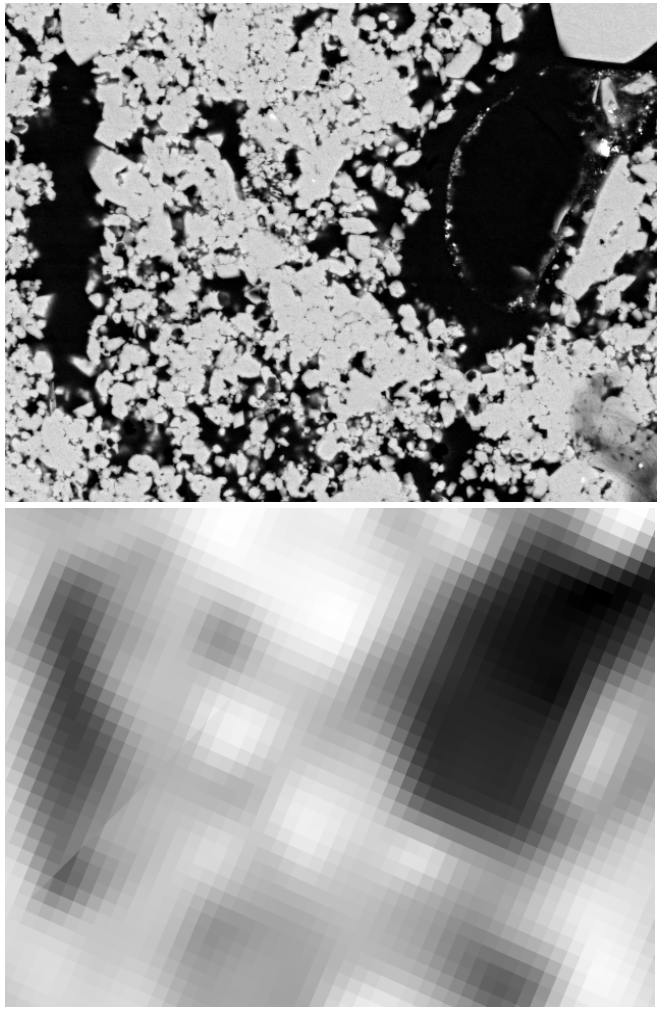

Figure 3: High-resolution 1000× magnification BSEM image (top) and registered region of micro-CT image (bottom) from the sample shown in the top row of Figure 2.

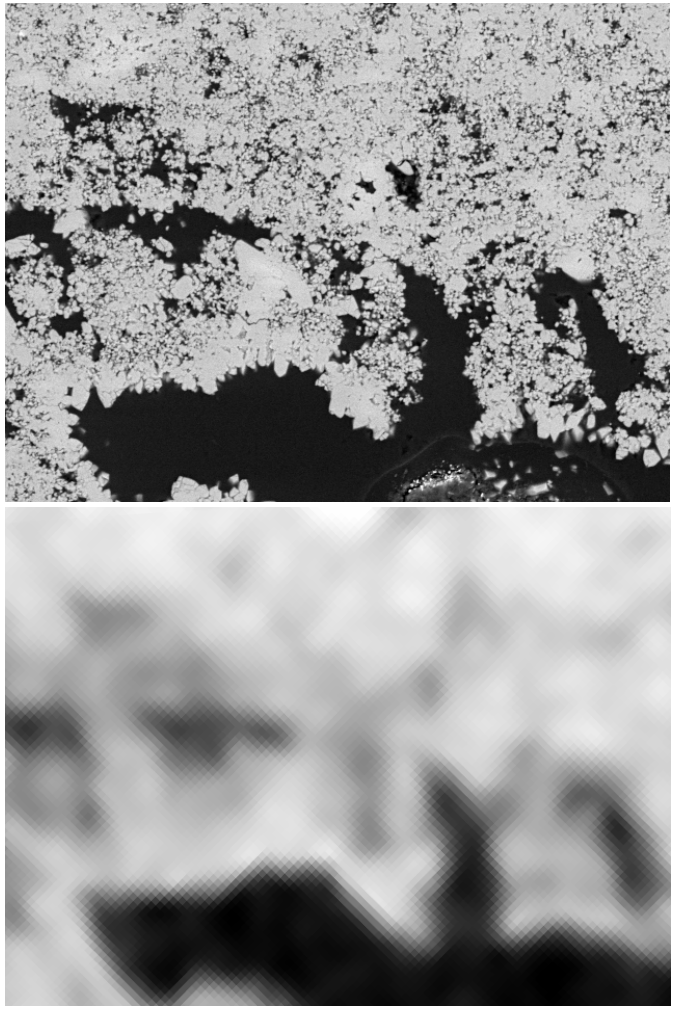

Figure 4: High-resolution 450× magnification BSEM image (top) and registered region of micro-CT image (bottom) from the sample shown in the bottom row of Figure 2.

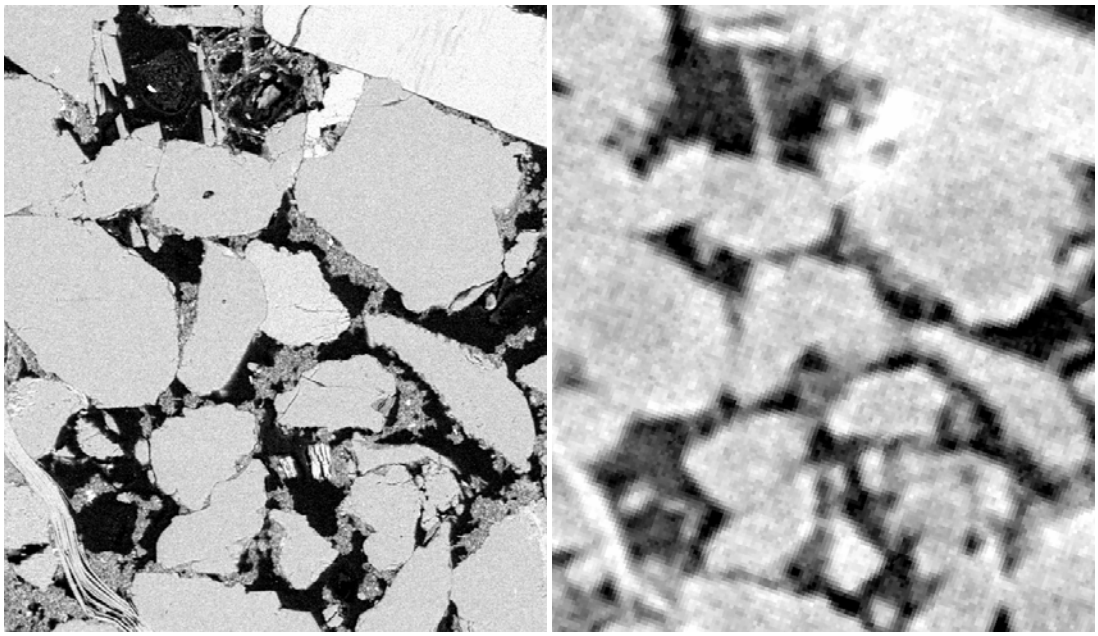

Figure 5: High-resolution 500x magnification BSEM image region (left) and registered region of micro-CT 3D image (right).

The distortion removal uses a BSEM image of a *known grid* to de-warp the sample images. The separate BSEM thin section images were then registered to form a single montage image of the thin section. This 2D registration was performed using a *phase-correlation frequency-domain* technique, which is highly efficient at determining translational shifts between images. The result of this registration is a mosaic image typically containing around 4000x4000 pixels. The montage image was then registered with the micro-CT 3D image by employing our multi-start multi-resolution algorithm, with the correlation coefficient used as the distance-metric. The higher-resolution BSEM images were registered with the montage image and then were overlaid with the corresponding region of the registered micro-CT image. Figure 2 shows BSEM images (left column) of two carbonate samples with the registered region of their corresponding micro-CT image (right column). The left image of the first row is the montage image formed by performing 2D registration on the "quadrants" of the de-warped images of Figure 1. The right image of the first row is the corresponding *slice* of the micro-CT 3D image (2.59 micron voxel size). The alignment between the two images is good, with a correlation coefficient of 0.641. A *perfect* correlation value of 1 indicates a linear relationship between the intensity values in the two images. In general, two random variables are said to have *large* correlation when the coefficient is between 0.5 and 1.0, *medium* correlation when the coefficient is between 0.3 and 0.5 and *small* correlation when the coefficient is between 0.0 and 0.3. Note that the BSEM images are visually similar to X-ray projection images, since both methods are primarily sensitive to the density of the electron cloud surrounding each atom.

There are noticeable differences between the two images in the top row of Figure 2 where some of the larger empty pore spaces in the micro-CT image appear to have non-empty counterparts in the BSEM image. Presumably, there has been some disruption to the sample during the thin-section preparation. The bottom row of Figure 2 shows good alignment between the montage BSEM image on the left (40× magnification, 1.24 micron pixel size) and the micro-CT image slice on the right (2.69 micron voxel size) with a correlation coefficient of 0.716. Figure 3 shows the 1000× magnification BSEM image (top), from the sample in the top row of Figure 2 and the corresponding region of the micro-CT 3D image (bottom). The high-resolution BSEM image clearly illustrates features (porosity) which are not easily classifiable in the micro-CT image. The upper image in Figure 4 is a 450× magnification BSEM image (0.110 micron pixel size) of the sample from the bottom row of Figure 2 and the lower image is the corresponding region of the micro-CT image. Figure 5 compares a 500× magnification BSEM image region (0.099 micron pixel size) with the corresponding micro-CT image slice for clay-rich reservoir sandstone. Again, there is structure present in the high-resolution BSEM image which is simply not present in the micro-CT image.

Alignment of Flooded with Dry Micro-CT 3D Images

The top row of Figure 6 shows slices from the registered micro-CT 3D images of a dry (left) and a flooded (right) 5 mm diameter sucrosic dolomite carbonate core. The dry sample tomogram was obtained by drilling the sample core in water followed by oven drying to constant weight and then imaging in the micro-CT apparatus. This dry core was then removed from the micro-CT apparatus to conduct the flooding experiment. The flooded sample was prepared by initially plasma cleaning (Kumar et al. 2008) the dry core and then bringing it into hydraulic contact with a capillary bed made from plasma treated glass wool, which was completely saturated with 1 M CsI solution. Spontaneous imbibition was carried out in a high humidity environment over a 24 hour period to produce a strongly water wet core. The flooded sample was then replaced into the micro-CT apparatus for 3D imaging. Both the wet and dry tomograms were acquired/reconstructed at image size of 1024^3 voxels with voxel side-length ≈6.72 microns.

In the flooded image, the contrast liquid shows as the brightest pixels. The correlation coefficient distance metric is a good indicator of correct alignment when there is a linear relationship between the attenuations of the fixed and moving images. For the dry and flooded samples there is a nonlinear relationship between the attenuations. The pore-phase attenuation in the dry image may be mapped to either the dry-pore-phase attenuation or to the liquid-phase attenuations in the flooded image. In addition, because a high-contrast liquid was used in this particular experiment, there is some *blurring* of the liquid-filled-pores in the flooded image, so that near the edges of the pores in the dry sample image the solid-phase attenuations may correspond to liquid-phase attenuations in the flooded image. Ideally, a lower contrast liquid would be preferred for fluid distribution analysis to reduce the blurring. However, this example highlights the ability of the registration method to deal with fundamentally different images by simple adjustment to the distance-metric.

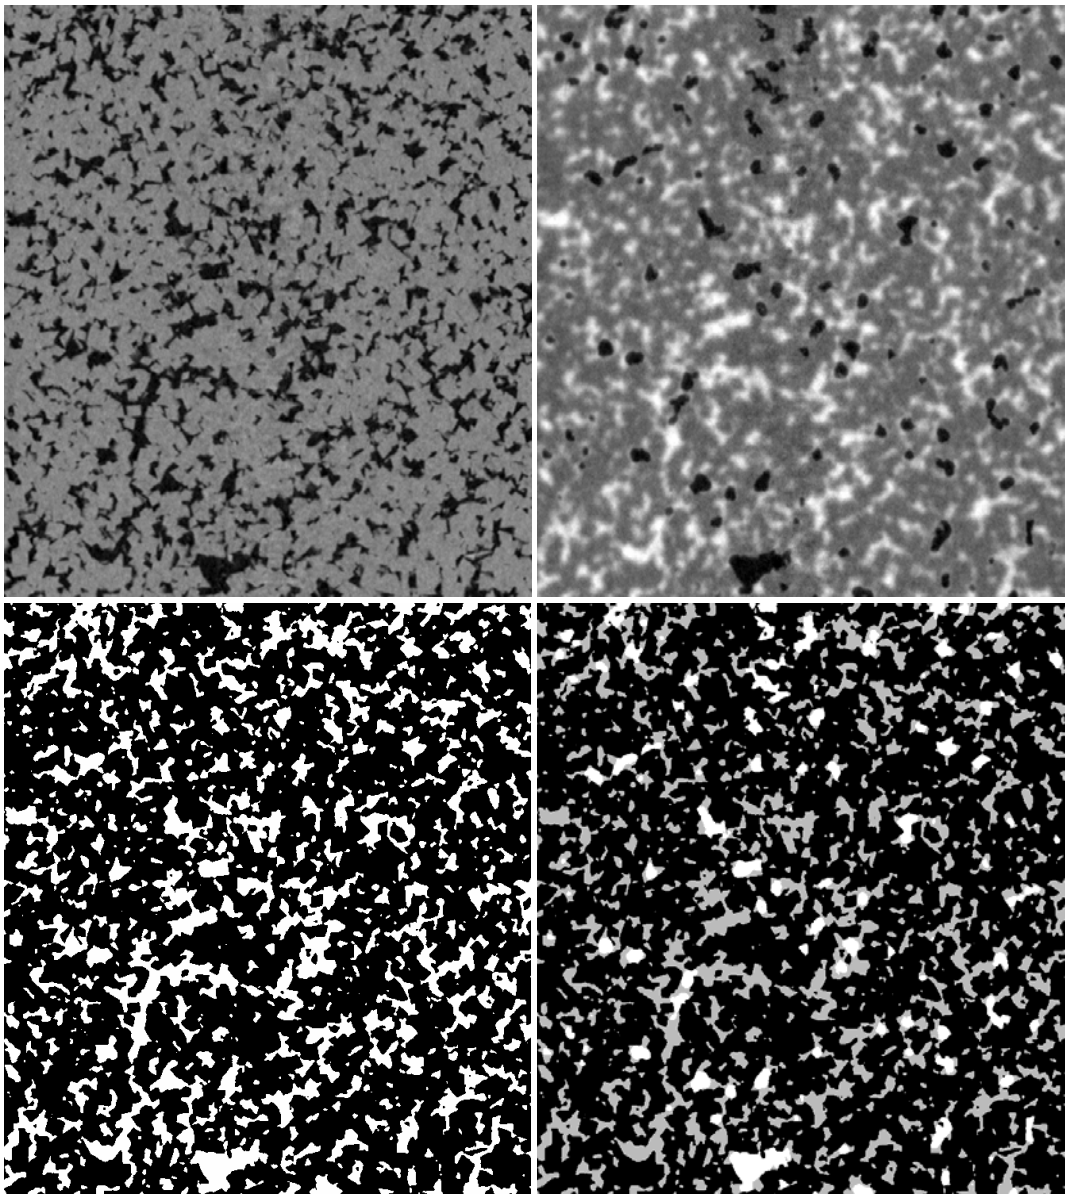

Figure 6: Micro-CT images: dry sample (top left), registered flooded sample (top right), segmented dry sample (bottom left) with pore space shown in white and segmented flooded sample (bottom right) with contrast-liquid shown in grey and dry pore space in white.

Because of the non-linear attenuation relationship, the correlation coefficient was not a good discriminator for correct alignment when calculated using the entire range of image attenuations. To overcome this difficulty, only voxels with dry-pore attenuation in the flooded image were considered when calculating the correlation coefficient. Spurious registrations were avoided, where only a few dry-pore regions may overlap but still yield a large correlation coefficient, a penalty term based on the percentage of non-overlapping

dry-pore voxels is subtracted from the correlation coefficient. The bottom row of Figure 6 shows a slice from the segmented dry 3D image (left) and the corresponding slice from the segmented flooded 3D image (right). In both segmented images the rock matrix is shown in black and the dry-pore space shown in white. The segmented dry image was used as a mask on the flooded image in order to delineate the regions containing contrast-liquid, which are shown as grey pixels in the bottom right image of Figure 6.

CONCLUSION

We have described a registration method for accurately aligning high-resolution BSEM 2D images of core thin sections with the corresponding region of the micro-CT 3D image of the core. Typically, the BSEM alignment takes less than 1 hour to run on 64 CPUs of an SGI Altix-3700-Bx2 system and consumes less than 32GB of RAM. Figures 3, 4 and 5 clearly illustrate structure present in the BSEM images which is below the resolution of the corresponding micro-CT image. Visual comparison of the aligned images gives a measure of the quality of the micro-CT image and a preliminary indication of whether network models generated from the micro-CT data will be able to accurately simulate fluid transport within the sample. Currently under development are quantitative methods for combining information from the BSEM image with the 3D micro-CT image-analysis.

Micro-CT imaging of contrast-liquid flooded core samples currently appears to be the only method for obtaining accurate fluid partitioning information. The registration method, with a modified distance-metric, can be used to align flooded-core micro-CT 3D images with the corresponding dry-core micro-CT images. The run-time to achieve alignment is typically less than 2 hours on 64 CPUs of an Altix-3700-Bx2 system and consumes less than 64GB of RAM. With the availability of suitable computational resources, the registration method is capable of processing large numbers of sample images. The flooded/dry alignment method allows for the possibility of conducting *out-of-apparatus* flooding experiments over indefinitely long time periods. During the experiment, the core may be placed back into the micro-CT apparatus (without the need to maintain constant orientation) for re-imaging in an arbitrary time sequence. All the resulting micro-CT 3D images can subsequently be aligned in a common frame of reference for analysis. Fluid distribution information, obtained in such a manner, has the potential to greatly enhance the discrimination capabilities of micro-CT and, in particular, to provide key validation and calibration data for multiphase flow simulations.

The author's have implemented automated calculations for most of the registration method parameters. A crucial parameter for achieving successful registration and convenient runtime is the maximum down-sample factor. Specifying too large a factor will cause the registration to fail to find the best match and too small a factor will cause the runtime for Step 2 (see Appendix) to become inconveniently large. There are automated methods for choosing a suitable maximum down-sample factor, however, the authors are yet to implement them. For the flooded/dry registration, the attenuation ranges of the dry-pore regions in each image must be specified, but these ranges can be be easily found during visual inspection of a conveniently sized portion of the micro-CT data.

ACKNOWLEDGEMENTS

The authors acknowledge the Australian Research Council and member companies of the Digital Core Consortium for providing funding support and the Australian Partnership for Advanced Computation for supplying HPC resources. The authors also thank reviewers Norman Morrow and Craig Lindsay whose comments improved the content of the paper.

REFERENCES

Arns, C. H., Knackstedt, M. A., Pinczewski W. V. and Garboczi E. J. "Computation of Linear Elastic Properties from Microtomographic Images: Methodology and Agreement Between Theory and Experiment". *Geophysics,* (2002) **67**, no. 5, 1396-1405.

Jenkinson, M. and Smith S. "A Global Optimisation Method for Robust Affine Registration of Brain Images". *Medical Image Analysis,* (2001) **5**, no. 2, 143-156.

Knackstedt, M. A., Arns C. H., Limaye A., Sakellariou, A., Senden T. J., Sheppard A. P., Sok R. M., Pinczwiki V.W. and Bunn G.F. "Digital Core Laboratory: Properties of Reservoir Core Derived from 3D Images". In *Proceedings SPE Asia Pacific Conference on Integrated Modelling for Asset Management*, Kuala Lumpur, Malaysia, Soc. Petrol. Eng., (2004), 1-14.

Kumar, M., Senden T. J., Latham S. J., et al. "Designing for Mixed Wettability". In *2008 SPE Improved Oil Recovery Symposium*, Tulsa, Oklahoma, Soc. Petrol. Eng., (2008), 1-9.

Padhy, G. S., Lemaire C., Amirtharaj E. S. and Ioannidis M. A. "Pore Size Distribution in Multiscale Porous Media as Revealed by DDIF–NMR, Mercury Porosimetry and Statistical Image Analysis". *Colloids and Surfaces A: Physicochemical and Engineering Aspects,* (2007) **300**, no. 1-2, 222-234.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. *Numerical Recipes in C (2nd ed.): The Art of Scientific Computing,* Cambridge University Press, New York, (2002), 412-420.

Prodanovic, M., Lindquist, W. B. and Seright R. S. "Porous Structure and Fluid Partitioning in Polyethylene Cores from 3D X-Ray Microtomographic Imaging". *Journal of Colloid and Interface Science,* (2006) **298**, no. 1, 282-297.

Prodanovic, M., Lindquist, W. B. and Seright R. S. "3D Image-Based Characterization of Fluid Displacement in a Berea Core". *Advances in Water Resources,* (2007) **30**, no. 2, 214-226.

Zitová, B., and Flusser J. "Image Registration Methods: A Survey". *Image and Vision Computing,* (2003) **21**, no. 11, 977-1000.

APPENDIX: REGISTRATION ALGORITHM DETAILS

The next two subsections give a more detailed description of the steps of the registration algorithm (a generalised version of the Jenkinson and Smith 2001 method for registering 3D brain-images) along with a discussion of the parallelisation strategy.

Algorithm Steps

1. Pre-process images. The reconstructed tomographic 3D images are pre-processed to identify any voxels which are not contained with the core-sample region. These voxels are assigned a *mask-attenuation value.* Voxels with mask-attenuation value do not contribute to the evaluation of the distance-metric. Lower resolution BSEM images generally contain distortion which is removed during pre-processing. All BSEM images are then combined via 2D registration into a mosaic. Pixels in the BSEM image which do not lie within the sample region are also assigned a mask-attenuation value.

2. Exhaustive search at lowest resolution. The fixed-image and the moving-image are down-sampled to a lowest resolution by a factor d (the maximum down-sample factor). The lowest resolution is determined by the size of the features present in the images. Typically, d=16 in each dimension so that a 2048×2048×2048 voxel image is down-sampled to generate a 128×128×128 voxel image. A search of the parameter space is conducted by evaluating the distance-metric at points which form a regular grid in the transformation parameter space. The translational step length is that of a down-sampled voxel side, and the rotational step-size is 2 degrees. In this step, the rotational DOF are restricted to the Z-axis as BSEM images are generated from thin-sections taken (approximately) from the *XY*-plane, and the micro-CT images are generated from samples which have their cylindrical axis aligned with the Z-axis. The scaling parameter is held constant at 1, as the tomogram voxel side lengths and BSEM pixel side lengths are generally known to within ±3% and this error can be ignored at the lowest resolution. Keep the N_d sets of transformation parameters, which give the N_d lowest values of the distance-metric, to pass to the next step as starting solutions for the iterative optimization.
3. Check if finished. If d<1 go to Step 7.
4. Perform local mathematical optimizations. Create new fixed and moving images by down-sampling by original images by a factor of d. For each of the N_{2d} best transformation parameters from the previous step, perform a local optimization with the solution from the previous step as the starting point.
5. Update down-sample factor. d=d/2.
6. Go to Step 3.
7. Finished. The best transformation parameters from the previous step are the final solution giving the optimal alignment of the original fixed and moving images.

Parallel Implementation

In order for the above global-minimization method to be computationally feasible, there are two parallelization strategies which are used when implementing the above steps in software on high-performance Non Uniform Memory Access (NUMA) architectures. The first is a *task-parallelism* strategy where individual tasks are performed independently on *computational-units*. The second strategy is a *data-parallelism* approach where each computational-unit contains only a subset of the discrete-image pair data. Task-parallelism is of great advantage when calculating the distance-metric values in Step 2. At low resolutions each computational-unit can evaluate the distance-metric for a subset of the transformation parameter grid points independently of other computational-units. Similarly, task-parallelism can be used to advantage in Step 4 for the lower resolution images, where each computational-unit can perform a subset of the iterative local optimizations, independently of other computational-units. For the higher-resolution images, the data-parallel strategy is preferred. In this strategy, each computational unit only contains a subset of the discrete-image data. Step 4 is then performed using a *Master-Worker* division of labour. The *master* computational-unit controls the iterative optimization algorithm while *worker* computational-unit calculates a portion of the total distance-metric result based on the subset of image data held in its local-RAM.